

# The Sustainable Adoption of Industry 4.0 & Machine Learning in the Automotive Industry

by

James Flynn

Thesis submitted to

Swansea University

In fulfillment of the requirements

For the Doctorate of Engineering

Eng.D

Department

Mechanical Engineering

2023

Copyright: The Author, James Flynn, 2023.

## Supervisors

The following acted as the academic and industrial supervisors for this thesis.

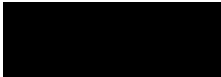
Academic Supervisor (1): Cinzia Giannetti,  
Professor, Dept. of Engineering,  
Swansea University

Academic Supervisor (2): Christian Griffiths,  
Associate Professor, Dept. of Engineering,  
Swansea University

Industrial Supervisor: Steven Buck,  
PTO Business Manager,  
Dunton Technical Centre,  
Ford Motor Company

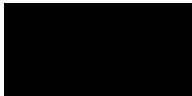
## Author's Declaration

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed.....  .....

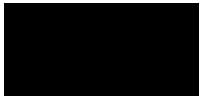
Date.....  
12.06.23

This thesis is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed.....  .....

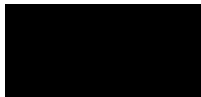
Date.....  
12.06.23

I hereby give my consent for my work, if relevant and accepted, to be available for photocopying and for inter-library loans after expiry of a bar on access approved by the University.

Signed.....  .....

Date.....  
12.06.23

The University's ethical procedures have been followed and, where appropriate, that ethical approval has been granted.

Signed.....  .....

Date.....  
12.06.23

## Abstract

The automotive industry is undergoing a major transformation. New environmental legislation, changing consumer requirements, Industry 4.0 technologies, and advancements in battery technologies, have contributed to an industry-wide shift towards electric powertrain. To remain competitive in this rapidly changing environment, automotive manufacturers must ensure high levels of technical and organisational innovation to transition towards digital and data-driven business practices.

This research aims to address these growth opportunities and manage ongoing change in three steps. First, the literature on machine learning applications in automotive manufacturing is critically reviewed and the barriers to developing and implementing machine learning are discussed. Secondly, a structured framework is developed to assess the industry 4.0 maturity of automotive manufacturing operations and guide digital transformation at the factory level. In the third and final step of this research, two machine learning projects identified by the assessment are presented in detail. The first case study presents an anomaly detection solution to identify process errors in engine assembly. This research introduces multiple advancements in anomaly detection in manufacturing, including the introduction of the Anomaly No Concern class. The second case study is a greenfield project to explore new digital value chains to add value to EV customers and explore data-as-a-service business models. This case study uses a combination of Google Street View data and GIS data to identify houses suitable for EV charging and represents a major advancement towards fully automated remote surveying of the built environment.

To conclude, multiple advancements are presented that contribute to the academic literature. Clear stepwise frameworks support the proposed industrial solutions to develop, implement and replicate these solutions across the business. These contributions have been used to support ongoing digitalisation efforts, implement state-of-the-art anomaly detection solutions, and explore new data-as-a-service business models in one of the world's largest automotive companies.



## Acknowledgements

I am deeply indebted to my academic supervisor, Professor Cinzia Giannetti, for her exceptional mentorship and unwavering support throughout the entire duration of this project. Professor Giannetti's dedication and guidance have been instrumental in shaping my research, academic pursuits, and personal growth. Her contributions to this project go above and beyond what is expected of an academic supervisor, and I cannot overstate the profound impact she has had on my personal development.

I would also like to express my sincerest gratitude to my industrial supervisor, Steve Buck, at Ford Motor Company. Without his unwavering support and guidance, this research would not have been possible. His mentorship and collaboration have ensured that this research has had a meaningful impact in the real world and will continue to do so in the future. Moreover, his mentorship has played a pivotal role in my professional development.

I extend my heartfelt thanks to Michael Higgins, Hessel Van Dijk, and Andreas Billstien at Ford Motor Company for their invaluable contributions to the delivery of the anomaly detection research presented in Chapter 4. Working under Michael's exceptional management skills has been an absolute pleasure. Hessel and Andreas have also provided exceptional support and mentorship, generously giving up their time to teach and guide me on this journey. I am forever grateful for all of their support and commitment to ensuring my involvement in the delivery of this research.

To Cinzia, Steve, Michael, Hessel, and Andreas, I am delighted to be able to convey my formal thanks to you all as part of this thesis. I am equally grateful for the friendships we have formed over the years, and I look forward to our continued collaboration in the future.

## Dedication

To my loving fiancé Ellen for helping me throughout this project and for making these past 5 years the best of my life.

# Table of Contents

<b>List of Tables</b>	<b>xiii</b>
<b>List of Figures</b>	<b>xv</b>
<b>1 Introduction</b>	<b>2</b>
1.1 A Brief History of the Industrial Revolutions . . . . .	3
1.2 Machine learning in the Factory of the Future . . . . .	6
1.3 Business Models of Industry 4.0 . . . . .	8
1.4 Objectives . . . . .	10
1.5 Thesis Layout . . . . .	10
1.6 Industrial Impact and Research Outputs . . . . .	11
1.6.1 International Research Journal Publications . . . . .	12
1.6.2 International Conference Papers and Presentations . . . . .	12
1.6.3 Industrial Projects . . . . .	12
1.6.4 Awards . . . . .	14
1.6.5 External Factors . . . . .	14
<b>2 Machine Learning in the Automotive Industry</b>	<b>16</b>
2.1 Summary . . . . .	16
2.2 Introduction . . . . .	17
2.3 Search method and strategy . . . . .	18

2.4	Search results . . . . .	21
2.4.1	Data Extraction Fields . . . . .	22
2.5	Discussion . . . . .	24
2.5.1	Enabling Technologies of Machine Learning in Manufacturing . . . . .	31
2.5.2	Digital Skills . . . . .	46
2.5.3	Management . . . . .	51
2.5.4	Computer Vision . . . . .	58
2.5.5	Key Findings of Machine Learning Opportunities and Barriers . . . . .	66
2.6	Conclusion . . . . .	69
<b>3</b>	<b>Industry 4.0 Readiness Assessment</b>	<b>72</b>
3.1	Summary . . . . .	72
3.2	Introduction . . . . .	73
3.3	Related Research . . . . .	77
3.3.1	Assessment Tools Limitations . . . . .	82
3.3.2	Summary of Key Findings . . . . .	86
3.4	Methodology . . . . .	88
3.4.1	Design Approach . . . . .	88
3.5	An Industry 4.0 Assessment tool for Automotive Manufacturers . . . . .	94
3.5.1	Planning . . . . .	94
3.5.2	Workshops . . . . .	95
3.5.3	Questionnaires . . . . .	98
3.5.4	Interviews . . . . .	98
3.5.5	Scores and Report Writing . . . . .	99
3.5.6	De-brief . . . . .	100
3.6	Areas of Assessment . . . . .	101
3.6.1	Manufacturing Technology Readiness . . . . .	101

3.6.2	Strategy, Organisation, and Culture . . . . .	112
3.7	Applying the Assessment Tool . . . . .	123
3.7.1	Summary of the assessments and resultant outcomes . . . . .	123
3.7.2	Time and Resource Involved in Deploying the Method . . . . .	127
3.7.3	Qualitative analysis of the assessment outcomes . . . . .	129
3.8	Conclusion . . . . .	132
<b>4</b>	<b>In-Process Anomaly Detection in Engine Assembly</b>	<b>136</b>
4.1	Summary . . . . .	136
4.2	Introduction . . . . .	139
4.3	Related Research . . . . .	144
4.3.1	Types of Anomalies . . . . .	145
4.3.2	Dimensionality Reduction . . . . .	147
4.3.3	Semi-Supervised Anomaly Detection . . . . .	149
4.4	Methodology . . . . .	151
4.4.1	Labeled Data . . . . .	151
4.4.2	Machine Learning Model Descriptions . . . . .	154
4.4.3	Metrics . . . . .	171
4.5	Results . . . . .	172
4.5.1	LSTM Results . . . . .	174
4.5.2	Semi-Supervised Clustering Results . . . . .	175
4.6	Discussion of Results . . . . .	176
4.6.1	LSTM . . . . .	176
4.6.2	Semi-Supervised Clustering Findings . . . . .	176
4.7	Industrial Case Study . . . . .	178
4.7.1	Opportunities for Future Work . . . . .	181
4.8	Conclusions . . . . .	182

<b>5</b>	<b>Identifying Houses Suitable for EV Charging</b>	<b>185</b>
5.1	Summary . . . . .	185
5.2	Introduction . . . . .	186
5.3	Related Research . . . . .	188
5.3.1	Machine Learning in Remote Sensing . . . . .	188
5.3.2	Streetscape Remote Sensing . . . . .	194
5.3.3	Geographic Data Sources and Digital Tools . . . . .	195
5.4	Image Processing Methodology . . . . .	197
5.4.1	What makes a property EV suitable? . . . . .	198
5.4.2	Data Acquisition . . . . .	199
5.4.3	Image Classification . . . . .	200
5.5	Proposed Workflows to Identify EV Suitable Houses . . . . .	204
5.5.1	Workflow 1 . . . . .	205
5.5.2	Workflow 2 . . . . .	208
5.5.3	Workflow 3 . . . . .	212
5.6	Discussion . . . . .	216
5.7	Conclusion . . . . .	219
<b>6</b>	<b>Conclusions</b>	<b>221</b>
6.1	Main Research Contributions . . . . .	223
6.2	Future Work . . . . .	227
6.2.1	Industry 4.0 Assessment Outcomes and Future Work . . . . .	227
6.2.2	Anomaly Detection Future Work . . . . .	230
6.2.3	Mapping Homes Suitable for EV Charging Future Work . . . . .	231
	<b>References</b>	<b>232</b>

<b>APPENDICES</b>	<b>232</b>
.1 Machine Learning Models and Descriptions . . . . .	265
.1.1 Machine Learning Approaches . . . . .	267
.1.2 Common Machine Learning Models . . . . .	271
.2 Applying the Assessment tool: A Case Study at Ford Motor Company . .	280
.3 Industry 4.0 Assessment Questionnaire . . . . .	281

# List of Tables

2.1	A list of all papers in the reviewed literature for which machine learning experiments are conducted, as well as an overview of the research topic and any critiques identified by the author. . . . .	25
3.1	A comparison between the various existing approaches to assessing the Industry 4.0 readiness of a company. *Limited information available. . . . .	82
3.2	The following suggested interview list was shared with senior management during the planning phase. . . . .	97
3.3	A simplified version of the scoring matrix for manufacturing technology. . .	102
3.4	Simplified version of the scoring matrix for Strategy, Organisation, and Culture. . . . .	114
3.5	A categorised list of outcomes of the Industry 4.0 assessment performed across all of Ford's UK manufacturing sites. These outcomes are expanded upon further in section 2.6 of the appendix. . . . .	130
4.1	Composition of training and testing datasets to evaluate performance of the models. . . . .	154
4.2	Comparison of ML approaches for test Dataset 1a where ANC's are considered as True Anomalies. . . . .	173
4.3	Experiment results for Dataset 2a where ANC's are considered as True Anomalies. . . . .	173
4.4	Experiment results for Dataset 1b where ANC's are NOT considered as True Anomalies. . . . .	173



4.5	Experiment results for Dataset 2b where ANC's are NOT considered as True Anomalies. . . . .	174
5.1	A comparison of the most commonly used CNNs in related remote sensing works in the reviewed literature. . . . .	194
5.2	The following definitions were used when labelling the training and testing data. . . . .	201
5.3	A breakdown of the image dataset used to retrain CNN 1. The Oswestry data is used for validation and hyper-parameter optimisation. . . . .	202
5.4	A breakdown of the image dataset used to retrain CNN 2. The Oswestry data is used for validation and hyper-parameter optimisation. . . . .	202
5.5	Geographical data sources . . . . .	210
1	A list of all machine learning methods used in the reviewed literature and the respective references arranged from most common to least common. . . . .	268

# List of Figures

2.1	A flow diagram outlining the main steps of the literature searches. . . . .	20
2.2	Graphs showing the year of publication of the 90 papers returned by the searches (left) and a visualisation of the ratio of journal articles to conference papers (right). . . . .	22
2.3	A pie chart (Top) shows the most common topics discussed in the reviewed literature. Note that one paper may include multiple key topics. A chord diagram (Bottom) shows the relationship between these key topics. The bar width between two topics is proportional to the number of times those topics appear in the same research paper. For example, papers discussing computer vision also often discuss quality assurance but rarely discuss unsupervised learning, as indicated by the wide and thin bars. . . . .	23
3.1	An iterative approach was taken when designing the assessment tool, using feedback from industrial partners to guide the design process throughout. .	89
3.2	A description of each of the six stages of assessment and the key tasks associated with each stage. . . . .	93
3.3	An timeline for the assessment to take place was proposed during the planning phase and agreed during the workshop. . . . .	95
3.4	Four examples of questionnaire responses related to data usage and availability.	124
3.5	An example comparing how some questionnaire responses vary between different levels of management. . . . .	125
4.1	An example of a line worker using a DC nut runner tool in engine assembly.	141

4.2	Datasets 1 and 2 with a random example of a single observation highlighted in red. Dataset 1 is a manual nut runner process with high variability. Dataset 2 is an automated process where the staging problem can be clearly observed.	142
4.3	Labelled PCA plots for Datasets 1 and 2. 'Normal' data are shown in blue, 'True Anomalies' are shown in red, and 'Anomaly No Concern' are shown in green. Here it can be seen that for both datasets, not all anomalies are outliers, and not all outliers are anomalies.	143
4.4	The data labelling dashboard allows users to label batches of 12 normal waveforms at a time.	153
4.5	UMAP uses combinations of simplicies to provide a simplified representation of the continuous topological space defined by the high dimensional dataset $X$ while retaining the global and local structures that define the space.	158
4.6	The same GMM approach applied using t-SNE and UMAP to reduce and cluster the data. Labelled data includes Nominal points (blue), ANC (green), and True Anomalies (red).	160
4.7	Outlier regions calculated using Gaussian mixture model trained on the reduced normal data. Any points that fall in the red area are identified as anomalies. Labelled data includes Nominal points (blue), ANC (green), and True Anomalies (red).	163
4.8	Graphs showing the Sigmoid and Tanh activation functions used in the LSTM network.	165
4.9	The high level LSTM architecture is similar to that of an RNN.	165
4.10	Inside the LSTM cell the three gates controlling the flow of information can be observed.	166
4.11	Examples of how LSTM forecasts vary between normal waveforms (Top) and anomalous waveforms (Bottom). The LSTM performs poorly when forecasting anomalous oscillations between 50 and 200 time steps, leading to a high reconstruction error above the anomaly threshold.	168
4.12	Two methods are compared at to find the error threshold for the LSTM forecasting approach: one setting the threshold as 1.5 times the interquartile range above the third quartile (Top), and the other using the elbow method on a plot with RMSE values sorted in ascending order (Top).	169

4.13	Results of the parameter optimisation experiments showing how anomaly detection accuracy varies with the models' performance at predicting normal data. . . . .	171
4.14	Four examples of waveforms included in the off-line trial. Any process data within the green boundary is classified as 'Normal' by the proposed PCA-GMM model. Green points indicate process data that are classified as 'Normal' by both the current PCA-DBSCAN model and the proposed PCA-GMM model. Blue points show process data that has been labelled as 'Normal' but the current PCA-DBSCAN model, but classed as a 'True Anomaly' by the proposed PCA-GMM model. Red points are classified as 'True Anomalies' by the current PCA-DBSCAN model. Because of the lack of validation, data that are suspected to be False Positives or False Negatives are highlighted. . . . .	180
5.1	The AlexNet architecture is one of the most simple CNN architectures with only 8 layers and 60 million paramters [1]. . . . .	190
5.2	The two types of inception modules introduced in the GoogLeNet CNN architecture [2]. . . . .	192
5.3	The VGG-16 network contains 13 3x3 convolutional layers and three fully connected layers. [3]. . . . .	193
5.4	The ResNet architecture uses an identity connection to feed information forward between layers to address the vanishing gradient problem [4]. . . .	194
5.5	Locations from which training, testing, and case study data were retrieved.	198
5.6	Example images from the 6 categories used to train the CNNs. . . . .	200
5.7	A box plot showing the range of F-scores resulting from the hyperparameter optimisation experiments for three different architectures. Googlenet performs the best for both CNN1 and CNN2. . . . .	203
5.8	Confusion matrices of the optimal networks selected for CNN1 and CNN2	204
5.9	Flow diagram for Workflow 1. . . . .	205
5.10	Survey area for the Birmingham case study. . . . .	206
5.11	Confusion matrices showing the performance of CNN1 and CNN2 on the Birmingham dataset. . . . .	206

5.12	A heat map of EV-suitable properties in the Birmingham survey area. . . .	207
5.13	Flow diagram for Workflow 2 . . . . .	208
5.14	Confusion matrices showing the performance of CNN1 and CNN2 on the Petersfield dataset. . . . .	211
5.15	A map showing all locations of suitable and unsuitable properties in the Petersfield survey area. . . . .	212
5.16	A flow diagram for Workflow 3. . . . .	213
5.17	Gloucestershire Map [5] . . . . .	215
5.18	Workflow 3 F-scores . . . . .	216
5.19	The confusion matrices of the Gloucestershire test locations for Workflows 2 and 3. . . . .	217
1	The area under the ROC curve is a metric used to evaluate the performance of classifiers. ( <i>Source: <a href="https://en.wikipedia.org/wiki/Receiver_operating_characteristic">https://en.wikipedia.org/wiki/Receiver_operating_characteristic</a></i> ) . . . . .	266
2	An artificial neuron used in neural networks. ( <i>Source: <a href="https://lanstonchu.files.wordpress.com/2018/08/artificial-neuron/">https://lanstonchu.files.wordpress.com/2018/08/artificial-neuron/</a></i> ) . . . . .	267
3	An example of a 3 layer ANN, sometimes called a Multilayer Perceptron. Each node in the network is an artificial neuron depicted in Figure 2 Adding at least one more layer to the network would make it a Deep Neural Network. . . . .	272
4	Plots showing the most common activation functions used in machine learning algorithms. . . . .	273
5	An Example of the CNN architecture [6]. . . . .	274
6	An example of an autoencoder architecture where each node of the network is an artificial neuron. . . . .	275
7	example of a hyperplane generated by a linear SVM algorithm to separate classes in 2D feature space ( <i>Source: <a href="https://www.analyticsvidhya.com/blog/2020/10/the-mathematics-behind-svm/">https://www.analyticsvidhya.com/blog/2020/10/the-mathematics-behind-svm/</a></i> ). . . . .	276
8	A graphic demonstrating the kernel trick in which the kernel function $\phi$ being used to map points from a 2D input space into a 3D feature space to learn non-linear relations between the classes ( <i>Source: <a href="https://medium.com/@KunduSourodip/finding-non-linear-decision-boundary-in-svm-a89a97a006d2">https://medium.com/@KunduSourodip/finding-non-linear-decision-boundary-in-svm-a89a97a006d2</a></i> ). . . . .	276

9	A simple decision tree with class labels of each node depicted by coloured dots. Notice the increase in class purity through the tree with nodes terminating when only 1 class label exists in the subset. . . . .	278
10	The random forest architecture showing bagging and ensemble tree classification. . . . .	279
11	Architecture of a Generative Adversarial Networks (GAN) . . . . .	280
12	Industry 4.0 Assessment Questionnaire page 1 . . . . .	281
13	Industry 4.0 Assessment Questionnaire page 2 . . . . .	282
14	Industry 4.0 Assessment Questionnaire page 3 . . . . .	283
15	Industry 4.0 Assessment Questionnaire page 4 . . . . .	284
16	Industry 4.0 Assessment Questionnaire page 5 . . . . .	285
17	Industry 4.0 Assessment Questionnaire page 6 . . . . .	286
18	Industry 4.0 Assessment Questionnaire page 7 . . . . .	287
19	Industry 4.0 Assessment Questionnaire page 8 . . . . .	288
20	Industry 4.0 Assessment Questionnaire page 9 . . . . .	289
21	Industry 4.0 Assessment Questionnaire page 10 . . . . .	290
22	Industry 4.0 Assessment Questionnaire page 11 . . . . .	291
23	Industry 4.0 Assessment Questionnaire page 12 . . . . .	292
24	Industry 4.0 Assessment Questionnaire page 13 . . . . .	293
25	Industry 4.0 Assessment Questionnaire page 14 . . . . .	294
26	Industry 4.0 Assessment Questionnaire page 15 . . . . .	295

# List of Abbreviations

ADAPT	Automated Detection of Anomalous Production Tests
AGV	Autonomous Guided Vehicles
AI	Artificial Intelligence
AIAC	Artificial Intelligence Advancements Center
ANC	Anomaly No Concern
ANN	Artificial Neural Network
API	Application Programming Interfaces
ARIMA	Autoregressive Integrated Moving Average
BEP	Bridgend Engine Plant
CAD	Computer-Aided Design
CMMI	Capability Maturity Model Integration
CNN	Convolutional Neural Network
Cobots	Collaborative Robots
COVID	Coronavirus
CPS	Cyber-Physical System
DAE	Deep Autoencoder
DBN	Deep Belief Network
DBSCAN	Density Based Spacial Clustering
DEP	Dagenham Engine Plant
DMAIC	Define, Measure, Analyze, Improve, and Control
DNO	District Network Operator
DREAMY	Digital Readiness Assessment Maturity
EM	Expectation-Maximization
EV	Electric Vehicle
FIS	Factory Information System
FMS	Flexible Manufacturing Systems

FMS	Flexible Manufacturing Systems
GAN	Generative Adversarial Network
GCP	Google Cloud Platform
GDIA	Global Data Insights and Analytics unit
GIS	Geographic Information System
GMM	Gaussian Mixture Model
GPU	Graphics Processing Unit
HPC	High-Performance Computing
ICT	Information and Communication Technology
IIoT	Industrial Internet of Things
IoT	Internet of Things
I-RPA	Internal Robotic Process Automation
k-NN	k- Nearest Neighbour
LSTM	Long Short Term Memory
LULC	Land Use and Land Cover
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MSE	Mean Square Error
MSOA	Middle Super Output Area
NDVI	Normalized Difference Vegetation Index
OEE	Overall Equipment Effectiveness
OS	Ordnance Survey
OSM	Open Street Map
PCA	Principal Component Analysis
PDF	Probability Density Function
PdM	Predictive Maintenance
PEV	Plug-in Electric Vehicle
PTME	Powertrain Manufacturing Engineering
R&D	Reserach and Development
RF	Random Forest
RFID	Radio Frequency Identification
RMSE	Root Mean Squared Error
ROI	Return On Investment
RUL	Remaining Useful Life



SPC	Statistical Process Control
SVM	Support Vector Machine
TDWI	Transforming Data with Intelligence
t-SNE	Distributed Stochastic Neighbor Embedding
UMAP	Uniform Manifold Approximation and Projection
UPRN	Unique Property Reference Number
WDAD	Well-Defined Anomaly Distribution

# Chapter 1

## Introduction

The automotive industry is undergoing a major technological revolution. Contributing factors to this disruption include emerging manufacturing technologies, environmental legislation, and changing consumer requirements. These factors, combined with improvements in lithium-ion battery technologies, digitalisation, and automation, have resulted in an industry-wide shift towards electric powertrain as an alternative to traditional internal combustion engines that have dominated the market for nearly a century. A major driver of this change was the Paris Climate Agreement, in November 2016 outlined ambitious targets to cut global greenhouse emissions. While many hybrid and EV products existed before the Paris Climate Agreement, this global call to environmental action resulted in many leading economies introducing legislation to ban the sale of new internal combustion engine vehicles over the coming decades [7].

More recently, automakers have faced new supply chain disruptions following the COVID-19 pandemic, the Russian invasion of Ukraine, and the subsequent geopolitical and global economic instability [8]. The resultant energy crisis has meant many countries are putting new policies and guidelines in place to motivate consumers to reduce their reliance on fossil fuels further [8]. These recent market disruptions introduce new risks to automotive manufacturers due to supply chain disruptions, global trade tensions, fluctuating sales, and political uncertainty [9, 8]. There remains some uncertainty around how regional policies will affect the global EV market, with EV sales slowing in some markets, including the US and China, in recent years [9]. However, as emissions regulations become more stringent, battery technologies advance and EV products become cheaper with longer ranges, experts predict that EV models will emerge as the dominant mode of powertrain in the near future [9].

In addition to the challenges of transition towards electric powertrain, the automotive industry also finds itself amid a wider technological revolution affecting the whole manufacturing sector, referred to as 'Industry 4.0'. The term Industry 4.0 was first used to demonstrate the impact of Smart Systems and the Internet of Things (IoT) on the German manufacturing sector and set the basis for the German government's "High-Tech Strategy 2020 Action Plan"[10, 11]. Since then, the term has evolved to describe various emerging technologies and business practices that are redefining the manufacturing sector as industries shift towards increased digitalisation and automation. Such technologies include big data analytics, IoT, Cyber-Physical Systems (CPSs), wearable technologies, additive manufacturing, cloud computing, advanced robotics, and machine learning. The application of Industry 4.0 is not confined to manufacturing technologies but also considers organisational innovation. As organisations digitalise and integrate all areas of business operations throughout the product life-cycle, new business models are enabled to deliver added value to customers [12].

## 1.1 A Brief History of the Industrial Revolutions

Industry 4.0 is an abbreviation of 'the fourth industrial revolution'. Before discussing this further, a brief overview of the previous industrial revolutions is presented. Entire books have been dedicated to understanding the causes of each separate revolution, and therefore the discussion is limited to a high-level overview in relation to advancements in manufacturing to provide some background and context for the reader. For further details on the first and second revolutions, the reader is referred to books by P.M.Deane and R. Joel Mokyr respectively [13, 14].

The three previous industrial revolutions were all triggered by various innovations, the first of which can be traced back to Manchester, Great Britain, in the 18th century [11]. Innovations during this era often had little to no scientific base, with discoveries often arising serendipitously through the practical application of the engineering, medical, and agricultural knowledge base of the time [15]. During the mid-18th century where the mechanisation of the textile industry was made possible by the introduction of water- and steam-powered manufacturing processes such as the Roberts Loom [11]. Before the Roberts Loom, workers manually powered machinery using foot pedals while repeatedly weaving through a taught matrix of tightly strung threads. By mechanising these processes, the physical limitations of the human body were overcome, allowing less skilled users to oper-

ate the looms, making the process considerably faster and cheaper. Similar machine tools and technologies began emerging throughout the late 18th and early 19th century, bringing with it a new era of industrialisation that quickly spread across the globe, transforming what was previously an agricultural society into a new industrial society [15].

The second industrial revolution can be traced back to multiple innovations in manufacturing between the 1870s and 1920s that addressed the changing nature of organisation and production, and their associated technologies [15]. During this period, technologies such as railroads, telegraph networks, and city infrastructures were expanded massively thanks to various technological advancements that enabled the mass production of steel, chemicals, and oil [15]. New systems, such as electric power, were also introduced that helped further improve mass production systems [11]. However, the second industrial revolution is not defined by these new technologies but rather by the new approaches to the organisation of production and more scientific approaches to innovation. The economy of scale became increasingly important as high-volume manufacturing became widespread and giant corporations began to rise, such as Carnegie Steel, Dupont, Ford Motors, and General Electric [15]. One of the most significant advancements in manufacturing during this era was the introduction of the moving assembly line, which was first introduced in 1870 to improve the process efficiency of meat production in slaughterhouses [11]. The concept was later popularised by Henry Ford, who proved that combining the concepts of division of labour, continuous flow processes, and interchangeable parts made it possible to produce highly complex products at low prices [15].

There is a large crossover between Industry 1.0 and Industry 2.0. Both revolutions involved improved manufacturing processes and technologies, spurring new eras of economic and societal growth. However, the second industrial revolution is distinguished by two main factors. Firstly, the new approaches to production organisation and a greater focus on systematic, scientific approaches to innovation and process optimisation. Secondly, industrialisation spread much quicker during Industry 2.0 thanks to improved transport and communication networks, and growth spread outside of the British Isles into the United States, Germany, and other countries across Europe [15, 16].

The major technological advancement that marked the beginning of the third industrial revolution was the development of the first Programmable Logic Controller (PLC) in 1968 [17]. This enabled users to digitally reprogram a computer to perform multiple operations and kickstarted the information and communication technology (ICT) age which quickly spread across the world, redefining the way that companies operated [17]. Rapid devel-

opment over the following decades led to smaller, more affordable computers aimed at both consumers and large companies. The introduction of personal computers drastically changed the way people worked and lived. As the world became evermore connected, personal computers quickly shifted from being a luxury item to a necessity in much of the western world.

Industry 4.0 is widely regarded by scholars and industry as an evolution from the third industrial revolution, in the same way, that the second was a continuation of the first [11, 15]. In terms of technological advancement, the rise of the Internet triggered the fourth industrial revolution, the subsequent development of Cyber-Physical Systems (CPS) and the creation of a digital value chain [17, 18, 19]. CPSs are the mechanism that combines elements of both the physical world and digital software, providing a means for components, objects, devices, and other things to communicate information [18]. During the beginning of the 21st century, the popularity of these wireless, networked devices grew significantly in both industrial and commercial applications. The widespread distribution of CPS is the main distinguishing aspect of the fourth industrial revolution. With further miniaturisation and improved infrastructures such as cloud computing, these technologies soon evolved into a wider system of connected physical CPSs termed ‘The Internet of Things’ (IoTs), a major enabling technology of the fourth stage of the industrial revolution [11]. The key difference between IoT and CPS is that CPS involves integrating software, computation, networking, and physical processes, not necessarily via the internet. In contrast, IoT refers to physical objects and systems connected through internet networks.

By integrating CPSs and IoTs into every aspect of a production line with the appropriate framework, manufacturers can work towards a fully networked factory with the potential to develop a self-organising factory environment, often referred to as a Smart Factory which represents the pinnacle of the Industry 4.0 manufacturing environment [20]. These highly automated manufacturing environments require high levels of digital skills and new management approaches and business models that challenge well-established manufacturing practices. To address these challenges, change management, business intelligence, and automated production management are among the most common research topics in the Industry 4.0 literature [21, 22, 23, 12]. Industry 4.0 is a human-centric philosophy in which digitalisation and automation are explored in a socially sustainable way as a means to empower the existing workforce to drive innovation while minimising the negative social impact. However, some researchers raise concerns that in real-world applications, the human-centric concepts of Industry 4.0 are not well-understood [24].

There are many similarities between the transition from Industry 1.0 and Industry 2.0 compared to that between Industry 3.0 and Industry 4.0. The first industrial revolution resulted from technological innovations that overcame the physical limitations of the human body in production environments and radically increased production rates. This, in turn, led to a new era of high-volume production across various industries demanding new approaches to business organisation and production management to manage this new supply chain. Similarly, in the third industrial revolution, technologies such as PLCs, advanced data analytics, robotics, and CPSs allowed users to overcome further physical limitations of the human body but, more importantly, overcome the analytical limitations of the human mind. Organisations are again seeing how technical innovations in automated data analytics are changing approaches to business organisation and production management.

## **1.2 Machine learning in the Factory of the Future**

In the modern automotive industry, data and human resources are two of the most valuable assets to a company. As the workforce is upskilled and manufacturing environments become increasingly integrated, data becomes increasingly valuable [25, 26]. This value creation lies in the ability to apply data analytics and generate insights into integrated systems and processes to users throughout the value chain, leading to new organisational and technical knowledge which can deliver competitive advantage and drive further innovation [25]. Value is assigned not only to data but also to the integrated systems, technologies, and people allowing for its collection, integration, and exploitation.

As manufacturers continue to improve data collection and integration by installing cheap sensors and networked IoT systems, new challenges arise in analysing these data. Manufacturers must adapt to deal with high volumes of high-dimensional manufacturing data that do not follow Gaussian distributions. Existing analytic toolsets commonly used in automotive manufacturing environments, such as Six Sigma and Excel, will become increasingly ill-suited to analyse these data. Manufacturing organisations must recognise the limitations of the current analytical tools and adopt new approaches to data management and analytics.

Machine learning has emerged as a powerful tool to analyse big manufacturing data with various applications discussed in the reviewed literature across all aspects of manufacturing production and the broader organisation [27, 28]. Research in machine learning in automo-

tive manufacturing has increased considerably in the last five years. While machine learning applications are already fairly widely employed in the manufacturing sector, much of this research into automotive applications has focused on proof-of-concept solutions, with few papers demonstrating application readiness. This highlights a major opportunity for both further academic research and industrial application of these technologies in the automotive sector.

Machine learning presents tremendous growth opportunities through increased automation and business intelligence. Autonomous production systems driven by machine learning models can support production managers to predict production progress, react quickly to issues and provide prescriptive analytics to support root cause analysis [29]. At the process level, various machine learning applications can deliver quality improvements, time savings, and cost savings. Algorithms can monitor in-built sensors to predict tool breakages in machining processes, to automatically schedule maintenance and reduce machine downtime [30]. Anomaly detection systems can monitor processes and automatically report when processes exceed limits, reducing the requirements for manual inspection and designate products for repair [31]. Integrating engineering processes throughout the product life-cycle allows production to be tailored to meet fluctuating consumer demand of multiple product families. This delivers added customer value through highly customisable products while also reducing lead times [27]. Flexible manufacturing systems enabled by machine learning can also react quickly to changes in the local manufacturing environment to deliver reduced labour costs, improved inventory management, reduced lead times, and reduced manufacturing costs [27]. As data produced by automotive products become increasingly integrated, added value can be delivered to customers through machine learning enabled services such as predictive and prescriptive maintenance. By monitoring vehicle data in-service, a customer can be notified if maintenance is required improving safety, reducing cost, and improving customer experience [32].

Managing the transition toward implementing machine learning technologies can be challenging. This transition requires a well-structured long-term strategy that challenges existing organisational cultures [33, 34, 35]. It requires investment in new technologies, practices for which Return on Investment (ROI) is difficult to quantify [33, 34, 35]. Research shows that the use of machine learning in automotive manufacturing applications has been limited until recent years [36, 37]. Researchers and industry practitioners are still working to understand how best to create value from the combination of machine learning and other Industry 4.0 technologies, given that many of these technologies are still in their infancy

[37]. The application of machine learning requires careful selection of the suitable use case, requiring both technical knowledge of advanced data analytics and domain-specific knowledge of the target process. Gaps in organisational knowledge in automotive manufacturers have resulted in a reluctance to embrace machine learning and its enabling technologies due to the challenges of complexity, technical expertise, and uncertainty of investment requirements [38]. Previous research has also explained the slow adoption of machine learning enabling technologies due to the limited availability of skills and poor change management [39].

To address these challenges and support organisational change, considerable research has been done in developing Industry 4.0 maturity assessments to measure progress towards industry 4.0 and develop roadmaps to guide this transformation [12, 40, 41, 42, 43, 44, 45, 46]. These self-assessment maturity models allow organisations to compare technological and organisational aspects of operations against a well-defined benchmark to quantify progress toward Industry 4.0 and highlight areas to improve. Industries with high technological readiness levels supported by well-communicated sustainable automation strategies will achieve higher process scores and Industry 4.0 readiness. Despite research in Industry 4.0 change management, research shows that the full realisation of machine learning and its enabling technologies is yet to be realised in the automotive manufacturing industry [47, 48]. Further research is required to understand how organisations should manage the transition towards increased digitalisation and automation and how to prioritize investments in emerging technologies to maximize value creation.

### **1.3 Business Models of Industry 4.0**

Innovation is critical to successfully managing this change and requires a strong understanding of the most recent challenges, and opportunities of Industry 4.0 technologies [49]. Innovation is required not only in technological research and development but also in organisational aspects as new business models emerge as a result of new integrated Industry 4.0 technologies [37, 50]. Digitalisation is a key enabler of new business models, where novel digital platforms create new digital markets and embrace consumer involvement in the product and service innovation process [51]. These business models include 'The Sharing Economy', 'On-Demand Services', 'Manufacturing-as-a-Service', 'Data-as-a-Service', 'Mobility-as-a-Service', and Circular Economy [52, 53, 54, 55, 56]. As new products and services are introduced, and data are integrated vertically and horizontally



throughout the organisation, machine learning and Big Data analytics will play a critical role in the business strategy [57]. Predictive analytics and automated process control support real-time prescriptive analytics mitigating risks such as unpredictable raw material quantity, quality, availability variations, and constantly changing market trends and consumer behavior [57].

Customer behavior analytics presents a significant opportunity for automotive manufacturers to improve their understanding of the potential value of different customer segments through analysing external data sources such as social media and internal sources [58]. This knowledge can be used to strategically target new customers as well as improve customer experience and ensure the loyalty of existing customers [58]. Big data analytics of consumer data also significantly improve manufacturing process agility and flexibility, allowing production and procurement to adapt to changing markets. By gathering and analysing real-time data from an automaker's fleet and warranty data, cloud-based big data processing will enable new after-market services for customers. Examples include predictive and preventative maintenance of products, various infotainment services, and self-driving support services such as autopilot and collision prevention that are already found in many modern vehicles [59].

Like many automotive companies, Ford Motor Company recently began offering Mobility-as-a-Service. Mobility-as-a-Service models involve customers paying an upfront fee, followed by weekly payments to access a vehicle with all-inclusive service, including insurance, MOTs, servicing, and maintenance [60]. Non-ownership and sharing of vehicles are not limited to passenger transport, with recent research in the Mobility-as-a-Service model also highlighting the potential for opportunities in freight transport. As the fixed cost of ownership is replaced with variable costs of travel use, electric vehicles and autonomous driving become major enablers of these services [56, 61].

A key challenge of business model innovation is understanding how to identify, select, and implement digital innovations [37]. This is particularly true in areas where innovations are neither a product nor a service but instead promise to add additional value for customers. These innovations are difficult to develop a business case and equally difficult to measure the impact of outcomes [37]. Further research is required to understand business model innovations at both the organisational and technical levels [37].

## 1.4 Objectives

The overall aim of this thesis is to present and critically review the application of machine learning in the automotive industry and develop various methodologies to support the further uptake of these technologies to deliver increased quality and value creation in various organisational settings within the automotive industry. To achieve this goal, we outline the following main objectives:

1. Identify the main machine learning technologies used in automotive manufacturing and identify the barriers and opportunities for further sustainable growth and value creation in this field.
2. Development of a strategic framework to support the future uptake of machine learning in the automotive industry with a focus on sustainability.
3. Using the proposed framework, develop machine learning solutions to create value from existing data sources in the automotive industry.

## 1.5 Thesis Layout

This thesis is structured as follows. Chapter 2 introduces and critically reviews the main body of the literature on machine learning applications in automotive manufacturing. It includes descriptions of the various machine learning approaches with detailed descriptions and examples of the most commonly used models and architectures. The barriers to developing and implementing machine learning and its enabling technologies are also discussed. With human-centered design and sustainability being one of the main pillars of Industry 4.0, particular focus is placed on understanding the social implications of machine learning and its enabling technologies and managerial practices.

Realising the value of machine learning in manufacturing requires investment in emerging technologies and adopting new business practices that challenge cultural norms. This transition must be supported by well-structured step-wise change management strategies communicated to the entire workforce. These challenges are addressed in Chapter 3, in which an Industry 4.0 maturity assessment tool is presented. This structured framework is aimed at automotive manufacturers to self-assess the technological, strategic, and cultural

maturity of production facilities. By reviewing these areas against a well-defined benchmark, growth opportunities can be identified to create added value from existing data. A roadmap toward the organisations' vision of Industry 4.0 can be then be developed. As organisations increase their industry 4.0 maturity level, further opportunities are created to add value through data analytics and machine learning solutions to improve flexibility, productivity, and efficiency.

Chapter 4 presents a machine learning-based In-Process Anomaly Detection system in an engine assembly plant. This project was identified as a result of an Industry 4.0 assessment at Ford Motor Company and demonstrates how machine learning solutions can automate tasks and deliver quality and process improvements following a sustainable approach. This research includes the first instance in the reviewed literature where anomaly detection is applied to time-series data gathered from manual production processes. Manual production data presents challenges such as process staging, human-induced variability, and the subjectivity and ambiguity of the anomalous class. Multiple novel concepts are introduced to overcome these challenges, including the first use of an 'Anomaly No Concern' anomaly class in the literature. Furthermore, to address the lack of publicly available datasets to develop anomaly detection approaches in production settings, the datasets used in this study are made public to support future research.

As automotive manufacturers reach the highest levels of Industry 4.0 readiness, new data-as-a-service business opportunities emerge, enabled by machine learning technologies. Chapter 5 presents a novel method of surveying the built environment using automated machine-learning approaches to identify houses suitable for EV charging. Automotive products are changing how we live and function as a society and play a major role in ensuring a sustainable future. By exploring novel methods such as these to gain insights into the future uptake of electric vehicles, organisations can become better connected with customers, react quicker to market changes, and explore data-as-a-service business opportunities.

## 1.6 Industrial Impact and Research Outputs

Some of the research outputs of this doctorate thesis have been disseminated by publishing papers in international research journals and international conferences. This research has also contributed to the completion of several industrial projects, delivering cost savings of over £10m per annum at Ford Motor Company.

### 1.6.1 International Research Journal Publications

The following papers have been peer-reviewed and published in international research journals:

- Flynn J, Giannetti C. Using Convolutional Neural Networks to Map Houses Suitable for Electric Vehicle Home Charging. *Ai*. 2021;2(1):135–49.
- Borghini E, Giannetti C, Flynn J, Todeschini G. Data-Driven Energy Storage Scheduling to Minimise Peak Demand on Distribution Systems with PV Generation. *Energies* 2021;14:3453.
- Flynn J., Giannetti C., van Dijk H., Anomaly Detection of DC Nut Runner Processes in Engine Assembly. *AI*. 2023 Feb 7;4(1):234-54.

### 1.6.2 International Conference Papers and Presentations

The following papers have been reviewed and published in conference proceedings:

- Flynn J, Brealy E, Giannetti C. Making Green Transport a Reality: A Classification Based Data Analysis Method to Identify Properties Suitable for Electric Vehicle Charging Point Installation. 2021 IEEE Int Geosci Remote Sens Symp IGARSS. 2021;(2018):6229–32.
- E. Brealy, J. Flynn and A. Luckman, "Multi-Criteria Approach Using Neural Networks, GIS, and Remote Sensing to Identify Households Suitable for Electric Vehicle Charging," IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium, 2022, pp. 283-286, doi: 10.1109/IGARSS46834.2022.9884517.
- J. Flynn, Machine Learning Anomaly Detection for In-Process Quality Assurance, 2022 M2A Annual Conference, 2022

### 1.6.3 Industrial Projects

#### Project ADAPT

The work carried out during this EngD has contributed to successfully delivering the ADAPT project, a machine learning strategy to address production anomalies and en-

hance quality in powertrain manufacturing. The author worked as part of a global cross-functional team to develop and implement the anomaly detection solution presented in Chapter 4. Project ADAPT has been successfully implemented in two trials at Ford's Dagenham engine plant and is estimated to deliver greater than £10m per annum savings per plant. Following the success of these trials, the anomaly detection method has demonstrated application readiness and is planned to be rolled out globally within Ford Motor Companies' manufacturing operations. As a subject matter expert in machine learning within Ford's Power Train Manufacturing Engineering team, the author continues to work on project ADAPT to explore further opportunities for cost savings and quality improvements elsewhere within the company. Further applications of the authors' contribution to the ADAPT project have since been identified in vehicle assembly operations, with ongoing collaboration with the Cologne Vehicle Assembly plant to deliver further process optimisation.

As part of this research, a major gap was identified in the company's machine learning development strategy. Prior to this research, no standardised method was in place to support data labelling tasks within the company. Data labelling is a critical stage in model development to produce high-quality training and testing datasets. To address this, software was developed to provide engineers with a dashboard interface to label data and minimise the time taken to label large amounts of time series data. This dashboard has since been adopted by Ford Motor Company to support additional internal projects, including a recent project exploring new approaches to deal with disagreement in labelled production data.

## **Industry 4.0 Assessment**

In Chapter 3, an Industry 4.0 maturity assessment tool is presented. This tool supports automotive manufacturers in identifying growth opportunities in the technological, strategic, and cultural aspects of current operations. As well as identifying specific industrial projects, this assessment tool helps develop a roadmap toward long-term Industry 4.0 goals. Following the success of two assessments at Ford's UK manufacturing sites, the company has adopted this assessment methodology to perform further internal assessments and manage ongoing change towards increased digitisation and automation. Most recently, an Industry 4.0 assessment was performed at Halewood Transmission Plant following £230m investment into new electric vehicle transmission lines. The Industry 4.0 assessment sup-

ported identifying and prioritising digitalisation projects in the warranty, logistics, and production departments. The outcome of this assessment led to new business strategies adopted on-site, including introducing new business metrics to measure progress toward digitalisation. The assessment findings also led to changes in training strategies to up-skill the existing workforce using socially sustainable approaches. In addition to guiding investments into human resources, subsequent investments were also made into emerging technologies as a direct result of this assessment. As Ford Motor Company continues to invest in electrification in EU markets and new production lines are launched, other Industry 4.0 Assessments are planned to manage this organisational change at manufacturing sites across Europe.

#### **1.6.4 Awards**

As part of a team of 5 data science and engineering researchers, the author achieved 3rd place in the 2021 International Open Data Challenge Series hosted by the Energy Systems Catapult and Western Power Distribution. The goal of the Open Data Challenge was to generate innovative and sustainable open-data solutions to social problems. The author's contributions include model construction, validation, experimentation, and result visualisation of the state-of-the-art machine learning approaches used in this research. These models are tuned and combined with ad hoc and convex optimisation techniques to maximize peak load shaving and power storage. This research led to a journal article published in *Energies* as part of the Special Issue Forecasting and Management Systems for Smart Grid Applications.

#### **1.6.5 External Factors**

There have been multiple external factors outside the control of the university and sponsor company that has had some effect on this EngD. This doctorate research topic was originally sponsored by Ford's Bridgend Engine Plant to explore the development and implementation of Industry 4.0 technologies on-site. In June 2019, the Bridgend plant announced its closure which resulted in considerable work on the development of machine learning implementation being unable to continue. As a result, the candidate's industrial supervisor and other industry contacts could no longer provide support for the project. For almost a year, the EngD project's future was uncertain. During this time, it was decided

to change the scope to explore a new topic that relied on open-source data and would not require the input of a sponsor company. This led to the development of the project presented in Chapter 5 to map houses suitable for EV charging using Google Street View data.

It was not until Feb 2020, when Dagenham Engine Plant agreed to continue funding the EngD that a new industrial supervisor was arranged. At this time, the scope of the EngD was to return to that outlined initially. Given the success of the Industry 4.0 Assessment Tool at Bridgend, plans were made to expand this project to replicate this methodology at Dagenham and other European plants. The ongoing works on mapping EV charging locations were socialised with Ford UK teams and work was also agreed to continue in this area. However, shortly after reestablishing regular contact with the sponsor company, the COVID-19 pandemic in March 2020 caused further delays to the project. Site visits were banned, forcing the project scope to be reviewed once more and the focus of work returned to EV charging. It was not until Feb 2021 that work with Ford was able to continue once more, working with new teams at Dunton Technical Center where the candidate could contribute to projects working remotely. Further work on the Industry 4.0 assessment was not able to continue until June 2022.

# Chapter 2

## Machine Learning in the Automotive Industry

### 2.1 Summary

This chapter introduces the main body of the literature related to the use of machine learning in the automotive industry. A systematic review is conducted using the PRISMA framework to answer the following research question: *”How are machine learning technologies and practices used to create value in automotive manufacturing environments, and what are the barriers to their development and implementation?”* . In answering this research question, a range of machine learning approaches are discussed, as well as the most common models researchers use to explore novel manufacturing solutions. Enabling technologies in manufacturing are also discussed, including a range of Industry 4.0 technologies such as IoT, Big Data analytics, Flexible Manufacturing Systems, Digital Twin, Cloud and Edge computing. In addition to discussing the technical aspects of machine learning and their associated challenges, this review also explores cultural and organisational challenges. These findings are used to inform the development of a strategic framework outlined in the following chapter to support the future uptake of machine learning and other Industry 4.0 technologies and practices. Furthermore, these findings are also used to guide research in Chapters 4 and 5 into specific case studies to implement machine learning in automotive applications.

The literature review is structured as follows. Section 2.3 presents the research question, describes the methods used to search the literature, and details the inclusion and exclusion



criteria. This section includes a description of the data extraction sheet used to gather the quantitative and qualitative information from the reviewed literature and is referred to throughout the study. Section 2.4 presents a summary of the results gathered from the data extraction sheet, including key topics, types of data used, data collection methods, common machine learning models, hyperparameter tuning methods, and evaluation metrics. Section 2.4 also includes an overview of the various machine learning training approaches discussed in the reviewed literature and the most commonly used models to provide a foundation of knowledge to which to refer later in answering the research question. The discussion of these results is presented in section 2.5. In this section, the emerging technologies of Industry 4.0, as presented in the reviewed literature, are introduced concerning a wide range of research into their industrial applications in manufacturing. The challenges and opportunities for value creation in manufacturing are discussed and how they are used to support machine learning applications. Three main topics are identified in the reviewed literature: digital skills, computer vision, and management. These main topics are discussed to structure the discussion on these various technological, organisational, and cultural challenges and opportunities and their respective implications on social sustainability. These key themes are digital skills, management, and computer vision and were chosen because they were the most common recurring topics throughout the reviewed literature. Finally, the research conclusions are presented in section 2.6.

## 2.2 Introduction

The automotive industry is undergoing significant technical, organisational, and cultural change, driven by global environmental policies, changing consumer demands, and a wide range of Industry 4.0 technologies. Innovation is critical to successfully managing this change and requires a strong understanding of the most recent challenges and opportunities of Industry 4.0 technologies [49]. Innovation is required not only in technological research and development but also in organisational aspects as new business models emerge as a result of new integrated Industry 4.0 technologies [37, 50].

As manufacturing environments become increasingly integrated, data becomes a significant asset [25, 26]. This value creation lies in the ability to apply data analytics and generate insights into integrated systems and processes to users throughout the value chain, leading to new organisational and technical knowledge which can deliver competitive advantage, and innovation [25]. Value is assigned not only to data but also to the integrated systems,

technologies, and people allowing for its collection, integration, and exploitation. Machine learning has emerged as a powerful tool to analyze big manufacturing data with various applications discussed in the reviewed literature across all aspects of manufacturing production and the broader organisation [27, 28].

Developing machine learning solutions requires people with advanced data analytics skill sets to analyze large volumes of integrated data easily accessible through networked infrastructures. To provide these high levels of data collection, integration and accessibility, multiple enabling technologies of machine learning are discussed in the literature, including: IoT [62], Big Data analytics [52], Flexible Manufacturing Systems [27, 29], Digital Twin [63, 64], Cloud and Edge computing [28, 65, 57]. Machine learning solutions can be implemented at the functional level without the widespread adoption of these technologies. However, as companies reach high maturity levels in these enabling technologies, the opportunities for value creation using machine learning increase significantly as new enterprise-level opportunities are created to deliver business intelligence [12].

Research shows that the full realisation of Industry 4.0 technologies is yet to be realised in the manufacturing industry [47, 48]. This is particularly true for machine learning, with research showing its use in manufacturing applications has been limited until recent years [36]. Consequently, there are gaps in organisational knowledge leading to a reluctance to embrace machine learning and its enabling technologies due to the challenges of complexity, technical expertise, and uncertainty of investment requirements [38]. Previous research has also explained the slow adoption of Industry 4.0 due to the limited availability of skills and poor change management [39].

This chapter aims to overcome these organisational knowledge gaps by providing a comprehensive overview of the machine learning technologies and practices currently used in automotive manufacturing environments. The barriers to developing and implementing machine learning and its critical enabling technologies are also discussed. With human-centered design and sustainability being one of the main pillars of Industry 4.0, particular focus is placed on understanding the social implications of machine learning and its enabling technologies and managerial practices.

## **2.3 Search method and strategy**

This section introduces the research question and describes the methods used to search the literature. The search methodology used in this chapter is based upon the PRISMA

framework, a commonly used systematic literature review approach widely used in the medical literature that has recently been adopted by software engineering researchers also [66].

## Research Questions

The research question for this literature review is defined as follows:

*”How are machine learning technologies and practices used to create value in automotive manufacturing environments, and what are the barriers to their development and implementation?”*

## Inclusion and Exclusion Criteria

This literature limits searches to peer-reviewed, published works. The searches are limited to publications from 2011 onwards, as this was when Industry 4.0 was first presented in the literature. Due to the scope defined by the industrial partners, research related to the development of autonomous robotics is excluded from this review. The automotive industry relies heavily on automated manufacturing robots and already has access to considerable talent in developing and implementing machine learning solutions for these systems. Given that this project’s scope is focused on opportunities for increasing the usage of existing data sources, research on augmented / virtual reality, additive manufacturing, supply chain, and cyber security are excluded. After the initial screening, papers that did not include sufficient detail on the methodology were also excluded. The complete exclusion criteria are as follows :

- Papers published before 2011.
- Papers that are not related to manufacturing production.
- Papers for which the primary topic is augmented / virtual reality, additive manufacturing, robotics, supply chain, or cyber-security.
- Papers that focused specifically on Small or Medium Enterprises.
- Papers not in English.
- Papers not peer-reviewed.

- Papers with no methodology included

While this review excludes papers not related to manufacturing, the author includes papers in the searches that include relevant developments in data analytics or machine learning that could be applied in manufacturing settings, even if this is not the main topic of the presented use case. These additional sources are mainly from the reference list of research papers included in the review. For transparency, topics that include additional sources outside the primary literature searches are highlighted to the reader.

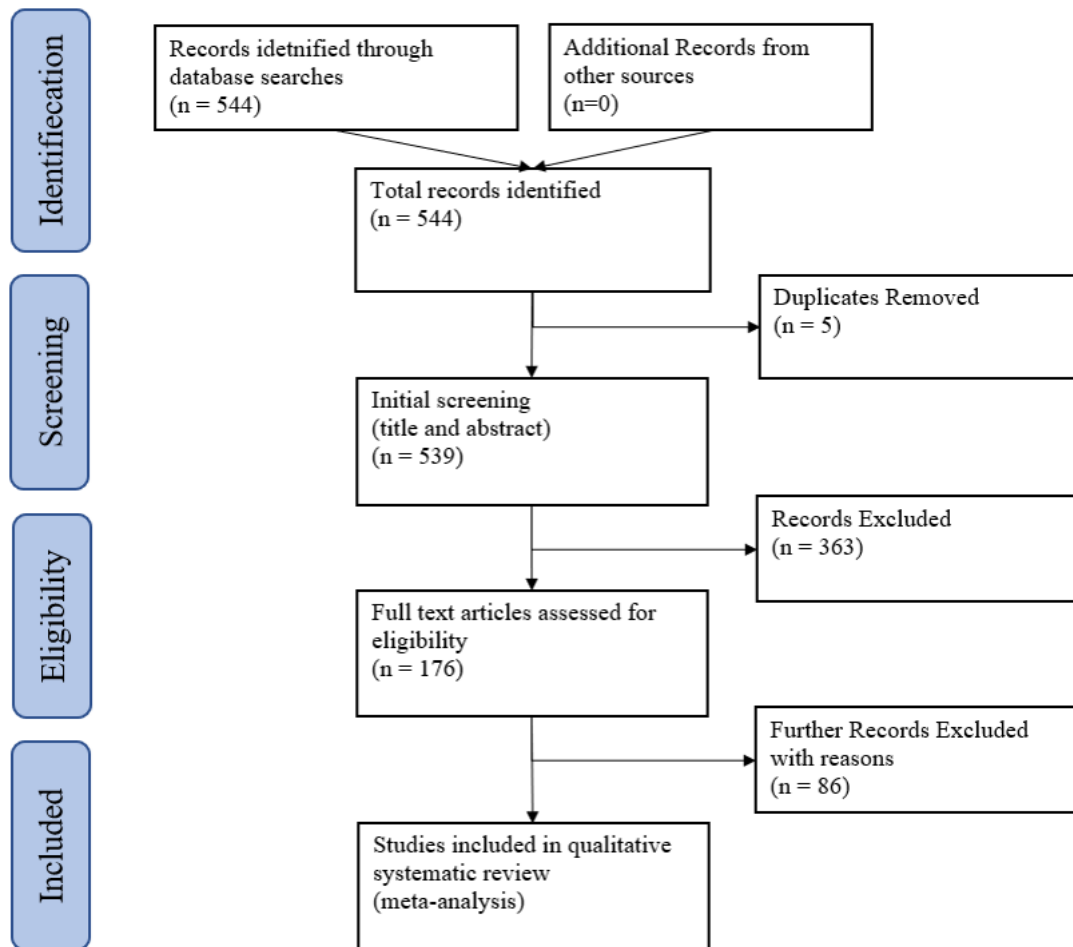


Figure 2.1: A flow diagram outlining the main steps of the literature searches.

### Searching the literature and the Screening Process

The literature search uses three academic databases: Semantic Scholar, Science Direct, and Scopus. The following search string is used to obtain results from these databases:

*"("Industry 4.0" OR "Industry 4" OR "fourth industrial revolution" OR "4th industrial revolution") AND (Automotive) AND (Manufacturing OR Production)) AND ("Artificial Intelligence" OR "machine learning" OR "Deep Learning")"*

The initial search returned 401 records from Science Direct, 110 from Semantic Scholar, and 33 from Scopus, giving a total of 544 records to be included in the screening process. After removing duplicates, the title and abstracts of 539 papers were read to determine their suitability according to our inclusion and exclusion criteria. Three hundred sixty-three papers were excluded during the initial screening. The remaining 176 papers were read entirely, and an additional 86 papers were excluded, leaving 90 records to be included in the final review. To gather some quantitative data for comparisons and visualisations, an Excel spreadsheet was used to extract data on several fields, including critical topics, types of data used, data collection methods, machine learning models compared, hyperparameter tuning methods, evaluation metrics, experiment results, and industrial impact. In addition to the literature search results, this study also includes some additional sources to provide further context into specific topics where necessary, i.e., citing the original paper on network architecture.

## **2.4 Search results**

This section includes a summary of the quantitative results gathered from the data extraction sheet, including: key topics, types of data used, and data collection methods. Details on common machine learning models, hyperparameter tuning methods, and evaluation metrics are discussed in section 2.5. By presenting these results and describing the machine learning technologies and practices presented in the literature and the primary data sources, this chapter aims to provide a foundation of knowledge for the thesis to refer to later in answering the research question. While the data extraction sheet also includes additional fields such as experimental results, industrial impact, economic appraisal, and methodology critique, these fields are not reviewed in this section as they are discussed in detail in the discussion in section 2.5.

## 2.4.1 Data Extraction Fields

### Key Topics

Figure 2.2 shows the publication year of each of the papers returned by the searches, as well as the type of publication, i.e., conference paper or journal article.

The key topics discussed in the reviewed literature include: quality assurance, computer vision, management, automation strategies, supervised learning, Big Data, predictive maintenance (PdM), and Flexible Manufacturing Systems (FMS). A chord diagram in Figure 2.3 shows how these key topics and others relate. This diagram highlights that one of the most common research topics in the reviewed literature is using computer vision systems for anomaly detection/fault detection as part of a quality assurance process in manufacturing. Management approaches to Industry 4.0 are also widely discussed, with social sustainability and workforce digital skills being the most common topic in the Industry 4.0 management literature [67, 39, 48, 68, 69, 47]. Digitisation change management in the automotive manufacturing sector is also widely discussed [70, 71, 72, 73].

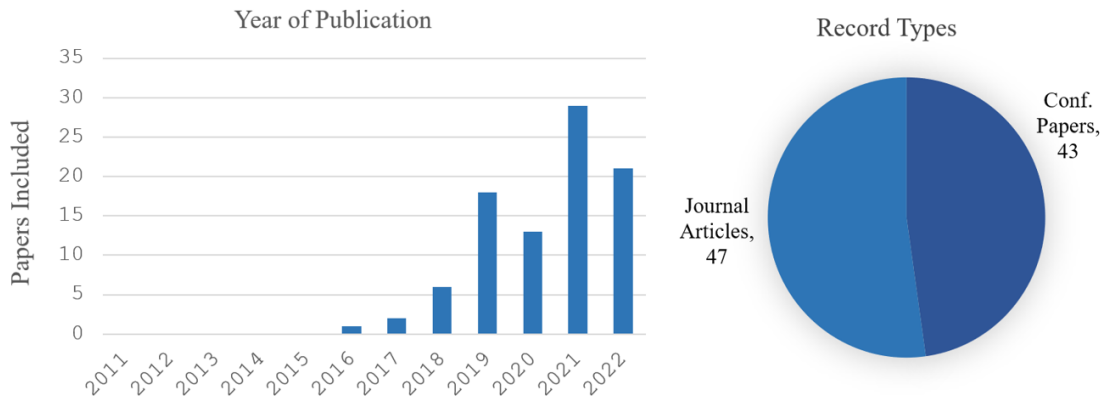


Figure 2.2: Graphs showing the year of publication of the 90 papers returned by the searches (left) and a visualisation of the ratio of journal articles to conference papers (right).

### Types of Data and Data Collection

Of the 46 papers that included machine learning experiments, 48% used image data as inputs, while the remaining 52% of papers used time series such as torque, vibration, or

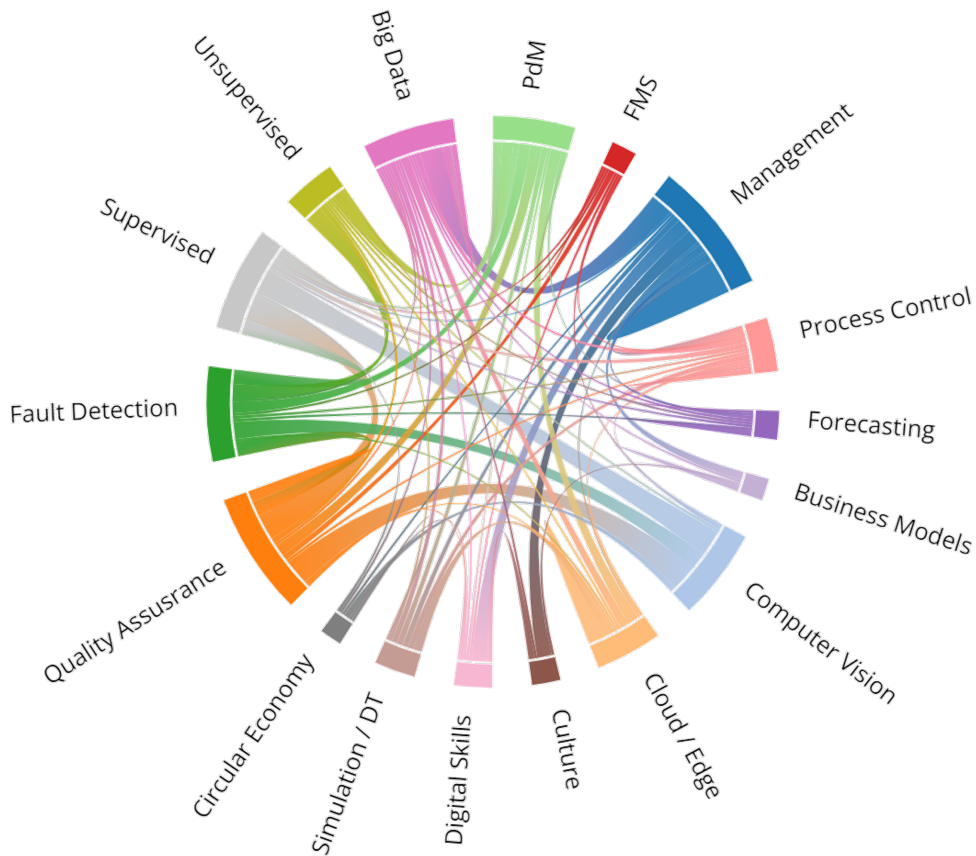
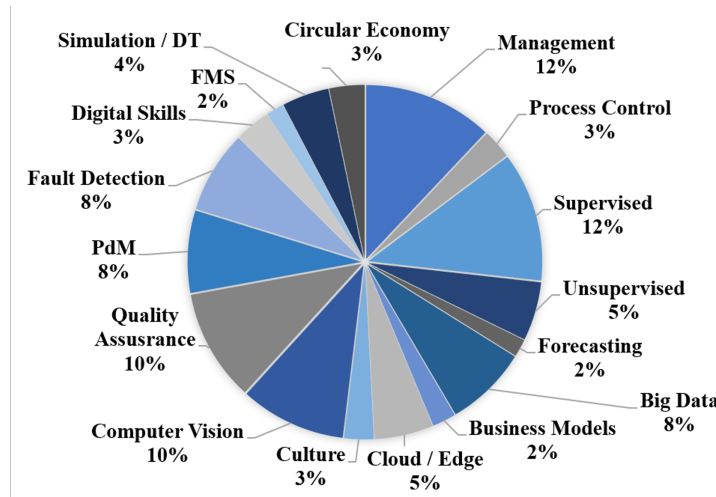


Figure 2.3: A pie chart (Top) shows the most common topics discussed in the reviewed literature. Note that one paper may include multiple key topics. A chord diagram (Bottom) shows the relationship between these key topics. The bar width between two topics is proportional to the number of times those topics appear in the same research paper. For example, papers discussing computer vision also often discuss quality assurance but rarely discuss unsupervised learning, as indicated by the wide and thin bars.

audio. Most time series data are uni-variate, with only a few papers exploring multivariate time series analysis [74, 75, 76, 77, 29]. The most common methods to evaluate time series data include SVM [78, 79, 80, 81, 82], k-NN[78, 80, 81, 76], RF [78, 75, 79, 81, 83, 76], ANN[78, 79], Auto-Encoders [74, 75, 84]. Given the significant research on computer vision for industrial quality inspection, some papers used techniques to transform time series data into 2D representations to exploit the use of CNN to classify the resultant images [81, 82, 84]. The data collection methods vary significantly depending on the type of data and machine learning models used. In most machine learning experiments in the reviewed literature, data are gathered from real-world production settings to demonstrate some proof of concept in off-line experiments. Many researchers highlight the challenge of collecting or accessing sufficient labelled anomaly data in industrial settings. In many cases, additional data are generated synthetically to increase the diversity of the minority class, and address data imbalance [85, 86, 87, 84, 88, 89] while others rely on publicly available data [6]. Only one instance was identified where researchers rely entirely on synthetic data to create both training and testing data [90]. Proof of concept experiments that require off-line development, such as data-driven production management systems or scheduling systems, often use historical data to train, and test models to avoid impacting production [29, 74]. Alternatively, digital twins can be used to simulate production data for these use cases [64, 91, 92].

## 2.5 Discussion

In this section, various enabling technologies of machine learning in manufacturing are discussed, and a critical review of the literature is presented. Fields from the data extraction sheet referred to in this discussion include experimental results, industrial impact, economic appraisal, methodology critique, and further details on other extraction fields discussed in the previous section. Throughout the review, reference is made to the research question, and the barriers and opportunities to machine learning in automotive manufacturing are discussed.

### Machine Learning Approaches in Manufacturing

Table 2.1 lists all papers in the reviewed literature for which machine learning experiments are conducted to provide an overview of how different machine learning methods are ap-



plied across various manufacturing use cases. This table also summarises any critiques of these research papers which are expanded upon further in the discussion in section 2.5.

Table 2.1: A list of all papers in the reviewed literature for which machine learning experiments are conducted, as well as an overview of the research topic and any critiques identified by the author.

Ref <sup>o</sup>	Description	Critique	Manufacturing Use Case	Data Analysis Methods
[93]	Real time analysis of RFID production data using a combination of unsupervised clustering and gradient decent.	Little information is given about the origins of the datasets, or the targeted end use of the proposed unsupervised clustering workflow. It's not clear how manufacturers might apply this technology to add value.	Big Data, Unsupervised	K-means, gradient decent
[94]	Aim to develop a zero defect manufacturing solution to identify defects in rubber parts.	Very small sample size for supervised training and testing datasets.	Computer Vision, Quality Assurance	ANN, Physics based methods
[85]	GAN used to create synthetic images of discontinuity and blob defects in an adhesive application process and improve the performance of a computer vision system for quality assurance.	Evaluation metrics not discussed.	Computer Vision, Fault Detection, Supervised	YOLOv4-Tiny, GAN
[86]	This paper expands on the same use case as [85], showing how including GAN generated images in a training dataset can increase performance.	This article goes into further detail on training approach, although details still missing on the testing and validation datasets used.	Computer Vision, Fault Detection, Supervised	YOLOv4-Tiny, GAN
[95]	Proof of concept using computer vision techniques and point cloud detection to develop an unsupervised approach for image based quality assurance.	Not compared with supervised approaches it aims to replace. Not clear that the proposed approach achieves better accuracy, or saves time.	Computer Vision, Quality Assurance	Domain-Adversarial Neural Network (DANN).
[96]	Vision based classification models to determine suitability of damaged automotive components for remanufacturing.	Limited information on data labelling methodology, such as time taken and required personnel.	Computer Vision, Quality Assurance, Circular Economy	CNN (AlexNet, CaffeNet), GMM, key point-based CNN
[90]	Aims to categorise the suitability of automotive frames for remanufacturing. Training and testing dataset was created using CAD data.	Tested using synthetic data and needs to be validated in real-world scenarios. Not clear how well the accuracies found in this study will translate to real world. Some similarity with [95] that should be considered in future research.	Computer Vision, Quality Assurance, Circular Economy	
[89]	Evaluates three object detection approaches and their performance in detecting workpieces and ability to recognise failure situations.	Clear detailed methodology on training / validation / testing / optimisation approach. Only two types of failures are simulated. Not clear how this will perform in real world scenarios, particularly where unknown errors might occur. No consideration of the economic benefits / cost benefit analysis of this project. No consideration of the required skills to implement this solution nor time taken to gather and label data.	Computer Vision, Quality Assurance, Fault Detection	YOLO, PANet, CSPResNeXt50,
[97]	Computer vision system to detect defects in automotive break components using a bespoke algorithm based on a sliding window approach.	The paper presents an accuracy of 92%, however, they do not account for the high dataset imbalance. F-score would have been a more appropriate evaluation metric. It's also unclear if the small sample used to demonstrate the method is representative of the anomaly actual distribution of the manufactured parts. No detail in the exact methods used to calculate the anomaly threshold value. Don't compare their bespoke algorithm with current state-of-the-art approaches i.e. CNN.	Computer Vision, Quality Assurance, Fault Detection	Bespoke Defect Detection Algorithm, Sliding Window
[98]	Supervised learning approach to inspect quality of soldering of electronic components. Augmented or simulated data used to overcome data imbalance.	Good example of how synthetic data can be used to improve performance.	Computer Vision, Quality Assurance, Fault Detection	CNN
[87]	Semantic segmentation is used to identify defects in images of a cylinder head manufacturing.	Very small sample of defective parts for cylinder head inspection (10 images). Defects were created synthetically using scrap parts. The paper seems to use GAN to generate a further 16000 images for training and testing. However, the exact method of testing and evaluation is unclear, with no results included in the final F-score of the proposed method.	Computer Vision, Quality Assurance, FMS, Process Control, Supervised	VGG, GAN
[99]	Deep learning approach to identify manufacturing defects in gears based on photos taken at multiple angles.	Good details of how datasets were built and good overview of the industrial impact of the solution. However, only use a single model to detect one type of defect.	Computer Vision, Quality Assurance, Supervised	Faster R-CNN

[100]	supervised computer vision system to ensure quality of informative labels on automotive parts in final assembly.	Good example of a human centric approach to ML implementation where the was to assist operators, not replace them. However, lacks technical information. Methods for labelling data not discussed. No details on model optimization.	Computer Vision, Quality Assurance, Supervised	YOLOv5, YOLOX
[28]	An iPad app is developed to document images of vehicles at the end of the line into a total of 100 categories. The author states that the main contribution of this paper is the creation of an automotive dataset, that allows users to learn and automatically recognize different vehicle properties.	The purpose of the visual inspection is not discussed with no end goal. No details on how training data was obtained / no methodology for the training data. Given that one of the main contributions of this paper was the dataset, without a clear methodology describing how it was built the quality of this data is in question.	Computer Vision, Supervised	Alexnet, Googlenet, Inception
[101]	The paper aimed to demonstrate that hand-crafted features in DL lead to improvements in performance when analysing fracture surfaces for root cause analysis of industrial processes i.e. tools, parts.	The benchmark selected for this study used databases have small samples with only 81 and 108 images per class. These datasets are too small to be reliable. The authors note the lack of large datasets for texture and fracture use cases. The results for the second experiment are more reliable given the increased dataset size and rigorous quality assurance processes. However, only explored one CNN (VGG).	Computer Vision, Supervised	VGG
[102]	Quality assurance of bodywork painting using spectrometers.	Lack of methodology throughout. Not clear what evaluation metrics used. Little information provided on how training and testing data was gathered and validated.	Quality Assurance	PCA, ANN
[81]	Quality assurance of wiring harness connector assembly in automotive manufacturing using ultrasonic microphone.	While the method was demonstrated to work with high degree of success in lab settings, when tested in real world scenarios there were no instances of failures to validate these findings. Further work required with a larger testing sample that includes anomalies to fully demonstrate the accuracy of this method.	Quality Assurance, Fault Detection, Supervised,	CNN, MLP, KNN, SVM, RF
[38]	Statistical process control and process monitoring in flexible machining processes to increase part quality and reduce scrap.	While no comparison with state-of-the-art ML techniques, this is one of the few papers that successfully applies the proposed solution in a real world setting.	Quality Assurance, Digital Skills, PdM, Quality Assurance, Supervised	Mean, Stdv, Median Absolute Deviation, Threshold Limits
[75]	Supervised classification task trained on eight separate torque testing processes to identify in process faults on highly imbalanced data.	Good overview of dealing with imbalanced data, however, no information on how the anomaly threshold limits are determined. This would make these approaches difficult to replicate. No economic appraisal.	Quality Assurance, Supervised	RF, AutoML, AE
[103]	Anomaly detection in discrete manufacturing using self-learning.	The paper presents an unsupervised method but provides little detail on how the test datasets were developed. Test datasets were not validated by domain experts.	Fault Detection	LSTM Auto-encoder with
[78]	Supervised fault detection on time series data in elective drive systems. Data is collected from a real world system for which various faulty components can be installed.	No cost benefit analysis. Not clear how this system will perform at identifying unforeseen failures. Some detail of engineering requirements to develop this solution beyond that considered by most other works, although further detail on specific requirements would be useful i.e. time, human resources, capital, software, data preparation, etc.	Fault Detection	SVM, K means, KNN, SVM, RF, XGBoost, ANN
[104]	A novel approach to analyse vibration data from a doser pump system to develop a predictive maintenance solution.	No consideration of the economic impact of this solution. Only explored 1 method: GMM unsupervised clustering. No testing data is used to assess the accuracy of the model and it's effectiveness to identify process failures is not discussed.	Fault Detection, Circular Economy, Management, Big Data, Fault Detection, Unsupervised Digital Skills, Supervised	k-means
[84]	Anomaly detection on refiguration units using low cost directional microphone arrays.	Unclear how the anomaly data was generated for the testing dataset.	Fault Detection, Unsupervised Digital Skills, Supervised	AE, VAE, GAN
[88]	Process monitoring of installing connectors in automotive assembly. Uses data collected from gloves with embedded sensors.	Experiments conducted in lab environment based on a rough estimate of the failure distribution, as actual distribution not known. No methods of how models are optimised. No details on time taken to collect and label data.	Culture, Supervised	CNN
[79]	16 data warehouses from various companies are used to train machine learning models to demonstrate data driven decisions related to Stock-Keeping Units profiling, inventory profiling, workload profiling, layout profiling.	Good comparison of big data analytics models for vertical and horizontal integration of manufacturing data sources.	Management, Big Data	ANN, Decision Tree, SVM, RF, regression, perceptron, Adaboost, gradient boosting, baggin Tree, Naïve Bayes, DAE, DBN,
[29]	Propose a method to predict production progress for IoT factory environments and support in identifying when production plans are executed incorrectly.	While the data set is very large, this method has only been tested in 1 real world scenario. Little detail was included on the manufacturing site in which this was implemented. It would be interesting to compare this across multiple locations, particularly those with FMS. Without a good understanding of the manufacturing process and the flexibility of production and how these data are collected, it's makes this method difficult to replicate.	Management, Process control, Supervised, Forecasting, Big Data	DNN, PCA
[105]	A 1D CNN used to predict the quality of an assembly screwing process.	Limited information provided on how the data was labelled / gathered and how he quality of this data was validated.	PdM, Quality Assurance, Supervised,	1D CNN

[106]	RUL prediction on conveyor chains at an engine assembly line for maintenance strategy identification and maintenance planning.	No information included on how failure data are collected, how many types of faults are analysed, and how these data are used to set the failure threshold limits. This lack of information would make the solution difficult to replicate in practice.	PdM, RUL	LSTM, ARIMA, Holt-Winter
[80]	Quality diagnostics approach for milling composite honeycomb cores. Two fault modes are considered: tears and uncut fibres.	No detail given on the amount of time taken to label data or other challenges of data acquisition. Training and testing data sizes not give, only percentages. Not clear on the imbalance of the testing datasets.	PdM, Supervised	KNN, DT, SVM
[74]	Proposes a conditional variational autoencoder to determine machine degradation. Unsupervised method is applied to two datasets: NASA Turbofan, and FMS production data.	This method requires unsupervised training data with a low number of failure events and therefore relies on the failure rates and distributions of historic unlabelled data being well understood. This presents a challenge in industrial use cases where long term failure data are often difficult to access.	PdM, Unsupervised	AE (CVAE)
[77]	Unsupervised clustering approach for PdM on a waterjet cutting process.	Good study with a range of methods compared and evaluated using visual inspection. Industrial impact of solution not discussed.	PdM, Unsupervised, Fault Detection	PCA, MLHL, CMLHL, CMDS, SM, FA, k-means, agglomerative clustering
[91]	Reinforcement learning for statistical process control for a manufacturing production line.	While this is a good overview of the technical barriers, the results were tested and validated using synthetic data. Further work required to understand the additional strategic, organisational and cultural barriers to implement this. Initial iterations of complex machine learning models like this have the possibility to significantly impact production in early stages of implementation. Careful economic appraisal of the long term ROI of needs to be better understood.	Process Control	RL, transfer learning
[107]	Authors work closely with domain experts to develop an unsupervised, hierarchical clustering approach to detect throughput bottlenecks on a automotive CNC machining line.	A good example of an unsupervised application using real world data with good description of the challenges this data presents.	Simulation / Digital Twin, Unsupervised Clustering	Hierarchical Unsupervised
[83]	Data on wiring harness manufacturing processes is used to predict key performance indicators using multiple ML models to allow managers to proactively address production issues.	The problem definition is also not well defined as it's not clear how the proposed prediction model can support root cause analysis. While the proposed ML methods do accurately predict OEE, it's to clear how this method improves on traditional methods of measuring cycle time distributions of specific processes and setting alerts if these processes exceed some limit. Given the fact that this data is normally distributed, it seems more simplistic statistical techniques would be easier to implement, even basic Six Sigma DMAIC technotes.	Supervised	SVR, RF, XGB, DL
[76]	A test bench of a permanent magnet synchronous motor was used to collect 139 hours of multivariate time series data across a range of operating parameters. A range of machine learning models are trained to predict the temperature of the magnets inside motors and compare results with predictions made using thermodynamic calculations. ML models gave similar results.	The author notes that this was a limited study on 1 type of motor. Further work required to explore the transferability of this approach.	Supervised	KNN, RF, SVR, ET, LPTN (thermal model), RNN, MLP, OLS, CNN

The reviewed literature discusses five main types of machine learning: Supervised, Unsupervised, Semi-supervised, Reinforcement learning, and Transfer learning. Only two papers in the reviewed literature include details on these learning approaches; even in these cases, information is limited. For further information on these machine learning settings, the reader is referred to section .1.1 of the appendix where these topics are described in detail.

In supervised learning settings, the most common machine learning approaches in the reviewed literature is the Convolutional Neural Network (CNN) due to its wide use in computer vision for quality inspection [94, 96, 89, 97, 108, 98, 100, 85, 109]. Other popu-

lar machine learning algorithms include Support-Vector Machine (SVM), Random Forest (RF), Artificial Neural Networks (ANN), Auto-Encoders (AE).

The most common unsupervised learning approaches in the reviewed literature apply dimensional reduction approaches such as PCA and t-SNE to transform higher dimensional data into 2D or 3D feature space, followed by cluster analysis techniques such as k-means to identify outliers in the reduced feature space [30, 64, 103]. Many researchers highlight that unsupervised methods lend themselves to anomaly detection in as they enable a system's normal behavior to be learned, and any deviations from the norm can be assumed to be anomalies [30, 63, 75]. However, in real-world settings, some labelled data are still required to validate and test the accuracy of any trained model against some known ground truth [30].

Reinforcement learning has been a key technology in the automotive industry over the past few decades as it has been an enabling factor in autonomous driving; however, applications of reinforcement learning in manufacturing and production environments are limited [30]. Only four papers are identified in the reviewed literature exploring reinforcement learning solutions in manufacturing [83, 91, 30, 110]. Mathematical details of the common machine learning models outlined above can be found in section .1.1 of the appendix, as well as details on the evaluation metrics and optimisation approaches presented in the reviewed literature. Reinforcement learning approaches require commitment by many parties to develop complex and robust reward functions relevant to the target domain [91]. Viharos et al. discuss the complexity of setting up these models, which requires an in-depth understanding of failure distributions of relevant processes and careful tuning of reward functions in initial model setup [91]. Despite these challenges, Theissler 2021 highlights reinforcement learning as a promising opportunity in PdM applications due to the minimal labelled data requirements compared to supervised and unsupervised approaches [30]. Gankin et al. and Viharos et al. demonstrate how reinforcement learning can be used to optimise production scheduling digital twin environments [91, 110]. However, no research in the reviewed literature discusses the use of reinforcement learning in real-world production settings. The author suggests future research explore frameworks to implement these systems in real-world production environments. Further research is also required to understand the additional strategic and cultural barriers to the users' trust in the system's ability and how best to define reward functions.

Regression is a supervised setting which aims to understand the relationships between dependent and independent variables of continuous data to predict future observations based

on these trends. Examples of regression in the reviewed literature include prediction of jobs remaining [29], predicting due-date of wafer fabrication [29], Remaining Useful Life (RUL) prediction of a cutting tool [64] and RUL of conveyor chains [106]. However, some of these papers lack sufficient description of the methods used to develop the training and testing datasets. For example, when using autoregressive models to predict RUL of conveyor chains, Einabadi 2022 include no information on how failure data are collected, how many types of faults are analysed, and how these data are used to set the failure threshold limits. This lack of information would make the solution difficult to replicate in practice. [106]

In their 2021 literature review, Viharos et al. find that autoregression, linear regression, and quasi-linear autoregressive models are the most popular approaches for statistical process control. However, this finding is based on only two cited examples [91]. The only significant example of regression used in manufacturing settings in the reviewed literature is by Hassani et al., in which the Support Vector Regression model was used to predict Overall Equipment Effectiveness (OEE) [83]. However, in this paper, the problem definition is not well defined, and it is unclear how the proposed prediction model can support root cause analysis and identify the actual problem [29]. While the proposed ML methods accurately predict OEE, it is unclear how this method improves on traditional methods of measuring cycle time distributions and setting alerts if these processes exceed some limit. Furthermore, there is a lack of detail on the data acquisition methods, the distributions of the quality and performance data, and how the extracted features of the data were used to make the final predictions.

Research by Theissler et al. finds a large number of research on regression in a range of PdM settings [30]. The notable difference between the two papers is that the vast majority of research into predictive maintenance relies on historical data to train models rather than synthetic data generated from Digital Twins [30]. van Dinter et al. highlight that Digital Twins are costly to develop and require expert domain knowledge to extract the required data. The author also argues that synthetic data's trustability will likely impact its use in real-world PdM cases.

Limited applications of regression models were identified within the scope of this review and therefore are not discussed at further length. However, regression is a valuable tool for manufacturing applications for which there is considerable research. The author refers readers to the following papers where various examples of regression are presented [111, 112, 113]. Further research should explore the use of regression models in the automotive industry.

In the reviewed literature, Generative Adversarial Networks (GAN) are often used to overcome a lack of high-quality labelled data by generating new synthetic data to expand existing datasets [85, 86, 87, 84, 90, 88]. This is particularly true for fault detection/anomaly detection problems where datasets are highly imbalanced due to limited availability of labelled failure instances [85, 86, 87, 84]. Peres et al. present a good use case where GAN are used to create synthetic images of discontinuity and blob defects in an adhesive application process [86, 85]. This data was then used to train a YOLO network to achieve accuracies of up to 93% for the two fault types [86, 85]. Including an additional 4000 synthetic image data on top of the 88 real training dataset improved the performance of the anomaly detection system significantly when compared to using only synthetic data or only real data [86, 85]. The research shows that GAN should not be used to create entire training and testing datasets, but rather as a means of enriching and expanding existing high-quality labelled datasets by generating additional images to increase the diversity of the training data further [86, 85].

Mazzetti et al. discuss how in addition to generating synthetic data, GANs can also be used as a means of anomaly detection for image data [87]. The proposed method first involves training a GAN network using images of non-defective parts to generate additional images of parts with no surface defects [87]. A new unlabelled image is then given as the real sample for the pre-trained network to compare with the generated normal sample. Calculating the difference between the images by measuring the residual loss in the image space allows users to create a heat map using the residual image to identify specific regions in the new image of possible anomalous regions. While Mazzetti et al. present a good theoretical description of this method, its application in the two case studies lacks key information on the methods used to label the normal training data [87]. Furthermore, no anomalous test data was included to evaluate the model's performance, and no quantitative results are included in the paper.

Hatanaka et al. present a similar GAN-based anomaly detection used to identify potential faults in refrigeration systems by training only on normal audio data transformed into a 2D spectrogram [84]. While this paper does include quantitative results to compare anomaly detection performance with other image-based anomaly detection methods, this paper lacks detail on how the testing data was gathered or produced [84]. The paper states in the introduction that no failure records are available for the refrigeration system and yet includes 275 "abnormal" images in the testing data without any detail on how this anomalous testing data was gathered or generated [84]. Based on these works, future research

directions should consider testing GAN in anomaly detection settings using high-quality training and testing data representative of real-world industrial data.

Another notable finding in GAN research is the contrasting nomenclature used by researchers. Both Mazzetti et al. and Hatanaka et al. present GANs as unsupervised approaches despite relying on large amounts of high-quality training data, with up to 16000 normal images used to train Mazzetti et al.'s discriminator [87, 84]. In contrast, Peres et al present GAN as a semi-supervised approach due to these training requirements [86]. The original paper in which GAN networks are first proposed does not address which category the network falls under [114]. However, the paper highlights its suitability for semi-supervised settings in suggestions for future works [114]. Given the growing popularity of these training approaches to overcome data imbalance in developing quality assurance systems, the author argues that establishing a clear definition for this is not a trivial matter to ensure training and testing best practices are well defined for each of the various approaches. The author finds that in the context of anomaly detection described above, GANs fall under the description of a semi-supervised, or clean-semi-supervised approach as outlined in section .1.1. Future research in this field should consider best practices on implementing semi-supervised solutions outlined in [115].

### **2.5.1 Enabling Technologies of Machine Learning in Manufacturing**

The reviewed literature discusses various applications of machine learning solutions, including Visual Inspection, social media analytics, autonomous driving, advanced robotics, process control, predictive maintenance and scheduling, prescriptive analytics, anomaly detection, supply chain optimisation [27, 28]. Developing machine learning solutions such as these requires people with advanced data analytics skill sets to analyze large volumes of integrated data that are easily accessible through networked infrastructures. Due to the complexity of developing and implementing these solutions, this requires high levels of collaboration between technical experts and domain experts supported by effective digital communication channels [107, 116]. To provide these high levels of data collection, data integration, data accessibility, and other ICT infrastructures, various enabling technologies of Industry 4.0 are discussed in the literature, including IoT, Big Data analytics, Flexible Manufacturing Systems, Digital Twin, Cloud, and Edge computing. Many machine learning solutions mentioned above can be implemented at the functional level

without the widespread adoption of these technologies. However, as companies reach high maturity levels in these enabling technologies, the opportunities for value creation using machine learning increase significantly as new enterprise-level opportunities are created to deliver business intelligence [12]. In this subsection, the author refers to various additional references as few papers returned from the literature search included details on IoT infrastructures and challenges in manufacturing environments.

### **The Industrial Internet of Things**

The Internet of Things (IoT) describes the latest evolution of the conventional Internet, where World Wide Web (Web) technologies are becoming integrated into everyday appliances, devices, and services. The Web and the internet are two different technologies. The Internet is a global computer network infrastructure with a wide range of information resources and services. The Web is one of these services and currently the main internet platform, providing browsing capability through technologies such as Hypertext Markup Language (HTML), Uniform Resource Locator (URL), and Hypertext Transfer Protocol (HTTP) [62]. The IoT expands the capabilities of Web-based protocols by enabling internet presence for a wide range of devices, appliances, places, people, and other 'things' [62]. This bridge between the physical world and the Web is referred to as the Physical Web, which represents a new way that consumers and businesses can use the internet to access, monitor, and control the world around them and the connected things within it [62]. IoT products enable real-time data collection, which companies can use to understand consumer behavior better and offer added value through data-driven services that are fully integrated with the customer [12].

IoT is one of the key elements driving digital transformation around the world. While IoT technologies can be found almost everywhere in modern society, they generally fall under two main categories: The Internet of Things (IoT), and The Industrial Internet of Things (IIoT). IIoT is the adoption of IoT within industrial and manufacturing applications with an emphasis on improving the connectivity and integration of machines and real-time data collection and utilisation. The data collected through the end-to-end integration of IIoT throughout the manufacturing process and in-service use is a key enabler of machine learning and other data driven manufacturing solutions such as business intelligence, Big Data analytics, cloud computing, edge computing, and flexible production systems [34, 117]. The main barrier to developing and integrating IoT technologies is the limited research



on safety and security, and the associated challenges of confidentiality and data protection [118, 119]. From a cybersecurity perspective, the IoT architecture consists of 4 layers: the sensing layer, network layer, middle-ware layer and the application layer [120]. While each layer has its own vulnerabilities and associated attacks, most of the discussion in the available literature surrounds those associated with the network layer. IoT-based applications are particularly vulnerable at the network layer as communication is often wireless. The network layer uses technologies such as Wi-Fi, Bluetooth, 4G to allow data transmission and routing between IoT devices over Internet and mobile networks [120]. Radio Frequency Identification tags (RFID) are another common networking technology used in manufacturing that provides a read-write capability to record and track the birth history of a product through the assembly process. These open network gateways are susceptible to a range of both passive and active attacks, including man-in-the-middle and denial-of-service attacks [119, 120, 121]. The network layer is especially vulnerable to man-in-the-middle attacks, which describe a range of ways a poorly secured network transmission can be intercepted and decrypted [120]. Eavesdropping is a man-in-the-middle attack, otherwise known as spoofing, that takes advantage of unsecured networks, allowing data transmission to be intercepted and stolen [121]. Advanced protocols, software, and hardware can be implemented in devices to detect network security threats and keep devices secure, such as HTTPS or Secure Shell. However, the more devices integrated within the same network increases the risk of attack [120]. For further details on these vulnerabilities and attack vectors, the author refers the reader to Lu 2018 and Roberts 2006 [120, 121].

It is agreed among the reviewed literature that there needs to be more focus by both industry and academia on the issues surrounding cyber security in IoT systems. In manufacturing, this requires increased collaboration between IT, manufacturing production, and innovation teams to ensure IoT manufacturing solutions consider the legal, security, and network infrastructure requirements during the early stages of pilot projects involving IIoT.

## **Big Data Analytics**

The term 'Big Data' originated from Silicon Graphics Inc. in the mid-1990s, although it was only at the turn of the 21st century that significant academic references began to emerge [122]. Distributed Smart Systems and the Internet of Things were becoming increasingly widespread, and the digitisation of systems, services, and processes resulted in vast amounts of data being generated daily. Existing computing and service modes were

struggling to meet the high demands for network bandwidth, response speed, and data storage, driving the development of cloud computing and data center networking, key enabling technologies of Big Data analytics [123].

In an increasingly digital world, data generation continues to be produced from various sources at unprecedented rates. Every website, digital system, process, sensor, and connected device produces data that, if analysed correctly, can be used to generate tremendous value for businesses. However, these huge volumes of data are often presented in semi-structured or completely unstructured formats making traditional data analysis techniques inefficient. Over the past decade, powerful new technologies and advanced algorithms have been developed to deal with the complexity of modern Big Data sets. Following huge amounts of academic and business interest, it was not long before Big Data became its own field of study.

Big Data refers to a broad field of interdisciplinary study and has no single definition due to its widespread usage. Different industries use different data sets, each of which will require very different analysis methods, for example, tabular data versus image data. Applications range from social media analytics to government monitoring to medical diagnosis, each of which will have its perspective on Big Data [124]. The product-oriented, process-oriented, cognition-oriented, and social-oriented perspectives will consider different features of their data, as well as the respective methods of analysis [52]. Furthermore, in terms of capacity, what may be considered 'big' by one industry may not be by others, and as technology progresses, the same data may not meet the threshold for Big Data in the future. While definitions vary between industries, most data scientists and experts describe Big Data by referring to 'the 3 V's': volume, velocity, and variety [125, 52, 126].

Big Data sets are typically in the range of tens of terabytes to multiple petabytes [127, 128]. The volume of data continues to increase at rates faster than current tools can process [129]. A common challenge of machine learning applications in the automotive industry is the storage and processing of large volumes of data and the associated challenges of dealing with unstructured Big Data [28].

Velocity measures the speed at which data is generated, streamed, pre-processed, and processed [130]. Depending on the data velocity, there are different processing methods: batch, near-time, real-time, and stream [52]. Some Big Data applications have strict time requirements, such as social media or weather forecasting, and require real-time or stream processing. High throughput batch-oriented processing of Big Data sets is not suitable for online processing demands as it may take several hours or days to process [131]. However,

batch processing provides an efficient way to analyze higher volumes of data collected over a time period [132]. This makes batch processing superior for applications such as training predictive models. A hybrid approach is often used to benefit from real-time and batch processing advantages. Data velocity management is much more than just a bandwidth issue, it is also an ingest issue relying on Extract, Transform, Load (ETL) processes to gather, organise and centralize data [130].

Data is generated from various sources and formats, such as images, text, audio, sensor signals, graph, logs, and many more. Big Data sets are often made up of incompatible data formats that may be structured, semi-structured or unstructured, public or private, complete or incomplete, etc. Structured data includes spreadsheets, relational databases, and others with fixed fields. Semi-structured data is harder to define and loosely describes data that is neither raw data nor strictly laid out in fixed fields [133]. Tags or other markers give semi-structured data some order, capturing some elements of the data such as Extensible Mark-up Language (XML) or Hyper Text Mark-up Language (HTML) [134]. Unstructured data sets such as raw sensor data or interactions of consumers on social media have no fixed fields or format. From an analytic perspective, unstructured data in various formats with non-aligned data structures or inconsistent data semantics pose a significant challenge to effectively using large volumes of data [134, 130].

While the 3 V's are often described slightly differently in the context of different industries, they are widely used to define Big Data. Some scholars and businesses extend this definition to include dimensions such as variability, value, veracity, and complexity [129, 130, 135]. With the 3 V's already well established, these additions are usually discussed separately, used to highlight the importance of other characteristics and properties inherent in Big Data that are more difficult to quantify.

The variability of Big Data combines both velocity and variety and is used to describe the constantly changing structure, meaning, and/or flow rate of incoming data [129, 130]. Analysing and predicting these changes has become increasingly important with the increased use of social media, the internet, and connected devices [124, 130]. Given that these systems are highly anthropocentric, the variety and velocity of data vary significantly with human behavior. For example, certain events may cause a sudden increase or decrease in social media usage, similar to how energy usage across the national power grid varies throughout the day as people go about their activities. Variability adds further variety to Big Data, providing unexpected, valuable information hidden within the data.

Veracity in the context of Big Data refers to its quality, particularly its uncertainty and

unreliability. Before the Big Data revolution and Industry 4.0, data in the scientific and academic community was often assumed to be clean and precise, a view also found in traditional data warehouses [136]. In modern Big Data projects, cleaning up the data can make up to 80% of an analytics project [137]. Uncertainty is often inherent in unstructured Big Data sets due to incompleteness, ambiguity, variety, and latency and the approximations of the techniques used to analyze it [134, 52]. These factors all reduce the veracity of data which in turn can significantly reduce the predictability of the data. This aspect is critical for applications such as epidemiology, medicine, and healthcare [138, 139]. In terms of the results of the outcome of Big Data analytics, many scholars and data scientists agree that veracity is becoming, if not already, the most important of the many 'V's' presented in the literature [136, 137, 140, 141].

The value of Big Data lies in the ability to analyze and extract useful business information and varies greatly depending on the application. The potential of analytics to generate value for automotive manufactures lies in areas such as: customer behavior analytics, marketing spend management, global supply chain management, and predictive analytics.

Complexity is inherent in unstructured Big Data due to the variety and variability of data [52, 142]. Identifying, linking and transforming relevant data between different systems and sources poses a huge challenge for practitioners [125]. To deal with this complexity, advanced data analysis techniques, such as neural networks are often required [125]. As well as requiring high levels of skills in data science disciplines, the development and implementation of such techniques also requires high computational costs and complexity that can be difficult to integrate into existing systems.

Expanding the common 3 V's definition to include these essential features provides a much clearer insight into the challenges and opportunities presented by Big Data. While a formal definition is yet to be agreed upon, based on these characteristics that are by far the most widely discussed in the literature the following definition coined by De Mauro et al seems most appropriate: "Big Data is the Information asset characterised by such a High Volume, Velocity, and Variety to require specific Technology and Analytical Methods for its transformation into Value" [143].

Data usage and Big Data analytics feature heavily in the Industry 4.0 discussion and the available literature presents many opportunities for many industries and organisations to increase productivity, reduce waste and reduce costs. These benefits are often discussed regarding the new innovative business models that Big Data and other Industry 4.0 concepts are enabled. These business models include 'The Sharing Economy, 'On-Demand

Services' and many circular economy models presented in the ReSOLVE framework by The Ellen MacArthur Foundation, which has been gaining much attention recently years [52, 53, 54]. These models focus heavily on collaboration and identify Big Data as a major enabling technology. Several external uncertainties can have adverse effects on manufacturers, many of which are difficult to identify without appropriate predictive analytics and control strategies [57]. Examples include unpredictable raw material quantity, quality, availability variations, and constantly changing market trends and consumer behavior. Customer behavior analytics presents a significant opportunity for automotive manufacturers to improve their understanding of the potential value of different customer segments through analysing external data sources such as social media and internal sources [58]. This knowledge can be used to strategically target new customers as well as improve customer experience and ensure the loyalty of existing customers [58]. Big Data analytics of consumer data also plays a major role in improving manufacturing processes' agility and flexibility, allowing production and procurement to adapt to changing markets. By gathering and analysing real-time data from an automaker's fleet and warranty data, cloud-based Big Data processing will enable new after-market services for customers. Examples include predictive and preventative maintenance of products, various infotainment services, and self-driving support services such as autopilot and collision prevention that are already found in many modern vehicles [59].

Big Data is still an emerging technology requiring significant research and development. The literature presents a multitude of barriers to the full realisation of Big Data benefits across all applications. These barriers generally fall under two categories: internal technological constraints and cultural barriers. The Big Data life cycle consists of data generation, storage, and processing, each of which faces significant technological challenges [144, 145]. As the volume and velocity of Big Data continue to increase, high demands are placed on networks, and servers [135]. Cloud computing is often used to outsource data storage and processing to avoid new data management systems' capital and organisational expenditure [145]. Not only does this lead to numerous privacy and security issues but as data generation increases in the order of exabytes, current daily networks cannot handle data sets of this scale. A 2018 report by IBM revealed that while 80% of senior technology executives recognised the competitive advantage of IT infrastructures, less than 10% of organisations said that their existing infrastructures were able to meet the demands of Big Data, cloud computing, social media, and mobile technology [146].

To industries yet to fully embrace Big Data, these technological challenges may seem of

little concern at the early stage of development. However, this highlights the additional cultural barriers to the implementation of Big Data analytics discussed across the literature. These cultural barriers are widely considered a greater challenge than those of a technical nature. Organisational culture is defined as "The visible and less visible norms, values and behaviour that are shared by a group of employees which shape the group's sense of what is acceptable and valid." [147]. In many industries, adopting Big Data represents a significant change to current organisational culture requiring new managerial approaches and structures. As organisations rely increasingly on data to drive business decisions, it is crucial that this information is made available in an accurate, complete, and timely manner to ensure business intelligence and productivity improvements [135]. This will require companies to share data in a standardised manner with other businesses, and government agencies outside of their organisational boundaries [12]. This level of external collaboration is an important part of the Industry 4.0 strategy.

### **Flexible manufacturing Systems**

The concept of a high volume, high variability manufacturing production environment has existed since Industry 3.0. Recent developments in Industry 4.0 technologies such as IoT, Cloud computing, machine learning, and digital twin, have resulted in these systems becoming increasingly widespread in industrial settings [27]. During industry 4.0, flexible automation emerged as an extension of programmable automation where reprogramming is done off-line, resulting in no downtime during reconfiguration [87]. By combining the above technologies in manufacturing environments through lean management practices, a new type of assembly line emerged known as the Flexible Manufacturing Systems (FMS). In a typical FMS, workstations are arranged more freely on-site in a modular arrangement, with AGVs automatically routing products to the required workstation. Tracking technologies such as RFID enable individual processes to be controlled by an automated production system to select the appropriate production process for the incoming part. This highly automated data-driven process allows production to be tailored to meet fluctuating consumer demand of multiple product families and deliver highly customizable product variants. As well as reacting quickly to market changes, FMS can react quickly to changes in the local manufacturing environment, such as processing changes, material changes, or new product variants [87]. FMS results in less space, reduced operational headcount, improved inventory management, reduced lead times, and reduced manufacturing costs [27]. The concept of a fully integrated factory with ubiquitous integration of Industry 4.0 tech-

nologies throughout the entire business is widely referred to as a 'Smart Factory', although other terms such as 'Intelligent Factory', 'Intelligent manufacturing', 'Ubiquitous Factory' or 'Real-Time Factory' are also used. Research finds that organisations investing in Smart Factory projects report increases of up to 12% manufacturing production, factory utilisation, and labor productivity [27]. Despite the numerous benefits of FMS presented in the literature, there are several challenges in developing, implementing, and maintaining these systems. FMS requires high initial investment both in physical assets and personnel with the required automation and digitisation skills to deal with the high complexity of these systems [27].

Researchers present novel machine learning approaches to support automation in FMS in the reviewed literature. Huang et al. propose a system to support production managers through a machine learning-based system to predict production progress for IoT factory environments [29]. A two-layer transfer learning approach using a combination of Deep Auto encoders and Deep Belief Network (DAE-DBN-TL) is trained on historical data using a bootstrap sampling approach. The proposed method was tested using real-world historical data over 15 orders with 1118 features. The experiment finds that the DAE-DBN-TL method achieves high performance ( $R^2 > 87\%$ ) in predicting production progress based on historical production data. The author argues that as well as monitoring and analysing production progress, this model can also identify instances where production plans are executed incorrectly and support root cause analysis of these abnormalities. While the data set used in this study is very large, this method has only been tested in one real-world scenario. Little detail was included on the manufacturing site in which this was implemented. By comparing the DAE-DBN-TL across multiple locations and understanding its applicability in FMS environments, it is easier to determine the validity of these results [29]. Future research should further examine the human-centric implementation of production process prediction systems in FMS and consider the effects on the workplace experience of production managers.

### **Horizontal, Vertical and End-to-End Integration**

Horizontal and vertical integration are terms that already have a variety of meanings in IT, marketing, and business applications. In the machine learning and Industry 4.0 literature, end-to-end integration is widely discussed as a key requirement of flexible manufacturing environments [148, 149].

External collaboration is an important part of the industry 4.0 strategy and is often discussed in terms of the horizontal integration of a company's value networks. By cooperating with suppliers and other organisations by exchanging real-time information, data, and resources, companies can drive innovation and growth while mitigating risks [148]. This external collaboration also extends to customers through technologies such as IoT and other smart devices that offer real-time data transfer throughout the value network allowing companies to react quickly to changes in the market and customised production. Moreover, connected devices can provide a foundation on which new service-orientated business models can be developed to deliver added value to customers [148].

The integration between ICT systems within an organisation is called vertical integration. This considers information exchange throughout all functional layers of the business hierarchy from the asset layer, which includes the physical subsystems such as robots, machines, personnel, etc., all the way through to the business layer, where production management and corporate planning take place [92]. While vertical integration requires a high level of digitisation, the result is a fully networked manufacturing system that promotes intra-company collaboration and improves process efficiency and flexibility through real-time availability of process data throughout the hierarchy [92]. The third kind of integration focuses on the digital integration of engineering throughout the product life-cycle, from product design and development through the entire value chain to end-of-life considerations. By engineering smart products to gather and transmit real-time data throughout manufacturing production and in-service use, individual products in some sectors will have the ability to automatically control the stages of their life cycle [134]. The Cloud-based design manufacturing model is a good example of end-to-end integration applicable to the automotive industry [134]. Cloud-based design manufacturing is a service-orientated product development model that uses in-service vehicle data to provide additional services for customers and gives engineers access to real-time data on individual product performance. By integrating this data with engineering teams, Big Data analytics can be used to predict faults and automatically schedule maintenance based on equipment conditions. This increases customer value in a horizontally and vertically integrated factory environment. This information can help improve the smart factory's design, logistics, production, and management processes.



## Digital Twin

Digital Twins are a key enabling technology of Industry 4.0 and are presented as one of the main tools to manage the complexity of FMS [63]. Digital twin is an evolution of traditional simulation methods that encompass increased data availability, ubiquitous connectivity, and new user requirements that enable manufacturers to model complex behaviors of production processes based on real-time and historical data. Digital twin plays an important role in planning when new lines are being designed and developed, using historical process data from previous systems to simulate and optimize the proposed layout. These accurate digital representations of the manufacturing environment can be used to analyze real-time behaviors of the system and can be reconfigured off-line to optimize throughput and address challenges that may not have been modeled in advance [63].

Vinter et al. provide a rigorous systematic literature review on Predictive maintenance using digital twins [64]. Researchers found that the most common application of these combined technologies was in the manufacturing sector on systems such as CNC machines and industrial robotics, where parts experience rapid wear due to components such as rolling bearings and gearboxes [64]. The main machine learning approaches used for PdM in digital twin environments are SVM, regression, decision tree, RF, K-means, and PCA [64]. The researchers highlight that the complexity and high computational requirements of both digital twin and machine learning models present a challenge, impacting the cost-effectiveness, delivery time, and energy requirements of these methods [64]. Another major challenge of applying machine learning in digital twin environments is the requirement for high quality and a wide variety of data [64]. Digital twins can be used to generate synthetic data to train machine learning models for many use cases, however, they require high-quality historical data and supporting data. For example, in a predictive maintenance setting, data is required on healthy, semi-healthy, and faulty machine performance and data of the failure distribution [64].

Although the initial set-up requirements for digital twin systems are high, once developed, they are a key enabling technology of applied machine learning to deliver a high level of automation through PdM and process control [150, 64, 110]. Research shows that adopting digital twin predictive maintenance frameworks can increase machine up-time by 10 to 20 percent [150]. Mourtzis et al. present an edge-based architecture that uses SVM model to classify time-series production data and update digital twin systems in real-time to predict the remaining useful life of critical components [150]. By processing data in edge environments, the researchers overcome the high bandwidth requirements of real-time prognostics

and predictive analytics.

Gankin et al. propose a production control system to address the scheduling challenges of AGV routing between FMS modules in a digital twin environment based on a real-world automotive case study of 25 workstations [110]. AGV routing is a complex problem that involves optimising vehicle management, workstation management, deviation management, routing control, and job release control [110]. In the author's own experience, simulation AGV routing in digital twin environments results in similar data acquisition challenges outlined by VanDinter et al [64]. Digital twin models for AGV routing require high-quality data on AGV systems, with a detailed understanding of speed distributions, effects of interference and collision detection systems and battery charging requirements. The researchers propose a Deep Reinforcement Learning Multi-Agent System approach that uses reward design incentivising agents to achieve maximal throughput in a digital twin environment [110]. The reward-based learning approach of deep reinforcement learning leads to optimising routing and scheduling problems, and the researchers find that combining this method with a multi-agent systems approach helps improve the robustness of the solution [110]. These methods have rarely been combined for modular process control. Further research should explore this topic in other digital twin environments as well as the frameworks to support the real-world implementation of these systems [110].

## **Cloud / Edge**

Delivering machine learning solutions and managing Big Data requires high levels of IoT integration and High-Performance Computing (HPC) platforms. While some manufacturers may already have HPC systems in place, significant work is often required to develop these existing architectures to process large volumes of data in real-time, particularly for the level required for iterative data-driven machine learning models [112]. Furthermore, it is expensive and difficult to scale these systems as demand increases overtime[112]. Cloud solutions can reduce the costs associated with training new personnel to manage, maintain, service, and scale these IT infrastructures [55]. Service-oriented networked computing platforms are a common solution to these challenges due to their scalability and flexibility. Many organisations have already adopted them to process IoT data [112]. These services are called Cloud platforms and provide a means for organisations to structure and manage the functionalities of distributed IoT systems [151].

Cloud-based manufacturing enables the highest integration levels that extend throughout the entire product life cycle. Production resources and capabilities can be monitored and

controlled in real-time through integrating tracking technologies such as RFID and sensors with Zhong et al. describe how Cloud manufacturing is also an enabler of new business models in which a fully integrated flexible manufacturing environment can provide manufacturing as a service to support cross-business applications [55]. Achieving this level of Industry 4.0 maturity requires high data integration and management levels to handle the complexity of service matching, planning, scheduling, and execution [55]. Furthermore, cloud platforms may also offer additional data analytics services such as predictive maintenance, condition-based monitoring, and risk-orientated production planning [151]. By using these additional services, organisations can reduce costs associated with licensing new software [55].

Cloud computing cannot meet real-time data analytics and decision-making requirements due to high Round Trip Latency Time [112]. Many IoT applications are latency sensitive, and therefore hosting analytics in the Cloud can sometimes compromise the performance [112]. In some instances, manufacturing analytics can be performed using only locally stored information, such as inventory management or logistic processes. In these cases, analytics can be performed locally by networked IoT computing devices such as laptops and PCs. Researchers have also explored using Raspberry Pi single-board computers to provide low-cost Edge solutions [150]. However, this is a relatively new concept that requires further research to demonstrate its application readiness [150].

Edge systems are well suited to support computationally intensive and latency-sensitive machine learning tasks such as speech recognition and face recognition, as edge processing generally reduces the power requirements and increases processing speed.[152, 112]. Decentralisation is a common theme in Industry 4.0 literature. By decentralising computation away from the Cloud and using local computing resources, manufacturers can deliver cost savings through reduced energy costs and service fees [150]. Edge computing can increase efficiency and decrease power consumption by over 40% more than conventional cloud systems [112].

In many manufacturers, outdated legacy systems present a major challenge in terms of data accessibility and data integration, particularly mainframe IT systems. In their review of key technological requirements for Industry 4.0, Chen et al. discuss Cloud technologies as a potential solution to mainframe systems [153]. However, they do not provide any further details or specifics on these solutions [153]. Chen et al. suggest that legacy systems are reevaluated or replaced for Industry 4.0 to overcome the limitations in data handling [153]. The author disagrees with these findings due to the high costs that would be as-

sociated with redeploying existing mainframe systems onto new Cloud-based platforms. Furthermore, the economic risk of doing so would be considerable, given the distributed nature of these systems throughout the entire organisation. Chen et al. include no search methodology, strategy, research aims or research question in their review and include a very limited number of sources given the wide range of topics considered [153].

Migrating to the Cloud is often a considerable task. Therefore businesses often adopt a hybrid approach where more innovative areas of the business lead the way and processes slowly migrate over time. Similar to the challenges of IoT, researchers raise concerns on the security of communication between Edge and Cloud-based solutions and the legal barriers of storing personal information on externally managed platforms [150]. The main limiting factor of Cloud computing is the high costs associated with high-bandwidth transmissions, a challenge in machine learning applications that require real-time processing [112].

## **Circular Economy**

Since the first industrial revolution, the global industrial economy has always been based on the same linear model. In this model, raw materials are extracted, transported, processed, and manufactured into products that serve until their end-of-life when they are disposed of, usually by incineration or landfill. The circular economy is an alternative economic model that focuses on the effective use of resources and the elimination of waste by designing products to be reusable, easily repairable, or upgradeable [154]. When raw materials are required, they must be obtained from sustainable sources, ensuring no damage is caused to the natural or human environment. Many scholars, governments, and climate experts believe that adopting a circular economic model is not only beneficial but critical to our survival as a species and is already being proven to be possible in many industries and cities [155, 156].

Industry 4.0 technologies and concepts are key in overcoming the barriers to the circular economy, and trends in the literature show that the topics have begun to converge in recent years [54, 157]. Both share the revolutionary vision that existing organisational and operational systems of production and consumption must change to reflect the changes of the modern world with a focus on digitisation, resource efficiency, productivity and collaboration [54]. Digitisation has been a key enabler of the Circular Economy thanks to Industry 4.0 technologies such as IoT, which facilitates better relationships and communication with customers and provides more sustainable consumer relationships through product-service

systems. End-to-end integration and engineering are key in facilitating a circular product life cycle and product-service systems as it enables smart products to communicate in real-time with engineers and the wider production management system. Other technologies such as Big Data, machine learning, digital modelling, and AI are critical to reducing waste throughout the supply chain by optimising and automating production processes, and supply chain management [54, 155]. Achieving a circular economy requires revolutionary approaches to production, product design, and organisational culture, which are directly impacted during the Industry 4.0 transition.

Digitisation is already revolutionising the way modern manufacturers and other industries conduct business as technologies and concepts such as IoT, vertical and horizontal integration, and predictive analytics provide opportunities for new innovative circular business models [37]. The Ellen MacArthur Foundation is a global charity that focuses on accelerating the transition to a circular economy. In a 2015 report, the charity presented the ReSOLVE framework, a tool for businesses and governments to use to develop circular strategies and initiatives [158]. The tool identifies six business actions that will accelerate a company's transition to circular economy:

- **Regenerate:** Firstly, businesses must shift away from relying on fossil fuels and increasing scarce raw materials towards renewable energy and more sustainable material usage through circular value chains.
- **Share:** In a sharing economy, the product life cycle is extended through designing products to be easily maintainable and upgradeable. Moreover, goods, assets, and services should be shared between consumers.
- **Optimise:** By utilising new technologies like Big Data, AI, IoT, and other Industry 4.0 technologies, production processes can be optimised to remove waste from the supply chain.
- **Loop:** Manufacturers must focus on designing for a product's end-of-life by engineering products and components to be re-manufacturable, repairable, or recyclable.
- **Virtualise:** Servitisation of products enables companies to remove material elements directly and indirectly. This can be achieved by virtualising the business through online models such as internet shopping.
- **Exchange:** Old non-renewable materials, processes, and goods must be identified and replaced with more advanced solutions and technologies.

Since its conception, this framework has been developed further by Lopes de Sousa Jabbour et al, who combine the six elements of the tool with a range of Industry 4.0 technologies that can guide manufacturers towards more sustainable and intelligent production [54]. Servitisation is a common theme, and throughout the Industry 4.0 literature, it is discussed that businesses should move away from product-orientated business models and towards providing products as a service [37, 158, 159]. In these Product service systems, the customer enters into a contract to rent or lease a product, such as a washing machine or a car to establish a long-term, mutually beneficial arrangement with the provider [160]. Product service system models help achieve the highest level of end-to-end integration to gain valuable insight into customer behaviour and the quality and performance of products throughout the life-cycle, including end-of-life. Because customers are paying regular instalments for a product as a service, this continues even if the product breaks down. By considering the end-of-life of a product at the development stage, products can be designed and engineered for assembly and reassembly, making repair and re-manufacturing a quicker and cheaper option to waste [160].

### **2.5.2 Digital Skills**

Data is collected at all levels of the ICT hierarchy, from the production process and control on the line to factory-level manufacturing execution systems and enterprise resource planning systems [161]. Manufacturing data are becoming increasingly valuable as improvements in data-driven systems create new opportunities to leverage and integrate this information to improve decision-making throughout all stages of the product life-cycle [161]. Traditional approaches to data-driven decision-making, such as Six Sigma, lean manufacturing, and discreet event models, are becoming increasingly ill-suited to create meaningful insights from these data [161, 107]. In Industry 4.0 and the age of Big Data, the main technical skill requirements to enable manufacturing are data science, which in the context of the manufacturing environment includes: data analytics, statistics, programming, data mining/data management, optimisation, computer science, and machine learning [47, 48]. The manufacturing sector has the third highest demand for these roles after the professional services and finance sectors, as managing the transition towards Industry 4.0 relies on workers who have an in-depth understanding of applied data science in manufacturing [47]. Among the main barriers to implementing data science in automotive manufacturing, particularly the implementation of machine learning, is the shortage of the aforementioned

data science skills [162, 48, 163, 36, 47]. The lack of workforce skills has limited the development and implementation of machine learning and other Industry 4.0 technologies in manufacturing. Despite this widely recognised skills gap, the reviewed literature highlights that manufacturers are yet to address this through education and training programs fully [47, 38].

### **Human-in-the-Loop Automation**

In their research into Statistical Process Control (SPC) of flexible manufacturing cells, Martinez-Arellano et al. suggest that a key factor in successfully implementing Industry 4.0 technologies is adequate training and support for users through internal education or internal advisors [38]. Research finds that the complexity of the analytics skill-sets required to implement machine learning solutions such as data processing, model selection, and training make these solutions inaccessible to most industrial practitioners [38]. Martinez-Arellano et al. state that despite the considerable research into Industry 4.0 technologies and their industrial applications, industrial practitioners remain hesitant to embrace these technologies [38]. This is largely due to a need for more organisational knowledge on the ROI, implementation requirements, human resource/skills requirements, and capital requirements [38]. The authors' own research in Chapter 3 supports this finding.

To address the social sustainability challenges of Industry 4.0, Martinez-Arellano et al. propose a series of machine condition monitoring dashboards and visualisations implemented in a flexible manufacturing cell at a BMW Group production facility [38]. The researchers highlight throughout the development of the solution the importance of human-centered design, focusing on delivering a solution that improves the workplace experience of the end user. Following the implementation of the dashboard and integrated statistical tools, the researchers report the immediate time savings on root-cause fault analysis and highlight the positive impact on the workers' workplace experience by freeing up time to focus on more intricate problems. Workers became an integral part of the SPC feedback process by focusing on a tool that improved workplace experience rather than headcount reduction. They developed skill sets in feature extraction processes which positively impacted trust in predictive maintenance systems [38]. In addition to this project's social and cultural benefits, the system also led to a 97% reduction in waste. This project is an example of human-in-the-loop innovation. Many scholars present this approach as a medium- to long-term solution for automotive companies to address various technical and cultural challenges

while managing the transition to increased automation and Digitisation [70, 36, 161, 107]. Fahle et al. discuss a lack of machine learning training opportunities in the manufacturing industry and highlight a research gap in the use of 'learning factories' to support the development of these skills in both industrial and academic settings [36]. A learning factory is a physical, realistic manufacturing environment at a very small scale used for training, education, and research that is well suited to understanding the application and implementation of machine learning [36]. This paper finds gaps in the research on frameworks and systems to support education and training using learning factory environments [36]. The researchers also state that these are important tools for increasing the uptake of machine learning in manufacturing [36]. These research findings, however, are solely based on the authors' review of the literature, and no primary evidence was gathered from manufacturers to support these findings [36]. The set-up costs, set-up times, and other factors that could impact production are also not considered in this review. The feasibility of learning factory environments as an effective means of training employees in advanced analytics requires further research.

Research by Suvarna et al. also discusses the benefits of human-in-the-loop cyber production systems in delivering improved flexibility, control, decision-making ability of the user, and improved reaction and response times [161]. Suvarna et al. also discusses the impact on workplace experience and how this synergy between advanced data-driven models and workers can lead to reductions in stress and improved safety [161].

Subramaniyan et al. propose a generic, unsupervised machine learning-based hierarchical clustering approach to detect throughput bottlenecks in a real-world automotive production line [107]. In this study, researchers work closely with domain experts to utilize their understanding of the underlying system requirements during model development. This close collaboration during the development phase led to more innovative solutions and resulted in highly applicable models that deliver improved workplace experience compared to traditional siloed approaches [107]. Furthermore, a two-way transfer of knowledge occurs between process domain experts and data scientists.

Domain experts learn how to participate in new data science projects and improve their understanding of the importance of gathering labelled data and developing testing and training datasets. Depending on their skill level, domain experts can participate in these processes and further develop their technical skills in feature extraction and statistical modelling. At the same time, data scientists gain a more in-depth knowledge of real-world systems and improve their understanding of human-centred design concepts that will in-



crease the usability of future interfaces to improve further data-driven solutions[107]. While the primary aim of this research by Subramaniyan et al. is to reduce bottlenecks and increase throughput, the researchers highlight the importance of developing a solution to improve the workplace experience by reducing the workload for the domain expert as focusing on developing organisational knowledge.

## **Social Sustainability**

Some researchers explain the slow adoption of Industry 4.0 due to the limited availability of skills, poor change management, and a lack of organisational knowledge [39]. Some researchers further suggest that this problem is more complex than just a skills shortage, but that the Industry 4.0 paradigm at its core needs to be better aligned with social sustainability goals [39]. In their research on social sustainability in the age of Digitisation, Grybauskas et al. find evidence to suggest that both high-skill and low-skill jobs are equally at risk in Industry 4.0 [39]. The review highlights studies that suggest 40%-60% of jobs are at risk from technological change [39]. These findings are supported by further research that warns 47% of jobs are at risk of computerisation and 24% of UK jobs are at high risk of automation [39]. While those most at risk are manual workers performing repetitive tasks, highly skilled engineering jobs are also at risk [39]. Machine learning systems have been proven to outperform human performance in image recognition, prediction, and diagnostics in certain domains [39]. There is some debate on this topic as Boavida et al. find in their literature review that automation creates jobs, and as capital investment increases, so does employment [68]. Parida et al. also find that digitalisation is an essential enabler for sustainable business practices in the long term by enabling new business models aligned with goals of circular economy [37]. Evidence from interviews at various automotive manufacturing companies in Portugal supports this, where human-in-the-loop systems are used to support workers to increase efficiency by leveraging the combined benefits of both the domain expert and data-driven systems [68]. On the other hand, multiple cases were found in these companies where workers were displaced from their jobs. As a result, there was some resistance to automation technologies from lower-level employees [68]. Management was also found to be resistant to change in some instances. Boavida et al. identify one example where automation solutions do not consider the end user, and resulting dashboards create more work for the manager due to poor user interfaces [68].

Grybauskas et al. argue that these social sustainability challenges can be addressed through

government intervention to regulate the implementation of Industry 4.0 through taxation, education, and labour relationships [39]. This hypothesis is well supported by the outcomes of legislation passed in South Africa, where organisations are scored on their workforce skills [48]. A low score may limit organisations with whom they can do business, therefore incentivising investment in workforce skills [48]. This legislation positively improved the workforce Digitisation skills in South Africa’s automotive industry and led to increased uptake of Big Data analytics, and AI solutions [48]. However, Bag et al. note that this research is limited to a single country where digitisation skill levels are generally low [48]. While there is some argument that government intervention may be an effective approach at the macro level, this is clearly outside the individual organisation’s control and not considered further in this review.

It is important to note that when discussing education, training, and up-skilling, this does not mean all employees across the board should re-train and gain new qualifications in programming, data science, and advanced analytics. This is both infeasible and impractical, as some employees may need the prerequisite technical or soft skill requirements, or this level of advanced training may not be relevant to their job function. However, all workforce members should be trained and educated to support automation and digitisation projects where required. For example, a common issue identified in the literature is the lack of high-quality datasets, a topic discussed further in section 2.5.4. This issue could be solved through basic training on the importance of gathering high-quality labelled data and how to identify and report opportunities where data collection could be improved.

## **Digital Skills Conclusion**

To conclude, research shows that the full realisation of Industry 4.0 is yet to be realised in the manufacturing industry and that one of the main barriers is a lack of data science skills. Consequently, there are gaps in organisational knowledge leading to a reluctance to embrace emerging technologies due to complexity, technical expertise, and uncertainty of investment requirements [38].

For manufacturers to manage the transformation towards increased automation through machine learning, organisations must shift hiring strategies to focus on data analytics and data science skill-sets and develop digital skills in the existing workforce [69]. This up-skilling of the existing workforce must be approached with careful consideration of the impacts on the workforce through the application of human-centered management approaches. The available research suggests that an effective solution to overcome these

barriers is to begin by exploring short-to-medium term innovation pilots to implement human-in-the-loop analytics projects that focus on improving workplace experience rather than headcount reduction. Human-in-the-loop innovation in data analytics has been shown to deliver high economic returns [38]. More importantly, it builds organisational knowledge and improves trust in these systems, thereby encouraging further innovation and replication of these technologies [38]. This human-centric approach to innovation addresses organisational and technological barriers associated with developing and implementing Industry 4.0 technologies. Companies beginning this journey or unsure where to focus efforts should explore systems for which considerable research is available and ROI can be easily quantified, such as: PdM, quality assurance, and anomaly detection.

### **2.5.3 Management**

#### **Management Approaches in Industry 2.0 and 3.0**

During the early 20th century, following the introduction of mass manufacturing, traditional management approaches operated in functional silos and strict hierarchical management structures where the main focus was to deliver high predictability, and efficiency [164]. This was achieved through the introduction, and gradual optimisation of the moving assembly line, where workstations are fixed in position and parts are moved between stations to deliver high production rates and low labor costs at the expense of low product variability. This type of assembly is now referred to as a transfer line. Over the years, the concept of a transfer line has been extended to modern-day production systems. It now refers to production systems where automated processes are hard programmed into machines and systems [87].

Throughout the latter half of the 20th century, as competition in the automotive industry grew, automotive companies began adopting more consumer-centric approaches and differentiating products to meet various consumer needs to maintain a competitive edge [165, 164]. [87] In response to the increasing demand by consumers for a diverse range of products, Toyota introduced the Toyota Production System between the late '60s and early '70s. The Toyota Production System was a production philosophy focused on delivering increased consumer value by minimising lead times in production and increasing the response times from suppliers to the customer [166]. This is achieved through the continuous improvement of production and business processes by adapting or eliminating aspects of that process that results in lost time or money. From the authors' own experience, a key

principle of lean manufacturing is the concept of creating a continuous workflow for each department and minimising bottlenecks. This requires effective cross-functional collaboration between departments and represents a significant change from the traditionally siloed workstreams of Industry 2.0. This 'just-in-time' philosophy encompassed by the Toyota Production System was the precursor to the modern-day lean management practices widely used throughout the automotive industry and other sectors [166].

Throughout Industry 3.0, following the introduction of consumer-centric and lean management approaches, bureaucratic-hierarchical structures emerged to define fixed lean processes that could be well coordinated throughout the organisation [164]. This bureaucracy also made large-scale planning and coordination of multiple product families more predictable and easily managed by a centralised control group through a top-down management approach [164]. Project management approaches during this period are characterised by systematic and rigid processes such as stage-gate and waterfall methodologies. Targeted goals must be reached at set project stages before progress can continue [167]. These phase-based models enable management to monitor and effectively control project progress throughout development and are still widely used today in the automotive industry [167]. Throughout the 1970s and 1980s, increased amounts of data became available through new digitalised production systems and improved ICT infrastructures. As companies began to recognise the value of this data, data analytics became an increasingly important part of improving individual production processes to meet performance metrics, improve product quality and reduce lead times. During this era, programmable automation systems were introduced where production systems could be reprogrammed to accommodate batched production. However, this reprogramming would usually result in machine downtime during changeover [87]. By the 1990s, the maturity of manufacturing systems had developed enough that programmable systems were becoming commonplace in automotive manufacturing environments. Manufacturers began introducing new project management approaches, such as Six Sigma, that focused on creating value from the large amounts of process data produced by these systems.

In the reviewed literature, the origins of Six Sigma are not discussed, and therefore refer the reader to an additional reference by Hahn et al. on the evolution of Six Sigma for further information [168]. Six Sigma is a top-down management philosophy originally introduced by Motorola in the 1980s to improve product and process quality using a disciplined data-driven process [168]. Traditionally, Six Sigma 'Champions' are appointed by senior leadership who facilitate implementing and deploying the strategy within their

business area. At the process level, 'Process Owners' are responsible for managing the Six-Sigma project, which involves SMART goal setting for quality objectives, monitoring the progress towards these goals, and training and mentoring other team members. Project team members are then tasked with implementing the strategy within their respective workflows. A Six Sigma project was originally designed to follow four key stages: Define, Measure, Analyze, Improve, and Control. This approach is called DMAIC, which is used to improve existing processes.

Six Sigma has since been adopted by many industries that have further developed these methodologies into other areas, including process design (DMADVO), product servicing, and non-manufacturing business applications. For a detailed description of the DMAIC and DMADVO stages, as well as the evolution of Six Sigma in business, see [168]. Even as the Six Sigma method was adapted and generalised, the focus remains on using disciplined, quantitative approaches to improve process quality. Six Sigma continues to be the leading approach for quality management at the world's largest manufacturers like Sony, Lockheed Martin, Nokia, Ford, and GE. For many of these companies remains a standard for data and analytic skills training to guide process improvements [168].

### **Management Approaches in Industry 4.0**

Like the revolutions that preceded it, Industry 4.0 brings new opportunities for organisations to increase productivity, flexibility, and customer value by embracing new technologies, business models, and management practices. High levels of machine learning-driven automation in FMS environments challenge well-established organisational structures and require increased collaboration between production, simulation, and data science teams to optimize production offline and explore new automated management and control processes. This not only requires a significant long-term investment in physical assets and human resources to build organisational knowledge but also cultural changes to ensure the workforce is empowered to support ongoing innovation. Various management strategies are presented in the literature to manage the ongoing cultural, and technological changes in the automotive industry at various levels of the organisational management hierarchy [70, 72, 166, 169].

The journey towards the digitisation and automation of production should be viewed as a continuous evolution as opposed to a revolution, one that is driven by operating costs and improving productivity [33]. This should not require a complete overhaul of outdated systems to be replaced with the latest technologies. Instead, the Industry 4.0 philosophy

is to optimise, digitalised and integrate existing processes using technologies such as IoT, Big Data analytics, and cloud-based systems to create further value from existing Industry 3.0 based technologies [33]. To adopt Industry 4.0 technologies and remain competitive in a world of digitalisation, considerable amount of long-term investment is needed at both the corporate level and the supply chain level to achieve Industry 4.0 goals. Management needs to be open to this investment and ensure the long-term business strategy is well communicated throughout all levels the workforce. Effective downward vertical communication of the ongoing organisational changes ensures supervisors and managers have the best possible understanding of the future technological and human resource requirements in their respective areas. This in turn enables managers and supervisors to put together clear business cases for required investments to be communicated back up the management chain, allowing businesses to act quickly to change.

A bottom-up approach to Industry 4.0 is ineffective, whereby an organisation invests in Industry 4.0 technologies used by their competitors or within internal innovation hubs and then implements them without having global goals or specific problems to solve [70, 33]. Instead, it is argued that digital transformation should begin with an organisation understanding its long-term vision and objectives. Once these goals are established, organisations should explore top-down solutions that focus on addressing business-level goals, prioritising those that offer the highest long-term ROI [70]. Albukhitan et al. suggests that during the final stages of implementing business goals relating to digital transformation, organisations should then consider adjustments in organisational culture and infrastructure requirements. This can be achieved by reviewing training, qualification requirements, or hiring strategies to meet the new demands of changes in high-level objectives or organisational change [70]. Albukhitan et al. argues that this top-down approach helps develop a holistic strategy to transform all aspects of an organisation that can be implemented while avoiding the inefficiencies of functional silos [70].

In contrast, Kulvisaechana et al. find in their research on change management that if a strategic change is to succeed, changes should initially take place in the cultural beliefs and assumptions of the organisation, thus leading to the cultivation of workforce commitment in later structural changes [116]. This finding seems to contradict Albukhitan et al.'s suggested approach. However, Kulvisaechana et al. later conclude that to effectively manage organisational change while promoting a culture of innovation, management must think and act holistically and make changes on several fronts in careful alignment [116].

Research shows that employees do not resist all organisational changes, only that which

has not been well communicated or is perceived as psychologically or economically threatening [116]. Communication strategies are identified in the literature as a key mechanism to facilitate organisational change and should support a systematic change management process that considers both changes in culture and organisational structure [70, 116]. This requires careful consideration of changing internal environmental pressures that dictate employee behavior, motivation, and performance of teams. Decentralised agile management approaches are regarded by scholars as one of the most effective management models in this regard, as it enables lower levels of management to be more embedded in their local environments and respond faster to changes [170, 171].

Agile management is a project management process originally popularised in software development but has recently become widely adopted in the manufacturing sector [47]. The agile methodology promotes continuous development and testing during the whole life cycle of a project to deliver benefits throughout the process rather than just at the end. This iterative approach to developing and implementing manufacturing solutions enables management to adapt quickly to changing environmental requirements and deliver highly flexible solutions [47, 69, 162]. This management approach is particularly applicable for machine learning, as the optimal solution will likely change throughout the product life cycle due to data availability. For example, consider an anomaly detection system for quality assurance of an automotive manufacturing process. Due to data availability limitations in the initial development phases, early model iterations may be fully unsupervised or rely heavily on synthetic data. This solution could be implemented in early pilots to demonstrate concept readiness while high-quality data are still being collected. Through later development stages, engineers may decide to transition towards new semi-supervised or supervised approaches that utilize these new data to deliver increased reliability. This decentralised approach to project management challenges well-established phase-based models and requires higher levels of communication and collaboration between teams [170, 171]. Some research shows that agile methodologies can be challenging to apply in manufacturing environments as this flexible management approach challenges traditional routines and work processes of Industry 3.0 that are culturally embedded in factories [69]. Despite these challenges, research in hiring trends in manufacturing shows that organisations recognise the benefits of agile methodologies as experience in agile management is among the sector's most sought-after skills [47].

Agile management approaches are widely regarded as the most effective management approach to manage the requirements for many Industry 4.0 projects due to the agile method-

ologies' ability to deliver high flexibility, adaptability, and faster time to market [69].

## **Maturity Models**

Throughout the transition towards increased levels of digitisation and automation, the specific challenges and opportunities of any given solution will continuously change as organisations build the required organisational knowledge to develop and implement these solutions successfully. This presents a complex management challenge due to the wide range of emerging technologies and organisational changes that must be considered carefully to deliver this transition successfully. Various digital maturity assessments are presented in the literature as a change management tool for organisations to manage the technological and organisational requirements throughout this journey. Various maturity models were identified in the reviewed literature to help assess an organisation's maturity relating to Industry 4.0 objectives [172, 72, 69].

In their 2005 paper titled 'Understanding the Main Phases of Developing a Maturity Assessment Model', Bruin and Freeze argue that any maturity assessment tool fits into 3 general categories: descriptive, prescriptive, and comparative [173]. A descriptive tool is used to define the current state of a business, providing no means of improvement or relationships between the current state and key performance metrics. A prescriptive tool provides further insight into how the current state of a business relates to key performance indicators, highlighting which areas can be improved to deliver value. Finally, a comparative model compares maturity across industries, providing insight into the differences in business practices and how this relates to value generation in different environments. A comparative model requires many assessments to be carried out across various industries to gather sufficient data to draw these relationships.

The author also argues that although a wide range of maturity models can be found in the available literature, despite their claims, these models are not well suited to give actionable feedback to management to guide specific process improvements and wider organisational change. Limited qualitative and no quantitative evidence supports the argument that organisations benefit from performing an Industry 4.0 maturity assessment. This claim is discussed further in Chapter 3.

The maturity models included in this review give little direction on how the assessment should be conducted, what personnel should be involved, and who should take responsibility for delivering the resultant roadmap. In their discussion on digital transformation and change management approaches, Albukhitan et al. find that successful digital transfor-



mation requires support from all stakeholders, including top-level executives, employees, and customers [70]. The author suggests that future research into maturity models focus on addressing the suggestions of Albukhitan et al. by developing new maturity models that consider the involvement of these key personnel throughout the assessment process. Further research should also consider the real-world application of maturity assessments within an organisation.

Regardless of the usability of existing maturity models to deliver useful roadmaps to Industry 4.0, the importance of developing a structured roadmap to support change management is not contested.

## **Management Conclusions**

The automotive industry is in the midst of major organisational change due to various emerging technologies and business practices that challenge well-established business models at both the strategic and operational levels. Research shows that to manage this organisational change successfully, automotive manufacturers must think and act holistically and make changes on several fronts in careful alignment [116]. Because of this holistic approach discussed throughout the literature, the barriers to developing and implementing machine learning solutions aligned with the more general challenges of Industry 4.0 change management. These include a need for increased communication both vertically and horizontally to facilitate effective collaboration between teams to deliver cross-functional manufacturing solutions. Automotive manufacturers must learn from the lessons of Industry 3.0 and recognise that investment should not be limited to manufacturing technologies, but also into developing good organisational structures and communication channels to support the effective management of human resources [33]. Senior management must take steps to communicate the business strategy to all levels of the workforce to identify future skill requirements across different functional areas and ensure they have sufficient personnel with the right digital skill sets to support the operation and future development of Industry 4.0 systems. Agile methodologies are presented as an effective way to manage these projects at the process level and particularly lend themselves to projects exploring machine learning solutions. Given that this project management approach challenges well-established methods such as lean, stage-gate, and waterfall, organisations should ensure that good communication structures are in place to address these new models' cultural barriers. Managing the large-scale organisational transformation of the workforce, processes, strat-

egy, and culture requires well-structured step-wise approaches [172, 72, 69].

## 2.5.4 Computer Vision

Computer vision is a field of machine learning in which algorithms are trained to interpret the image or video data and perform tasks such as object detection or classification. Computer vision applications in manufacturing settings include quality testing, safety supervision, inventory management, and process monitoring [27]. If implemented correctly, machine vision systems can optimize quality assurance processes as they are quicker, more objective, and continually functioning compared to human inspectors [27]. This leads to lean improvements, reduced labour costs, improved part quality, waste reduction, and improved traceability [24]. Combined with the predictive and prescriptive process management approaches, integrated vision systems collect data on important production metrics for analysis [27]. Computer vision can be combined with other machine learning technologies to optimize processes and deliver business intelligence [27].

In the reviewed literature, 18 papers were identified for which the main topic is the development of a novel computer vision approach or a novel application of an existing system. Of these papers, only 3 papers were found to have been actually implemented in a production setting [28, 86, 97], while 12 were a proof of concept or a literature review [94, 96, 174, 95, 27, 24, 89, 108, 175, 87, 101, 99], and 3 were unclear whether a solution was implemented or not [98, 100, 85]. Most of the reviewed literature focuses on computer vision systems for quality assurance in which inspection tasks are automated or semi-automated and deliver a pass/fail result to the production control system, guiding routing and rework requirements accordingly. Other use cases for computer vision in manufacturing include bar-code checking [27], supervision [27], tracking and reporting [27], condition monitoring [27] and inventory management [27].

### Computer Vision Environmental Conditions

A common challenge in the computer vision literature is the variation of environmental conditions such as changing light conditions, colour, background, object placement, and orientation [24, 87]. Flexible manufacturing environments present a challenge for computer vision systems which work optimally in environments where variation in light conditions, texture, scale, and position are minimised. FMS may require reconfiguration of the line,

affecting these lighting conditions, and variation in the produced parts require more training datasets and higher levels of integration to identify the incoming part and perform the correct analysis.

Schluter et al. use weight and image data as input to vision-based classification models to determine the suitability of damaged automotive components for re-manufacturing [96]. To overcome the challenges of lighting conditions, the researchers suggest using an enclosed lighting rig with a uniform white background to place the part inside for the highest possible accuracy of the computer vision system, achieving accuracies of 98%. In this study, training and testing data are collected from the automotive manufacturer, considering the variability of how different workers might place objects in the lighting set-up. This paper also includes the only identified use case where additional data on the weight of the components, in addition to images, as part of the classification input. However, the paper's methodology lacks detail on the datasets as well as the industrial impact of the final system. No information is provided on the set-up cost of this system, nor the amount of time taken to label the datasets, however, the researchers do discuss that future work will focus on automating this process of data acquisition.

Courville et al. use a similar approach to overcome lighting challenges when identifying manufacturing defects in gears [99]. The paper uses a novel image acquisition approach where photos of the gear are taken at multiple angles. The vision system only flags the part as defective if the defect is identified on multiple images. A predetermined threshold of the number of required defective images to return an anomalous result. Image data are collected using a high-resolution camera and within a similar enclosed inspection cell to minimize the impact of changing environmental conditions. A single-axis gripper rotates at a rate that ensures every tooth of the gear is scanned multiple times. This process ensures consistent rotation angles and constant lighting to ensure high-quality data and repeatable results. This paper includes good details on the data labelling approach, in which a ground truth dataset was built by having a domain expert manually label the parts' state before scanning to assign labels. Courville et al. state that each gear scan took about 90 seconds using this method, and further optimisation of this process could reduce this to 20s per gear [99]. Parts were randomly selected from the line to be scanned, resulting in a dataset of 193 gears, 93 of which had some defect present, and no instances were found where two defects occurred on the same part. Because of this random selection, this suggests that this is representative of the actual anomaly distribution. From all images of the scanning process from the 193 parts, 3172 images of defects were labelled by the

domain expert. This paper is a good example of the required level in developing training and testing datasets that many papers fail to include. Labelling data is time-consuming, and understanding the estimated time commitment is important for reproducing computer vision researchers in both academic and real-world manufacturing settings. Courville et al. also give an estimate of the solution’s impact, finding that their approach could reduce the requirements for manual quality inspection by 66% [99]. Few researchers consider the real-world impact of machine learning solutions and no papers in the reviewed literature comment on the estimated economic impact of their solutions.

Papavasileiou et al. also discusses the importance of stable lighting when computer vision detects manufacturing defects in automotive brake components. In this instance, the challenge was overcome without using an enclosed lighting rig, instead using a bright white LED strip with careful orientation of the inspected part controlled using a robotic arm [97]. In addition to the lighting conditions, the resolution of the camera is also important [99, 87]. Mazzetti et al. find that this causes challenges when using video to detect quality defects in brake components as the resolution is too low to detect surface defects [87].

## **Data Availability and Computer Vision**

The challenge of limited available data is common among the reviewed literature [97, 87, 101, 95, 98]. Some researchers overcome this through generating synthetic images using GAN networks as discussed in section 1.2, although this approach is complex and requires high levels of machine learning skills as well as a high-quality sample of training data. Some researchers rely on publically available datasets to validate their models, however, these open-source datasets are also often limited in size [101]. In their research on visual analysis of fracture surfaces for root cause analysis of industrial processes, Bastidas-Rodriguez et al. aim to demonstrate that handcrafted features in deep learning lead to improvements in performance [101]. The benchmark selected for this study used databases with small samples of only 81 and 108 images per class in the KTH-TIPS and KTH-TIPS2-B datasets, respectively, resulting in very small sample sizes with only 70 for training and 11 for testing. The authors note the lack of large datasets for texture and fracture use cases.

In their research on quality inspection of soldering of electronic components, Schwebig et al. study the impact of adding augmented training data to overcome training data imbalance. This rigorous study compares multiple training datasets on multiple testing datasets with various fault types. The common image data augmentation approach is to perform geometric and colour variations through rotations and translations, filtering techniques [98].

This approach is much simpler than using GANs to generate synthetic data, as discussed in section .1.2. By comparing the F-score of training datasets that include augmented data, researchers could significantly increase computer vision performance compared to non-augmented datasets [98]. The researchers found that in all test cases, high-quality and diverse training data are the most significant factor that impacts successful implementation [98].

Malburg et al. evaluate three CNN object detection methods and their performance in detecting workpieces and recognising failure situations [89]. This proof of concept shows how this technology could be applied to identify workpiece misalignment. However, only two failures are simulated by manually misaligning the parts and collecting image data. It is unclear how the accuracy of this model trained using manually simulated data will perform in real-world scenarios, particularly where unknown errors might occur. This limited information on the real-world use of this system means the researchers do not consider this project's economic benefits/cost-benefit analysis. Once again, no consideration is given to the required skills to implement this solution nor the time taken to gather and label the synthetic data.

Papavasileiou et al. present an image-based system to detect manufacturing defects in automotive brake components [97]. A bespoke unsupervised algorithm uses a sliding window process to measure brightness variation in a greyscale image and classify parts as either OK, Minor defects, or, Not OK [97]. While some papers consider different fault occurrences, this is the only example in the reviewed literature where an algorithm indicates the severity of a single fault. During model development, domain experts observed the sliding window results to determine anomaly threshold values manually, however, little detail on the exact methods used to calculate these thresholds is presented. The solution was implemented at an automotive manufacturer with processing done locally and presented to the user in an online dashboard. The paper presents an accuracy of 92%, although this is not compared with other methods such as CNN or ANN [97]. F-score would have been a more appropriate evaluation metric given the high imbalance of the testing dataset. A small sample size of only 25 parts for the testing dataset is unclear if this sample is representative of the actual anomaly distribution of the manufactured parts.

Luckow et al. present an automotive use case of a computer vision system in supporting walk-around quality inspection of the final product [28]. The paper includes considerable detail on the integrated system architecture used to enable data collection using hand-held IoT devices before a fine-tune transfer learning approach performs classification in the

Cloud [28]. The researchers state that the main contribution of this paper is the creation of an automotive dataset to learn and recognise different vehicle properties [28]. Despite this primary aim of the paper, the end use of the IoT-supported computer vision system is not discussed beyond its use for visual inspection. The paper goes beyond other works in describing how long it took to build the test dataset, which took multiple workers weeks to build [28]. Without a clear insight into the labelling process, the definition of labels, and who was assigning the labels, it's difficult to replicate this. Without a clear methodology, the quality of this data and the results are in question [28].

In addition to 2D inspection approaches, many companies now also use 3D scanners in combination with computer vision to deliver quality improvements [27]. One such example is presented by Zhu et al., who propose a 3D scanning system for quality assurance using point cloud detection [95]. Point clouds of manufactured parts are generated using a cobot-mounted camera system and are compared with point clouds generated using CAD data to identify manufacturing defects. The paper shows that unsupervised approaches can be used to generate point clouds that are representative of those produced by CAD data, negating the need to develop supervised training and testing datasets. However, the proposed solution is complex to set up in comparison to the supervised alternative. The author speculates that developing a supervised dataset would take less time than what would be required to replicate the complex point cloud detection approach. Further research is needed to compare the time requirements of both solutions and directly compare the performance of both methods in detecting a wider range of defects. Stavropoulos et al. propose a remanufacturing cell to robotized the remanufacturing of automotive frames for which training and testing dataset was created using CAD data [90]. The proposed approach is only tested using synthetic data, with further work required to validate the results in real-world scenarios as it's not clear how well the accuracies found in this study will translate. Given the similarities between this case study and that of Zhu et al. there are opportunities to improve both methods by combining these findings.

Runeson et al. discuss the data management challenges of machine learning in industry and the requirements for large volumes of high-quality data [176]. Open Data Exosystems (ODEs) is presented as a potential solution to this challenge [176]. An ODE is a network of community actors consisting of both organisations and individuals that collaborate on developing datasets and related resources to foster innovation, create value, and support new business [176]. This research is presented as a wider solution to enterprise-level challenges, however, this collaborative sharing of both knowledge and data could also be applied to

solve limited data availability in computer vision settings. ODEs challenges the existing culture of competition and secrecy within the automotive industry and presents new legal, organisational, and technical challenges. Researchers conducted 5 focus groups with 27 experts across 22 organisations to survey ODEs [176]. Participants raised concerns that sharing data requires giving away business value and would rather collaborate with businesses that are not competitors. Furthermore, these organisations recognise that many challenges can be overcome through collaboration but raise concerns about security, legal barriers, authenticity, standardisation, quality, and trust. Kauffman et al. discuss these legal challenges in their research on intellectual property in Industry 4.0, highlighting that in the newly interconnected manufacturing environments, data is of high value and worthy of protection [25]. This value creation lies in the ability to gain insights into integrated smart objects through data analytics, thus generating new organisational and technical knowledge which can deliver competitive advantage and innovation [25].

While there remains some debate on ODEs, the author would add that in the automotive industry, this concept should also be considered when collaborating with organisations or other manufacturing sites within the same company. In large-scale automotive companies, different sites may view each other as competitors when bidding for new contracts and, therefore may be resistant to sharing information on innovation projects. Runeson et al.'s research is limited in highlighting how Open Data Exosystem culture varies between industries and what specific barriers need to be overcome in individual sectors, and further research is needed to understand the value proposition of ODE. Theissler et al. include further information on the organisational challenges of data-driven topics in the automotive sector [30]. Research finds that engineers often have access to data in automotive settings due to political and bureaucratic barriers [30]. Furthermore, researchers find that a lack of understanding and support in gathering labelled data on faults in production environments makes developing machine learning models difficult [30].

## **Social Sustainability in Computer Vision Research**

Social sustainability in the context of manufacturing refers to the ability of the organisation to promote social well-being and equity, while also minimizing negative social impacts across all aspects of operations. This includes considering the well-being of workers, the community, and society at large, as well as the ethical and responsible use of resources.

Babic et al. present a systematic review of computer vision systems for quality inspection in which they find the majority of the literature focuses on comparing the accuracy between

human operators and autonomous systems [24]. Only 4% of the surveyed results compare existing systems with improved methods [24]. This is supported by the author’s research findings and suggests that these systems are not currently widely used in manufacturing, with many processes still largely relying on some manual process for quality assurance.

Babic et al. find that in Computer vision quality inspection, most researchers either developed bespoke software or did not mention what software was used [24]. Of those that did, CAD and MATLAB was the most popular choice for computer vision inspection. CAD was used in instances where 3D representations were required, while common machine learning methods for computer vision include ANN and object segmentation [24]. These findings differ from the results of this study which finds the most common machine learning method for computer vision to be CNN’s such as YOLO, Alexnet, GoogLeNet, VGG. When considering variants of the YOLO network, such as YOLOv4, YOLOv5, and YOLOX, this network is by far the most common CNN variant in the reviewed literature.

Within the context of Industry 4.0, Babic et al. also raise concerns that the human-centric concepts of Industry 4.0 are not well understood in applications where computer vision is being used for quality inspection [24]. Researchers found that the knowledge of industry 4.0 needs to be more widespread with regards to human inclusion but do not include examples of where this has been done well, nor guidance on how this could be improved [24]. The authors’ findings agree with this. Some studies mention workplace experience, automating non-value added work processes, and improving safety [100, 108, 175]. However, these examples do not explore the qualitative or quantitative impact of automation on workplace experience[100, 108, 175].

Rosa et al. propose an IoT network architecture for computer vision systems to inspect the quality of wiring harness parts, however, the proposed solutions still need to be tested or implemented [174]. The research also discusses using other IoT technologies to support increased data collection and integration, such as RFID cards to monitor employee behavior. In this example, little consideration is given to the IoT security challenges of safely storing this personal information, nor consideration to how this data collection might conflict with union requirements in real-world applications.

Torres et al. demonstrate a use case in which a supervised learning approach is used to train a CNN to classify different types of wiring harness connectors received from a supplier in an unsorted box [108]. The system then classifies and sorts the connectors while separating faulty instances to reduce waste and improve the workers’ experience by automating a highly repetitive and time-consuming task. In instances such as these, where an entire



process or sub-process is automated, the desired outcome of the end system must be well communicated to the workers following guidance on change management, as outlined in section 2.5.3.

Ferreira et al. present a computer vision system to inspect informative labels in vehicle assembly, a process previously performed through manual inspection using a paper-based checklist [100]. With increased product variety, quality inspection of these labels becomes increasingly complex, and researchers find that mental fatigue, physical fatigue, and the lack of experience can lead to a reduction in performance that could lead to safety-critical production errors [100]. Instead of fully automating this process, Ferreira et al. present a solution to assist operators in the quality inspection process [100].

Javaid et al. suggests that such systems can be implemented in a socially sustainable and human-centric manner without resulting in job losses [27]. Instead, automation should address high-priority work while workers have more opportunities to learn more demanding skills from their cognitive point of view [27]. The author would add that workers may not want to be upskilled or may not have the prerequisite skill requirements to learn more advanced skills to support computer vision systems.

## **Computer Vision Conclusions**

Given the rapid advancement of these technologies in academic and industrial settings in recent years, there are significant opportunities for further research and industrial application of these systems. Multiple instances are identified where previous researchers have faced environmental challenges relating to computer vision systems, such as changing light conditions, colour, background, object placement, and orientation [24].

In the reviewed literature, ODEs are identified as a potential organisational solution to this challenge, in which automotive industries need good digital communication channels to collaborate between internal and external organisations to develop high-quality datasets. This concept of ODE extends beyond computer vision and offers a solution to improve other datasets, such as time series for condition-based monitoring as well as enterprise data [176]. ODEs are a relatively new concept and challenge existing cultures and organisational practices of the automotive industry, with the main barriers to ODEs being political and bureaucratic processes [30]. Furthermore, researchers find that a lack of understanding and support in gathering labelled data on faults in production environments makes the development of machine learning solutions difficult [30]. Organisations must recognise these cultural barriers and address them through structured change management strate-

gies supported by high levels of communication of these systems' social, environmental, and economic benefits to the workforce [116].

A common technical solution to lighting and environmental challenges is to use a robotic arm to accurately control the orientation of the inspected part inside an enclosed lighting rig [97]. Environment enclosure for vision systems is required even where lighting conditions seem constant as additional light from opening shutter doors, windows, etc may affect performance [87]. Even in highly controlled conditions, research shows that the most significant factor that impacts successful implementation is the quality and diversity of training data [98]. In instances where position and lighting solutions are not practical due to limited available capital, or physical space, larger training datasets may be required to account for the increased variability of the inspected parts. Very bright lights in a fixed position may also be used to minimise the impact of changing environmental lighting. In all cases, researchers highlight the importance of using cameras with sufficiently high resolution to identify the respective defects.

A research gap is also identified in exploring the social sustainability of these systems. A key consideration of Industry 4.0 is the impact that innovations in digitalisation and automation technologies have on the workforce, and how any negative impacts should be managed. Further research is required to develop frameworks to support the implementation of computer vision systems in a socially-sustainable way that considers these aspects. This finding is supported by prior research in computer vision systems in manufacturing [24]. Future research focusing on human-centric applications of computer vision should use questionnaires to explore these systems' cultural and social impact on management and line workers in automotive manufacturing settings.

### **2.5.5 Key Findings of Machine Learning Opportunities and Barriers**

For the reader's convenience, this subsection includes a brief summary of the key findings of the key topics discussed above. Various machine learning solutions in automotive manufacturing are discussed in the reviewed literature, with the most common applications including computer vision for quality inspection, automated process control, predictive maintenance, automated scheduling, and anomaly detection. In all these cases, limited organisational knowledge of machine learning requirements at the managerial and process

levels is one of the main barriers found in the reviewed literature. Researchers found that the limited understanding of the potential value creation and return on investment of machine learning technologies results in a lack of support in gathering high-quality labelled data on faults in production environments [30]. This is a common theme throughout the literature that makes developing machine learning solutions difficult [30]. Organisations must recognise these cultural barriers and address them through structured change management strategies supported by high levels of communication of these systems' social, environmental, and economic benefits to the workforce [116]. Maturity models are presented as a tool to support this change management by helping to develop a roadmap towards improving the technological readiness of key enabling machine learning technologies, however, this research finds little evidence to show that maturity models lead to measurable change in organisations Industry 4.0 maturity. Further research is required to quantify the impact of these tools.

Only a few papers in the reviewed literature discuss the potential economic impact of the proposed machine-learning solutions. Communicating the business case for industrial applications of machine learning is critical to help address the knowledge gaps within the automotive industry. Future works should focus more on the economic appraisal of real-world implementation of machine learning solutions to support replication in industrial settings.

This oversight on communicating the business case of proposed solutions is largely because most papers in the reviewed literature present a proof of concept rather than an application-ready solution. As shown in Fig 2.2, research in this field has increased considerably in the last 5 years. This suggests that many of these technologies are still in their infancy and suggests a wide range of opportunities for both further academic research and industrial application of these technologies.

Computer vision systems for quality assurance are the most common research topic in the reviewed literature. Multiple instances are identified where previous researchers have faced environmental challenges relating to computer vision systems, such as changing light conditions, colour, background, object placement, and orientation [24].

Many researchers highlight the challenge of collecting or accessing sufficient labelled anomaly data in industrial settings. In many cases, additional data are generated synthetically to increase diversity of the minority class and address data imbalance [85, 86, 87, 84, 88, 89] while others rely on publicly available data [6]. These approaches present further challenges relating to the reliability, accuracy, and replicability of results [30, 85, 86, 87, 84, 88, 89]. De-

spite the challenges of gathering high-quality labelled data, Theissler et al. find that most machine learning research explores supervised solutions, a finding supported by our own research [30]. Future research should therefore explore unsupervised or semi-supervised solutions to existing challenges and the best approaches to engage domain experts and process owners in collecting high-quality labelled data.

Social sustainability and workforce digital skills being the most common topic in the Industry 4.0 management literature [67, 39, 48, 68, 69, 47]. Some researchers explain the slow adoption of industry 4.0 due to the limited availability of skills, poor change management, and a lack of organisational knowledge. Other researchers further suggest that this problem is not as simple as just a skills shortage but that the Industry 4.0 paradigm at its core is not well aligned with social sustainability goals [39]. There is some evidence to support this claim, and while social sustainability and workforce digital skills are one of the most common topics in the Industry 4.0 management literature, little consideration is given to the social impact of machine learning solutions when developing and implementing real-world manufacturing. No papers in the reviewed literature explore the qualitative or quantitative impact of automation on workplace experience, and a lack of structured frameworks exist to support the implementation of these systems in a socially-sustainable way. This finding is supported by prior research in computer vision systems in manufacturing [24]. Future research should consider using questionnaires and face-to-face meetings to understand the social impact of machine learning and its associated technologies on both management and line workers.

The available research suggests that an effective solution to overcome both the technological and organisational barriers of machine learning innovation is to begin by exploring short-to-medium term innovation pilots to implement human-in-the-loop analytics projects that focus on improving workplace experience rather than on headcount reduction. Human-in-the-loop innovation in data analytics has been shown to deliver high economic returns [38]. More importantly, it builds organisational knowledge and improves trust in these systems, thereby encouraging further innovation and replication of these technologies [38]. This human-centric approach to innovation addresses organisational and technological barriers associated with developing and implementing Industry 4.0 technologies.

Companies beginning the journey towards increased automation and digitisation who are unsure where to focus efforts should explore human-in-the-loop innovation in systems and technologies for which considerable research is available and the value proposition is easily quantified. Predictive maintenance of manufacturing systems and computer vision for

quality assurance are the most widely researched topics in the reviewed literature with a wide range range of use cases presented. A research gap is identified in understanding the social impact on line workers who may be displaced by automation and digitisation. The author suggests that future research in automotive manufacturing settings focus on implementing real-world solutions of predictive maintenance and computer vision. This research should focus on understanding the end solution's economic impact and measuring the social impact on the workforce through surveys.

## 2.6 Conclusion

This systematic literature review applies a rigorous search methodology to provide a comprehensive overview of the machine learning technologies and practices currently used in automotive manufacturing environments and the barriers to their development and implementation. In answering this research question, a range of machine learning approaches are discussed, as well as the most common models researchers use to explore novel manufacturing solutions. Enabling machine learning technologies in manufacturing are also discussed, including a range of Industry 4.0 technologies such as IoT, Big Data analytics, Flexible Manufacturing Systems, Digital Twin, Cloud and Edge computing.

Computer vision is the most widely researched topic in the reviewed literature. A key finding is that while various papers discuss the importance of high-quality training data [98, 100, 108, 87], few papers include insufficient information on how data are collected, labelled, cleaned and validated, which often makes these studies difficult to replicate. Research studies on anomaly detection also often lack information on how threshold limits are set, making it difficult to replicate these solutions [106, 75, 97]. Future research must include this information to ensure research is easy to replicate. These findings help guide our labelling methodologies, threshold limits, and validation approaches in chapters 4 and 5, ensuring our methods could be easily replicated and that our data are validated by domain experts to ensure the data are of sufficiently high quality.

In addition to discussing the technical aspects of machine learning and their associated challenges, the cultural and organisational challenges are also discussed. Research shows that the full realisation of Industry 4.0 technologies is yet to be realised in the manufacturing industry, with one of the main barriers being a lack of data science skills [47, 48]. Consequently, there are gaps in organisational knowledge leading to a reluctance to embrace emerging technologies due to complexity, technical expertise, and uncertainty of

investment requirements [38]. The majority of papers in the reviewed literature focus on proof-of-concept solutions, with few instances where machine learning solutions are implemented in real-world manufacturing settings, with the exception of computer vision for quality assurance. Future research should give more consideration to the economic impact of machine learning solutions in order to support the uptake of these technologies in industry to drive further sustainable growth. This is considered in Chapters 4 and 5 where two machine learning case studies are presented, each of which gives careful consideration to the economic aspects of the projects.

Agile methodologies are presented as an effective way to manage machine learning projects at the process level, however, these project management approaches challenge well-established methods such as lean, stage-gate, and waterfall methods that are embedded in the automotive manufacturing culture [47, 69, 162]. In order to maximize the value creation of machine learning in automotive manufacturing, research shows that well-structured human-centric change management approaches must be in place at the early stage of an organisation's digital transformation to address these cultural barriers and create an environment that promotes innovation and develops trust in new technologies [172, 72, 69]. These findings are used to inform the development of a strategic framework outlined in Chapter 3 which encompasses change management best practices presented by Kulvisaechana 2001, and focuses on human-centric innovation to support the future uptake of machine learning and other Industry 4.0 technologies and practices [116].

A key finding is that limited knowledge of the machine learning requirements at the managerial and process levels are one of the main barriers in the reviewed literature. Researchers find that the limited understanding of the potential value creation and return on investment of machine learning technologies results in a lack of support in gathering high-quality labelled data on faults in production environments [30]. This is a common theme throughout the literature that makes developing machine learning solutions difficult [30]. Organisations must recognise these cultural barriers and address them through structured change management strategies supported by high levels of communication of these systems' social, environmental, and economic benefits to the workforce [116]. These research findings are an important consideration in the assessment methodology presented in Chapter 3 which ensures the involvement of management throughout all stages of the assessment process in order to transfer knowledge on machine learning and Industry 4.0 technologies. Furthermore, these findings also led to the development of 'Bite Sized Training' as a secondary outcome of Chapter 3, in which short training content is developed

for senior management to further address the knowledge gaps that present a barrier to machine learning uptake.

In addition to education and training efforts, the available research suggests that an effective solution to overcome cultural and organisational knowledge gaps of machine learning is to explore short-to-medium term innovation pilots that focus on improving workplace experience rather than headcount reduction. Human-in-the-loop innovation in data analytics has been shown to deliver high economic returns [38]. More importantly, it builds organisational knowledge and improves trust in these systems, thereby encouraging further innovation and replication of these technologies [38]. This human-centric approach to innovation addresses organisational and technological barriers associated with developing and implementing Industry 4.0 technologies. Companies beginning this journey or unsure where to focus efforts should explore systems for which considerable research is available and ROI can be easily quantified, such as: PdM, quality assurance, and anomaly detection.

# Chapter 3

## Industry 4.0 Readiness Assessment

### 3.1 Summary

The transition to Industry 4.0 and the adoption of machine learning technologies is a complex process, requiring organisational changes that challenge well-established business practices. Increased levels of digitalisation and automation will present new social challenges and require careful change leadership to maintain an innovative culture that supports ongoing digital growth. Ford Motor Company recognises these changes, however, there are gaps in the current Industry 4.0 strategy to guide low-level changes at the factory level.

To address these challenges, this chapter presents a strategic framework to support management at the factory level in guiding digital growth. The proposed framework builds on previous research on Industry 4.0 maturity assessments. It considers technological, organisational, strategic and cultural aspects, and presents user-friendly tools to measure progress in these areas against a well-defined benchmark. To design the tool, a wide range of maturity models are critically reviewed and research is conducted into best practices in maturity model design and questionnaire design.

The assessment tool is used to perform three assessments at Ford's UK manufacturing sites. These findings are compiled in this chapter to provide a comprehensive assessment of digitalisation and automation strategy across Ford's UK operations. The key findings of this research highlight growth opportunities in various aspects of the business strategy. Skills gaps are identified in IT and data analytics which present a barrier to IoT development and implementation and prevent the site from maximising value creation of existing data sources. A lack of metrics surrounding digitalisation and automation is



discussed, making it difficult to measure progress towards the company's own long-term digitalisation objectives. It is also highlighted that while most sites are quick to identify and replicate technological successes, less consideration is given to organisational innovation opportunities, and further cultural barriers prevent strategic-level innovation changes. These findings are discussed in detail, and various innovation projects, project management tools, and strategic changes are proposed to overcome these barriers to innovation. The assessments resulted in various outcomes that have had a direct impact on the company. Some of these proposed innovation projects have since been implemented at the company, and are also described in detail as well as their impact on the company. For example, one of the assessments identified an opportunity to create further value from AGV monitoring data and error logs. This resulted in a short-term data science project that provided on-site innovation teams with new insights into the root causes of AGV errors, as well as direct actions that could be taken to mitigate these instances on-site and reduce cycle times of material handling processes on site. Furthermore, these findings were also used to update the company's simulation models to improve the accuracy of AGV simulations when planning future replication of these technologies. This is but one example of the numerous projects that resulted from the research in this chapter.

The value of these assessments to measure and guide Industry 4.0 strategy as the company continues to expand into EV markets is being recognised by Ford Motor Company who have since adopted the assessment methodology as an internal tool to perform regular assessments at European manufacturing sites. The Simulation and Process Optimisation team at Ford's R&D Center in Dunton, UK has taken ownership of this tool with plans to perform further assessments at Valencia and Cologne manufacturing sites. In addition to full assessments, the company also uses the questionnaire as a means of measuring ongoing progress following the initial assessment by comparing responses change over time. This provides senior management with a new way to assess cultural changes at the factory level, something that had not yet been explored within the company.

## 3.2 Introduction

Industry 4.0 brings tremendous growth opportunities through increased automation, digitalisation, and business intelligence. Research shows that companies focusing on data-driven solutions deliver higher productivity than other companies [26]. However, managing the transition towards implementing highly automated data-driven manufacturing

solutions can be challenging. This is particularly true during the early stages of digital transformation as organisations may not yet understand the value add of emerging digital solutions. Prior research into measuring progress towards Industry 4.0 in the manufacturing sector finds a lack of clear metrics devoted to key Industry 4.0 technologies such as digitalisation, automation, cyber security, talent acquisition, and data analytics [177, 178, 179]. This lack of metrics and limited organisational knowledge makes it difficult to estimate the ROI of emerging Industry 4.0 technologies, and as a result, organisations are often reluctant to invest in digitalisation as the economic benefits of individual investments are often unclear [180]. This has resulted in a strategic dichotomy in Ford’s UK manufacturing sites, where long-term business objectives envision a highly productive and automated Smart Factory, while risk-averse investment strategies at the factory level prevent the high levels of innovation and digital growth required to realise this vision. The shift towards Industry 4.0 practices is not only driven by investment in new technologies, but also requires a changes in business strategy, organisational structure, and workplace cultures. This presents further challenges, long-term objectives in these areas are more difficult to define and progress is difficult to measure.

At Ford Motor Company, there are multiple business strategies to guide digitalisation and automation strategy. The ‘Factory of Tomorrow’ is Ford’s equivalent of the ‘Smart factory’ discussed in Chapter 2, and outlines the company’s technological vision of a highly automated manufacturing site. In the EU, Fords Powertrain Manufacturing Engineering (PTME) team oversees the launch of new production lines with a dedicated Industry 4.0 team to guide and support digitaliaiton and automation solutions such as IoT, AGVs, and collaborative robotics. To guide changes in workplace culture ‘Ford+’ presents a comprehensive list of the behaviours and values that are required to ensure the success of organisational changes. Despite these various high level strategies, visions, and business units, this research finds that Fords UK manufacturing sites are not well aligned with these objectives. Senior management report that while they are aware of these long term goals, they lack the low-level step-wise strategies and roadmaps to drive changes at the functional level.

These challenges are not limited to Ford Motor Company. Considerable research has been done to address these challenges across the manufacturing industry to support organisational change, with much of this research focusing on developing Industry 4.0 maturity assessments [12, 40, 41, 42, 43, 44, 45, 46]. Industry 4.0 maturity models are presented as a tool that businesses can use to quantify digitalisation progress and understand their

current maturity compared to a well-defined benchmark. Assessment tools also present opportunities for management to justify investment portfolios, improve project management, better understand strengths and weaknesses, learn from past mistakes, and review business strategies [181]. However, upon reviewing existing Industry 4.0 maturity tools and exploring how previous maturity models have been used at Ford Motor Company, it was found that existing tools are not well suited to deliver these goals in automotive settings. Existing tools lack clear guidance on the step-wise process of performing a self-assessment. For example, current maturity models lack information on which personnel should carry out the assessment, the scope of the assessment, what levels of management should participate, how many people should participate, and how long the assessment should take. Without this information, applying these assessment tools in large multinational organisations is challenging, particularly in organisations at the early stages of their Industry 4.0 transformation. Previous research has shown that companies tend to assume their Industry 4.0 maturity is higher than it actually is [46], a finding supported by our research at Ford Motor Company. Therefore, without clear, detailed guidance on correctly performing a given maturity assessment, results may be skewed by confirmation bias in favour of the optimistic perception of current readiness, leading to roadmaps that miss key growth opportunities. Furthermore, existing assessment tools cover a very wide scope, often aiming to be used by multiple organisations and industries. While this approach may be useful in providing high-level strategic guidance, it fails to provide specific low-level actions to enact change and develop a roadmap to Industry 4.0.

To address these challenges and address gaps in Ford’s Industry 4.0 strategy, this chapter presents an Industry 4.0 assessment tool aimed at automotive manufacturers. This tool addresses gaps in the company’s current strategy to deliver Industry 4.0 and drive digital growth at the factory level following Ford+ objectives. The assessment tool was developed by critically reviewing a wide range of existing maturity models and following current best practices in maturity model design and questionnaire design. The assessment is designed to be performed over 6 stages, with information gathered primarily through questionnaires and 1-to-1 interviews with employees and management. Assessors are provided with user-friendly supporting documents with detailed guidance on each process stage. Three assessments were performed at Ford Motor Companies UK manufacturing sites, including two engine manufacturing and assembly plants and one transmission manufacturing and assembly plant. An iterative design process was followed to develop the finalised design working closely with industrial partners throughout. The company is now using this assess-

ment tool to guide Industry 4.0 strategy across Fords European manufacturing operations. The chapter is structured as follows. Section 3.3 includes a review of existing Industry 4.0 assessment tools highlighting key research gaps in this field and providing a foundation for which to develop our own tool. This section also introduces the latest research in maturity model design and questionnaire design, in sections 3.3.1 and 3.3.1 respectively. Section 3.4 describes the methodology used to develop our assessment tool, as well as details of the iterative design process that led to the final application-ready version.

In Sections 3.5 and 3.6 the assessment tool is presented. Section 3.5 includes a description of the 6 stages that assessors should follow to perform an assessment. Each stage includes information such as required personnel, estimated timelines, scope, legal requirements, and other key considerations for industrial applications that previous assessment tools have failed to consider. This section also includes references to supporting documents such as questionnaires, scoring tables, and workshop guides.

Section 3.6 includes a detailed description of each of the 11 areas considered in the assessment, summarising the industry benchmark for each respective area. This section is also intended to be used as a source of further information for assessors to refer to in the scoring process to assess their findings against the current Industry 4.0 benchmark defined by the current state-of-the-art literature. These 11 areas of assessment are split into 2 general focus areas: 'Manufacturing Production', and 'Strategy, Organisation, and Culture'. Section .2 presents the results from three assessments carried out at Ford's UK manufacturing sites. Section ?? presents the various innovation projects and other outcomes of these assessments and their impact to the company. Based on these combined findings, section ?? outlines the Industry 4.0 roadmap for Fords UK operations to guide further digitalisation and automation efforts. Finally, research conclusions are presented in section 3.8.

### 3.3 Related Research

As outlined in the conclusion of Chapter 2, Industry 4.0 is a new concept with wide-reaching scope for which academics are yet to agree on a specific definition. Industry 4.0 continually evolves as emerging technologies, practices and concepts are refined, evaluated and reassessed. Furthermore, Industry 4.0 is widely described as a vision that varies between companies based on size, available capital, and industries. Because of this wide-reaching scope, developing a single tool for any company to use to assess Industry 4.0 readiness is challenging, and different authors take various approaches.

Some assessment tools address the wide scope of Industry 4.0 by presenting highly generalised tools that any company can use to provide limited, high-level direction towards digitalisation and automation [12, 42, 44]. In their 2010 paper, Steenbergen et al. describe these types of assessments as fixed-level maturity models, which typically cover approximately 5 key areas of assessment [182]. This approach is beneficial as an assessment can be performed by a small team, or even an individual in a matter of hours or days, making it an approach accessible to a wide range of industries and companies. However, because this approach is so generalised, many areas of assessment will inevitably not apply to all companies, which can lead to confusion when performing the self-assessment and any resultant guidance will be limited, and subject to interpretation. Steenbergen et al. argue that while many of these fixed-level tools are presented as a means of developing a roadmap to guide future improvements, these tools fall short of these promises [182]. This idea is supported by our own research, which finds that when Ford Motor Company has tried to use generalised tools such as these in the past, the tools have had to be adapted to such an extent that the comparison with the benchmark is no longer valid, and no actionable outcomes were identified. This is discussed further in section 3.3.1.

Steenbergen et al. also discuss the second category of maturity models with a narrower scope, referred to as focus area maturity models. These assessments are specific to a particular functional domain within an organisation and are more useful for identifying specific actions to make improvements [182]. Since this research was published, a wide range of maturity models have been introduced that blur the lines between Steenbergen's description of fixed-level and focus area maturity models, however, in general, those more narrow in scope tend to provide clearer guidance to develop a roadmap to increase the maturity of any given area [182].

An alternative approach is to develop an assessment methodology that can be tailored to the desired scope and Industry 4.0 vision of the company that is being assessed. Only 2

such assessments have been identified in the literature, both produced by private companies for which limited information is available [40, 41]. Assessments following this approach will take considerably longer to perform due to the added complexity of defining the Industry 4.0 vision, and providing specific guidance on multiple areas of operations and how to overcome existing barriers. This requires the team of assessors to develop a strong understanding of the Industry 4.0 vision of the respective departments, as well as support from senior management to commit these resources and review findings. Some external providers that offer assessments of this type dedicate 3-5 months to gathering this information [40], while others provide self-assessment guidance on how to effectively gather this information internally through workshops and guided questionnaires [41].

An example of a fixed-level tool is the Warwick University Industry 4.0 Readiness Tool developed in collaboration with Crimson & Co and Pinsent Masons adopts an intuitive self-assessment approach that requires the assessor to review tables of varying readiness criteria for various aspects of business and rank these aspects from 1 to 4 [12]. The assessment covers 6 different areas of business including products and services, manufacturing and operations, strategy and organization, supply chain, business model and legal considerations. Each of these sections is in turn, broken down further into sub-dimensions resulting in a total of 37 different areas for a business to rank its Industry 4.0 readiness. The tool itself is simple to use and includes the results of 53 other assessments performed on a wide range of companies for assessors to compare the results of their assessment to other similar companies [12]. However, a major challenge of this tool lies in its simplicity. The sub-dimensions in this tool each include single-sentence descriptions of the requirements for a company to achieve level 1, 2, 3 and 4 readiness. The vague nature of these descriptions makes it challenging to accurately rank a business as a whole as there are likely to be some areas of the business that meet level 4 descriptions, while others that are yet to achieve level 1. This is particularly true for Small and Medium Enterprises (SMEs) for whom many of the sub-dimensions are irrelevant [183]. Removing sub-dimensions in the assessment makes it difficult to assign scores and compare maturity with the benchmark. Researchers have addressed the difficulties of Industry 4.0 assessment in SMEs and presented alternative tools specifically for these smaller organisations [183], and [184]. Despite its limitations, the Warwick assessment tool has been highly cited and other studies have since adopted this style of self-assessment [185]. In contrast to the Warwick assessment tool, Engineering USA present an Industry 4.0 assessment that is designed to be more flexible, tailored to the client's 3-5 year vision of Industry 4.0 and designed to be carried out over a 3-5 month

period [40].

Engineering USA is a private company that offers Industry 4.0 consulting services. Although this assessment was developed to be provided as a private service, and therefore does not share the full details of the assessment process, it does contain useful information on their general methodology, outlined in their white paper and website [40]. The assessment is designed to maximise production metrics such as quality and throughput through the use of digital tools, as well as focusing on the effective utilisation of data on the manufacturing floor, and integrating these data throughout different levels of the organisation. Furthermore, emphasis is placed on digital manufacturing technologies such as simulation, IoT, process monitoring and control, and system integration. From the information provided, the assessment does not appear to address other key aspects of Industry 4.0 outside of production settings such as human resource management, management practices, and supply chain. This approach is extensive, requiring significant time and personnel commitments by both parties to plan, scope, write up, and review findings to deliver a clear actionable roadmap.

A tool that sits in between the simple and user-friendly Warwick assessment, and comprehensive Engineering USA assessment is the Acatech Industry 4.0 Maturity Index [41]. Acatech is a national academy in Germany which specialises in delivering science and engineering advice to policymakers, businesses, and society through independent research. Acatech's tool for assessing Industry 4.0 readiness is described in a six-stage "Maturity Index" focusing on 4 main areas of the business: Resources, Information Systems, Organisational Structure, and Culture. The process is split into 6 stages designed to be carried out step-by-step over several years to guide companies towards implementing some level of autonomy and self-optimisation in production settings. Similar to the Warwick University assessment, Acatech's tool provides a single, generalised document that any company can use as a self-assessment to support the development and integration of digital technologies in production. Acatech takes a highly qualitative approach, describing in detail their vision of Industry 4.0 and how each stage of deployment will create value for the company using real-world examples and case studies to demonstrate how these stages can be applied in various environments.

The six stages towards Industry 4.0 maturity described in the Acatech tool are related to data strategy and include: Computerisation, Connectivity, Visibility, Transparency, Predictive Capacity, and Adaptability. As described by these stages, the main focus of this report is on the effective collection, integration, and utilisation of data. The report

also includes limited guidance on the necessary approaches to change management in this journey towards flexible production environments, highlighting the need for agile project management strategies, structured communication channels, revised training programs, and collaboration.

While Acatech's report is the most comprehensive tool reviewed in this section, there is a large amount of detail that is not included such as initial assessment questionnaires, workshop guides, estimated time, and a list of required personnel to perform the assessment. The lack of this information makes the maturity index difficult to replicate in industry.

Many maturity models give few details on the design methodologies used to develop their tool, with the exception of the DREAMY toolkit [42]. The Digital Readiness Assessment Maturity (DREAMY) toolkit was developed using the maturity assessment design philosophies presented by Bruin et al. [173] and incorporates maturity principles proposed by Carnegie Mellon University in their Capability Maturity Model Integration (CMMI) framework [186]. The CMMI maturity model was originally designed to assess the maturity of a businesses software development practices against 5 levels of maturity, with a particular focus on maintenance, flexibility, and innovation and has been widely used by maturity models focusing on assessing the maturity of software process capability [116, 182, 181]. The DREAMY toolkit re-purposes the five-levels of CMMI to provide a generalised ranking of a company's digital readiness by evaluating both the technological and organisational aspects of the business. Similar to the original work by CMMI, the DREAMY assessment also focuses on identifying growth opportunities in process control and maintenance.

The DREAMY maturity lacks detail on the specific process a user should take to perform the self-assessment and what aspects of a business's operations should be considered in the assessment. Instead, the DREAMY assessment focuses on addressing an important research finding from their literature review, in which they state that the majority of prior maturity models fail to provide detail on the theoretical basis and methodology that was used to develop the tool.

Another popular and highly cited tool in recent years is the Industry 4.0 Maturity Model [43]. This maturity model for assessing Industry 4.0 readiness and maturity of manufacturing enterprises covers a wide scope, with a total of 62 areas of assessment grouped into nine core dimensions: Strategy, Leadership, Customers, Products, Operations, Culture, People, Governance, and Technology. A self-assessed questionnaire is used to rank each of the 62 areas on a scale of 1 to 5, e.g. *"On a scale of 1 to 5, one being strongly disagree and 5 being strongly agree, do you use a roadmap for the planning of Industry 4.0 activities in*



*your enterprise?”.*

The assessment notes that the model’s accuracy is largely dependent on the participants having some understanding of Industry 4.0 to understand the questions. This type of ambiguous questioning is found across many Industry 4.0 maturity models and goes against the best practices of questionnaire design outlined by Kronsnick et al., which is discussed further in section 3.3.1 [187]. Schumacher et al. discuss how the ambiguity of key terms and other knowledge gaps between the assessor and participants can be addressed through either external consulting or group sessions prior to the questionnaire [43].

The Industry 4.0-MM provides a high-level description of the 5 levels of Industry 4.0 capability by presenting a range of technologies and best practices. Similar to the Warwick university assessment tool, without a clearly defined scope and guidance to apply this assessment assessors may unintentionally cherry-pick examples from more innovative areas of a business, particularly in larger organisations, which may make it difficult to identify opportunities to further develop such topics within less mature areas of an organisation.

Assessment tools such as the TDWI assessment that are more narrow in scope tend to give more consideration to supporting materials, such as personnel requirements, user guidelines, ISO documentation [44, 188].

The TDWI maturity assessment uses a series of questions to rank five general areas of assessment: Organisation, Data Infrastructure, Resources, Analytics, and Governance. This assessment has a narrow scope, focusing specifically on these five areas within the context of business intelligence. By narrowing the scope, the assessment ensures that those taking ownership of the tool will likely have the required domain knowledge to act upon any opportunities identified in the process.

The TDWI maturity tool improves upon other assessment tools by providing some information and guidance on who should take the assessment, stating that it should be used by business and IT professionals involved in both new programs for analytics and older programs. This is beneficial when compared to other more general tools where findings might lay outside of the area of expertise of the assessor which can make it difficult to interpret results and develop a roadmap to guide organisational improvements. The TDWI approach also improves on other questionnaire-based assessments by including a short assessment guide to help assessors interpret results and highlight some potential areas for improvement.

Table 3.1: A comparison between the various existing approaches to assessing the Industry 4.0 readiness of a company. \*Limited information available.

Assessment Name	Areas of Focus	Assessment Method	Advantages	Disadvantages
<b>Warwick Uni</b>	Production BI Supply Chain Strategy Management	Self-assessment Scoring matrix	Simple, yet holistic coverage of the key themes of Industry 4.0	Lack of information in each of the criteria makes it difficult to distinguish between levels in many areas.
<b>Engineering USA*</b>	Production BI	External Certified Practitioner	Comprehensive review to define a roadmap to implement digital manufacturing solutions with long-term vision	Limited to manufacturing aspects of the business and doesn't address many other key themes of Industry 4.0
<b>Acatech</b>	Development Production Logistics Services Marketing	Third Party Assisted	Highly detailed supporting material with examples and case studies.	Key details for self-assessment not included i.e. questionnaires, required personnel
<b>DREAMY</b>	Design Production Quality Maintenance Logistics	Self-assessment	Detailed design methodology	Highly generalised. Lacks detail on the specific process to perform the assessment
<b>Industry 4.0 -MM</b>	Management Customers Products Culture Technology	Self assessment Questionnaire	Detailed scoring methodology	Lack of supporting materials, i.e. questionnaires, assessment process
<b>TDWI</b>	BI Analytics	Self assessment Questionnaire	includes an assessment guide to help assessors interpret results	Lacks detail on the specific process to perform the assessment
<b>SIMMI 4.0</b>	IT Systems	Self-assessment Scoring matrix Questionnaire	Includes supporting material	general and abstract questions

### 3.3.1 Assessment Tools Limitations

The tools outlined above assess different aspects of an organisation's Industry 4.0 maturity, some focusing entirely on specific Industry 4.0 goals like business intelligence, while others cover a much broader scope. Of these reviewed maturity models, the Acatech tool provides the most detailed direction to the reader on the specific steps to take to apply the tool and perform an assessment of a business [41]. However, these directions still lack key information such as which specialities should carry out the assessment, what levels of management should participate, how many people should participate in each of the stages, how long the assessment should take, ect. Furthermore, the methods used to score a plant's maturity are complex compared to other tools. Although this is designed to give a more tailored roadmap than other tools, this complexity and the lack of supporting information make this tool not user-friendly. Research by Schumacher et al. finds that maturity models tend to fail if they are too complex [43]. On the other hand, Steenbergen et al. discuss

that while simpler tools with generic fixed maturity levels can be useful to benchmark companies, they are not well suited to highlighting opportunities for improvement [182]. Overall, the evidence suggests that although Industry 4.0 maturity models are widespread, these models do not lead to improved project success or other quantifiable improvements within the respective organisation to gain a competitive advantage [181].

In the past 5 years, two instances were identified where Ford Motor Company has used maturity models to guide Industry 4.0 strategy: Warwick University Assessment, and TDWI. In both cases, the application of these tools failed to result in any actions taken. Upon investigating this with those involved in the assessment, two main reasons were identified for this: 1. The feedback from the tools was too general, 2. The assessment did not have input from senior leadership to support proposed changes. For example, the findings of the TDWI assessment found opportunities for growth in the 'Organisation' section. This was presented in a score out of 20, and the information in the assessment guide was limited and too general to draw identify specific actions. As a result, the assessment relied heavily on the experience and knowledge of the assessor to interpret this score and identify actions through their own extensive literature review. This process is not only time-consuming, but the action items identified required input by persons in more senior management positions than suggested by the TDWI guidance. For these reasons, it is critical that senior management who have the power and influence to enact some level of organisational change, however small, are included at some level in the assessment process. This involvement is required to ensure the assessment is performed within a clear scope for which changes can be implemented, as well as delegating key tasks within the relevant teams to deliver this.

Similar results resulted from using the Warwick Uni assessment at Ford Motor Company's Bridgend Engine Assembly plant. The tool was again found to be too general to provide a clear roadmap for change and also failed to address key areas of interest, a finding which supports research by Steenbergen et al. [182]. To overcome this, the assessor attempted to add their own areas of assessment based on the design principles of the tool, as well as removing several areas that did not apply. This resulted in biased and skewed scores and the results of the assessment presented to senior management were very high compared to the Industry 4.0 benchmark, despite our own assessment identifying multiple areas where on-site technology and organisation fell below the current Industry 4.0 outlined in Chapter 2. In this instance, the lack of guidance and direction of the Warwick University tool led to a clear selection bias and confirmation bias by the assessor. This is supported by previous

research finds that companies have a tendency to assume that their Industry 4.0 maturity is higher than what is calculated using questionnaire-based maturity assessments [46].

## **Types of Assessment Tools**

In their 2005 paper titled 'Understanding the Main Phases of Developing a Maturity Assessment Model', Bruin and Freeze review several popular assessment tools and present a single methodology outlining the main phases of developing these models [173]. Since then, this methodology has been used by researchers to develop a wide range of Industry 4.0 assessments [42, 116, 182]. Bruin and Freeze argue that any maturity assessment tool fits into 3 general categories: descriptive, prescriptive, and comparative [173]. A descriptive tool is used to define the current state of a business, providing no means of improvement or relationships between the current state and key performance metrics. The Warwick university tool is an example of a purely descriptive tool. A prescriptive tool provides further insight into how the current state of a business relates to key performance indicators, highlighting which areas can be improved to deliver value. The Acatech tool is an example of a prescriptive tool, focusing on developing a roadmap to add business value through Industry 4.0. Finally, a comparative model compares maturity across industries, providing insight into the differences in business practices between industries and how this relates to value generation in disparate industries. A comparative model requires many assessments to be carried out across various industries to gather sufficient data to draw these relationships.

Bruin and Freeze suggest that these different types of tools also describe the evolution and design phases of a maturity model. As a descriptive tool is used to gather an in-depth understanding of maturity within a specific domain, this knowledge can provide guidance and insight and evolve towards a prescriptive model. As this model is improved it can be used across domains to become comparative. Each of these different types of tools has its uses. For example, a single comparative maturity rating is useful to benchmark and track progress across organisations, but in large multinationals, different business functions will vary in their organisation, strategy, culture, and technological readiness. Therefore, at the functional level, the same comparative tool may give very different results. This is a common challenge with generalised assessment tools where opportunities for localised improvements may be overlooked or roadmaps may only apply to certain functional areas [182].

## Questionnaire Design

Most assessment tools rely on some form of questionnaire as a primary information gathering source [189, 44, 46, 190, 45], often relying entirely on score-based questionnaires to provide final scores on various aspects of operations. Despite the wide use of questionnaires, few maturity models include details on the design frameworks or methodologies used to guide the design of these questionnaires. Without full details of the processes used to communicate the questionnaire to participants, and without the questionnaire itself, the accuracy of these maturity models is difficult to assess.

Only one data analytics maturity model was identified where the authors describe how their questionnaire was designed in which researchers studied 4 other questionnaires from related maturity models to develop their self-assessment [45].

In 2018, Kronsnick et al. published 'Questionnaire Design', a book detailing best practices on designing questionnaires for research purposes and has since been highly cited by researchers [187]. Kronsnick et al. state that questionnaires should use simple, familiar words, avoid ambiguous meanings that different respondents may interpret in different ways, and avoid general and abstract questions [187]. Of the reviewed maturity models for which questionnaires are openly available, some do not follow these best practices [46, 44]. This is especially true for fixed-level maturity models for which questionnaires are highly generalised so they can be applied to a wide target audience across multiple industries. Research also shows that as they become increasingly fatigued, respondents are more likely to select answers presented early in a multiple-choice list which can impact the reliability of a questionnaire [187]. The effects of fatigue become more prominent in lengthy questionnaires or as questions become more difficult [187]. However, this can be minimised by ensuring items at the very beginning of the survey bear a strong connection to the topic and purpose of the overall study [187]. The TDWI questionnaire is an example of a very long and complex questionnaire with over 100 questions with many detailed multiple choice answers [44].

In practice, previous Industry 4.0 research has shown that high questionnaire response rates can be achieved by ensuring that the questionnaire is sent from an institutional email, with follow-up emails to chase up results after 2 weeks [191]. Frank et al. also suggest liaising with industrial partners to ensure the technical language of the assessment is aligned with that of the company, as well as getting multiple responses from each functional area that is considered to avoid common method bias [191].

A common finding among Industry 4.0 assessment questionnaires is the use of closed questions with ranked answers from 1 - very low / strongly disagree, to 5 or 7 - very high/

strongly agree [191, 44, 190] or multiple choice answers with detailed descriptions for each answer [44, 189]. Krosnick et al. find that the optimal number of points on questionnaire rating scales is most often a 7-point scale for visually administrated questionnaires [187]. Most assessment tools and Industry 4.0 readiness studies avoid using text input descriptive answers or rely on face-to-face interviews or workshops to gain further insights into questionnaire responses [181, 68, 48].

### **3.3.2 Summary of Key Findings**

To conclude, a wide range of Industry 4.0 maturity models are presented in the literature ranging from highly specific focus area assessments that focus on specific functional areas [44, 46], to highly generic fixed-level assessments that cover Industry 4.0 as a whole [189, 190, 42, 12, 43]. Regardless of scope, all of these tools consider both the adoption of emerging technologies and how these technologies should be coordinated strategically throughout the businesses to deliver Industry 4.0 objectives. This focus on strategy, organisation, and culture is a reoccurring theme throughout all tools referring to various themes discussed in chapter 2, including: vertical and horizontal integration, change management, shifts in culture, and agile management approaches.

Another common theme throughout existing tools are challenges of measuring short-term quantitative progress in areas relating to increased automation, digitisation, and data utilisation. Therefore these tools rely heavily on qualitative assessment often through self-assessed questionnaires. The use of questionnaires also ensures the simplicity of the assessment tool, as researchers suggest that maturity models tend to fail if they are too complex [43]. However, in striving for simplicity and generalisation, most existing assessment tools lack the key information required to replicate results and apply these assessments in industrial settings. This is particularly the case for the lack of information on the exact process that should be followed to ensure the right personnel are included in the assessment to ensure the scope of the assessment is appropriate to deliver actionable results. In some cases, researchers also do not include key resources such as questionnaires that are required to apply their proposed assessment methodologies in practice [41, 43]. Assessment tools such as the TDWI assessment that are more narrow in scope tend to give more consideration to supporting materials, such as personnel requirements, user guidelines, ISO documentation, [44, 188].

Overall, evidence suggests that although Industry 4.0 maturity models are widespread,

these models do not lead to improved project success or other quantifiable improvements within the respective organisation to gain a competitive advantage [181]. This finding is supported by the author's own internal research at Ford Motor Company where previous attempts to apply these assessment tools have been unsuccessful, and assessors suggested that increasing the involvement with senior management would have helped overcome this. To address these challenges, a new assessment tool should tread a careful balance between being general enough to cover a wide range of growth opportunities relating to technological and organisational aspects of digitisation and automation, while also being specific enough that the quantitative results can be interpreted by the user to provide clear direction for improvements within these areas. This balance must also be achieved through a simple, user-friendly design and should include supporting documentation to provide details on the exact process that should be followed to complete the various stages of the assessment. Based on the available literature, the best approach to deliver this is a flexible methodology that can be tailored to meet the goals and requirements of the organisation, similar to those used by private consulting firms [40]. The assessment should consider using questionnaires sent to a large and diverse sample, including multiple people in each functional area to be assessed with clear direction to the assessor on selecting the appropriate participants [191]. While they are a useful information-gathering tool, questionnaires should not be the only method used to gather information and assessors should also make use of workshops, site visits, and face-to-face interviews to discuss opportunities and previous successes relating to technology adoption, business strategy, and culture [41, 40, 191].

## 3.4 Methodology

The end goal of this research is to develop a tool for Ford Motor Company and other engine manufacturers to self-assess their Industry 4.0 readiness. The proposed tool should compare current technologies, and practices against a pre-defined Industry 4.0 benchmark so that areas falling short of this benchmark can be identified. Once identified, the proposed tool must also provide specific actions and guidance to senior management to help develop a roadmap to support the transition to increased digitalisation and autonomy. The proposed tool should be easy to use and adaptable so that it can be tailored for use at different Ford sites and other companies where the scope and Industry 4.0 vision may vary. With these aims and objectives in mind, a detailed prescriptive assessment tool is presented that focuses on qualitative assessment to be carried out over multiple stages. This research was carried out through close collaboration with senior management at Ford Motor Company to ensure the tool is well suited for industrial application.

The final tool is designed to be conducted in five stages of assessment, including a kick-off workshop, questionnaires, one-to-one interviews, assignment of scores, and a de-brief workshop. Each of these stages is discussed in detail in section 3.5 in which the full assessment tool is presented. This section describes the design process that was used to develop this proposed assessment tool and outlines the motivation for this 5 staged structure based on research findings from both the academic literature and through collaboration with industrial partners. It is discussed how this unique assessment structure overcomes challenges of previous maturity models to result in an assessment that results in a specific roadmap towards Industry 4.0.

### 3.4.1 Design Approach

The proposed assessment tool is developed using design principles by Bruin and Freeze, and involves splitting the design phase into 5 stages: inception, elaboration, construction, deployment, and adaptation [173]. These design stages are the foundations on which the stages of assessment are developed.

#### **Phase 1: Inception**

In the inception phase, an in-depth literature review of Industry 4.0 is conducted, as presented in Chapter 2 and section 3.3. A second shorter review was also conducted to identify



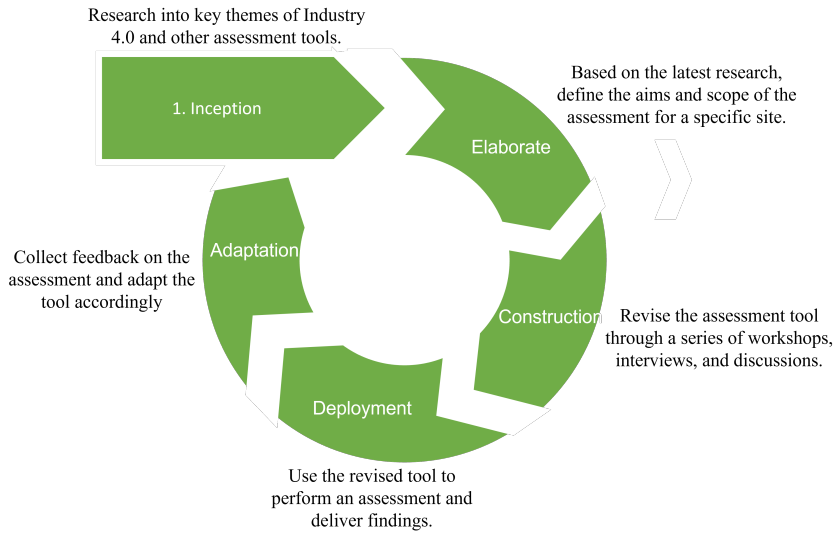


Figure 3.1: An iterative approach was taken when designing the assessment tool, using feedback from industrial partners to guide the design process throughout.

other Industry 4.0 assessment tools, discussed in section 3.3. The information gathered on common themes, research gaps, current state-of-the-art technologies, and proven maturity model methods are then considered in the initial development of the first iteration of the assessment tool.

During the inception phase, the authors toured two of Ford Motor Company’s engine manufacturing and assembly plants at which the assessment was to be carried out. The site tours were led by production engineering managers and other senior managers to understand the current technologies and practices on the line, identify any pilot innovation projects, and get some understanding of the long-term vision that managers had for different areas within the plant. In addition to these site tours, meetings took place with production managers and senior managers to understand the existing technologies, strategies and culture within the engine plant. These informal conversations provided considerable information that not only helped guide the design of the tool but also was used to a large extent in the final report following the assessment. This was also true for further conversations with industry partners in the construction and adaptation phases.

Given the value of these conversations in identifying Industry 4.0 readiness, it was decided early on to rely on one-to-one meetings with key personnel as our primary information-gathering stage of the assessment. While this method of information gathering is much more time-consuming than other assessments, this approach enables a large amount of specific information to be gathered to provide detailed insight into the differing cultures

between departments, organisational barriers, investment opportunities, and novel innovation projects identified during the interview process. This approach of using meetings and interviews has been successfully used by other maturity models that use meetings in selected parts of the assessment process [41, 68]. Schuh et al. suggest including group sessions prior to the information-gathering stages to address knowledge gaps, define key terms, and work with management to guide the scoring process using guided handouts [41]. Boavida et al. use interviews with experts in automotive R&D to understand trends in automation technologies in the automotive industry [68].

## **Phase 2: Elaboration**

In the elaboration phase, the knowledge gained in the inception phase is applied to develop a first iteration of the assessment tool. It was decided to incorporate multiple elements of different maturity models such as questionnaires, qualitative ranking tables, and group sessions into a single assessment that is designed to be carried out in multiple stages, each stage diving deeper into information gathered from those previous. Each of the stages include detailed supporting material to provide guidance and direction for assessors throughout the assessment process. These supporting materials also provide definitions of key terms with real world case study examples to clarify any ambiguity. In order to narrow the scope of the assessment it was decided to focus mainly on data analytics, digitisation, and automation.

To ensure this interview stage can be conducted within a reasonable time frame, it was decided to use a questionnaire to identify key participants for the interview stages. This selection process was done by reviewing responses to questions designed to identify participants who are involved in innovation projects, those who regularly use advanced data analytics tools, those that work with key Industry 4.0 technologies such as cloud computing, discrete event simulation, machine learning, IoT, predictive analytics, automation, change management. The questionnaire also includes questions relating to the participant's views towards innovation and automation to gain insight into aspects of workplace culture. Management may also identify key personnel to include in the interview stage in the kick-off meeting.

This approach differs from most maturity models. Our questionnaire is sent out to a wide range of employees throughout the site rather than a single individual tasked with performing the assessment. By increasing the sample size and diversity of the questionnaire, participants assessors can see how responses vary between departments and functional

teams and highlight specific opportunities for improvements through direct comparison of results. The questionnaire is developed by applying best practices in questionnaire design outlined by Krosnick et al. [187] as well as more specific guidance on Industry 4.0 questionnaires presented by Frank et al. [191]. Furthermore, inspiration is drawn from other assessment tools for which questionnaires are available [189, 44, 46, 190, 45].

### **Phase 3: Construction**

In the construction phase, the author worked closely with Ford Motor Company to review initial tool designs through a series of workshops and discussions, continually iterating the tool's design. This design phase focused less on developing and incorporating research findings and theoretical frameworks, but instead on the tool's usability in an industrial setting. For example, it was decided to include a planning stage at the beginning of the assessment process to provide time for administrative activities such as scheduling meetings with stakeholders and communicating the project to the relevant union representatives. A suggested participant list is also included in the supporting documentation that includes departments and individual job roles that should be included in the questionnaire and interview stages. This ensures that the assessment considers participants across a range of departments and throughout the management hierarchy. This list of suggested job roles is outlined in Table 3.2.

To ensure the assessment covers both technological and strategic aspects of Industry 4.0, the assessment is split into two main areas: 'Manufacturing Technology' and 'Strategy, Organisation and Culture'. These two areas are broken down into multiple sub-dimensions addressing key Industry 4.0 topics to assess what technologies are being implemented on-site to deliver digitisation and automation and how the skills of the organisation's workforce are being utilised to create the greatest value from these technologies. To score these areas, a similar 4-level description to that presented by Warwick University was used, as it was found to be the most user-friendly [12]. A detailed supporting document is also provided to elaborate on each sub-dimension to provide assessors with direction and guidance to assign scores through case studies and supporting research.

In addition to the planning, questionnaire and interview stages, it was also decided to include two workshop stages at the beginning and end of the assessment that would consist of senior leadership and other relevant stakeholders. This ensures the involvement of senior management throughout the assessment process while minimising time commitments. These workshops also provide an opportunity for senior leadership to provide direction on

specific areas to focus assessment efforts, i.e. investments in new technology, organisational changes, and workplace skills shortages. This is not to say these are the only areas that assessors should focus on, but alignment with senior management and key stakeholders on specific areas of interest and ongoing strategies increases the likelihood that assessment findings will result in actions taken. This addresses the limitations of other assessment tools that need to specify who should take ownership of the assessment findings or suggest how these findings should be used to enact organisational change. A diagram summarising each of these stages of assessment is included in Figure 3.2.

#### **Phase 4: Deployment**

Once the tool is refined, we enter the development stage in which a full assessment is carried out and our findings are delivered in the required format. This begins with an assessment of Bridgend Engine Plant (BEP), an engine manufacturing and assembly plant that produced two engine families. This plant was chosen as these two lines vary significantly, one being a brand new installation with some of Ford's most advanced manufacturing systems, while the other was decades old and approaching the end of its operation. The contrast of the two lines provided an opportunity to test the maturity model in identifying known investment and growth opportunities in the older lines, as well as exploring state-of-the-art manufacturing technologies and the extent of which current organisational practices and cultural norms were suited to create the greatest value from these systems. Further details of the deployment is included in section 3.5.

#### **Phase 5: Adaptation**

Upon discussing the results of the BEP assessment findings with senior management and department managers various changes were made to the tool in order to streamline the assessment process, make information gathering more efficient, and ensure the results were presented in a suitable format for senior management to use to guide future changes. A discussion on suggested tool modifications is included in the de-brief workshop in which all workshop participants are encouraged to critique the methodology in order to make improvements for future assessments elsewhere in the organisation. The author notes that limited information on the actual implementation of the proposed changes will be available at this stage and therefore the conversation should be focused on the assessment process and not the proposed roadmap. The assessors should take notes on this feedback and

included a section in the report for future assessors to consider during their own respective planning phases.



Figure 3.2: A description of each of the six stages of assessment and the key tasks associated with each stage.

## 3.5 An Industry 4.0 Assessment tool for Automotive Manufacturers

This section introduces the assessment tool and describes in detail the 6 stages of assessment. This section also provides supporting documentation for assessors to use when performing an Industry 4.0 readiness assessment. The first stage of the assessment is the planning stage which takes place in the weeks leading up to the start of the assessment process and mainly includes administrative tasks. Stage 2 is the workshop, in which the designated assessors meet with senior management to define the scope and objectives. In stage 3, questionnaires are sent out to the relevant teams to gather key information. In stage 4, a series of meetings are scheduled to gather further information on any opportunities identified from the questionnaire responses. In stage 5, those leading the assessment assign scores based on the findings of all previous stages before providing a de-brief of these scores and findings to senior management in the 6th and final assessment stage. This assessment process is summarised in Figure 3.2.

### 3.5.1 Planning

Once a manufacturing facility is identified to perform an Industry 4.0 assessment, some planning and preparation are required before performing the assessment. Administrative tasks at this stage may include travel planning for site visits, project scheduling, and selecting the assessment team, and socialising the project with the senior management team. It is recommended that at least two assessors are required to deliver the assessment with some knowledge of Industry 4.0 and change management. The literature review presented in Chapter 2 covers many of these topics which may be useful to revise the current state-of-the-art.

For companies and organisations performing their first assessment, it's important to make sure the assessment process is aligned with GDPR (or equivalent) rules on fair data usage and personal information that are required to be gathered in the questionnaire phases. Consideration should also be given to ensure the assessment process is compliant with workers union requirements.

If a previous assessment of this type has been carried out by the organisation, assessors should request to read the appendix of the most recent assessment carried out. This will

include any details of proposed changes to the assessment tool made by previous assessors. This step is important to ensure the continued development of the methodology as new technologies and best practices are identified, as well as ensuring the tool is up to date with any changes in company strategies.




Figure 3.3: An timeline for the assessment to take place was proposed during the planning phase and agreed during the workshop.

### 3.5.2 Workshops

This assessment is designed to be extensive, covering 11 different aspects of manufacturing operations. When performing an assessment, time and resource constraints will make it challenging to provide highly specific low-level feedback in all of these areas to improve technological maturity, Industry 4.0 strategy, organisational structure, and culture. Therefore, a workshop session is included in the initial stages of the assessment. Workshops and group sessions are a popular approach in previous research to engage senior management and align the scope with current long-term goals [43, 41].

The aim of the Kick-Off Workshop is to brief the management team on the aims and scope of the assessment, identify particular areas of interest, and gain support in organising the questionnaire and interview stages. Including departmental managers and other senior leadership in the assessment process is critical to ensure any proposed actions can be reviewed and acted upon, as well as communicating requirements for wider organisational change vertically upwards.

The Kick-Off Workshop is an informal workshop chaired by the primary assessor and can be done remotely or in person and is estimated to take 1 hour. The following discussion points are suggested to encourage a discussion on current readiness:

- What comes to mind when we think about Industry 4.0?

- What are we excited about when thinking about Industry 4.0?
- What are we apprehensive about when thinking about Industry 4.0?
- What are the main opportunities for Industry 4.0 on-site?
- What is the current long-term vision of digitalisation and automation on-site? How are we currently measuring progress towards these goals? How could we improve this?
- In which areas do we lack metrics and how could we improve this?

In addition to understanding management’s current perception of Industry 4.0 readiness, this workshop allows the assessors to understand any limiting factors that may affect the assessment, such as unavailable personnel or time constraints, and adjust the scope accordingly. The suggested outcomes for the workshop are as follows:

- **Clarify Scope:** Leadership will already have some understanding of the maturity of different departments. Understanding this may help assessors focus their efforts on a particular area of the site from the start. Areas with the lowest Industry 4.0 maturity represent the greatest opportunity for growth.
- **Specific Objectives:** This workshop provides an opportunity to identify and review current Industry 4.0 strategies, or Industry 4.0 goals. For example, the site may be aiming to establish itself as a center of excellence in a specific technology, investing heavily in a specific emerging technology, or experimenting with new business models. Understanding these goals of senior leadership provides direction and specificity for the assessors. If specific goals are identified, assessors should use this workshop to understand what steps are being taken to realise these goals, how these goals are communicated on-site, how progress is measured, and how this progress compares to other sites.
- **Suggested Timeline:** The delivery time for the assessment largely depends on the number of assessors, the number of participants to be included, and their availability during the interview stage. To minimise this time, we suggest that senior leadership ensure that the requirements of the assessment are well communicated to all participants to ensure interviews are scheduled within a short space of time. For reference, an assessment carried out using the finalised tool included 54 participants who were



interviewed over a 3-week period. Two assessors working 3 days a week took 4 weeks to finish the report.

- Next Steps:** Following the workshop, questionnaires are sent out to all departments before the interview stage. Input is required from senior leadership to distribute questionnaire links to the participants. At this stage, key personnel to be included in the interview stage are likely to be already identified e.g. innovation leads, production managers, data analysts, and cloud engineers. These interviews can be scheduled immediately, and it should be decided how the remaining interviews are to be arranged and who takes responsibility for scheduling these meetings. A list of suggested job roles for various departments is presented in Table 3.2.

Table 3.2: The following suggested interview list was shared with senior management during the planning phase.

<b>Production and Assembly Engineering</b>	<b>Productivity and Forward Planning</b>
Senior Process Engineer, Senior Production Engineer, Senior Manufacturing Engineer, Production Team Manager, Senior Controls Manager, Tooling Engineer, Systems Engineer, Maintenance and Scheduling Organiser, Test Engineer	Engineering Change Coordinator, Launch Manager, Forward Planning Manager, Senior Productivity Engineer, Senior Industrial Engineer, Simulation Engineer / Manager, Process Optimisation Engineer
<b>Supply Chain and Inventory Control</b>	<b>Emerging Technologies</b>
Supply Chain Manager, Material Handling Manager, Inventory Audit Supervisor, Logistics Manager	Innovation Manager, IoT Engineer, Emerging Technologies Engineer, Data Analyst, Software Engineer, Senior Data Analyst
<b>Human Resources</b>	<b>IT</b>
Hiring Manager, Training Coordinator	Software Engineer, IT Manager, Senior IT Engineer, Network Engineer, Cloud Engineer
<b>Quality and Product Development</b>	<b>Finance</b>
Senior Quality Engineer, Quality Engineer, Quality Manager	Plant Controller, Finance Manager, Finance Analyst

### 3.5.3 Questionnaires

With the aims and scope of the assessment defined, assessors can begin the information-gathering stage by distributing the questionnaires. The questionnaire for this assessment is included in section 3 of the appendix. The questionnaire includes a data protection statement, followed by a series of questions on key Industry 4.0 topics, including automation, data analytics, data management, software usage, communication and collaboration, role and responsibilities, ongoing projects, ect. This information is quick to gather and provides useful quantitative information to support the qualitative feedback in the final report. Furthermore, with a sufficient sample, assessors can provide feedback on how cultures vary between departments.

A well-designed questionnaire is a useful tool that many maturity models have relied on to assess companies against existing benchmarks [43, 44, 45], however, in this assessment, the questionnaire is also used to prepare for the interview stage.

### 3.5.4 Interviews

As questionnaire responses are gathered, assessors should review each of the responses and identify participants who may be able to give insight into specific areas in the scoring matrix. Due to time constraints, interviews in this research are limited to 30 minutes, so assessors should also use the questionnaire responses to prepare discussion points and questions to gather information as efficiently as possible. Any additional questions to the participant can then be sent in an email after the interview, or follow-up meetings can be arranged. Interviews may be carried out either face-to-face or remotely.

Given that many participants will be interviewed, it is important for assessors to take notes throughout the interview stage to make information retrieval easier when writing the report. After each interview, it is recommended that assessors use the simplified scoring matrix to assign scores for each individual participant. When writing the report, these will allow assessors to quickly identify departments and individuals demonstrating high or low readiness. The simplified scoring matrix is included in the following section.

It is important to note that not all participants will be able to provide information on all topics on the scoring matrix. Throughout the interview stage, the 'Interview Record Sheet' can be used to identify areas of assessment where further information is required, and identify participants to fill these knowledge gaps based on the questionnaire responses. Assessors should also use the spreadsheet to note down examples of high maturity, low

maturity, ongoing innovation projects, and opportunities for innovation projects.

### **3.5.5 Scores and Report Writing**

When reviewing the scores of the individual interviews after the interview stages, there will likely be some variation in the results. Some departments may have higher maturity in certain areas of assessment compared to other departments. There may also be variations within departments. Assessors should identify these variations and include this information in the report to highlight areas of excellence as well as opportunities for organisational growth, technological innovation, and replication.

Scores are not calculated but rather determined through a qualitative assessment process. However, averaging the scores across all interviews may be useful to provide an initial guide for scoring. Using these initial scores, assessors should use the scoring matrix to determine if this score is appropriate or not for the whole site. In some cases, further information may be required that what is provided in the scoring matrix. Therefore, this chapter includes further guidance in section 3.6 where further detail is provided for each area of assessment including references to additional academic resources and case studies. In practice, even with this supporting information, it is unlikely that for a specific area of assessment, all areas of an organisation's operations will be perfectly described by a single maturity level. Therefore, assessors are recommended to select a score which they decide is the most appropriate fit for the site as a whole and use any discrepancies to highlight growth opportunities. For example, in instances where a higher score is given, but some areas of operations fall below this level, this highlights growth opportunities that can be included in the roadmap. Similarly, areas that demonstrate a maturity level higher than the final score can be highlighted as replication opportunities to advance other areas of the business.

The final report is targeted at the management team, and therefore assessors should use clear, concise language making use of bullet points, graphics, and visualisations of the results. If any modifications are suggested based on changes in company strategy, or new technologies, these changes should be included in the appendix. Additional documentation identified during the assessment process that may be useful for management to enact change should be included in a separate folder to be distributed with the report. Examples may include internal documents, process standards, new corporate strategies, and scientific papers.

### **3.5.6 De-brief**

Once the report is finished and distributed to the management team, the final stage in the assessment process useful-brief meeting. The same people who attended the initial workshop should attend the debrief to discuss findings and identify the next steps. This is an opportunity for assessors to expand on the findings presented in the report and answer any questions about the roadmap. Any additional documentation should also be presented and described. Potential innovation projects should also be discussed and if feasible, process owners should be identified or assigned. This meeting should also explore opportunities to improve the assessment tool. Assessors should request feedback from management teams on how the assessment was conducted and document any feedback to be used by future assessors.

## 3.6 Areas of Assessment

This section includes a detailed description of each area of assessment. Assessors may find this information useful when preparing interview questions, assigning scores, and identifying growth opportunities. This section is split into two main areas of assessment: 'Manufacturing Technology', and 'Strategy, Organisation, and Culture'.

### 3.6.1 Manufacturing Technology Readiness

Industry 4.0 is a technological revolution characterised by highly integrated digital manufacturing technologies, including Big Data analytics, machine learning, IoT, Cloud computing, intelligent robotics, [92]. This section provides detailed information to assess the current readiness levels of an organization's manufacturing technologies. This information can also be used to help identify examples of current practices that are aligned with Industry 4.0 goals. This assessment aims not to suggest a complete overhaul of existing technologies to align with the current Industry 4.0 standard but rather to identify the key opportunities for growth to support further digitalisation and automation efforts utilising the existing technologies. This should be achieved by maximising the value creation of existing data value chains through improvements in data collection, data integration, and data analytics. Depending on the organisation's current equipment readiness, some investment may be required to upgrade existing systems to enable these data requirements. Manufacturing and Operations are broken down into a total of 6 sub-dimensions: Automation, Equipment Readiness, Automation of Material Handling, Data Collection and Integration, Data Analytics, and Cloud Solutions.

#### Process Automation

Process automation considers both hardware and software automation opportunities. Hardware solutions will often be in relation to automating manufacturing production and assembly processes and the line. These opportunities may be identified during the site visit. When considering hardware opportunities, assessors should look for examples where operators perform highly repetitive tasks that could be performed by six-axis robots, collaborative robots, vision systems, machining systems, or other hardware solutions. These systems require high investment, however, they are well-established technologies in automotive manufacturing with clear standards to guide purchasing, infrastructure setup, and

Table 3.3: A simplified version of the scoring matrix for manufacturing technology.

Readiness	Level 1	Level 2	Level 3	Level 4
<b>Level Automation</b>	Some automated machines and processes. A large number of processes are still manual, with no plans in place to address this.	Widespread automation of manufacturing process with evidence of automation pilots in more complex areas i.e. inventory management, and quality assurance.	Automation across the majority of production processes, data management, quality assurance, product testing. workforce largely supports innovation and automation.	Flexible production processes are managed and controlled by high levels of automation with widespread data collection to deliver business intelligence.
<b>Equipment Readiness</b>	Significant overhaul of all systems and processes is required to meet industry 4 requirements.	Few systems have capability, and multiple instances of outdated legacy systems prevent widespread data integration. No plans to update or integrate these systems.	M2M across most machines and systems. Some legacy systems are still in place but workarounds have been developed to integrate these systems and further updates are planned.	The vast majority of machines and systems ready meet all future requirements of Industry 4.0 and are regularly viewed to explore further integration.
<b>Automation of Material Handling</b>	Manual material handling throughout the business.	Some automation in selected areas, but routing is linear and delivers no product flexibility. Low digitalisation of inventory management.	Material handling is mostly automated using AMHS. Inventory management is supported by IoT systems.	AMHS support flexible production through automated material handling that adapts to changing supply demands to deliver a high variety of products with high efficiency.
<b>Data Collection and Integration</b>	Multiple instances of manual data collection across departments with few plans to digitalise these areas. Data are not widely integrated.	Multiple instances of manual data collection and integration opportunities. These barriers are well understood, with work ongoing to automate and integrate data collection.	Comprehensive digital data collection throughout production and backroom processes with plans to address any gaps. Data are widely integrated, including those stored and collected on legacy systems.	Comprehensive data collection across the entire business with data integrated across Cloud platforms to enable real-time analysis and business intelligence.
<b>Data Analytics</b>	Data is rarely analysed other than for quality and regulatory purposes.	Data collection is very high, but most data are not well used to create value. Multiple barriers exist at departmental levels preventing data analytics.	Widespread data analytics throughout departments. Departments understand what data are not being well utilised with pilots in place to address these gaps.	Cloud-based data analytics delivers high levels of process automation and business intelligence. Multiple pilots are underway to address underutilised data.
<b>Cloud Solutions</b>	Some innovation pilot projects exist using advanced areas of the business.	Migration to Cloud-based services is part of the current strategy, with multiple examples of Cloud-based solutions in place across the business.	Cloud solutions are widely used to deliver real-time data analytics and high levels of data integration, automation and business intelligence. Work is ongoing to explore Edge solutions.	Both Cloud and Edge solutions delivered high levels of data integration, automation and business intelligence.

installation. These automation solutions lend themselves to delivering cost savings through headcount efficiencies and improving production performance metrics. Other hardware solutions include IoT devices that can digitalise current processes on the factory floor. This includes systems like handheld tablets or AR systems for improved access to data and information. These IoT hardware solutions can deliver a wide range of process improvements. IoT solutions generally have a lower initial investment, however, the ROI is often difficult to calculate as many benefits of these systems are not quantitative. For example, the IoT systems mentioned above include benefits such as improved workplace experience, improved data collection, and improved data access, improved ergonomics. To drive digitalisation and automation, businesses must have clear digital growth metrics to guide investment in these systems. This is discussed further in the Investments sub-section. Hardware automation opportunities to improve material handling, and inventory management through technologies such as AGVs, gantries, conveyors, and RFID are considered in the Automation of Material Handling area of assessment.

In addition to hardware solutions, process automation can also be delivered through software-based digitalisation and automation solutions. Examples may include automating reporting processes, automating data-driven Excel workflows, automating data transfer and data management between software, and automating maintenance scheduling using PdM and CMB. Software-based automation solutions often add value through time-savings and reducing non-value-added administration time. Software-based solutions can be challenging for businesses in the early stages of digital transformation as departments may not have access to the digital skills in the relevant software to identify, prioritise and deliver these solutions. The Internal Robotic Process Automation (I-RPA) activity is presented as an opportunity to support this digital growth. I-RPA is an exercise to be carried at the department level in which teams review the current digital workflows carried out regularly and evaluate the potential time and cost savings opportunities and establish the feasibility of automating these processes. I-RPA is discussed in further detail in section ?? .

For both hardware and software-based solutions, the available research suggests that an effective solution to overcome organisational and technical barriers of Industry 4.0 is to explore short-to-medium term innovation pilots to implement human-in-the-loop innovation [70, 36, 161, 107]. Human-in-the-loop innovation focuses on socially sustainable solutions driven by an aim to improve workplace experience. It has been shown to deliver high economic returns as successful solutions not only deliver short-medium term cost savings and productivity improvements but also contribute to long-term digital growth through

building organisational knowledge and establishing a workplace culture that supports digital growth and seeks out innovation [38]. For examples of human-in-the-loop innovation to deliver automation in manufacturing, refer to Chapter 2, Section 2.5.2.

In level 1 organisations, many machines and processes will be controlled using automation including machining processes, material handling and gantries, however, a large number of processes are still manual with no plans in place to address this. In level 1 organisations, both management and employees are resistant to change. At Level 2, assessors should find widespread automation of manufacturing processes with evidence of automation pilots in more complex areas such as inventory management and quality assurance. Management demonstrates a clear drive to automate manufacturing processes but this is mainly motivated by cost reduction and focus on headcount reduction to deliver this. Because of this, employees are more likely to resist further automation and innovation in their respective workstreams. Management's vision of automation does not extend to administrative tasks and other off-line manufacturing areas. Level 2 organisations are likely to find that poor data management and limited access to data and analytics skills make it difficult to extract and gain insights from more complex sources to deliver business intelligence. As a result, analytics may be predominantly descriptive rather than diagnostic, predictive or prescriptive. Examples of Level 2 automation practices include instances where simple repetitive tasks are performed such as data being manually transferred between spreadsheets to analyse data or to generate reports on a weekly basis.

Level 3 organisations should demonstrate a good understanding of the value of data throughout all departments with the ROI of various automation projects and pilots understood with multiple pilots in place to automate existing processes. Data-driven decisions should be used to automate most instances of production processes, scheduling, data management, quality assurance, and product testing. In more challenging cases, opportunities to develop predictive and prescriptive automation solutions should be identified with pilots in place. Some instances of human-centered automation with lower levels of employees more open to the idea of automation in their workstreams. In Level 4 organisations, flexible production processes are managed and controlled by high levels of automation with widespread data collection to deliver business intelligence. Management maintain a human-centric view of automation throughout the organisation with multiple examples of Human-In-The-Loop automation identified. The social sustainability of automation processes should be a key consideration when beginning new pilots. As a result, level 4 organisations will find that most employees throughout the organisation support automation and view it as an oppor-



tunity to improving workplace experience in addition to the economic benefits to the wider company.

## **Equipment Readiness**

Manufacturing solutions will continue to evolve as new processes and systems are developed, and systems become increasingly integrated. To ensure businesses can maximise the value creation of integrated data-driven solutions, outdated machines and systems that present barriers to the digital value chain should be well understood with actions in place to mitigate the impact of these barriers. Organisations should also be quick to recognise emerging technologies within the industry and make efforts to better understand the potential ROI in these technologies and how they can be best utilised within the existing business strategy.

Examples of outdated systems may include software or hardware. For example, outdated machining systems with no capability to record the birth history of the outputted part, or systems where data cannot be integrated and must be recorded manually from human machine interfaces on the machine. Instances of outdated technologies should be identified and their impact on the ongoing digitalisation strategy should be well understood. If possible, workarounds should be put in place to mitigate these impacts.

A common example of such barriers found in automotive companies established before the 1990s are Mainframes and Legacy Systems. Many core business applications continue to run on these mainframe systems as they are often reliable, secure, and represent decades of investment. However, companies often have multiple mainframe systems across different areas of business running on outdated hardware, languages, and frameworks making them very expensive to maintain [192]. Other technological barriers include old manufacturing machines running on outdated operating systems, or manufacturing machines and systems where vendors have ended support. These examples present serious operational risks related to cyber-security as software updates may no longer be available making systems venerable to new exploits. Furthermore, mainframe systems often present major challenges to data access and data integration. Many of these challenges can be overcome by developing middleman software to automate workflows requiring regular interactions with mainframe systems. Solutions may range from simple data management software like PowerBI, to more advanced software development solutions in C# or Javascript. When identifying examples of outdated software and hardware, assessors should look for previous solutions to address these challenges as well as any ongoing works.

Level 1 organisations are reluctant to invest in new technologies and only update existing systems and processes when absolutely necessary. In these cases, a significant overhaul of existing systems would be required to meet industry 4 requirements. Level 2 organisations demonstrate some level of investment in updating equipment to deliver automation in selected areas of production. Multiple outdated systems are identified that present major barriers to Industry 4.0 objectives being met. In cases where assessors identify Level 2 readiness, assessors should make use of the 'Industry 4.0 Metrics' and Industry 4.0 Investments' sub-dimensions to better understand these barriers and identify opportunities to improve the current Industry 4.0 strategy. To achieve level 3, departments should not only demonstrate a good understanding of any current machines and systems that are outdated but also be able to quantify the impact that these outdated systems have on key business metrics. In instance where outdated systems cannot be updated, actions should be in place to mitigate the impact through use of alternative means. For example, modernising an outdated mainframe systems using cloud-based services to integrate these data-sources across the business.

To achieve level 4, the vast majority of machines and systems should be integrated to deliver high levels of business intelligence. Level 4 business should have an excellent understanding of how any outdated systems are impacting KPIs and other metrics related to Industry 4.0.

In instances where organisations demonstrate a level 3 or 4 equipment readiness, assessors should consider how the organisations current human resources and digital skill-sets are managed to create the greatest value from these technologies and how the current businesses strategy coordinates these assets to deliver Industry 4.0 objectives.

### **Automation of Material Handling**

By combining the above technologies in manufacturing environments through lean management practices, a new type of assembly line emerged known as the Flexible Manufacturing Systems (FMS). In a typical FMS, workstations are arranged more freely on-site in a modular arrangement, with AGVs automatically routing products to the required workstation. Tracking technologies such as RFID enable individual processes to be controlled by an automated production system to select the appropriate production process for the incoming part. This highly automated data-driven process allows production to be tailored to meet fluctuating consumer demand of multiple product families and deliver highly customizable product variants. As well as reacting quickly to market changes, FMS can react quickly

to changes in the local manufacturing environment, such as processing changes, material changes, or new product variants [87]. FMS results in less space, reduced operational headcount, improved inventory management, reduced lead times, and reduced manufacturing costs [27].

The concept of a fully integrated factory with ubiquitous integration of Industry 4.0 technologies throughout the entire business is widely referred to as a 'Smart Factory', although other terms such as 'Intelligent Factory', 'Intelligent manufacturing', 'Ubiquitous Factory' or 'Real-Time Factory' are also used. Research finds that organizations investing in Smart factory projects report increases of up to 12% manufacturing production, factory utilization, and labor productivity [27]. Despite the numerous benefits of FMS presented in the literature, there are several challenges in developing, implementing, and maintaining these systems. FMS requires high initial investment both in physical assets and personnel with the required automation and digitization skills to deal with the high complexity of these systems [27].

Researchers present novel machine learning approaches to support automation in FMS in the reviewed literature. Huang et al. propose a system to support production managers through a machine learning-based system to predict production progress for IoT factory environments [29]. A two-layer transfer learning approach using a combination of Deep Auto encoders and Deep Belief Network (DAE-DBN-TL) is trained on historical data using a bootstrap sampling approach. The proposed method was tested using real-world historical data over 15 orders with 1118 features. The experiment finds that the DAE-DBN-TL method achieves high performance ( $R^2 \hat{=} 87\%$ ) in predicting production progress based on historical production data. The author speculates that as well as monitoring and analyzing production progress, this model can also identify instances where production plans are executed incorrectly and support root cause analysis of these abnormalities. While the data set used in this study is very large, this method has only been tested in 1 real-world scenario. Little detail was included on the manufacturing site in which this was implemented. By comparing the DAE-DBN-TL across multiple locations and understanding its applicability in FMS environments, it is easier to determine the validity of these results [29]. Future research should further examine the human-centric implementation of production process prediction systems in FMS and consider the effects on the workplace experience of production managers.

This sub-dimension involves the in-plant management of material, plattents, and other stock. A key aspect of Industry 4.0 is the concept of Flexible Production Systems (FMS)

where high levels of equipment readiness, systems integration, and business intelligence allow production to adapt quickly to changes in supply and demand of multiple product variants [193]. Automated material handling systems (AMHS) manage how materials are transported and guided through the use of tracking technologies such as RFID, barcoding, and near-field communication, in combination with material transport systems such as conveyors, elevators, gantries, and Autonomously Guided Vehicles (AGVs).

Many manufacturing sites will fall short of the Industry 4.0 benchmark defined by FMS capabilities. To assess and guide progress towards this vision, assessors should identify the extent to which inventory management is digitalised and automated. Assessors should explore the following:

- How material and stock are ordered.
- How inventory and stock levels are updated when the material arrives on-site and when outgoing stock leave the site. Both instances should be supported by IoT scanners to automatically update databases.
- How inventory is stored and the extent to which stores are digitally managed.
- How stock levels are updated throughout production and the extent to which stock levels are updated automatically in real-time.
- How material is transported between stations. This should be delivered using gantries, conveyors, elevators with RFID or other tracking systems collecting all information on birth history.
- How material is transported between different lines. This should be done using AGVs where possible.

Throughout each of these areas, assessors should look for examples where inventory management is done using paper-based processes indicating digitalisation opportunities. Similarly, instances where IoT devices are used indicate higher levels of readiness and highlight opportunities to replicate solutions elsewhere on-site.

In Ford Motor Company, AGVs are a key part of the company's innovation strategy. These systems are well-established technologies with clear standards for supporting infrastructures and well-defined ROI. Assessors should explore. Assessors should identify opportunities where forklifts and tugs could be automated using AGV systems.

Level 1 businesses will have considerable opportunities to deliver lean improvements. Inventory management is updated manually with a considerable number of paper processes still in place. To achieve level 2, some digitalisation and automation are identified in some areas. RFID or other digital systems track the birth history of individual parts. Material flow is automated by conveyors and gantries on the lines, however, material flow between lines is still largely manual. Stock levels are updated in real time throughout production. Pilots in place to digitalise inventory management towards Just-in-Time delivery, and further digital solutions plans improved inventory management throughout production. In order to achieve level 3, inventory management is largely digital with material handling in key areas of production automated with AMHS to enable reconfiguration of the production line without disruption. AMHS delivers some level of product flexibility. Furthermore, Level 3 businesses should be able to respond rapidly to changes in demand. Opportunities to further automate material handling are identified with pilots underway. In order to achieve a level 4, business should demonstrate FMS production capabilities using advanced simulation and digital twin environments to manage and control production and in-plant logistic processes.

### **Data Collection and Integration**

As discussed in the previous chapter, in the modern automotive industry two of the most valuable assets to a company are data and human resources. Modern OEMs will already be reliant on sophisticated data collection systems standardised across the wider organisation to measure various aspects of production relating to products, machines, production line metrics, quality metrics, human resources, inventory, and enterprise resource planning systems [161]. Much of this data will be distributed across multiple databases, including outdated legacy systems. This sub-dimensions aims to identify opportunities for organisations to automate existing data collection processes as well as opportunities to integrate existing data streams into data lakes where there are increased opportunities for value creation.

Areas that are often more difficult to automate data collection are stations relating to quality assurance. Assessors should focus on these areas to identify barriers to data collection as well as exploring ways the organisation may have overcome these challenges. These barriers and challenges may be technological, cultural or both. Technological barriers may include outdated ICT infrastructures, lack of tracability at certain process stages, or physical challenges of integrating sensing technologies in complex processes. Cultural

barriers may include a lack of understanding of the importance of integrated labelled data, limited availability of the required digital skills to deliver data collection technologies, or difficulty of accessing data due to highly bureaucratic processes. Assessors should also enquire about any legacy systems that are being used on-site and what steps are being made to ensure these data are well integrated with the wider businesses.

Examples of level 1 and 2 practices may include instances where employees have to manually input data to populate excel spreadsheets when taking inventory, recording quality metrics, or when producing reports. What distinguishes a level 2 business is that these instances of manual data collection are identified and plans are in place to digitalise and automate these processes. In level 2 organisations, the vast majority of personnel should understand the importance of collecting high quality data in their respective work functions. Level 3 organisations are expected to have comprehensive data collection across most areas of the business, with management recognising the value of data and demonstrating a clear drive to automate and integrate data collection. A good identifier of level 3 businesses are examples where automated reporting has been effectively used to reduce the administrative burden on departments. Level 3 businesses must also have solutions in place to integrate data collected and stored on any legacy systems if those systems exist. Level 4 businesses have a very high level of data integration using Cloud-based platforms to integrate all relevant data to deliver real time analytics and business intelligence.

## **Data Analytics**

Digitisation is strongly linked to data management and analytics. Not only does effective data management and analytics provide a means on measuring current processes, but also sets a foundation for new data-driven digital processes to emerge [178]. Production teams should demonstrate a good understanding of the various data sources available to them and how these data provide insights into production metrics and KPIs. All departments should demonstrate a good understanding of what data are not being utilised, with pilot studies in underway to address these gaps. Assessors should aim to identify opportunities for data analytics where data are not well utilised and understand the barriers preventing this. If no technological barriers are identified, assessors may use the strategy, Organisation, and Culture section to help identify additional barriers preventing the widespread uptake of data analytics. All teams should have the ability to easily access the data relevant to their job function and have access to people with data analytics skills to support with delivering further insights using these data.

In level 1 businesses, data analytics is still widely performed using basic tools but is mostly focused on improving quality. There is a low understanding of the potential added value that could be generated from improving analytic capability and as a result, innovation efforts do not focus on addressing data analytics gaps. Level 2 businesses are identified as having very high volumes of data available, but the majority of these data are not well utilised to create value. This is not to say that no data analytics occurs on-site, but rather that there is no clear data strategy in place to support and measure progress of ongoing data analytics projects, nor is there strategy in place to support departments in pursuing their own data analytics projects. In level 2 businesses, employees may recognise at a departmental level that data are not being well utilised and identify barriers preventing this realisation i.e. lack of human resources, lack of data analytics skills, poor communication, ect.

Level 3 is distinguished from level 2 by having a clear data analytics strategy in place to identify opportunities for data analytics on-site and provide the necessary support for departments looking to increase data utilisation. Level 3 must also demonstrate that data analytics is not just limited to manufacturing production, but that these support structures extend to all departments and backroom processes. High maturity of Cloud-based solutions, automated reporting, and real-time automated scheduling are indicators that suggest level 3 readiness. That being said, without a clear data analytics strategy a business cannot achieve a level 3, regardless of any individual cases where advanced tools and analytics are applied.

Level 4 businesses should demonstrate a scaleable data analytics strategy that uses Cloud-based solutions to deliver high levels of process automation and control, as well as a high business intelligence. Employees in all departments in a level 4 businesses should demonstrate a good understanding of what data in there departments could be better utilised with examples of projects in place to overcome these gaps.

## **Cloud-based Solutions**

Cloud computing is a service based business model that provides online data infrastructures to efficiency and securely store IoT data as well as a scalable platform for Bid Data analytics. Many researchers highlight Cloud computing as one of the most essential technologies to deliver Industry 4.0 due to its inter-dependent relationship with IIoT and scalable Big Data analytics [194, 183, 65]. Cloud computing is a key technology that enables the highest levels of systems integration and automation found in Smart factory environments

[194, 195, 196]. At these high levels of Cloud technology maturity, organisations are able to deliver advanced service-oriented business models such as Cloud manufacturing. Assessors should explore the current Cloud-based solutions used throughout the site and any ongoing plans to migrate different aspects of manufacturing operations to Cloud models. On site IT teams should demonstrate high levels of understanding of these plans. All departments for which Cloud migration is ongoing or planned should understand how this will effect their future day-to-day work, and supervisors should demonstrate a good understanding of the future training requirements in IT and data analytics to create value from new data sources and digital toolsets available through Cloud platforms.

Migrating to the Cloud is often a considerable task, and therefore business often adopt a hybrid approach where more innovative areas of the business lead the way and processes are slowly migrated over time. Assessors should aim to understand where the company currently is on this migration journey. Level 1 businesses will be on the start of this journey, with pilots underway in more advanced areas of the business Cloud solutions but Cloud migration is not a key target for senior management. Companies yet to begin this journey with no implementation of Cloud solutions throughout any area of the business cannot achieve a level 1. To achieve level 2, migration to Cloud-based solutions should be part of the current business strategy, with multiple examples of Cloud-based solutions identified within the organisation. Assessors should be able to find multiple examples of where Cloud solutions are already being used to integrate data, deliver data analytics, or where opportunities have been identified and places to migrate are ongoing. For Level 3, Cloud solutions should be widely used throughout the organisation to deliver high levels of automation and data insights. Assessors should draw on examples in both the production environment as well as in backroom processes and how these combine to deliver business intelligence. Level 4 is reserved for businesses who are capable of delivering some level of service-oriented business models aligned with Cloud manufacturing.

### **3.6.2 Strategy, Organisation, and Culture**

This section focuses more generally on the organisational culture of the business. Assessors may find that multiple people within a department need to be interviewed before sufficient information can be gathered to assign a score. Although this section is focused on managerial practices, it is important that assessors assign scores based on responses by all employees throughout the hierarchy. Assessors should consider how employees responses



vary throughout this hierarchy to determine how effectively Industry 4.0 strategies are communicated vertically up and down the chain of command.

When reporting on this section, assessors should identify similarities and differences in strategy and organisation between departments, drawing upon specific examples from the Manufacturing Technology Readiness section to communicate these findings.

## **Industry 4.0 Metrics**

As companies digitalise existing processes they are often faced with the challenge of prioritising investments in digital solutions, as well as how to estimate the ROI of digitisation [178]. This is especially challenging due to a lack of clear metrics devoted to key Industry 4.0 technologies such as digitalisation, automation, cyber security, talent acquisition, and data analytics [177, 178, 179]. This lack of metrics is one of the main reasons why so many Industry 4.0 maturity models have emerged in recent years in an attempt to address this gap. However, maturity models are time consuming to perform making them more suited to informing long term strategy as opposed to short term measurement of digital growth. Prior research finds that companies are often aware of the importance of data to deliver process insights and data-driven decisions, but often lack metrics to measure the value potential of this information asset, as well as the efficiency of data value chains [26]. This is particularly difficult during the early stages of digital transformation as organisations may not have understanding of the ROI of digital solutions. To build this organisational knowledge, ongoing innovation efforts require detailed measurement of any required inputs and resultant impact of innovation projects. Inputs may include, Human resources, skills, training requirements, data requirements, external support, IT requirements. Project outcomes may include technical barriers, organisational barriers, ongoing maintenance requirements, time savings, new data available, cost savings impacts, impact on production metrics. These findings should be communicated across departments regardless of their success to support replication and understand common barriers to innovation.

A major requisite to successfully manage the transition towards increased levels of digitalisation and automation is to deliver training requirements, acquire new talent to upskill the workforce to maximise value creation of new highly integrated systems and the data they produce. This requires departmental managers to have a good understanding of future skill requirements and clear metrics to assess ongoing efforts to address skills gaps.

Table 3.4: Simplified version of the scoring matrix for Strategy, Organisation, and Culture.

Readiness Level	Level 1	Level 2	Level 3	Level 4
<b>Industry</b>	4.0 KPIs and other	A structured set of busi-	Some departments have	New business metrics
<b>Metrics</b>	business metrics are inconsistent, not reviewed on a regular basis, and do not relate to Industry 4.0 objectives.	ness metrics are well understood but metrics relating to Industry 4.0 objectives are not included at a departmental level.	targets or metrics related to I4 concepts, i.e. digital training, data collection.	have been adopted to understand the impacts of digitalisation and automation at a department and organisational level and are regularly reviewed.
<b>Investments</b>	The business is reluctant to invest in new technologies.	Multiple investments in Industry 4.0 but limited to manufacturing production areas. Management are reluctant to invest in long term innovation projects.	Management are open to long-term ROI innovation projects but investments are generally limited to manufacturing production. Barriers to I4 that require large investment are well understood.	Multiple examples of both short- and long-term Industry 4 investments across a range of departments with dedicated innovation teams focused on the development and integration of emerging technologies.
<b>Human sources and Digital Skills</b>	Re- Many employees on-site have little or no experience with advanced digital technologies or data analytics and have limited opportunities to develop such skills.	Most teams have the required skills to deliver Industry 4.0 goals although multiple barriers still need to be overcome. Reactive recruitment and limited professional development plans to develop these skills.	All areas of the business have good access to people with the skills required to deliver digitalisation and automation. Proactive recruitment approaches and well-structured professional development plans to address skills gaps.	Employees with leading edge digital and data analytics skills are found throughout the business with a focus on collaborative innovation.
<b>External Collaboration</b>	Collaboration is poor between other sites within the company. Little or no external collaboration helps drive innovation.	Collaboration is good between other sites within the company with fast replication of innovation cross the business, however, limited collaboration with external organisations.	Departments are open to cross company and external collaboration to drive improvements and innovation.	Departments are open to all aspects of cross company and external collaboration to help meet Industry 4.0 objectives and metrics.
<b>Communication</b>	Poor communication of company strategy throughout the business with the company operating in functional silos.	There is good communication between departments supported by IT systems.	There is some use of IoT to support departments working on similar projects and to communicate company strategy.	All necessary horizontal and vertical communication channels are supported with a wide range of IoT technologies enabling good collaboration and communication across all areas of the business's operations.
<b>Change Management and Leadership</b>	All levels of management are resistant to any change within the business.	Splintered internal cultures, with some departments more open to ongoing organization change efforts that others.	All levels of management understand the current strategy to deliver Industry 4.0 objectives, although some resistance to change by employees is identified.	Long term strategic goals are well understood throughout the business. Agile practices support a highly innovative workforce that supports ongoing organisational change.

This sub-dimension is somewhat dependant on the Industry 4.0 vision defined in the workshop stage, and aims to ensure that appropriate metrics are in place to regularly assess the ongoing transition towards Industry 4.0. For example, in the workshop phase at BEP, senior and executive management stated that one of their current goals for the medium- to long-term were: 1) to increased utilisation and value creation of existing data sources, 2) increased automation of production and business processes. When asked what metrics were currently being used to track progress of these goals, it was found that no quantitative measures were in place. As a result, this meant that the ROI of any successfully instances of automation and digitalisation of this vision were not reported on, making it difficult to justify future projects with longer term ROI. In this case, it was proposed that new metrics be introduced to measure the share of tasks that are digital, and the share of jobs that are digital, the volumes of data produced, and percentage of these data being analysed. By measuring how these measurements change over time, these metrics can be used to demonstrate digital growth to stakeholders as well as demonstrating organisational innovation through the introduction of new digitalisation metrics. The author recognises that some of these areas are challenging to measure, and further work is required by the company to explore exact methods to collect and analyse these metrics.

Level 1 organisations will have inconsistent metrics between departments with individuals in teams not having a good understanding of their departmental KPIs. Metrics in level 1 organisations do not relate to Industry 4.0 objectives. Level 2 organisations should have a clear set of business metrics across all departments that are well understood and regularly reviewed. However, these metrics do not relate to industry 4.0 objectives and multiple improvements are identified to improve measurement of innovation efforts. In order to achieve a level 3 readiness, management should be able to give examples of key metrics that are used to measure the impact of Industry 4.0 solutions and measure digital growth. Furthermore, department managers should understand areas of business where Industry 4.0 metrics are difficult to define and the plans in place to address this. However, areas are identified where progress towards some Industry 4.0 goals are not being measured. Further opportunities are also identified in communicating these Industry 4.0 goals throughout the site. In order to achieve a level 4, businesses should have clearly defined metrics to quantify the impact of digitalisation, automation, and any other key Industry 4.0 metrics outlined in the workshop stage. These metrics being effectively used to guide business decisions at both the department and organisational level and are well understood by all members of the workforce.

## I4 Investments

The journey towards the digitization and automation of production should be viewed as a continuous evolution as opposed to a revolution, one that is driven by reducing operating costs and improving productivity while humancentric change management approaches. This doesn't necessarily require a complete overhaul of outdated systems and extremely large investments. Instead, the Industry 4.0 philosophy is to optimise, digitalised and integrate existing processes by using technologies such as IoT, big data analytics, and Cloud-based systems to create further value from existing Industry 3.0 based technologies. This can create a challenge for businesses as many of these emerging digital technologies require significant long-term investment. This makes it difficult to estimate the ROI, and as a result, business are often reluctant to invest in digitalisation as the economic benefits of individual investments are often unclear [180]. Organisations must recognise at an organisational level that some level of risk must be accepted to enable innovation and experimentation with emerging technologies. Raj et al. suggest that in order to overcome this barrier, managers should have a clear digital strategy in place and prioritize investment in digital infrastructures to support future implementation of Industry 4.0 technologies and mitigate the risks of these systems failing [197]. Department managers must have a strong understanding of this long term strategy to support and direct innovation local innovation efforts to align with these goals. Sufficient risk capital should be available to support local innovation efforts.

This investment process can be split into 3 stages: local innovation, Centers of Excellence, widespread implementation. These stages are derived from work by Camuffo et al. based on data from Fiat Auto Co in the mid 1990's [33]. Stage 1, describes innovation investment in selected parts of the manufacturing process. At this stage, some aspects of innovation efforts may still be done in silos. Opportunities may be identified to improve project management, improve communication of siloed progress, and combine ongoing efforts through collaborative innovation projects to implement solutions at a larger scale. This stage is associated with the highest risk as technologies will not always be successful and the ROI of new technologies can be difficult to estimate. Organisations must accept some level of financial risk to provide room for experimentation and innovation. Following multiple successes of innovation in stage 1, and the site develops a wealth of organisation knowledge of the respective technology, the site will progress to stage 2. Stage 2 describes large-scale investment of a technology at the factory level such that the site becomes a center-of-excellence in developing and implementing the given solutions. At this stage, the ROI of

these technologies is better understood as well as clear standards and strategies to identify, prioritise, develop and implement these solutions. Growth opportunities at stage 2 include collaborating with other sites to replicate these processes. This may highlight wider organisational barriers where competition between sites may prevent replication across the business. Other barriers may be identified in bureaucratic processes that also present a barrier to rapid replication of technological successes. The final stage builds on knowledge gained from stages 1 and 2 where individual sites are responsible for implementation of the new technologies and organisational changes, resulting in technological and organisational homogeneity across the business [33]. These changes are to be carried out with caution, subject to thorough economic appraisal and evaluation to minimise risk [33].

In reality, these evolutionary stages will coexist as multiple investments and organisational changes are pursued at once. Technological changes, economic factors and market changes will vary throughout the evolution of an organisation that may effect these long-term strategies and lead to inconsistencies between newly implemented technologies and those previous. A key finding from Industry 3.0 is that companies must recognise that investment is not limited to manufacturing technologies, but also required to effectively management human resources to ensure they have sufficient personal with the right digital skill-sets to support the operation and future development of these systems [33]. Assessors should identify any particular areas where investments are being prioritised and understand at which innovation stage the site is. This can then be used to guide next steps to accelerate the innovation efforts and replication of successes across the company.

Organisations where management are reluctant to invest in new technologies should achieve a level 1 readiness. To reach level 2, Multiple investments in Industry 4.0 are identified, but generally limited to short term investments on the factory floor. Management are reluctant to invest in long term innovation projects and little consideration is given to investing in digitalising backroom processes. To achieve a level 3, a clear investment strategy should be in place that supports the transition to increased digitalisation and automation. This data driven strategy should support the development and integration of new technologies by placing focus on improved utilisation of human resources through well structured training options as well as external talent acquisition. Senior management in level 3 organisations should understand the risks associated with investing in emerging technologies with plans in place to better understand and mitigate these risks. A level 4 site can be described as a center of excellence for one or more Industry 4.0 technologies, with high levels of investment and knowledge in that area including understanding of the ROI of these investments. This

investment is not limited to technological assets but also assess the investment in human resources from internal and external sources. Management in level 4 organisations should recognise the importance of their human resources in creating the greatest value from their digital investments. Level 4 sites should also be an active part of a wider collaborative effort to replicate their technologies and transfer their knowledge to other sites.

## **Digital Skills**

In Industry 4.0, the two most valuable assets to a company are data and human resources. A key finding from research into Fiat Auto Company's automation strategy by Camuffo et al. found that pushing the automation of manufacturing processes to the technological leading-edge does not necessarily result in better quality, higher flexibility and improved efficiency [33]. Instead the key success factor for competitive manufacturing was instead reliant on good Organizational structures and processes, as well as the competencies and commitment of a firm's personnel. A successful transition to Industry 4.0 is therefore the result of organisational learning based on internal development, external acquisition, replication, and selection of technological know-how relating to key digital skills required to meet Industry 4.0 goals [33]. Without the right distribution and management of data analytics and other digital skill-sets throughout the company, organisations will fail to identify and act upon opportunities to enhancing productivity, innovation, cost savings, and business intelligence made possible by digitisation and data insights. businesses should anticipate the skill-sets that will be required to remain competitive in the rapidly changing manufacturing landscape.

This section focuses on assessing how effectively human resources are being utilised to create value through automation and data analytics as well as other Industry 4.0 objectives outlined in the workshop phase. Assessors may find it useful to consider any common barriers from the 'Data Collection and Integration' and 'Data Analytics' sub-dimensions. In organizations where technological maturity is low, examples of digital skills required to deliver processes improvements are likely to include advanced Excel knowledge, macros, visual basic, SQL, and automated reporting and data management tools. These softwares can be highly effective in delivering quick wins through integrating existing excel spreadsheets between departments and migrating from paper based data collection to digital systems. As higher levels of maturity are achieved, digital skills required to digitalise, automate, and integrate data may include: programming languages such as Python, software development, automated reporting software, knowledge of data mining or machine learning

within the context of Big Data, experience with Cloud technologies, and experience with advanced manufacturing systems such as robotics or additive manufacturing.

Assessors should also consider the extent to which digital skills are available from outside of the immediate business. For example, some manufactures adopt a centralised approach to data analytics, digital asset management, innovation, and other Industry 4.0 related areas of business. Assessors should consider how these centralised, decentralised, or hybrid business models are effecting the departments to deliver Industry 4.0 goals and ensure that these departments can effectively access the necessary resources and communication channels to work collaboratively and deliver digital solutions.

Employees in level 1 businesses have little or no experience with advanced digital technologies or data analytics and have limited opportunities to develop such skills or access these skills elsewhere. At level 2, Most teams have access to the right skills to deliver Industry 4.0 goals although multiple barriers still need to be overcome. To achieve level 2, businesses must have professional development plans in place that offer some opportunities to develop these skills. At level 3, all areas of the business have good access to individuals with a high level of digital and data analytics skills required to address gaps in data collection, automation, and integration. Well structured, personalised, professional development plans are in place to develop these skills. Level 4, is distinguished by a focus on on-site innovation made possible by leading edge digital and data analytics skills are found throughout the business. At level 4, employees have time to apply their skill-sets to explore novel solutions to problems, and work collaboratively to develop and integrate digital value chains across the organisation.

## **Communication**

Communication strategies are identified in the literature as a key mechanism to facilitate organisational change and should support the change management processes by providing clarity on the ongoing change efforts throughout the business [116]. Research shows that employees do not resist all organisation changes, only that which has not been well communicated or that which is perceived as psychologically or economically threatening [116]. During the interview stages assessors should understand the extent to which business strategy is communicated throughout and how participants opinions on the effectiveness of vertical and horizontal communication channels. Assessors should also consider how responses vary between different levels of management and identify any opportunities for improvement. In cases where employees are dissatisfied with communications, assessors

should understand why i.e. adequacy of communication, timeliness of communication, consistency of communication. Assessors should also explore how senior managers at different management levels perceive the internal communication process surrounding the strategic goals outlined in the workshop phase. Innovation successes should be well communicated throughout the business with good digital channels to support collaboration efforts to develop and replicate digital solutions.

The communication sub-dimension is not limited to internal communications but also extends to the efficiency of external communication across the wider business. Communication strategies should be supported with IoT systems such as web conferencing services, email, or similar communication platforms to enable engaging and effective communication with other sites to collaborate, share knowledge and support innovation. Management must recognise the importance of an effective communication strategy to deliver digitalisation and automation and must invest in the necessary infrastructures to support this in order to maximise the value of human resources and team performance. Instances where information is communicated through paper based approaches should be identified and solutions to digitalise these processes should be proposed.

In level 1 organisations, poor communication of company strategy is found throughout the business. Instances are identified where multiple departments or individuals have produced the same work, or where work has been delayed, due to lack of communication. Innovation efforts are done in silos with poor communication of these efforts between departments. At level 2, there is good communication between departments supported by IT systems in most areas, however, some opportunities are identified to digitalise communication channels. In level 2 organisations long term business strategy is not well understood in most departments. At level 3, IoT is widely used to support good communication across departments and with other sites. Company strategy and on-going organisational change is well understood at the departmental levels with clear communication of innovation efforts throughout the site and how these efforts relate to long term business strategies. At Level 4, All necessary horizontal and vertical communication channels are supported with a wide range of user friendly IoT technologies enabling effective collaboration and communication across all areas of the business's operations.

## **Change Management and Leadership**

Industry 4.0 brings with it a number of external forces for change, including market disruption, legislation, new technologies, and changing consumer behaviors. In response to



these forces, organisations must act holistically, making changes on multiple fronts in a continuous process of experimentation and adaptation aimed at matching the organization's capabilities to the needs and dictates of the changing environment. Organisational changes must be managed with a systematic change management process that considers both changes in culture and in organisational structure [116]. However, scholars argue that if strategic change is to succeed, changes should initially take place in the cultural beliefs and assumptions of the organization, thus leading to the cultivation of employee commitment in later structural changes [116]. This requires careful consideration of changing internal environmental pressures that dictate employee behavior, motivation, and performance of teams. Decentralised agile management approaches are regarded by scholars as one of the most effective management models in this regard, as it enables lower levels of management to be more embedded in their local environments and respond faster to changes [170, 171].

Other scholars highlight the need to consider adjustments in organisational culture and infrastructure requirements during the final stages of implementing business goals relating to digital transformation. Examples may include reviewing training, qualification requirements, or hiring strategies to meet the new demands of changes in high level objectives or organisational change [70]. Assessors should aim to identify the reasons for any resistance to change i.e. self-interest, fear of job security, group pressures, poor communication. Employees should be motivated to actively seek change, challenging the status quo and looking for opportunities to improve their own working environments through innovative means of digitalisation and automation. As mentioned in section 3.6.2, department managers must have a strong understanding of this long term strategy to support and direct innovation local innovation efforts to align with these goals. Furthermore, high levels of vertical downward communication are required to ensure the workforce are kept up to date with ongoing organisation and structural changes.

In level 1 organisations, resistance to change is identified across all levels of management and innovation and proposed changes are actively discouraged. At level 2, the assessment identifies splintered internal cultures, with some departments more open to ongoing organisation change efforts than others. No metrics exist to incentivise management to drive changes in their departments to achieve long term goals aligned with Industry 4.0 objectives. In Level 2 organisations, assessors should aim to identify departments where the current Industry 4.0 strategy has been well communicated and well received and explore what management practices and cultural differences exist between these teams and others

who are resistant to change. At level 3, all levels of management understand and support the current long term strategy to deliver Industry 4.0 objectives, although some resistance to change by non-managerial employees. Department managers and supervisors recognise this with change management strategies in place to address these cultural barriers. At level 4, long term strategic goals are well understood throughout all levels of the business. Agile management approaches support a highly innovative workforce that are empowered to support ongoing organisational change and digitalisation.

## 3.7 Applying the Assessment Tool

Three Industry 4.0 assessments were carried out using the proposed framework. The first assessment was carried out at Bridgend Engine Plant (BEP) in April 2019 using an early iteration of the tool. The second assessment was performed at Dagenham Engine Plant (DEP) in March 2020. The third was performed at Halewood Transmission Plant in October 2022. At each plant, a report was provided to highlight growth opportunities relating to Industry 4.0 and machine learning. Each report also includes a roadmap to develop and implement these solutions as well as suggesting opportunities for digitalisation and innovation projects. Given that the results of these assessments are highly specific to Ford Motor Company and lie somewhat beyond the scope of this chapter, these results have been compiled in section 2 of the appendix. This section provides a summary of the insight gained by applying the assessment tool at these locations, as well as the resultant outcomes and a critical evaluation of the value-add of the complete exercise. Details of the implementation requirements are also discussed such as resource and time requirements in order to support further replication efforts. This section also includes a summary of any shortcomings of the method and plans for improvement.

### 3.7.1 Summary of the assessments and resultant outcomes

In the initial iteration of the assessment, the aim of the questionnaire was to gather general information on the participant's roles and responsibilities and identify discussion points for the interviews. Although the questionnaire used in the DEP assessment was an early iteration, it proved extremely useful to the assessors to identify interview participants, preparing questions, and structuring the interview process. Furthermore, management found the quantitative data valuable to gain insights into key Industry 4.0 topics such as communication, data usage, and collaboration. For example, Figure 3.4 shows that when asked how important the usage of data analytics is to a participant's job, 90% stated it was very important. However, when asked how well their department creates value from this data, 33% of participants stated value creation was very good, 54% good, and 13% poor. These data are immediate indicators of opportunities for value creation through the further application of data analysis on-site.

Given the added value of the initial versions of the questionnaire, a more rigorous design approach was used to produce the final version, following best practices in questionnaire design outlined in Section 3.3.1. This finalised questionnaire is presented in section 3 of



Figure 3.4: Four examples of questionnaire responses related to data usage and availability.

the appendix, as well as a summary of the results of the survey. The questionnaire sent out at Halewood Transmission Plant was also found to be valuable to the senior management team in understanding data science knowledge gaps. By increasing the sample size at Halewood, further insights were able to be drawn on how responses varied between different levels of management. Figure 3.5 shows that those in more senior management positions have increasingly less access to data when required. This highlights an opportunity to improve business intelligence which is a major goal of Industry 4.0 [12]. This insight into the lack of data access for senior management became a major discussion point during the de-brief meeting, and led to direct changes in how data science training is delivered to senior management as well as the content of these training pathways. Further questionnaire findings related to innovation and data science also provided important quantitative information and data visualisations to support discussions in the de-brief meeting to amend existing data science hiring strategies. Within 2 weeks of delivering the report, the proposed changes were implemented by the senior management team.

Each interview was conducted by two assessors who independently gave scores based on the information obtained in each interview. At the end of the interview, these scores were compared. Any discrepancies in scores were discussed between the assessors, notes were




Figure 3.5: An example comparing how some questionnaire responses vary between different levels of management.

taken on why there was any disagreement and decided between them what the final score should be. This dialogue became an important part of the assessment to gather information on: the Industry 4.0 maturity, the effectiveness of the tool, and discuss future changes that may be required. For these reasons, it is highly recommended that at least 2 assessors are assigned to perform this self-assessment.

Many participants in the interview stage were quick to share their ideas on digitalisation and automation opportunities in their departments. In these instances, further questions would be asked on the value proposition, resource requirements, technological barriers, organisational barriers, and other considerations of the potential business case. It was found that many participants had promising innovation ideas related to Industry 4.0 technologies that demonstrated a clear business case. Many of these ideas had been proposed to supervisors and management but had not been pursued for three main reasons: Lack of data science skills, limited knowledge of resource requirements and potential ROI of digitalisation and data science solutions, and resistance to change.

Firstly, sites didn't have access to the required data science skills and programming skills to support the development of digitalisation solutions, such as PdM, Computer Vision, dashboard development, and low-level administrative process automation. This supports previous findings by Li 2021 and Bag 2021 who find that a lack of data science skills is one of the main barriers to Industry 4.0 [47, 48].

Secondly, management often lacked sufficient understanding of these emerging technologies which made it difficult to estimate the resource requirements and calculate the ROI for the proposed solution. These findings are aligned with prior research by Theissler 2021, who discusses how limited knowledge of the machine learning requirements at the managerial

and process levels is one of the main barriers to technology implementation [30]. Our findings suggest that this goes beyond machine learning solutions and also applies to complex data-driven solutions that require programming.

Finally, it was found that many managers were resistant to digital innovation solutions, with multiple interview participants stating there was a mindset of ‘If it’s not broken, why fix it’. Upon exploring the cause of this resistance to change, the author finds that one of the main reasons was due to the frustration of employees dealing with highly bureaucratic processes required to deviate from well-established company norms. Multiple instances were identified where bureaucracy presented a barrier to innovation, resulting in considerable wasted time for those involved with the respective proposals.

For proposed innovation ideas aligned with Industry 4.0 and where sufficient information could be gathered to form a business case, these ideas were compiled in the report to communicate the value proposition to senior management and central R&D teams. As a result, multiple new projects were identified during the interview stage that went on to be successfully implemented at the company. For example, one of the questions in the questionnaire asks participants to state to what extent they agree with the statement: ‘My department / team/ business function is good at identifying instances where data are not well utilised’. In cases where participants answered ‘strongly disagree’, the assessors would follow up on this during the interview stage, asking participants to identify which data they believed were not well utilised. When posing this question to a member of the logistics department, it was found that vehicle monitoring data and AGV error logs were not being utilised by AGV teams, off-site simulation teams, or external vendors to help address AGV downtime. Interviews found that on-site teams were aware that this data was available, but teams lacked the required skills in SQL and other programming languages that were required to analyse these data. These findings, as well as a proposed solution, were included in the report to senior management highlighting this as a potential innovation project. This project has since been completed and led to direct improvements in AGV cycle times on-site, as well as providing additional data on AGV routing to improve simulation models of these systems across the company’s European operations.

In addition to AGV data insights, five other projects were identified from the interview stages across all three sites. These projects include: digitalising Kanban in warranty departments, new management tools to support process automation, new training initiatives, Cobot voice command systems, and an anomaly detection solution for nut runner processes. The details of all six projects are discussed further in section 2.6 of the appendix.

To summarise, when reflecting upon the information gathering approaches, both the questionnaire results and interview findings provided valuable insights into various aspects of Industry 4.0 related to technology, organisation, and culture. In terms of delivering actionable outcomes at the end of the assessment, the questionnaire was a quick and easy approach to collect clear quantitative information on the organisational and cultural aspects of the factory. These data were valuable to support any proposed changes. While much more time consuming and resource intensive, the interviews allowed assessors to assess Industry 4.0 maturity in much more depth as well as gain insight into potential innovation opportunities at the department or process level.

Following the information gathering stage, the de-brief meeting was found to be the most important stage in the assessment process. By engaging with the senior management team and presenting a list of proposed actions, these items were able to be assigned to on-site personnel to take ownership of delivery. By having members of centralised manufacturing departments, any deliverables beyond the capability of on-site teams were also able to be communicated to the relevant teams to support any required next steps. This supports the authors' initial hypothesis that many assessments fail to deliver actionable results due to the lack of engagement of senior management.

In addition to these direct strategic changes related to Industry 4.0 strategy, management are also able to use these data to demonstrate cultural improvements on-site. By sending out the questionnaire every six months and comparing the results, management are able to measure how cultural values are changing over time. The resource requirements for this are very low, with the questionnaire taking 4 minutes to complete on average, with results immediately accessible through an online dashboard. This makes this a simple and cost-effective approach for senior management to demonstrate how corporate-level cultural objectives are being met.

### **3.7.2 Time and Resource Involved in Deploying the Method**

The following subsection details the time and resource requirements from the Halewood assessment. A team of 3 people was required to deliver the full assessment. A Ford Motor Company manager in the Powertrain Manufacturing Engineering was required to support with organisational aspects during the project planning phase, such as liaising with union representatives, scheduling meetings, and ensuring GDPR compliance. The remaining two team members included the author and a data science apprentice at Ford Motor Company

to act as the second assessor during the interview stage. Throughout the planning phase, the necessary arrangements were made via web conferencing via email. From project start, it took approximately 3 weeks to identify, contact, and schedule meetings with the required personnel prior to the de-brief meeting. Much of these 3 weeks was waiting for scheduled meetings and email replies, which required minimal time input from those managing these arrangements.

Questionnaires were sent out via email using Microsoft Forms. Of the 80 salaried staff at Halewood that were sent the assessment, a total of 54 responses were returned in total, all within 5 days. Most quantitative data were available through the MS Forms platform, however, an estimated two hours of additional work was required to explore further insights into the responses such as the comparison of responses between different levels of management shown in Figure 3.5.

Using the questionnaire results to identify discussion points and preparing interview questions was key to structuring the individual interviews to ensure a fast and efficient interview process. By scheduling and preparing for interviews while the questionnaire responses came in, the 30 interviews at Halewood were able to be completed by two assessors within a 2 week period alongside other full-time projects. While the interviews were only scheduled for 30 minutes via web conferencing, following the interview an additional 30 minutes was often required by the assessors to write up notes and discuss any findings. In some cases, only 1 assessor was able to be present, in which case the interview was recorded and reviewed separately by the absent party. In many cases, interviewees would send supporting documents, or email additional information relevant to any discussion which would also require time to review and disseminate.

Some key findings of the interviews were able to be compiled into a draft version of the report as the questionnaire responses and interview findings were being collected. In total, it is estimated that the final 26-page report took approximately 70 hours to complete, spread over a 5 week period. If necessary, additional web conferencing meetings were scheduled during this time to clarify any information.

Based on these findings, the full assessment is estimated to take between 8 and 12 weeks to complete. These findings were used to produce the proposed project timeline presented in Figure 3.3. This takes considerably longer than most other assessments in the reviewed literature, with the exception of the Engineering USA report for which the time estimate is similar.

This investment of human resources and time is not insignificant. However, given that this



is the only example in the reviewed literature where an assessment tool can demonstrate the impact in a manufacturing environment, the author argues that the proposed approach improves upon prior research in delivering Industry 4.0 objectives. Further work is required to quantify the long-term value add of the assessment framework and compare this with other assessment tools.

### **3.7.3 Qualitative analysis of the assessment outcomes**

The objective of the Industry 4.0 assessment tool is to support automotive manufacturing sites in advancing aspects of the Industry 4.0 strategy with a particular focus on enabling technologies of machine learning. In order to achieve this, based on findings from Chapter 2, the assessment ensures to engage management throughout the assessment process and, upon completion, provides a concise report and strategic roadmap to deliver quick wins and guide long-term strategic change.

Applying the proposed assessment tool using the methodology outlined in this chapter resulted in multiple outcomes to address a range of human-centric, technology-centric, and organisational growth opportunities at Ford Motor Company. Table 3.5 presents a categorised list of these outcomes. As discussed above, multiple short-term innovation projects have resulted from performing these assessments at multiple manufacturing sites. These short-term projects have already yielded tangible improvements in simulation models, production efficiency, and upskilling initiatives. Additionally, ongoing projects aim to achieve further cost savings and enhance product quality. Further information on the Ford-specific projects highlighted in Table 3.5 are discussed in section .2.6 of the appendix. Due to the time scale of this research and the impact of COVID-19 during the course of study, further work is required to continue research in the projects listed as long-term in Table 3.5 as well as quantifying the long-term value-add of the short-term solutions.

Although the long-term impact of the Industry 4.0 assessment at the factory level is difficult to quantify at this stage, this research has significantly influenced Ford's long-term strategy in other areas of operations. In 2022, based on the successful outcomes of the Halewood assessment, Ford's Powertrain Manufacturing Engineering Department adopted this assessment tool as part of its Industry 4.0 strategy to support innovation efforts across European manufacturing operations. Further assessments are currently being planned for other European manufacturing sites throughout 2023. By quantifying the impact of these projects identified through the Industry 4.0 assessment tool, the value of conducting such

assessments can be effectively communicated to other manufacturing plants. This provides further opportunities to promote digital growth on a global scale.

Table 3.5: A categorised list of outcomes of the Industry 4.0 assessment performed across all of Ford’s UK manufacturing sites. These outcomes are expanded upon further in section 2.6 of the appendix.

	<b>Human Centric</b>	<b>Technology Centric</b>	<b>Organisational / Strategic</b>
<b>Short-Term</b>	<ul style="list-style-type: none"> <li>Automation of Kanban to reduce non-value added admin time.</li> <li>Bite-Sized data science training for management.</li> <li>Updating data science hiring strategies at Halewood.</li> </ul>	<ul style="list-style-type: none"> <li>AGV Data Insights</li> <li>Nut Runner Anomaly Detection</li> </ul>	<ul style="list-style-type: none"> <li>Industry 4.0 Assessment Tool</li> <li>Quantifying workplace culture using questionnaire results.</li> <li>I-RPA - A quantitative assessment approach to identify and prioritise RPA activities at the department level.</li> </ul>
<b>Long-Term</b>	<ul style="list-style-type: none"> <li>New training initiatives related to emerging AI and data science tools.</li> </ul>	<ul style="list-style-type: none"> <li>Cobot Voice Command System.</li> <li>IoT scanners in logistics for inventory management.</li> </ul>	<ul style="list-style-type: none"> <li>Quantifying cultural change using repeated questionnaires.</li> </ul>

When considering the value of the Industry 4.0 assessment tool, management should consider not only the economic appraisal of engineering projects but also the added value of additional data access, knowledge growth, and the social benefits of the proposed solution. This chapter presents various metrics to evaluate organisational and strategic change that may be used to measure this impact. This addresses a major gap in the industrial sponsors’ current business strategy, thus providing a novel means of supporting future sustainable growth. Further work is required to explore the extent of these strategic gaps in other automotive manufacturers in order to better understand the potential impact of the proposed assessment tool beyond Ford Motor Company.

As mentioned previously, the assessment process identified multiple innovation projects by expanding upon questionnaire responses and collaborating with participants during the interview stage to develop robust business cases. While this approach has led to the development and implementation of numerous projects, it relies on assessors possessing a detailed understanding of Industry 4.0 and experience in emerging technology implementation for relevant solutions. Additionally, assessors need familiarity with existing technologies in the company to determine resource requirements and identify opportunities for replication. These knowledge requirements restrict the pool of individuals capable of successfully

conducting the assessment. However, given that this assessment methodology targets large automotive manufacturers, it is reasonable to expect that individuals with the necessary skills can be found within centralized engineering teams dedicated to exploring emerging manufacturing solutions. These limitations could also be addressed by further improving the structure of the interview process to provide clear guidance for a non-expert in Industry 4.0 technologies to support the identification of innovation opportunities.

As mentioned above, the full assessment requires a significant investment of time and resources, with a team of at least three working over approximately 8 and 12 weeks to complete. A large portion of this time is taken up by the interview stages which could be reduced considerably through improvements in methods to select participants and a more structured interview process. Future research may consider developing a flow chart of questions to structure the line of questioning and reduce interview time.

A further limitation of the research is that the results and findings presented in the chapter are specific to Ford Motor Company. While the proposed iterative design approach can be applied by any automotive company to refine the assessment tool, further research is required to understand the additional resource requirements this may require and the extent to which the proposed scoring matrix is applicable across organisations. Further research may also explore how this assessment methodology may be adapted for use by other industries beyond the context of automotive manufacturing, although this would likely require major changes to the scoring matrix.

One limitation of the assessment questionnaire pertains to questions concerning workplace culture, where there currently exists no benchmark for expected responses from sites with high Industry 4.0 maturity. This limits the insight that can be gained from these findings, and makes it difficult to highlight growth opportunities. To address this issue, it is suggested that the Industry 4.0 questionnaire be distributed to sites within the respective company known to have the highest Industry 4.0 readiness. The results from these questionnaires can then serve as an internal benchmark for comparison. Additionally, it is recommended that the same questionnaire be administered to all salaried employees every six months following an assessment. This would demonstrate the extent to which management has narrowed the gap between the current status and that set by the internal benchmark. By comparing responses over time, progress towards Industry 4.0 objectives related to workplace culture can be measured effectively.

## 3.8 Conclusion

The transition to Industry 4.0 and the adoption of machine learning technologies is a complex process, requiring organisational changes that challenge well-established business practices. Previous research explains the slow adoption of Industry 4.0 in the automotive industry due to the limited availability of skills, poor change management, and a lack of organisational knowledge [39]. Furthermore, researchers suggest that this problem is more complex than a skills shortage and that the Industry 4.0 paradigm at its core needs to be better aligned with social sustainability goals [39]. The research in this thesis supports these claims, and therefore, the maturity model focuses on social sustainability and human-centric innovation approaches. This research finds that Ford Motor Companies UKs manufacturing operations are in the early stages of the Industry 4.0 transition and, despite widespread opportunities for digital growth, the company lacks a clear digitalisation and automation strategy to guide technical and organisational growth at the factory level. To overcome these challenges, this chapter presents a methodology for automotive manufacturers to produce a prescriptive maturity assessment tool to measure ongoing progress towards Industry 4.0 objectives and develop a roadmap to guide progress.

This research argues that the approach of previous maturity models is too generalised, often targeted for use by a wide range of industries and organisations. These models provide limited, high-level guidance towards digitalization and automation, but many areas of assessment may not apply to all companies. As a result, the self-assessment process becomes confusing, and the resulting guidance is limited and subject to interpretation. There is currently no research to suggest that the roadmaps produced by these assessments have any impact on a companies Industry 4.0 strategy. Our own research supports these findings, as Ford Motor Company's experience with such generalized tools has shown that they require extensive adaptation, rendering benchmark comparisons invalid and failing to identify actionable outcomes.

The proposed assessment framework aimed to address shortcomings of previous Industry 4.0 maturity models by narrowing the scope. Firstly, the assessment is specifically aimed at automotive manufacturers. Secondly, the assessment is aimed to be delivered at the factory level, rather than across an entire organisation. This narrow scope, combined with a clear step-wise information gathering approach enables assessors to score various aspects of the business strategy based on primary evidence. This narrow scope also makes it easier to identify, prioritise, and deliver digitalisation and automation projects at the department level.

A critical review of existing maturity models is used to develop the initial design for the proposed assessment tool, as well as follow best practices in questionnaire design. The assessment tool is then refined through an iterative design process following multiple assessments at the sponsor company and addressing feedback from senior management at all three of the Fords UK manufacturing sites. In addition to providing a qualitative analysis of current progress, each assessment also provides quantitative measures that can be used to measure long-term growth.

In addition to assessing Industry 4.0 strategy and quantifying progress, the assessment tool is also designed to identify potential innovation projects to add value to current work streams in the short term. Multiple projects were identified from these assessments to drive technological and organisational innovation at the respective sites. Some of these projects have already yielded tangible improvements in simulation models, production efficiency, and upskilling initiatives. Additionally, ongoing projects aim to achieve further cost savings and enhance product quality. However, due to the time constraints of this research and the restrictions of the COVID-19 pandemic, it was not possible to thoroughly examine the long-term impact of all the proposed innovation projects. Ongoing work on the proposed innovation projects requires further research to quantify the ROI and evaluate their contribution to achieving Industry 4.0 objectives. Understanding the economic impact of these solutions is crucial to quantify the value of the assessment methodology and justify the time and resource expense. Future research should also explore the transferability of the assessment tool in different contexts and provide insights into how it can be utilized to facilitate digital transformation in various industries.

The findings of the three assessments carried out using the proposed framework highlights three main barriers to Industry 4.0: a lack of data science skills, limited knowledge of the resource requirements and potential ROI of digitisation and data science solutions, and resistance to change. This research confirms the significance of data science skills in delivering Industry 4.0, supporting previous research that identifies a lack of data science skills as a key barrier to implementing Industry 4.0 technologies [47, 48]. This emphasizes the need for organizations to amend current training and hiring strategies to acquire the necessary data science and programming skills to support the continuous development of Industry 4.0 solutions.

Furthermore, this research highlights the knowledge gap among management regarding emerging technologies, particularly machine learning. The limited understanding of these technologies makes it challenging for management to estimate the resource requirements

and calculate the ROI for proposed solutions. This finding aligns with prior research by Theissler 2021, who emphasizes the limited knowledge of machine learning requirements at the managerial and process levels as a significant barrier to technology implementation [30]. This research expands on these findings, suggesting that these barriers also apply to complex data-driven solutions that require programming knowledge. This broadens the scope of challenges faced by organizations in implementing Industry 4.0 technologies, further highlighting the need for comprehensive skills training and talent acquisition in areas related to data science.

Regarding the impact on the industrial sponsor, the strategic framework presented in this chapter and the resultant guidance to align the business strategy with Industry 4.0 goals is an important step towards addressing these skills gaps and increasing the analytical capabilities of the company. To continue digital growth and maximise value creation of existing technologies, leadership teams must recognise data as a major asset and should continue to develop their training and hiring strategies to ensure it has the necessary resources to digitalise processes and create value from new and existing data sources. As these skills are developed and teams have increased access to data through new Cloud-based solutions, considerable opportunities to explore machine learning solutions will be possible.

A key finding of this research is how the different approaches to information gathering resulted in different outcomes. The questionnaire findings were valuable in guiding strategic and organizational changes, such as amending hiring strategies, and training delivery. This was due to the quantitative information and clear visualisations that could be presented to management to support these proposed changes. The interview findings were more useful in identifying innovation and replication opportunities for Industry 4.0 technologies. While each of these information gathering stages added value in different ways, the de-brief meeting was the crucial aspect that ensured these findings were acted on. By presenting these findings in this workshop-style format and facilitating a conversation between senior management, various tasks to deliver the proposed opportunities were quickly outlined and delegated among the management team. The value of the proposed assessment framework has been recognised by Ford Motor Company which has since adopted this methodology as an internal tool to perform regular assessments at European manufacturing sites. The Simulation and Process Optimisation team at Fords' R&D Center in Dunton, UK has taken ownership of this tool with plans to perform further assessments at Valencia and Cologne manufacturing sites. In addition to full assessments, the company also uses the questionnaire as a means of measuring ongoing progress following the initial assessment by

comparing responses change over time. This provides senior management with new way to assess cultural changes at the factory level, something that has not yet been explored within the company.

Overall, this research contributes to the academic literature by identifying and validating barriers related to data science skills, knowledge gaps among management, and other organisational barriers to complex data-driven solutions in the context of Industry 4.0. These insights can inform organizations and policymakers in addressing these barriers, facilitating the successful implementation of innovative ideas and leveraging the potential of Industry 4.0 technologies for improved business outcomes.

# Chapter 4

## In-Process Anomaly Detection in Engine Assembly

### 4.1 Summary

Throughout automotive engine assembly, various in-process testing is carried out at key stages for quality inspection, such as leak tests, torque tests, DC nut runner and final testing. Each test produces either uni-variate or multi-variate time-series data, which results in over 22 Billion time-series data per year within Ford power-train plants making it infeasible to review these time-series data manually. Research has shown that automating anomaly detection in production settings can be challenging due to requirements for high-quality labelled data, advanced data analytics skills, and long-term investment in novel solutions for which ROI is difficult to determine [35, 178]. Because of these challenges, many anomaly detection processes in engine testing at Ford Motor Company rely on simplistic approaches, such as static limit thresholds, that fail to detect a large volume of anomalies. This research finds that based on vehicle warranty data, a successful solution could save an estimated \$3m per year, globally. Additional cost savings would be delivered by reducing repair teardowns on the line, increasing output, as well as providing further opportunities to implement anomaly detection in other sites such as vehicle assembly.

To address these challenges, based on the current state-of-the-art approaches to anomaly detection two methods are presented to improve anomaly detection for one of the most challenging in-process tests at Ford Motor Company, DC Nut Runner. This test represents as much as 10% of the warranty losses mentioned above. The process involves a human



operator using a handheld tool to fasten a series of nuts onto the engine, and an inbuilt torque transducer in the nut runner tool measures torque against time data which can be analysed in real-time. Some nut rundown processes have been automated whereby a machine controls the same tool to secure the nuts.

The variability introduced in the time series by the human error and manufacturing variability of the incoming parts presents novel challenges for anomaly detection that have not been explored in previous research. Most notable is the 'staging' problem. Because the nut runner process involves multiple stages, a human operator may pause for some short duration between stages, resulting in characteristic torque measurements being shifted in time due to these intermittent pauses. Not only does this introduce further variability to the data, but also removes the cyclicity and seasonality on which many methods rely to identify outliers. These characteristics of the time-series data have made unsupervised approaches difficult to implement.

In this study, multiple semi-supervised approaches are presented to overcome the staging problem and identify outliers in nut runner data. The presented solutions fall under two categories: semi-supervised clustering and semi-supervised forecasting. The semi-supervised clustering approaches apply dimensionality reduction before training a Gaussian Mixture Model on normal data to produce threshold regions to identify anomalies in near-real time. Three dimensionality reduction methods are compared: principal Component Analysis (PCA), t-Distributed Stochastic Neighbor Embedding (t-SNE), and Uniform Manifold Approximation and Projection (UMAP). While PCA and t-SNE are well-established dimensionality reduction approaches, UMAP is a relatively new technique that is emerging as a promising tool that has yet to be explored for application in manufacturing anomaly detection using real-world data. For the time series forecasting approach, a Long Short Term Memory (LSTM) model is trained using normal data, and the prediction error is used to identify anomalies. This approach has been successful in a limited number of studies in manufacturing applications but has yet to be applied to manual process data.

The success of these methods is demonstrated on two bespoke datasets collected from two nut runner processes, one manual and one automated. In developing these datasets, multiple challenges are faced relating to real-world data. A major gap was identified in the company's machine learning development strategy, as no standard method was in place to support data labelling tasks, a key stage in model development to produce high-quality training and testing datasets. To address this, a simple user interface was developed to minimise the time taken to label large amounts of time series data. This dashboard has

since been used to label over 15,000 time series data to support additional internal projects, including a recent project exploring new approaches to deal with disagreement in labelled production data.

This research also introduces the novel concept of the 'Anomaly No Concern' (ANC) category when labelling data. An ANC describes a time-series waveform where test engineers observe a clear process anomaly but based on their experience the anomaly can be explained and would not impact the quality of the part, meaning no action would be required. The ANC category was initially introduced to address the knowledge gap between test engineers and data scientists, however, throughout model development and testing it was found to give new insights into model performance and became a key consideration during model selection.

High accuracies are achieved in model testing by both the clustering and forecasting approaches. Both approaches achieve F-scores of 0.8, a tenfold improvement on previous attempts to identify anomalies in nut runner data. The success of the anomaly detection approaches presented in this study has resulted in these methods being included in an off-line trial at Dagenham Engine Plant. Details of the trial are discussed in detail as well as the implementation challenges for the two methods. These challenges include TensorFlow GPU implementation, dataset quality, data labelling limitations, and training data selection. These lessons learned to highlight the main opportunities for further research in nut runner anomaly detection and should be addressed before further trials. Regardless of these challenges, the anomaly detection solutions are found to outperform the existing solutions and plans are in place to include these methods in a live trial at Halewood Transmission Plant and Cologne Vehicle Assembly. If successful, this method will be rolled out globally as part of the company's ADAPT program to automate anomaly detection in engine manufacturing and assembly.

This chapter is structured as follows. In section 4.2 the research topics are introduced at a high level. Details on the nut runner process are presented and the current anomaly detection methods used at Ford Motor Company are described. In section 4.3 previous research into time series anomaly detection is reviewed, focusing specifically on semi-supervised approaches that lend themselves to our application and previous applications of LSTM, GMM, and dimensionality reduction approaches. The methodology is presented in section 4.4, which includes low-level details on the machine learning models compared in this study as well as the experimental setup and details on the metrics used to evaluate their performance. The methodology also discusses the challenges of collecting labelled data in

real-world manufacturing settings, and how these challenges were overcome by introducing the ANC category. The results of the experiments are presented in section 4.5 and the discussion is included in 4.6. Following the success of the development and testing, the anomaly detection methods were included in a trial at a Ford manufacturing site. The details of this trial are presented in 4.7 as well as the challenges of implementing and evaluating the proposed solutions. Finally, our research conclusions are summarised in section 4.8.

## 4.2 Introduction

Developing a system capable of detecting anomalies in production settings is challenging for several reasons. Access to labelled anomaly data is often difficult in production settings where there are often many potential failure modes, each of which is usually rare and difficult to interpret in time-series data [167]. Furthermore, there is a lack of publicly available datasets to develop and test anomaly detection methods in industrial settings.

Manufacturers must therefore develop their own training and testing datasets and solve complex processing and feature engineering challenges that require technical expertise in both data science and the target domain. Not only is this research and development time-consuming, but any given solution may not be transferable to other processes, even if the processes seem similar in nature. These challenges often make it difficult to estimate a ROI of such data analytics projects. As a result, the value of machine learning solutions is yet to be fully realised in the automotive industry, which typically focuses on short-term ROI projects.

Machine learning technologies are key enablers of Ford Motor Company's long-term vision of transforming its manufacturing sites into highly automated 'Factories of Tomorrow'. As discussed in Chapter 2, an effective approach to expand the use of advanced data analytics solutions is through collaborative innovation pilots between centralized R&D teams and manufacturing production teams. Such projects are important to deliver high levels of technological innovation by delivering human-centric solutions by combining the domain-level expertise of production engineers and the technical expertise of data scientists. This socially sustainable approach helps transfer knowledge of data science and analytics to production engineers to improve analytic competency and build wider organisational knowledge in understanding the business case to identify, justify, and prioritise investments in future automation solutions.

At Ford Motor Company, project ADAPT is a good example of a human-centric innovation pilot aligned with these sustainable objectives. Project ADAPT was introduced to improve quality in engine assembly by introducing an anomaly detection dashboard to identify process anomalies at test stations. Throughout the engine assembly line, there are various in-process and end-of-line tests to ensure the quality of the final product. Many of these tests use static process limits to identify potential fault modes. These limits are set by experienced testing engineers with considerable knowledge of the process and are reviewed and updated regularly manually based on recent test data. In many processes, this visual inspection of process time-series data through a series of dashboards is also an important step in identifying potential process errors. This method has been proven effective in many tests for which the data is clean, well-structured, and highly regular, and the failure modes are well-understood.

Process owners recognise that there are opportunities to improve the current anomaly detection processes, for which there are multiple inefficiencies. Firstly, this approach is not well suited to identify new, previously unseen anomalies where faults may occur within the specified limits. In these cases, anomaly detection is completely reliant upon visual inspection. Complex processes that result in highly variable test data require considerable human input to identify potential process errors. Secondly, current methods require regular tuning whenever the operating parameters of the test or machinery are changed. Automating these processes would reduce the burden on test engineers to evaluate and maintain the current anomaly detection methods. Furthermore, by using statistical approaches based on historical data, there is an opportunity to significantly increase anomaly detection rates in processes that exhibit high variability and deliver quality improvements by reducing false negative rates.

Nut runner is an assembly process in which anomaly detection is particularly challenging. The process involves a series of nuts being fastened onto the product, either by a manual operator or a machine. An inbuilt torque transducer in the nut runner tool measures torque against time data which can be analysed to detect process anomalies. The process occurs at multiple stations throughout the automotive engine assembly line involving various types of nuts, threads, required torque, process duration, and other process variables. The automated and manual processes produce highly variable data due to the manufacturing variation of the incoming parts and the staged nature of the rundown process. Because the nut runner process involves multiple stages, a human operator may pause for some short duration between stages, resulting in characteristic torque measurements



Figure 4.1: An example of a line worker using a DC nut runner tool in engine assembly.

being shifted in time due to these intermittent pauses. Similarly, an automated process may pause between processes for tool changes or geometrical differences between product variants. This staging can be observed in the torque time plots in Figure 4.2 where torque is applied at different stages of the process, separated by periods of 0 torque, which vary in length. Not only does this staging introduce further variability to the data, but it also removes the cyclicity and seasonality that many methods rely on to identify outliers. This high variability in both the normal data and the anomaly data makes traditional unsupervised clustering approaches such as one-class SVM and PCA ineffective, as anomalies are not always outliers. Figure 4.3 shows the first two principal components of two nut runner datasets, showing how not all anomalies are outliers, and not all outliers are anomalies. Reconstruction methods such as encoder-decoders are also infeasible due to the data being shifted in time at multiple stages, making it difficult to draw a probability distribution from initial data.

When considering the economic impact of the current process, quality issues result in significant economic losses associated with in-service warranty claims. Within Ford Motor Company, warranty claims traced back to manufacturing faults in engine assembly, an estimated \$1.3 million per plant per year is lost due to process anomalies being missed by




Figure 4.2: Datasets 1 and 2 with a random example of a single observation highlighted in red. Dataset 1 is a manual nut runner process with high variability. Dataset 2 is an automated process where the staging problem can be clearly observed.

current anomaly detection methods. Of these faults, as high as 10% are due to errors in nut runner processes. Based on the available warranty data, a solution to reliably detect process anomalies in nut runner has the potential to save \$130,000 in warranty claims per plant per year, and over \$3m across all 24 engine lines globally. If successful, there are further applications for nut rundown anomaly detection in vehicle operations that have been identified at the Cologne vehicle assembly plant to provide further cost savings. In addition to this economic impact is the additional impact on customer satisfaction, customer loyalty, and brand reputation.

Project ADAPT aims to address these challenges by developing machine learning algorithms to automate, or semi-automate, anomaly detection across multiple tests in engine assembly. Currently, a single unsupervised algorithm is used to detect anomalies for all processes. This approach uses PCA to reduce the dimensionality of the time series data and perform a cluster analysis using Density Based Spatial Clustering (DBSCAN) under the assumption that any noise points are anomalies. This has been shown to be successful on a range of end-of-line tests, outperforming the current static limit approach. However, the PCA method is ineffective at identifying anomalies in nut runner data.

In anomaly detection, some amount of labelled data is required to evaluate machine learning models, and because anomaly data are rare, the available labelled data are usually mostly normal data. Many machine learning solutions use these surplus normal data for training [30]. Anomaly detection approaches that train machine learning models using

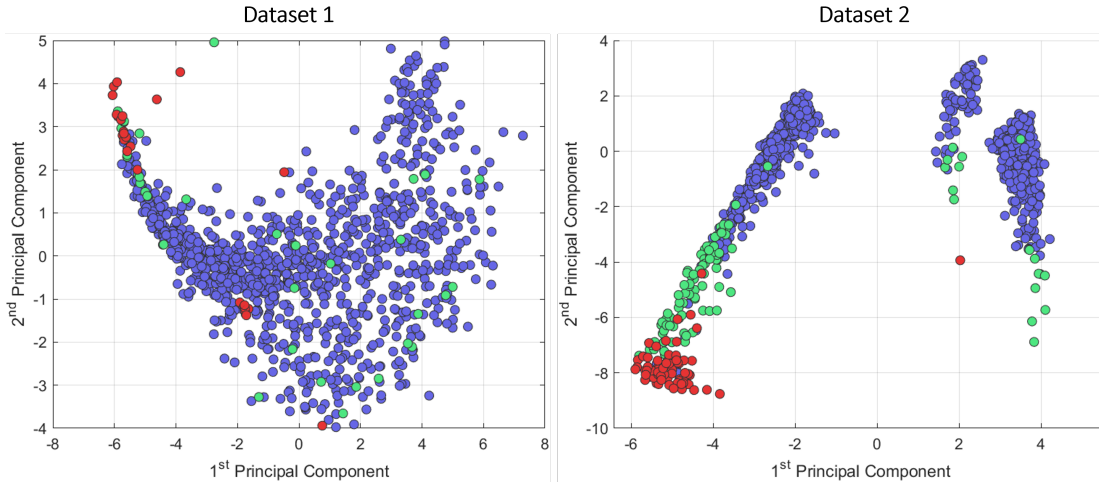


Figure 4.3: Labelled PCA plots for Datasets 1 and 2. 'Normal' data are shown in blue, 'True Anomalies' are shown in red, and 'Anomaly No Concern' are shown in green. Here it can be seen that for both datasets, not all anomalies are outliers, and not all outliers are anomalies.

only normal data are known as semi-supervised models, also called a "clean" approach or "clean-semi-supervised" approach [198]. In this chapter, two types of semi-supervised approaches are presented to identify outliers in nut runner data: semi-supervised clustering and semi-supervised forecasting.

Dimensionality reduction and clustering approaches are common approaches to identify anomalies in time-series applications [199, 200, 201, 202, 203]. Dimensionality reduction methods aim to represent high-dimensional data in a lower-dimensional space in order to visualise data in two or three dimensions and apply cluster analysis approaches that are more suited to lower-dimensional datasets. In this study, three dimensionality reduction methods are compared: PCA, t-SNE, and UMAP. PCA and t-SNE are well-established dimensionality reduction approaches, however, UMAP is a relatively new technique that is emerging as a promising tool but has yet to be explored for application in manufacturing anomaly detection using real-world data.

After applying dimensionality reduction, a GMM is trained using a semi-supervised approach. The GMM model is a common approach to clustering data that assumes the generative processes to produce the dataset can be described by a mixture of isotropic Gaussian probability density functions. By training the GMM on normal data, threshold regions can be defined, assuming any process that generates anomalies will fall outside of these regions and be identified as an anomaly.

For the time series forecasting approach, a semi-supervised Long Short Term Memory (LSTM) is trained using normal data, and the prediction error is used to identify anomalies. Machine learning-based forecasting methods such as LSTM have been proven to be powerful tools for time series forecasting in recent years, and perform better than traditional statistical techniques in more challenging cases where new or unknown anomalies may occur, or when there is no clear anomaly distribution [204, 205, 206]. LSTMs are a type of Recurrent Neural Network that has been widely used for time-series forecasting in industrial settings due to the ability of the architecture to retain both long- and short-term dependencies [207, 208, 209, 210]. For anomaly detection use cases, forecasting methods use some initial portion of the time series to predict the remainder of the process. If the prediction deviates significantly from the actual readings, an anomaly is assumed to have occurred.

### 4.3 Related Research

Anomaly detection is the process of finding, removing, describing, or extracting observations in a dataset that are generated by a different generative process than that of the majority of normal data [211]. Anomaly detection in time-series has been studied by data science researchers for over 50 years in various domains, including fraud detection [212, 213, 144], cyber security [214, 215, 216], stock market prediction [217], cardiology [211, 218, 219, 220], engine monitoring [221, 218, 219], fault detection and condition monitoring [199, 222, 223, 224], and manufacturing [209, 225, 226, 203, 31, 227, 228]. Approaches to anomaly detection vary greatly on the context of the task, however following several advancements in neural network architectures and computational statistics in the late 80's and early 90's, combined with the increased access to the required computational power to apply these methods, the majority of researchers have since focused on some form of machine learning to solve anomaly detection [229, 230, 231]. There are three typical approaches for anomaly detection:

- Supervised: Training data is labelled and includes both the nominal and anomalous data.
- Clean Semi-Supervised: Training data only includes nominal data, while test data also includes anomalies.



- Unsupervised: Training data is unlabelled and includes both the nominal and anomalous data.

Supervised methods frame the anomaly detection task as a binary classification problem (Normal vs. Anomaly) and use labelled data to train classifiers that distinguish between nominal and anomalous data. This can be effective in situations where the percentage of anomalies  $\alpha$  are high ( $\alpha > 1\%$ ). However, in most cases anomalies are typically very rare ( $\alpha < 1\%$ ) making supervised approaches infeasible as it is both difficult and time-consuming to obtain sufficient labelled anomalous data. Furthermore, the supervised approach makes the assumption that the distribution of anomalous data can be well-defined, and that this distribution can be used to train a statistical model [211]. This assumption is known as the Well-Defined Anomaly Distribution (WDAD) assumption [211]. In manufacturing, this assumption can be utilised to detect repeated machine failures for which the problem space is well understood and sufficient data are available to define the distribution. This is the theoretical basis for Six Sigma practices for which time-invariant data are modelled to fit a well-defined Gaussian distribution and if some measurement exceeds  $\pm 6\sigma$  from the mean, those instances are flagged as anomalous. Prior research finds that the WDAD assumption is rarely applicable in the real world as few approaches make the assumption that the anomaly and nominal distribution can be accurately modelled by the analyst [211]. This is especially true in manufacturing environments due to the increasing complexity and variance of data produced by modern manufacturing systems.

In cases where WDAD assumption does not hold, and the fraction of training points that are anomalies are very small ( $\alpha < 1\%$ ), unsupervised or clean semi-supervised methods can be used to detect outliers, although these methods may also fail if anomalies are not outliers or if the distribution of the nominal data has long tails [198].

### 4.3.1 Types of Anomalies

Anomaly detection of manufacturing systems deals with time-series data and requires different statistical approaches to those used on time-invariant data that assume constant variance and independence of variables. Time series data is defined as a sequence of observations taken by continuous measurement over time, with observations usually collected at equidistant time intervals [207]. Time series data can have properties such as trend, seasonality, cycles and level which can be used to make predictions on future trends and identify anomalies that deviate from the norm.

Much of the existing literature focuses on three types of anomalies in time-series data: point anomalies, collective anomalies, and contextual anomalies [201, 207, 205]. Point anomalies are instances where a single point in time deviates significantly from the majority of the time series. An example of a point anomaly in historic weather patterns could be a single day of heavy snowfall in British springtime. Point anomalies have been studied most extensively with most approaches making the assumption that anomalies are scarce and occur independently of each other [211]. Neural Networks [232], tree-based approaches [233], SVM [216, 234], and LSTM [235] have been successfully used to identify point anomalies. Collective anomalies are where multiple data points in the time series may be considered to be normal when analysed individually, but when viewed as a collective they demonstrate a pattern of unusual characteristics. Continuing with the weather example, a collective anomaly would be if the snowfall continues for multiple days. Contextual anomalies are cases where data may deviate from the majority of the data set but are dismissed as normal due to the context. Contextual anomalies are defined by two attributes [212, 211]:

1. a spacial attribute that describes the local context of a data-point relative to its neighbours,
2. a behavioural attribute that describes the normality of the data point.

Point et al. provide a detailed mathematical description of contextual anomalies, and how clustering algorithms can be used to identify contextual anomalies in a range of real-world and synthetic data [211]. A common example of contextual anomalies is described using credit card data [212, 213, 144]. For example, if an individual's credit card expenditure is significantly high over the course of a week in April it might be considered a collective anomaly and flagged as fraudulent activity. The same transaction behaviour the week before Christmas however may be considered normal behaviour given the context.

In the example of credit card transactions, we can see that there can often be an overlap between the different types of anomalies. Therefore, it is sometimes necessary to develop a solution that identifies all three types of anomalies. Hundman et al. demonstrate how LSTMs can be used to identify all three types of anomalies in a multivariate time-series dataset to identify spacecraft anomalies in telemetry data [205].

### 4.3.2 Dimensionality Reduction

The first step in any anomaly detection task is to use domain knowledge to extract meaningful features from the raw data using feature engineering techniques. These features can then be analysed using a wide range of statistical tools to highlight outliers, which are potential anomalies. The number of meaningful features a dataset has determined whether it has high dimensionality or low dimensionality. As the dimensionality of data increases, it becomes more difficult to draw relationships between these features. This not only requires more training data and more processing power to train models to learn these representations but also makes the trained models more susceptible to overfitting due to noise being present across all dimensions [236].

Dimensionality reduction methods aim to represent high-dimensional data in a lower-dimensional space to visualise data in two or three dimensions and apply cluster analysis approaches that are more suited to lower-dimensional datasets. The most common cluster analysis approaches that have been applied to anomaly detection in time series include: k-means clustering [237, 238, 193, 239, 240], Fuzzy C-Means clustering [241, 239], Gaussian mixture model [237, 242, 234], and hierarchical clustering [237, 193, 240, 107].

K-means and Fuzzy C-means clustering involve making initial guesses on the centroid position of a given number of clusters before applying stochastic approaches to iteratively optimise the centroid locations by minimising the distances to points that lie within each centroid's respective clusters. K-means clustering is a hard clustering approach in which each point is assigned to a specific cluster. C-means is a soft clustering approach that assigns individual probabilities to each data point so that data can be assigned to multiple clusters. Diez-Olivan et al. show how k-means clustering can be used for diesel engine condition-based monitoring by detecting anomalies in sensor data [239]. For CBM applications such as this, the normal operating conditions and the anomaly distributions can be well-defined, making cluster analysis a highly effective solution.

Gaussian Mixture Model (GMM) is a similar clustering approach that assumes that the process can be described by several sub-processes, each of which may generate a Gaussian component in the lower dimensional representation [237]. GMM is a probabilistic approach for which maximum likelihood estimation algorithms such as Expectation Maximisation are used for model fitting [237, 242]. Previous research has shown that GMMs perform well at semi-supervised anomaly detection in time-series data where the anomaly distribution is not known [243]. Amruthnath et al. compare unsupervised machine learning models to identify anomalies in machine vibration data for predictive maintenance. Of the various

clustering methods compared in the study, a combination of PCA and GMM was found to give the best result [237]. However, for this application, the normal operating parameters are well-defined, and only one fault instance was considered. GMM is often applied to analyse biometric time series as it is well suited to handle data with large sample distributions [243, 244]. Reddy et al. demonstrate how GMM can be used in unsupervised settings to identify outliers in network traffic data [245]. Reddy et al. apply a semi-supervised approach and discuss the importance of high-quality training data, as the model is sensitive to outliers.

In hierarchical clustering, the initial number of clusters  $K$  equals the number of data points. At each iteration, each point is merged with neighbouring clusters until a single cluster is formed. This bottom-up approach is called agglomerative hierarchical clustering and can also be performed in a top-down approach which is called divisive hierarchical clustering [107]. This process is then used to construct a dendrogram where branches are joined or split at a depth equal to the number of iterations at which those clusters were merged or split. The resulting dendrogram explains the relationship between all the data points in the system and can be horizontally sliced at any point to return the required number of clusters, where small clusters may indicate anomalous system behaviour [107, 237].

Dimension reduction techniques can be split into two main categories: Matrix Factorisation and Neighbour Graph approach. Matrix factorisation includes algorithms such as Linear Autoencoders, Generalised Low-Rank Models, and PCA. PCA is one of the oldest and most commonly used methods for dimensionality reduction across a range of scientific disciplines, dating back to work by Pearson in the early 1900s [246]. PCA uses the eigenvectors and eigenvalues of the dataset's covariance matrix to construct linear representations of the data in latent space. These linear representations are called principal components, and those with the highest variance capture the most information of the original data and can be retained for further analysis or plotting while components with low variance can be discarded. PCA has been widely applied in a range of time series anomaly detection tasks by researchers over the past few decades [202, 247, 238, 239, 237]. One limitation of PCA is that if the correlations between features are non-linear or unrelated, the resultant transformation may result in false positives or fail to draw any useful relationships [236]. Various tools and add-ins are included in common industrial toolsets, such as Microsoft Excel, that make PCA accessible to engineers. As discussed in Chapter 3, researchers propose that PCA should be included in industrial training programs such as Six Sigma to address the increasing complexity of manufacturing data for which the WDAD assumption

does not hold [248].

In recent years, there have been multiple advancements in the development of learning-based neighbour graph algorithms such as t-distributed Stochastic Neighbor Embedding (t-SNE) [249], and Uniform Manifold Approximation and Projection (UMAP) [250].

t-SNE is a variation of Stochastic Neighbor Embedding first proposed by Hinton and Roweis in 2002 [251]. While PCA retains global structure through eigenvectors with high variance, t-SNE reduces dimensionality by modelling high dimensional data neighbour points as a probability distribution in low dimensional space, thus retaining a more detailed local structure at the loss of some global information. This makes t-SNE favourable in producing visualisations where understanding this local structure is important and has been used in anomaly detection to visualise bearing faults [252], and superconductor manufacturing errors [253]. Furthermore, t-SNE can reveal non-linear relationships of the data that may be missed using PCA.

UMAP is a recent advancement in dimensionality reduction that has drawn much attention since its publication in 2020 in which McInnes et al. propose a topological mapping approach for dimensionality reduction [250, 202, 254]. UMAP has been shown to improve on t-SNE in preserving both local and global structure of data while also achieving superior run time performance [250, 202, 254]. UMAP outperformed PCA in clustering time-series data based on cyclical and seasonal characteristics [199, 254] and has been used in combination with density-based clustering approaches to highlight periods of anomalous behaviour in time-series data [199, 254]. Given the complexity and novelty of UMAP, further research is required to understand the performance of UMAP in industrial settings, with researchers suggesting opportunities for future works in comparing its 2D reduction performance with other distance methods [254].

### 4.3.3 Semi-Supervised Anomaly Detection

Unsupervised anomaly detection is a commonly used method of anomaly detection and is often beneficial as it can avoid the need to build high-quality labelled datasets to develop and implement the solution. However, in real-world applications, testing datasets will need to be developed to test and compare models during development to prove their effectiveness before implementation. In cases where the fraction of training points that are anomalies is very small ( $\alpha < 1\%$ ), any testing datasets will be highly imbalanced, with significantly more normal data than anomalies. In these cases, it is practical to utilize this

surplus normal data as part of a semi-supervised approach.

In time-series anomaly detection, LSTMs have been a popular area of research focusing on semi-supervised training approaches in which a model is trained on nominal data and anomalies are identified through various strategies of comparing prediction errors [205, 218, 220, 255]. Research by Malhotra et al. demonstrates how LSTMs can be used to successfully predict anomalies in four time-series datasets using a clean semi-supervised approach under the assumption that prediction errors fit a Gaussian distribution [218]. Mapping new error vectors onto this distribution is then used to identify anomalies in the contaminated dataset. In this study, LSTMs outperformed RNNs at semi-supervised anomaly detection. Chauhan et al. apply the same approach to identifying anomalies in ECG readings [226].

Hundman et al. use a similar semi-supervised training approach to anomalies in spacecraft telemetry data [205]. Hundman’s approach differs slightly in that error values are calculated at each discrete point in time, and instances where error exceeds a given threshold are highlighted as anomalies. By measuring the time above the error threshold, this method enables the distinction between point, contextual, and collective anomalies. A similar thresholding approach to Hundman et al. is used to identify specific time points for which collective anomalies occur in computer networks to highlight possible network intrusions [214]. Both Hundman et al. and Cao et al. highlight that the successful application of this threshold approach is sensitive to the careful selection of training data that must be clean and well-structured [214, 205]. Later research by Maleki et al. describes how outlier rejection can be used to select training data in an unsupervised manner using Central Limit Theorem, however, this requires the WDAD assumption, which makes this difficult to apply in practice [235]. This is discussed in their 2021 paper on how LSTM autoencoders trained using a clean semi-supervised approach can identify point anomalies in gas turbine measurement and CPU utilisation [235].

LSTM has been proven to be a powerful time series forecasting in recent years, outperforming traditional statistical techniques in more challenging cases where new or unknown anomalies may occur, or when there is no clear anomaly distribution [205, 207, 208, 209, 210, 218, 220]. A limitation of prior research in LSTM for anomaly detection is that research lacks detail in describing how anomaly thresholds are calculated once they have been modelled to fit a specific distribution [218, 226]. Some researchers focus on model comparisons and include no detail at all on how thresholds are set. In instances where the anomaly distribution is not known [255]. Prior research has explored LSTM for anomaly detection

in instances where either the nominal data distribution is well-understood [235, 218, 226], or vast amounts of historic training data are available [214, 205]. Because of this, less consideration has been given to methods that can be applied to minimise training data requirements and therefore make these methods more applicable in industrial settings where accessing high-quality training data is difficult. Based on this prior research, opportunities for further research are identified in the application of semi-supervised LSTM for anomaly detection in manufacturing environments. Given the success of setting anomaly thresholds for applications where testing data is limited and may not be sufficient to be modelled to fit a Gaussian distribution.

## 4.4 Methodology

In this section, the methodology to develop a solution for nut runner anomaly detection is presented. The two datasets were used to develop and test various machine-learning models. The methods used to collect and label these datasets in collaboration with domain experts are discussed, as well as the challenges of these real-world datasets.

Two semi-supervised approaches are presented to detect anomalies in nut runner data: 2D semi-supervised clustering, and time series forecasting. The former train a Gaussian Mixture Model on normal data to generate outlier thresholds in a reduced feature space. The latter trains an LSTM on normal data to make predictions on the test data and the resultant prediction error is used to detect outliers. This section describes the GMM and LSTM models at a low level, and how these mathematical foundations can be applied to detect anomalies in real-world manufacturing data. Descriptions are also included on the multiple dimensionality reduction techniques used in combination with the GMM. Details on the experimental setup to test and train each of the proposed solutions are also presented, as well as the metrics used to evaluate their performance.

### 4.4.1 Labeled Data

Nut runner anomalies are rare, and historical process data is not always stored long-term. This makes it challenging to get sufficient data on historical machine faults to develop training and testing datasets. If historical fault data does exist, this will still need to be reviewed by a domain specialist to ensure sufficiently high-quality datasets. The task of

labelling data is therefore the first major hurdle. Even fully unsupervised methods require high-quality datasets to validate model accuracy. In fault detection applications, training datasets will likely be highly imbalanced as fault instances and anomalies are usually rare. Therefore, large amounts of data may need to be reviewed by domain specialists to gather sufficient data to validate such models.

For this research, a dashboard was developed to speed up the labelling process. The dashboard presents a domain specialist with 12 on-screen examples of time-series nut rundown data to label. Given that anomaly occurrences were presumed to be very rare, the user was informed that all data they were being shown were examples of 'normal' operating conditions. If the user saw any instances that could be considered anomalous, they were asked to label this by using a series of push buttons to categorise the observation into one of three categories:

- **True Anomaly:** True anomalies are instances where either a known process anomaly has occurred that has compromised part quality, or an unknown anomaly has occurred that requires further inspection before the part is released.
- **Anomaly No Concern (ANC):** An ANC is defined as an anomalous observation for which no action is needed. This may be because the anomaly can be explained by a known process error that is unlikely to have compromised the quality of the outgoing part.
- **Re-hit:** A re-hit are instances where no data was recorded and the amplitude of the time series remains constant at 0.

If none of the 12 examples on-screen fall within one of these categories, a refresh button is used to label all observations as 'normal' and the display is refreshed with a new batch of 12 images.

This labelling approach proved very fast, as less than 1% of processes included True Anomalies and therefore most on-screen batches were all normal and labelled twelve at a time. Using this system, domain specialists were able to consistently label data at a rate of approximately 1000 observations per half-hour of labelling. This approach is designed to be used on personal computers to label historical data and is not integrated into any production process or data collection systems.

ANC class was introduced to overcome some confusion around what constitutes an anomaly. Production line test engineers consider anomalies to be any observation that would result in



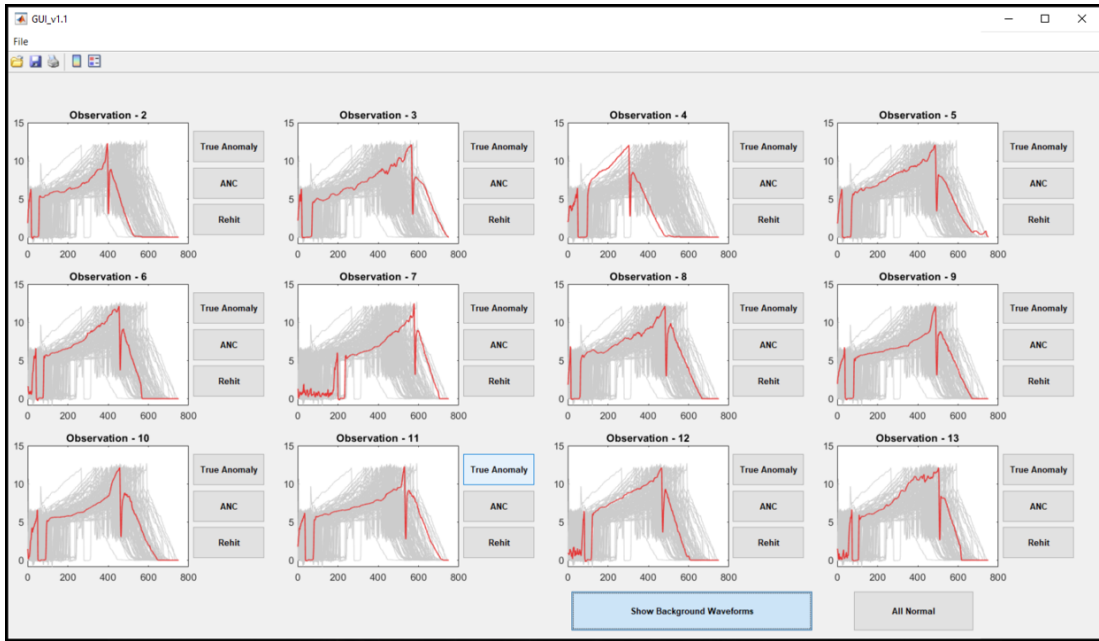


Figure 4.4: The data labelling dashboard allows users to label batches of 12 normal waveforms at a time.

a part being rejected for inspection/repair. In contrast, the data analysts viewed anomalies as having features or characteristics not found in the majority of the data. To address this contrasting definition of terms, the Anomaly No Concern (ANC) category was included in the labelling process. The judgement between a True Anomaly and an Anomaly No Concern is based on experience, and therefore there is some level of uncertainty in this class. For this reason, the task of labelling is given only to engineers who has a high level of understanding of the process and the test. While this does introduce some uncertainty into our testing datasets, it is notable that engineers will typically err on the side of caution, as product quality is prioritised over all other production metrics. For this reason, it is desirable for any proposed anomaly detection system to flag both a ‘True Anomaly’ as well as an ‘Anomaly No Concern’. This being said, a network’s ability to detect ‘Anomaly No Concern’ should never be improved at the expense of reducing the ‘True Anomaly’ detection rate.

By using this labelling dashboard we are able to overcome two of the major hurdles of developing and implementing machine learning approaches for fault detection: the time taken to label good quality data, and the knowledge gap between domain experts and data analytic experts. This labelling approach was used to build two testing datasets, one for each process. All labelled anomalies were utilised in this dataset, as well as a random

sample of normal data. Details of each dataset are shown below in Table 4.1.

Table 4.1: Composition of training and testing datasets to evaluate performance of the models.

Dataset	True Anomalies	ANC	Normal	Total
Test Dataset 1	26	37	1000	1063
Test Dataset 2	67	100	1000	1167

#### 4.4.2 Machine Learning Model Descriptions

Based on the related research, three dimension reduction approaches are compared prior to training a GMM: PCA, t-SNE, and UMAP. During the model development phase, the visualisations produced by these approaches proved useful in communicating the results and findings of the nut runner analysis to other team members. This was particularly useful when discussing the importance of high quality data, and revealed early on that test engineers would often disagree on data labels. Visualising the results through the early development phase made it easier to identify and communicate potential labelling mistakes and get feedback from test engineers. For these reasons, it was decided only to explore 2D representations for all dimensionality reduction approaches.

#### PCA

PCA is a dimensionality reduction technique that aims to preserve the global structure of the data by preserving pairwise distance among all data samples. This is achieved by applying linear mapping using matrix factorization. Consider a dataset  $\mathbf{X}_o$  comprised of  $m$  observations, and  $k$  variables. For ease of computation, the first step of the PCA is to centre the  $k$  dimensional data using the centering matrices  $\mathbf{C}_n$  and  $\mathbf{C}_m$ , given by:

$$\mathbf{C}_m = \mathbf{I}_m - \frac{1}{n} \mathbf{1}_m, \quad (4.1)$$

and

$$\mathbf{C}_n = \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n, \quad (4.2)$$

where  $\mathbf{1}_n$  and  $\mathbf{1}_m$  are the  $n \times n$  and  $m \times m$  matrices of all 1's respectively. The zero-mean centered matrix  $X$  can then be calculated by

$$\mathbf{X} = \mathbf{C}_m \mathbf{X}_o \mathbf{C}_n. \quad (4.3)$$

The simplest approach to derive a dataset's principal components is by finding the eigenvectors  $\mathbf{u}_i$  of the symmetric  $k \times k$  covariance matrix  $C$ , where [247]:

$$\mathbf{C} = \frac{1}{n-1} \mathbf{X} \mathbf{X}^T. \quad (4.4)$$

These eigenvectors are the principal components of  $\mathbf{X}$ . The inverse of the eigenmatrix  $\mathbf{U}$  is used to map the dataset onto the reduced feature space  $\mathbf{X}'$ , where

$$\mathbf{X}' = \mathbf{U}^T \mathbf{X}. \quad (4.5)$$

The first principal component of  $\mathbf{X}$  is the eigenvector  $\mathbf{u}_i$  that gives the largest eigenvalue, as this describes the direction in  $p$  dimensional space along which the data has the highest variance. The second principal component is the eigenvector orthogonal to the first principal component, and so on. For further information on the derivation of the covariance matrix, the reader is referred to [256].

### Application of PCA

The current anomaly detection method at Ford Motor Company identifies anomalies using Density-Based Spatial Clustering of Applications with Noise (DBSCAN) to identify noise points in the first two principal components. Figure 4.3, shows a labelled plot of the first two principal components for datasets 1 and 2. These plots show that when applying the PCA transform on nut runner data, not all noise points are anomalies, anomalies can form tight clusters, and not all are outliers. In dataset 1, many anomalies lie within the nominal distribution of the data, sometimes forming small clusters of anomalies within the nominal distribution. In dataset 2, most anomalies form tight clusters, and the ANC class overlaps significantly with the nominal data.

Because PCA is sensitive to the variance, the data are first normalized. Typically, when applying PCA, the first two or three components are selected, as these components retain the most information on the original structure of the data [257]. However, initial experiments with nut runner data found that using different combinations of principal components resulted in more distinct clusters of anomalies. For this reason, the principal components were also varied when optimizing the hyperparameters for the semi-supervised approach.

## t-SNE and UMAP

Following the state-of-the-art dimensionality reduction, t-SNE and UMAP are also compared with PCA at clustering the data in 2D before applying the semi-supervised GMM. t-SNE and UMAP are neighbour graph approaches that determine the similarity between the data points before projecting the data onto the lower dimensional space.

### t-SNE

Consider some training dataset comprised of  $T$  training vectors and  $n$  dimensions, given by  $X = x_1, \dots, x_T$  that we wish to map onto a low-dimensional space given by  $Y = y_1, \dots, y_T$ . The most basic method of calculating the similarity matrix is the k-NN approach, where for each point in the high-dimensional space, the euclidean distance between every other point is calculated, given by:

$$D_{ij} = \sqrt[p]{\sum_{i=1}^T |x_i - x_j|^p}. \quad (4.6)$$

where  $p$  is the user-defined power parameter. The resultant  $T \times T$  is the similarity matrix  $D$ . This k-NN approach is one of the fundamental clustering models used in machine learning [258].

Stochastic Neighbor Embedding (SNE) converts the euclidean distances  $D_{ij}$  to probabilities, such that the similarity between  $x_i$  and  $x_j$  is the conditional probability  $p_{ij}$  that  $x_i$  would select  $x_j$  as its neighbour if neighbours were selected in proportion to their probability density given by a Gaussian centred at  $x_i$  [249]. The conditional probabilities that make up this new similarity matrix  $D'$  are given by:

$$p_{j|i} = \frac{\exp(-C_{ij}^2)}{\sum_{k \neq l} \exp(-C_{kl}^2)}. \quad (4.7)$$

Similarly, a second similarity matrix  $C''$  is also calculated for the low-dimensional space, given by:

$$q_{j|i} = \frac{\exp(-\|y_i - y_j\|_2^2)}{\sum_{k \neq l} \exp(-\|y_k - y_l\|_2^2)}. \quad (4.8)$$

where the similarity between  $y_i$  and  $y_j$  is the conditional probability  $q_{ij}$  that  $y_i$  would select  $y_j$  as its neighbour. SNE aims to minimise the difference between  $p_{i|j}$  and  $q_{i|j}$  by minimising the sum of the Kullback-Leibler divergences over all data points using gradient decent [249]. The Kullback-Leibler divergences define the cost function  $C$  for the SME

algorithm as:

$$C = \sum_i KL(P_i || Q_i) = \sum_i \sum_j \log \frac{p_{i|j}}{q_{i|j}}, \quad (4.9)$$

and the gradient,

$$\frac{\delta C}{\delta y_i} = 2 \sum_j (p_{i|j} - q_{i|j} + p_{j|i} - q_{j|i})(y_i - y_j). \quad (4.10)$$

For a full derivation of  $\frac{\delta C}{\delta y_i}$ , as well as other approaches to determine the SME gradients can be found in the original paper by Maaten and Hinton [249].

t-SNE is a variation of the SNE algorithm that uses a heavily tailed Student-t distribution in the low-dimensional map, rather than a Gaussian distribution to determine the similarities. For the t-SNE algorithm, conditional probability  $q_{ij}$  is given by:

$$q_{i|j} = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}} \quad (4.11)$$

By applying a probability distribution with heavy tails,  $(1 + \|y_i - y_j\|^2)^{-1}$  approaches an inverse square law for large pairwise distances in the low-dimensional representation of the data. This helps retain global structure by separating clusters that are far apart, while retaining local structure within the respective clusters [249]. These new values for  $q_{i|j}$  give a gradient of the Kullback-Leibler divergence as:

$$\frac{\delta C}{\delta y_i} = 4 \sum_j (p_{i|j} - q_{i|j})(1 + \|y_i - y_j\|^2)^{-1}. \quad (4.12)$$

The gradient descent is initialised by randomly sampling from an isotropic Gaussian centred on the origin with a small variance. The initial data points  $Y$  are then shifted in this low-dimensional space such that the conditional probabilities  $Q$  converge on  $P$ . For further details on how the gradient decent process is optimised to avoid poor local minima, see [249].

## UMAP

Similar to t-SNE, UMAP is also a neighbour graph approach that uses stochastic processes to map  $X$  onto  $Y$ . UMAP is by far the most complex approach discussed in this section, with the theoretical foundations based on manifold theory and topology [250]. At a high level, UMAP applies manifold approximation together with local set representations to map the data onto lower dimensional space. These high-dimensional set representations, known as simplicial sets, provide a combinatorial approach to describe the high-dimensional




Figure 4.5: UMAP uses combinations of simplicies to provide a simplified representation of the continuous topological space defined by the high dimensional dataset  $X$  while retaining the global and local structures that define the space.

feature space by combining multiple simplices defined by the data points  $X$ . Figure 4.5 shows a visualisation of these simplicies and how they can be combined into a simplicial complex to describe a multi-dimensional feature space. The reader is referred to McInnes et al 2020 for further description of the mathematical description of simplicial complex and how it is used to describe the high-dimensional manifolds [250].

There are many similarities between UMAP and t-SNE. Numerous algorithms are presented in the original UMAP paper based on complex topological theories, however, these algorithmic steps can be expressed in a form similar to the t-SNE notation. For example, UMAP also uses the similarities in both the high-dimensional and low-dimensional space to define the cross entropy loss function  $C$ , given by:

$$C_{UMAP} = \sum_{i \neq j} v_{ij} \log \left( \frac{v_{ij}}{\omega_{ij}} \right) + (1 - v_{ij}) \log \left( \frac{1 - v_{ij}}{1 - \omega_{ij}} \right) \quad (4.13)$$

where  $v_{ij}$  are the local simplicial set memberships based on the nearest neighbour distances, and  $\omega_{ij}$  are the low-dimensional similarities. The local simplicial set memberships are given by:

$$v_{ij} = \exp[(-D_{ij} - \rho) / \sigma_i] \quad (4.14)$$

where  $\rho_i$  is the distance between  $x_i$  and the nearest neighbour,  $D$  is the similarity matrix, and  $\sigma_i$  is a normalising factor calculated using a binary search. For the algorithm used to

calculate  $\sigma_i$  see [250]. The similarities in low-dimensional space are given by:

$$w_{i|j} = (1 + a\|y_i - y_j\|_2^{2b})^{-1} \quad (4.15)$$

where  $a$  and  $b$  are user-defined positive values. Setting  $a = 1$  and  $b = 1$  give the Student t-distribution used in t-SNE. For information on how the points are initialised for the stochastic process, the reader is referred to McInnes et al [250].

### Application of t-SNE and UMAP

Because t-SNE is a probabilistic approach and both t-SNE and UMAP use stochastic processes, we must combine the training and testing data and perform dimensionality reduction on both datasets to ensure they are mapped onto the same lower dimensional feature space. The 2D outputs of the dimensionality reduction are then split back into the training and testing sets before applying GMM for semi-supervised clustering. Figure 4.6 shows t-SNE applied to the training and testing nut runner data.

For data with very high dimensions t-SNE has challenges associated with high computational requirements when compared to PCA. The initial construction of the k-NN graph to determine the similarity scores is a computational bottleneck for very high dimensional data, and the performance of the k-NN step deteriorates as the dimensionality is increased [259]. Furthermore, the t-SNE method becomes increasingly sensitive to parameter selection as the dimensionality is increased. This also requires users to exhaustively search for optimal parameters which become computationally expensive for very high-dimension datasets [259]. However, in the original paper by [249], the method was shown to have a low error at on a 784-dimensional dataset. This is a higher dimensionality compared to the 750 dimensions in the nut runner time series. Details of the optimization for the t-SNE and UMAP are discussed in 4.4.2.

### GMM

Consider a Gaussian process for which some output  $X$  is a continuous random variable. It is impossible to define a probability distribution function for all  $x$ , as there are an uncountably infinite number of potential values. To overcome this, a closely related function can be used to describe the probabilities associated with a continuous random variable [260]. This

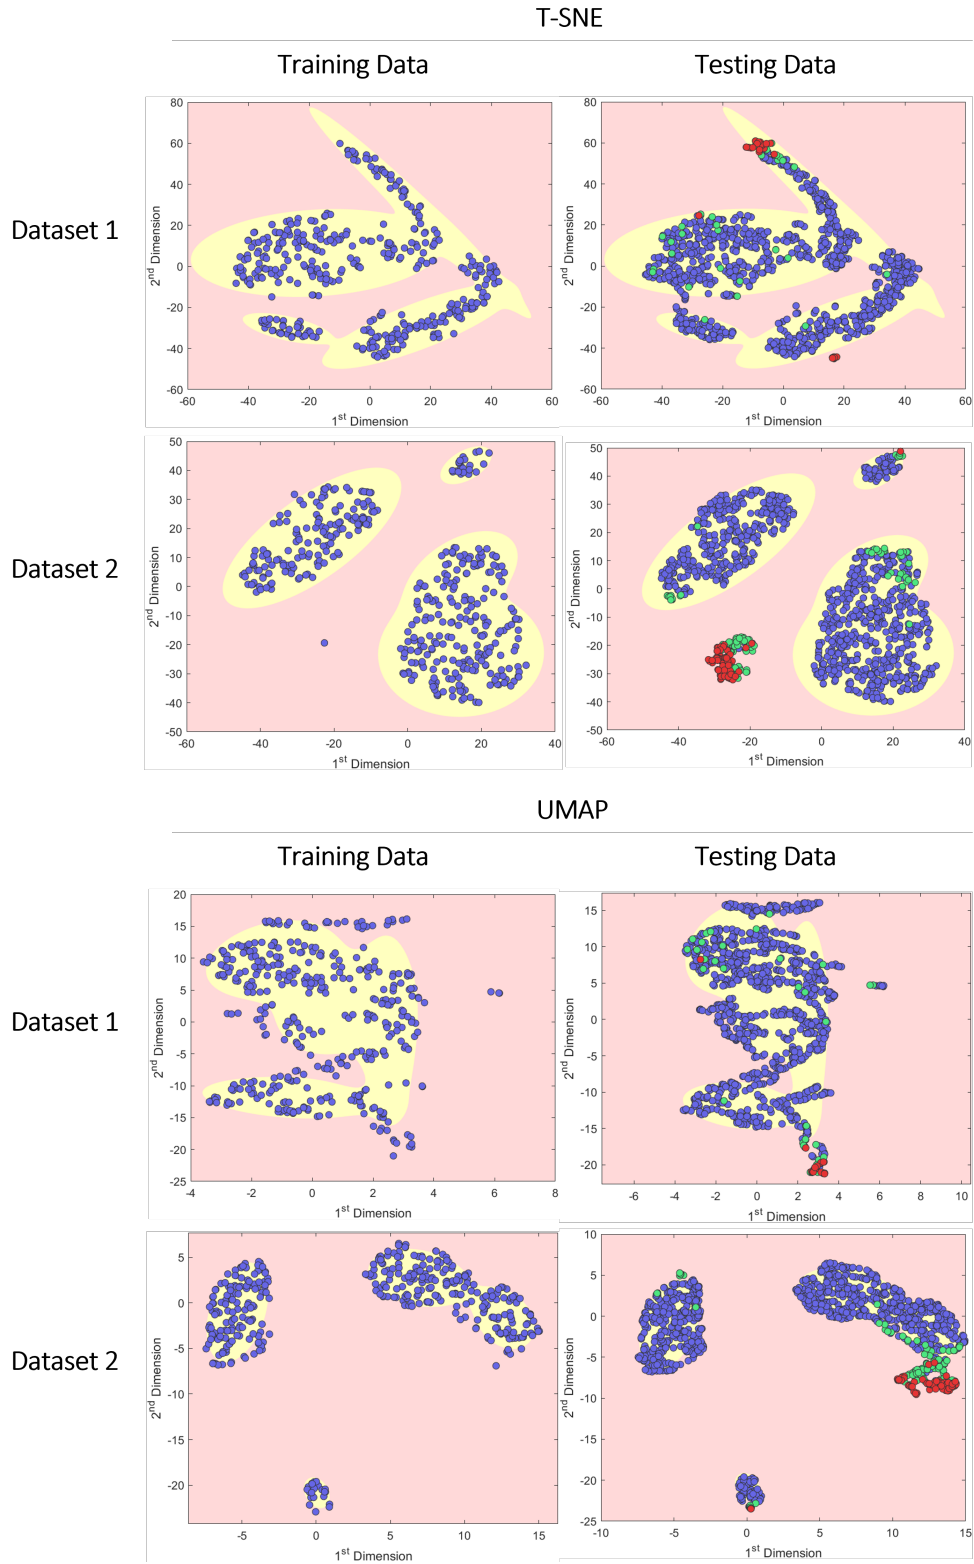


Figure 4.6: The same GMM approach applied using t-SNE and UMAP to reduce and cluster the data. Labelled data includes Nominal points (blue), ANC (green), and True Anomalies (red).



is called the Probability Density Function (PDF), given by:

$$p(x) = \frac{d}{dx}F(x) \quad (4.16)$$

$$F(x) = \int_{-\infty}^x p(x)dx. \quad (4.17)$$

A scalar Gaussian component has two parameters that can be used to describe the PDF: the mean  $\mu$ , and the variance  $\sigma^2$ . This gives a PDF in the form:

$$p(x|\mu, \sigma^2) = \mathcal{N}(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(x - \mu)^2}{2\sigma^2}\right). \quad (4.18)$$

Gaussian Mixture Model (GMM) assumes that the process can be described by several sub-processes, each of which can be described by a Gaussian probability density with a mean  $\mu$ , and the variance  $\sigma^2$ [237]. However, it is often the case when applying GMM that there are multiple features and high dimensionality [244]. For a multivariate Gaussian with  $n$  features and  $D$ -dimensions, a multivariate Gaussian PDF with the same quadratic form is used to describe these components, given by :

$$p(\vec{x}|\vec{\mu}, \Sigma) = \mathcal{N}(\vec{x}|\vec{\mu}_i, \Sigma_i) = \frac{1}{(2\pi)^{D/2}|\Sigma_i|^{1/2}} \exp\left(-\frac{1}{2}(\vec{x} - \vec{\mu}_i)' \Sigma_i^{-1}(\vec{x} - \vec{\mu}_i)\right), \quad (4.19)$$

where  $\vec{\mu}$  is the vector mean of length  $n$ , and  $\Sigma$  is the  $n \times n$  covariance matrix [244].

The GMM also introduces a scalar weight  $w_i$  for each Gaussian component, where  $\sum_{i=1}^M w_i = 1$ . Therefore, a GMM can be described as a weighted sum of  $M$  Gaussian components, given by:

$$p(\vec{x}|\{w_i, \vec{\mu}_i, \Sigma_i\}) = \sum_{i=1}^M w_i \mathcal{N}(\vec{x}|\vec{\mu}_i, \Sigma_i), \quad (4.20)$$

where  $i = 1, \dots, M$ . To apply the GMM to make predictions on new data, the model must first be fit to some training dataset comprised of  $T$  training vectors, given by given by  $X = \{\vec{x}_1, \dots, \vec{x}_T\}$ . This is achieved by making initial estimates for the mixture weights  $w_i$  mean vectors  $\vec{\mu}_i$ , and covariance matrices  $\Sigma_i$  before optimising these values. The most common approach to optimise the GMM parameters is to use an iterative Expectation-Maximization (EM) algorithm [245, 261, 244].

For a  $M$  components and with initial estimates for the mixture weights  $w_i$  mean vectors  $\vec{\mu}_i$ , and covariance matrices  $\Sigma_i$ , the next step in the EM algorithm is to calculate the probability that  $\vec{x}_T$  is assigned to component  $i$ , given by:

$$P_r(i|\vec{x}_t, \gamma) = \frac{w_i \mathcal{N}(\vec{x}_t|\vec{\mu}_i, \Sigma_i)}{\sum_{k=1}^M w_k \mathcal{N}(\vec{x}_t|\vec{\mu}_k, \Sigma_k)}, \quad (4.21)$$

where  $\gamma = \{w_i, \vec{\mu}_i, \Sigma_i\}$ . This probability  $P_r(i|\vec{x}_t, \gamma)$  is known as the *A Posteriori* and is used to calculate the next iterations parameters  $\gamma'$  using the following equations [244]:

$$w'_i = \frac{1}{T} \sum_{t=1}^T P_r(i|\vec{x}_t, \gamma) \quad (4.22)$$

$$\vec{\mu}'_i = \frac{\sum_{t=1}^T P_r(i|\vec{x}_t, \gamma) \vec{x}_t}{\sum_{t=1}^T P_r(i|\vec{x}_t, \gamma)} \quad (4.23)$$

$$\Sigma'_i = \frac{\sum_{t=1}^T P_r(i|\vec{x}_t, \gamma) x_t^2}{\sum_{t=1}^T P_r(i|\vec{x}_t, \gamma)} - \mu_i'^2. \quad (4.24)$$

The result of the EM process is dependent on the initialisation points for which to begin the EM optimisation process. This makes the user's selection of the number of Gaussian components important in achieving optimal results. Researchers commonly use methods such as the Bayesian Information Criterion or the Akaike Information Criterion to optimise  $M$  [261, 262, 261]. Similarly, the result of the GMM is also dependent on the training vectors, which often require careful pre-processing and feature engineering to reduce the dimensionality and cluster the data before applying the GMM. Because of the dependence on the initialisation points, GMM will converge on the local optimum, which may not necessarily be the global optimum. For this reason, multiple runs are required to compare model performance with different. Because the model is deterministic, the best result obtained over multiple runs can then be used to generate future Gaussian mixture components for classification.

## Application of GMM

A GMM is trained in a semi-supervised manner, using 400 normal training data and the number of Gaussian components. For the GMM model, two hyperparameters are optimised using a random search approach: The number of Gaussian components  $M$ , and the scalar weights  $w_i$ . The number of Gaussian components  $M$  are searched in the range of integers 1 to 6, and the initial estimates for the scalar weights  $w_i$  are multiplied by a value in the range 1 to 3 with intervals of 0.5. When optimising these hyperparameters for the GMM model, hyperparameters are also optimised for the respective dimensionality reduction methods. When applying PCA, the optimum principal components are also included in the random search, considering all possible pairwise combinations of the first ten components. For t-SNE, perplexity was studied in the range 5 to 50 in steps of 5,

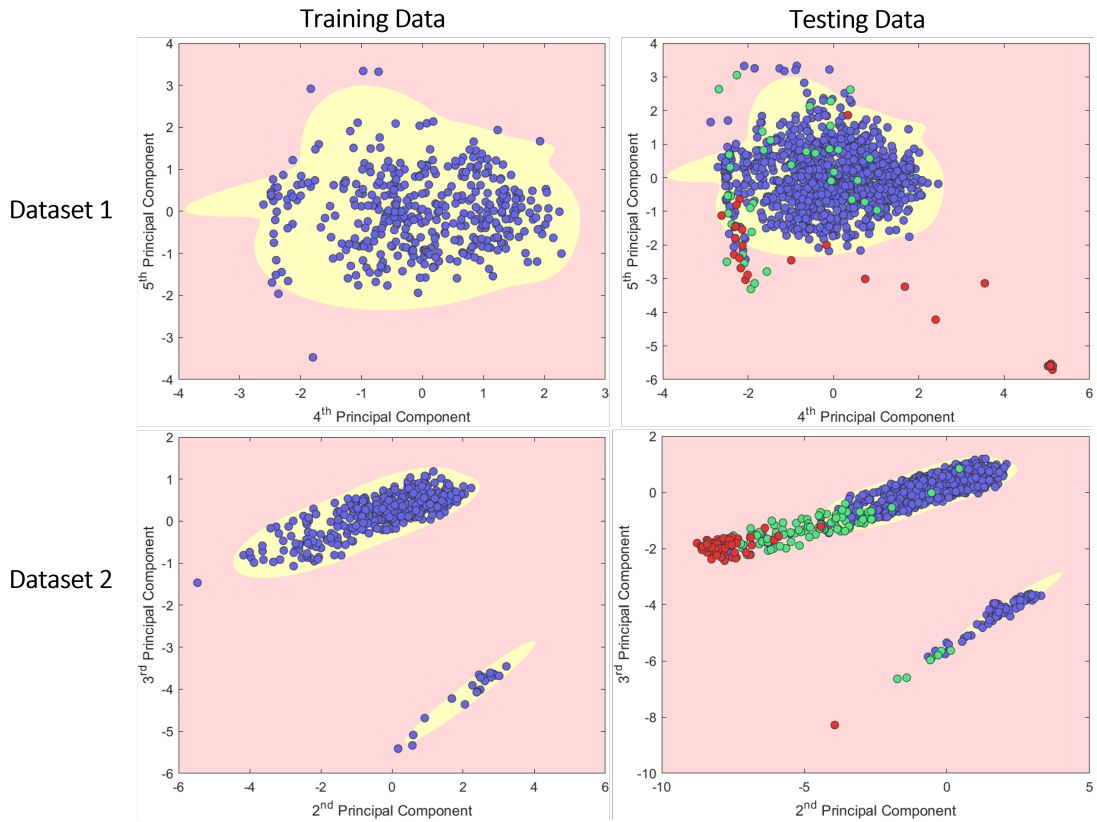


Figure 4.7: Outlier regions calculated using Gaussian mixture model trained on the reduced normal data. Any points that fall in the red area are identified as anomalies. Labelled data includes Nominal points (blue), ANC (green), and True Anomalies (red).

and learn rate in the range of 100 to 1000 with steps of 100, all other values are set as MATLAB defaults. For UMAP, the minimum distance was studied between 0.1 and 0.5 in steps of 0.1, and the number of nearest neighbours was studied between 5 and 25. All other parameters were kept as default in the modified code originally sourced from MATLAB File Exchange [263]. For each combination of hyperparameters, the experiment is repeated 3 times and the average f-score is calculated. For each model, the random search is stopped after 2 hours. This limitation on the optimisation time was decided by those managing the Cloud-based architecture of the anomaly detection pipeline. Because of bandwidth limitations, this 2-hour estimate would ensure that in a worst-case scenario, optimisation runs would still be able to be successfully completed during weekend non-productive time. This would avoid unnecessary downtime of the anomaly detection solution during the productive time. All experiments were run on an NVIDIA GeForce GTX 1050 GPU. By training the model using only normal data, the resultant Gaussian approximate components are used to define a threshold boundary to highlight outliers in the testing dataset. Figure 4.7 shows a plot of datasets 1 and 2 with the outlier thresholds visualized.

## LSTM

Long Short-Term Memory (LSTM) have been widely used in semi-supervised applications for anomaly detection. LSTM was introduced in the late '90s as an architectural modification to Residual Neural Networks (RNN) to address limitations of RNNs in their ability to learn and predict long-term dependencies in time series [230]. This limitation resulted from the propagation mechanisms in RNNs where the gradients used to update the adjacent weights of each hidden unit would either increase or decrease exponentially in proportion to the length of the time series  $n$ . For very large values of  $n$ , if the weights  $W$  are constant and greater than 1 then the gradients tend to infinity. This is called the exploding gradient problem and results in subsequent updates to these large gradients having a negligible effect on the output. In contrast, if  $W$  is constant and less than 1 this results in vanishing gradients, where the gradients tend to 0 and learned information is lost.

LSTM overcomes the gradient problem by replacing the tanh activation layer of RNNs with a memory cell containing multiple gates shown in Figure 4.9. The tanh activation function is depicted in Figure 4.8, where the output of the activation function is equal to  $\tanh(x)$ . Another important activation function that is used in the LSTM cell is the

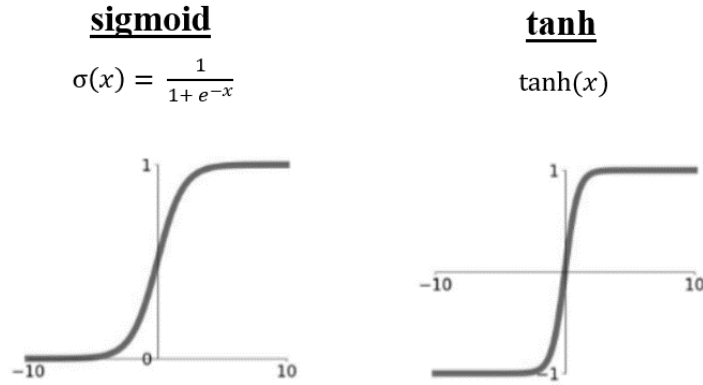


Figure 4.8: Graphs showing the Sigmoid and Tanh activation functions used in the LSTM network.

sigmoid function, for which the output  $\sigma(x) = 1 / (1 + e^{-x})$ . These activation functions and other mathematical operators make up the gates in the memory cell which regulate the amount of information the network passes between hidden units. A cell state vector  $C_{(t)}$  is also introduced to pass information between hidden units. This architecture allows the cell state  $C_{(t)}$  to act as the long-term memory component while the hidden state  $h_{(t)}$  acts as the short-term component. Figure 4.10 shows the 8 different learnable weights and 4 biases used to tune the network.




Figure 4.9: The high level LSTM architecture is similar to that of an RNN.

The three control gates in the memory cell are called the input, output, and forget gates, which can respectively perform write, read, and reset functions at each time step. The forget gate takes the input of the previous hidden state and the new input and passes

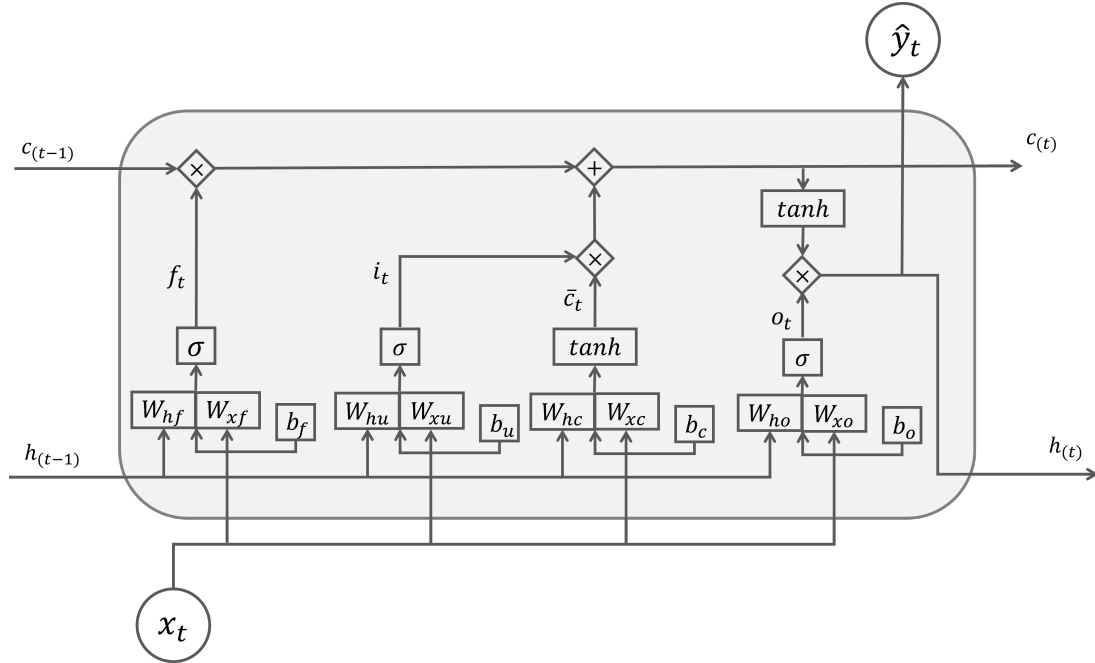


Figure 4.10: Inside the LSTM cell the three gates controlling the flow of information can be observed.

it through a sigmoid activation function. Values are returned in the range between 0 and 1, where values close to 0 are 'forgotten' while values closer to 1 are used to update the cell state. The forget vector  $f_t$  is then multiplied by the previous cell state output, where once again, forget vectors close to 0 will significantly reduce the cell state values. The forget vector is given by:

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f). \quad (4.25)$$

The input gate then applies the sigmoid activation to the previous hidden state and the new input to produce the input vector  $i_t$ , given by:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i). \quad (4.26)$$

The cell state vector  $\bar{c}$  is calculated by applying the tanh activation to the previous hidden state and the new input, given by:

$$c_t = \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c). \quad (4.27)$$

The tanh activation function distributes the gradients across a 0-centered range from 1 to -1, which helps regulate the cell state information over higher values of  $n$ , preventing the onset of the exploding and vanishing gradients problem. The input vector  $i_t$  is then used

to write information to the previous cell state by multiplying the result with the cell state vector  $\bar{c}$ . Finally, the output gate updates the hidden state  $h_t$ . This is done by applying a tanh activation to the new cell state and multiplying the result with the output vector  $o_t$ , where

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o), \quad (4.28)$$

and

$$h_t = \bar{y}_t = [o_t \times \tanh(c_t)]. \quad (4.29)$$

The new hidden state is then passed on to the following time step and is also outputted as the prediction  $\bar{y}_t$  for that timestep.

### Application of LSTM

An LSTM forecasting model is trained on normal data to make sequence-to-sequence predictions on the time-series data. This semi-supervised approach means that when making predictions on new 'Normal' data with no process errors, this trained model is expected to make accurate predictions. However, if the data contains anomalous readings, the error of the predicted sequence will be high as the model hasn't seen similar instances of these data before. This is shown in Figure 4.11 where LSTM predictions between a normal and anomalous observation are compared. This figure highlights a key benefit of this method over the clustering approaches in that the LSTM forecasting approach shows the specific parts of the time series that result in high prediction errors, therefore suggesting points in the process where anomalies may occur. This is useful in production as it can help identify instances of specific faults.

For this study, it was found that for any given waveform  $X$  with a sequence denoted by  $X_t = [x_{t1}, x_{t2}, x_{t3}, \dots, x_{tn}]$  a sequence-to-sequence approach was most effective. This is implemented using an X training input of  $X_{train} = (x_{t1}, x_{t2}, x_{t3}, \dots, x_{tn-1})$  and a Y training input of  $Y_{train} = [x_{t2}, x_{t3}, x_{t4}, \dots, x_{tn}]$ . All data are normalised prior to training.

### Anomaly Detection Threshold: Gaussian Distribution

Figure 4.4.2 shows a histogram of the RMSE values for predictions made on test datasets 1 and 2. Given that these values are normally distributed, we can use this to calculate the anomaly threshold similar to prior research [218, 226]. Previous research does not provide specific details on how an anomaly threshold should be calculated using this method




Figure 4.11: Examples of how LSTM forecasts vary between normal waveforms (Top) and anomalous waveforms (Bottom). The LSTM performs poorly when forecasting anomalous oscillations between 50 and 200 time steps, leading to a high reconstruction error above the anomaly threshold.

[218, 226, 255], therefore in this study we turn to a common heuristic of setting the threshold as 1.5 times the interquartile range above the third quartile [264]. This gives a threshold of 0.0602 for dataset 1, and 0.0468 for dataset 2.

### Anomaly Detection Threshold: Inflection Point

We can then use a test dataset contaminated with labelled anomaly data to determine an error threshold. Any data resulting in a prediction error above the given threshold can be assumed to be an anomaly.

The error threshold is calculated by sorting the RMSE values in ascending order and finding the corresponding RMSE value of the inflexion point. The inflexion point is defined by the maximum distance  $d$  on a plot of the sorted RMSE values, shown in Figure 4.4.2, for which  $d$  is defined as:

$$d = |p_2 - (p_1 \cdot \hat{\mathbf{b}})\hat{\mathbf{b}}| \quad (4.30)$$

where  $p$  are points on the graph, and  $\hat{\mathbf{b}}$  is the unit vector between the lowest and highest RMSE values.

This approach gives slightly lower threshold values compared to the Gaussian method, with 0.0562 for Dataset 1, and 0.0437 for dataset 2.

### Model Optimisation

As discussed in the literature review, prior research has explored LSTM for anomaly detec-



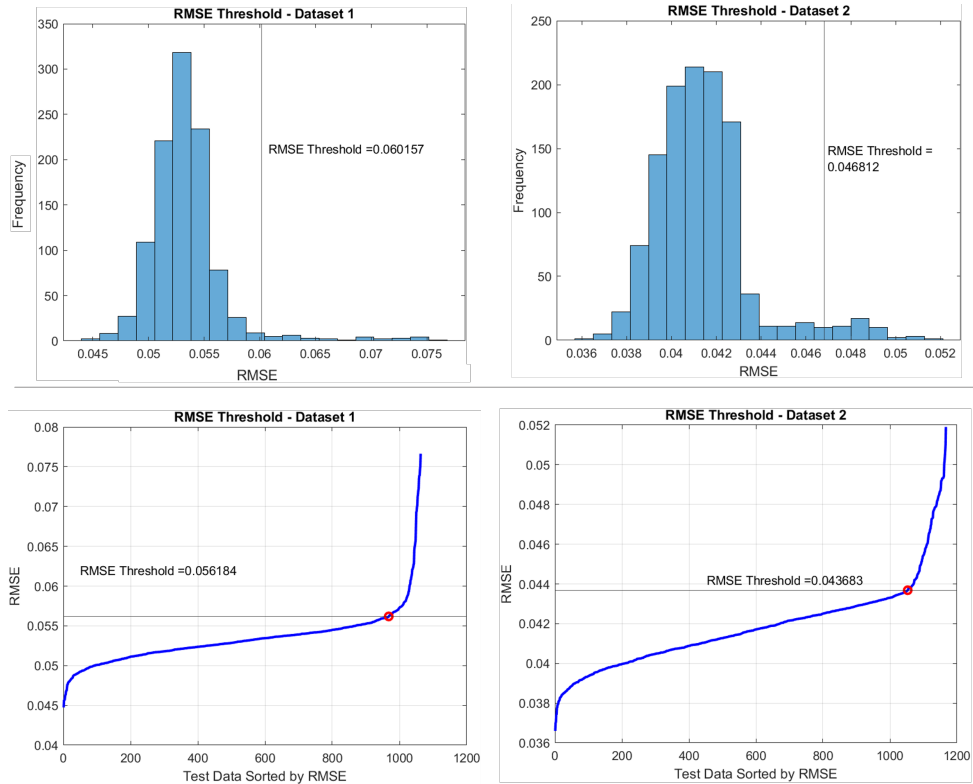


Figure 4.12: Two methods are compared at to find the error threshold for the LSTM forecasting approach: one setting the threshold as 1.5 times the interquartile range above the third quartile (Top), and the other using the elbow method on a plot with RMSE values sorted in ascending order (Top).

tion in instances where either the nominal data distribution is well-understood [235, 218, 226], or vast amounts of historic training data are available [214, 205]. In this application, due to the complexity of the process and the limited labelled data available, the distributions of our data are difficult to define, and limited data can be used for training.

To select optimal LSTM parameters and minimise training data requirements, a series of experiments is carried out to explore how LSTM prediction accuracy on normal data correlates with anomaly detection performance. In these experiments, the training data size, LSTM architecture, and model hyper-parameters are varied using a grid search approach with the goal of minimising the RMSE on a testing dataset containing 200 normal data. Each model setup is also used to make predictions on a second test dataset containing 200 additional normal data as well as 50 ANC and 50 True Anomalies. Figure 4.4.2 shows a plot the RMSE values on the normal test data against the F-score on the contaminated dataset.

In both datasets, there is little correlation between prediction accuracy on the normal test data and the models' ability to accurately detect anomalies. However, in both datasets when RMSE on the normal test data is low, two clusters are observed with one at an F-score of around 0.7 and the other above 0.9. These clusters are not the result of any particular combination of parameters and upon repeating the experiment, these clusters still form when all parameters are fixed and the only change between runs is the random sampling of the training data. This suggests that the training data sample is contaminated with some anomalous data, leading to a higher number of false negatives and therefore a lower F-score. This is reflected in the confusion matrices of the experimental results.

The optimisation experiments also help identify the amount of training data that can be used. For our testing datasets, we find that the optimal training data sizes are 200 and 400 for datasets 1 and 2, respectively. For both datasets, we find that a single LSTM layer with dropout and 200 hidden units gives the best results, as well as L2 regularisation set as 0.0001, learn rate = 0.0001, and max epochs as 20. The inclusion of L2 regularisation, dropout, and the low number of training epochs are common approaches to reduce overfitting, a common problem in training with noisy data in which models learn information on irrelevant features of the noise and fail to generalise important features to make accurate predictions on new data [265, 266, 267, 257].

The results of these experiments highlight the importance of high-quality test datasets to validate the accuracy of models prior to implementation.




Figure 4.13: Results of the parameter optimisation experiments showing how anomaly detection accuracy varies with the models’ performance at predicting normal data.

### 4.4.3 Metrics

In this study, anomalies are treated as the positive class. With this in mind, the following aims are outlined for this study:

- **Minimise False Negative Rate:** The end goal of this study is to develop an anomaly detection system to improve overall product quality. Therefore, our main objective is to reduce the number of True Anomalies incorrectly identified as Normal.
- **Identify a High Percentage of True Anomalies:** Reducing the false negative rate should not come at the expense of identifying a low percentage of True Anomalies.
- **Near Real Time:** Any solution must be able to identify a potential anomalous reading before the part continues onto the next process in the production line. While this time varies between processes, we set a target of under 5 seconds to perform the analysis.
- **Adaptable and Transferable:** As processes change over time, any provided solution must be re-trainable with minimal additional development by engineers. Furthermore, any solution must be demonstrated to be effective on multiple nut runner datasets to demonstrate its transferability to multiple use cases.

Given these objectives, we use the F-score false negative rate as our main metric to

measure the performance of our methods. F-score is defined as:

$$F - score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (4.31)$$

where precision is the ratio of between true positives and all positives,

$$Precision = \frac{TP}{TP + FP} \quad (4.32)$$

And recall is the measure of the method to correctly identify true positives

$$Recall = \frac{TP}{TP + FN}. \quad (4.33)$$

An optimal solution will maximise F-score while minimising the false positive rate.

These metrics are complicated slightly through the introduction of the ANC class. As discussed in section 4.4.1 it can be useful to view ANCs as either an anomaly or a nominal data point depending on the circumstances. For this reason, in any results, we state explicitly whether ANCs are being treated as nominal or as anomalies and discuss the findings within the relevant contexts.

## 4.5 Results

Typically, when evaluating the performance of models, data scientists at Ford Motor Company treat 'Anomaly No Concern' as 'True Anomalies'. This approach makes sense, as ANC's are still outliers and should be reviewed by test engineers to err on the side of caution and ensure the highest output quality. However, the author argues that this ANC information can provide additional insights into the performance of the models and should be further considered when analysing performance. This is particularly true for nut runner data where the quality of the datasets is difficult to assure given that there is some level of subjective judgement required when labelling the data that affects the overall quality of the datasets. For this reason, two situations are considered:

- a) ANC's are considered as True Anomalies. (Tables 4.2 and 4.3)
- b) ANC's are considered to be normal. (Tables 4.4 and 4.5)

Method	F-score	Precision	Recall	TP	FN	FP	TN
PCA+DBSCAN	0.14	0.14	0.14	9	54	55	945
PCA+GMM	0.49	0.56	0.43	27	36	21	979
t-SNE+GMM	0.25	0.35	0.19	12	51	22	978
UMAP+GMM	0.27	0.17	<b>0.68</b>	<b>43</b>	<b>20</b>	216	784
<b>LSTM Thresh1</b>	<b>0.56</b>	<b>0.82</b>	0.43	27	36	<b>6</b>	<b>994</b>
LSTM Thresh2	0.42	0.36	0.51	32	31	57	943

Table 4.2: Comparison of ML approaches for test Dataset 1a where ANC’s are considered as True Anomalies.

Method	F-score	Precision	Recall	TP	FN	FP	TN
PCA+DBSCAN	0.23	0.49	0.15	25	142	26	974
<b>PCA+GMM</b>	<b>0.83</b>	0.84	<b>0.81</b>	<b>136</b>	<b>31</b>	26	974
t-SNE+GMM	0.73	0.73	0.74	124	43	47	953
UMAP+GMM	0.59	0.48	0.77	128	39	137	863
LSTM Thresh1	0.43	<b>0.96</b>	0.28	46	121	<b>2</b>	<b>998</b>
LSTM Thresh2	0.66	0.78	0.57	95	72	26	974

Table 4.3: Experiment results for Dataset 2a where ANC’s are considered as True Anomalies.

Method	F-score	Precision	Recall	TP	FN	FP	TN
PCA+DBSCAN	0.09	0.06	0.15	4	22	60	977
PCA+GMM	0.43	0.31	0.70	18	8	40	1007
t-SNE+GMM	0.29	0.22	0.42	11	15	38	999
UMAP+GMM	0.02	0.01	<b>0.96</b>	<b>25</b>	<b>1</b>	234	803
<b>LSTM Thresh1</b>	<b>0.54</b>	<b>0.48</b>	0.62	16	10	<b>17</b>	<b>1020</b>
LSTM Thresh2	0.31	0.20	0.70	18	8	71	966

Table 4.4: Experiment results for Dataset 1b where ANC’s are NOT considered as True Anomalies.

Method	F-score	Precision	Recall	TP	FN	FP	TN
PCA+DBSCAN	0.07	0.08	0.06	4	63	47	1053
PCA+GMM	0.58	0.41	<b>1.00</b>	<b>67</b>	<b>0</b>	95	1005
t-SNE+GMM	0.61	0.45	0.97	65	2	78	1022
UMAP+GMM	0.35	0.22	0.87	58	9	207	893
<b>LSTM Thresh1</b>	<b>0.77</b>	<b>0.92</b>	0.66	44	23	<b>4</b>	<b>1096</b>
LSTM Thresh2	0.71	0.55	0.99	66	1	55	1045

Table 4.5: Experiment results for Dataset 2b where ANC’s are NOT considered as True Anomalies.

### 4.5.1 LSTM Results

For the LSTM solution, two methods are compared. Threshold method 1 (LSTM-Thresh1) sets the anomaly threshold 1.5 times the interquartile range above the third quartile, while threshold method 2 (LSTM-Thresh2) uses the elbow method on the sorted RMSE data. Because the calculated threshold is lower for LSTM-Thresh2, a higher number of anomalies are predicted, resulting in a higher true positive rate and a lower false negative rate, at the expense of a higher false positive rate.

For dataset 1, LSTM-Thresh1 achieves the best F-score of 0.56 due to its high precision. With the exception of GMM-UMAP, which can be ignored due to its very high false negative rate, the LSTM-Thresh2 method achieves the lowest false positive rate. However, this increased recall comes at the expense of a reduction in precision, resulting in a lower overall F-score than the best semi-supervised clustering approach, PCA-GMM. Here the user can tune the scalar constants to set the threshold value to achieve the desired balance between these metrics for the given situation. However, consider Table 4.5 for comparison. Here, the F-score for LSTM-Thresh1 is comparable, but there is a disproportionate drop in the F-score for LSTM-Thresh2. This shows that as the anomaly threshold is reduced, a large proportion of the new anomalies detected are ANCs. This would result in considerable added work for test engineers with minor gains in quality. For these reasons, the LSTM-Thresh1 method is preferable for dataset 1.

For dataset 2, when ANCs are considered True Anomalies, both LSTM achieve high precision, especially for LSTM-Thresh1 which only returns 2 false positives. However, both methods return a very high false negative rate, resulting in a lower F-score when compared to the clustering approaches. However, Table 2.4 shows that when ANCs are considered normal data, the LSTM methods achieve the highest F-scores. This example highlights

the importance of considering the result of situation b where ANCs are treated as normal data. Without considering this condition, initial experiments may have dismissed the LSTM solution due to a high false negative rate. However, in reality, the vast majority of these incorrect classifications are process errors that would not require any actions taken. LSTM-Thresh2 only misses 1 True Anomaly, making it the more preferable of the two methods.

#### 4.5.2 Semi-Supervised Clustering Results

The results of the semi-supervised GMM method are largely dependent on the dimensionality reduction approach used to prepare the data. For dataset 1a, containing manual nut runner data where ANCs are considered True Anomalies, the PCA-GMM performs the best of the clustering approaches, achieving F-scores of 0.55 compared to 0.25 and 0.27 for t-SNE-GMM and UMAP-GMM respectively. The PCA-GMM method achieves an identical recall to the LSTM-Thresh1 method, each returning 36 false negatives and 27 true positives, however, the PCA-GMM achieves a lower precision. The t-SNE and UMAP approaches are less desirable in comparison, with t-SNE returning a low recall while UMAP returns a low precision. However, all methods outperform the current PCA+DBSCAN used for anomaly detection.

When considering Dataset 1b, where ANCs are considered as normal, similar findings can be found compared to those discussed above, in section 4.5.1. The PCA-GMM method sees a comparable drop in true positive rate and false negative rates compared to the LSTM approaches. The t-SNE approach is the only method for which F-score increases when ANCs are treated as normal data. This aligns with the findings during the model development phases, where t-SNE was found to be particularly useful in identifying true anomalies contaminating the normal training data. However, the method was not as good as distinguishing between ANC and normal data. UMAP-GMM achieves the highest recall of all methods on dataset 1, however, the very low precision makes the approach undesirable in practice as it would result in considerable added work for test engineers to review these false positives.

For dataset 2a, PCA-GMM performs the best of all methods, achieving the highest recall and second-highest precision. Furthermore, when considering dataset 2b, it can be seen that this method identifies 100% of the true anomalies. t-SNE-GMM also performs well on dataset 2, with similar results to the PCA-GMM method. UMAP-GMM is once again

the least best method due to high false positive rates, however, it still achieves a higher F-score than the original PCA-DBSCAN approach.

## 4.6 Discussion of Results

### 4.6.1 LSTM

While no threshold selection method emerges as an optimal approach for the LSTM method, both semi-supervised LSTM solutions prove to be effective methods for predicting anomalies in nut runner data. These initial experiments suggest that the LSTM solutions are best at dealing with manual nut runner processes, although further experiments are required to confirm this. In real-world scenarios, the selection of an appropriate thresholding approach largely depends on the application. Given that the elbow method gives a slightly lower threshold, this is suggested for applications where a reduction in F-score is acceptable in order to reduce the false negative rate. The higher thresholds given by the Gaussian distribution approach results are more suited to minimising false positive rates. As mentioned in section 4.4.2, the model optimisation experiments resulted in very high F-scores on the validation datasets. These high F-scores are not replicated when using the same optimised pre-trained networks on test datasets 1 and 2. It is expected that this reduction in accuracy is again due to the high variance in anomalies. Further research is required to confirm this as more data are available.

A major benefit of the LSTM approach is that it's very easy to show at what point in the time series leads to the highest overall error. This presents an opportunity for future applications of this method to move towards a prescriptive analysis approach. This adds further value for test engineers to provide additional support in understanding the root cause of any errors. This is more difficult to achieve with other semi-supervised clustering methods.

### 4.6.2 Semi-Supervised Clustering Findings

For both datasets, PCA-GMM performs well and is the most consistent of all methods. This approach achieves similar precision and recall to the LSTM methods for Dataset 1, and performs the best on dataset 2 identifying 100% of the true anomalies. Unlike t-SNE



and UMAP, PCA is deterministic meaning the eigenvectors of the initial transform that gives an optimal result can be easily used to project any new data into the same feature space with minimal computational requirements. Furthermore, the ADAPT project already uses PCA in its current anomaly detection solutions and has a good understanding of the processes required to further develop and optimize this solution after concept readiness. For these reasons, it is decided to focus continued efforts on the LSTM and PCA+GMM solutions for nut runner anomaly detection.

Because t-SNE and UMAP are stochastic processes results will vary between runs, and specific results can be difficult to reproduce. Repeated experiments found that the f-score varied significantly, however, when successful t-SNE produced the most useful visualisations. For example, in Figure 4.6, the t-SNE results produce a distinct cluster of True Anomalies that were correctly identified and also resulted in a low false positive rate. Furthermore, this run also reveals a mislabelled data point in the training data that appears to be a True Anomaly at [-20, -20]. This highlights a major benefit of dimension reduction clustering approaches to produce 2D visualisations. By visualising the data in this abstracted feature space, a quick visual inspection can highlight potential labelling errors in training and testing datasets. These findings are aligned with previous research in which t-SNE was found to be the best method to visualise anomalies in fault detection and manufacturing production data [252, 253].

Despite the added value of the visualisations produced by t-SNE and UMAP, the re-training requirements for these algorithms present challenges when considering real-world implementation due to the variability of results and high computational requirements when compared to PCA. Given that the proposed architecture for the end solution uses Cloud-based services, any additional computational requirements will lead to higher processing costs and may affect the ability of the solution of delivering analysis in near real-time. There are opportunities for future reserach to explore non-random initialisation options for t-SNE and UMAP that reduce the variability of the final mapping after training. However, this can be complex to implement and may require additional optimisation steps to ensure a solution converges on global minima, rather than local minima. Future reserach should also explore more efficient Bayesian optimisation approaches that consider more hyperparameters, especially for higher dimensionality datasets greater than 784 dimensions, where the performance of the t-SNE approach is likely to decrease [249].

Although further work is required to apply t-SNE and UMAP for near real-time anomaly detection in nut runner data, our results show that t-SNE and UMAP are still useful tools.

Labelling production data is a difficult task, and during our research, it was found that even the most experienced test engineers disagree on True Anomaly and ANC labels, and mistakes are not uncommon when using our labelling tool. Throughout this project, visual inspection of the 2D t-SNE, PCA, and UMAP plots played an important role in cleaning labelled data and highlighting potential labelling errors. The false negative mentioned above in Figure 4.6 at [-20 -20] has since been confirmed by the test engineer to be an error and is indeed a True Anomaly. Following these results, these methods have since been adopted by Ford Motor Company within teams at the Dunton Technical Centre to validate data labelling efforts as part of this wider project. It is suggested that future research into real-world production data also use t-SNE and UMAP to help clean data before model training if sufficient labelled data is available.

## 4.7 Industrial Case Study

Following the successful results in the development stage, it was decided to include the semi-supervised anomaly detection methods in a live trial at Dagenham Engine Plant. The trial was focused on delivering anomaly detection for data in QualityWorks (QWX). QWX is a purchased data storage solution created by Sciometrics capable of storing Component Quality Data. QWX normally receives data from plant floor production cell machinery, such as Atlas-Copco, and Sciometrics and is commonly used in engine assembly and testing. For the full production volume at DEP, around 2,200 waveform IDs from 322 operations are written to QWX databases. For reference, a ‘waveform ID’ describes a collection of individual time-series data of fixed length produced by a specific process. Each individual time series is referred to as a ‘waveform’. The scope of the trial focused on waveform IDs for which at least 5000 historic waveforms were available to train the unsupervised PCA+DBSCAN model. Diagnostic waveforms are ignored. This resulted in 1,303 waveform IDs being considered in the trial, 256 of which are nut runner data.

Over the two-week trial, any detected anomalies are presented to test engineers via a dashboard. Test engineers review each of the flagged anomalies one-by-one, categorising each instance as “True Anomaly”, ‘ANC’, or ‘Normal’. If test engineers find a specific waveform ID is returning a high number of false positives, these labelled data are used to retrain the model and optimise performance.

Initially, it was decided to implement the LSTM-Thresh2 approach, due to the low false negative rates and high F-scores. While the overall performance seemed comparable to the

PCA-GMM method, the opportunity to explore prescriptive solutions in the future also influenced this decision. Furthermore, if the model returned false positive rates that are too high, a scalar constant can easily be applied to increase the anomaly threshold without requiring full retraining of the model.

Two main challenges were faced in the lead-up to the trial. Firstly, given the large number of waveforms to be included in the trial, it was infeasible to label normal training data for each case. Instead, it was decided to randomly sample training data from the 5000 historic waveforms, which may be contaminated with some ANCs and True Anomalies and may affect the performance. This was deemed acceptable by those leading the trial given the extent to which the solution outperformed the existing PCA-DBSCAN solution. The second major challenge was that the solution needed to be executed in Python, and the complex architecture of the anomaly detection pipeline made it difficult to set up a Tensorflow GPU environment required to train the model.

The complexity of the LSTM implementation presented a major challenge to the team, and as a result, the solution was not successfully included in the live trial. Due to time and resource constraints, it was decided to use the PCA-DBSCAN in the live trial for nut runner anomaly detection. An off-line trial would then be carried out afterwards to compare the semi-supervised approach with the results of the PCA-GMM approach. At the end of the trial, for the 256 nut runner data, a total of 1871 anomalies are identified by the system. Test engineers validate these anomalies for use in the off-line study. Due to the complexity of the LSTM solution, it was decided to use the PCA-GMM model given its similar performance during model development.

There are several challenges with this off-line approach. Firstly, only the anomaly data is validated by test engineers. This means normal data are likely to be contaminated with false negatives given the high false negative rates of the PCA-DBSCAN method when applied to nut runner data. Secondly, test engineers will often disagree on whether nut runner data are anomalies or not. This means the labelled data is also likely to contain some false positives. These factors result in low-quality testing datasets that are not representative of the ground truth.

The plots in Figure 4.14 show examples of why low-quality data presents a challenge for machine learning development. Each of these plots applied the PCA-GMM model, trained on 200 data sampled from the normal data, and optimized hyperparameters using the same approach described in section 4.4.2. Anomalies validated by test engineers are shown in

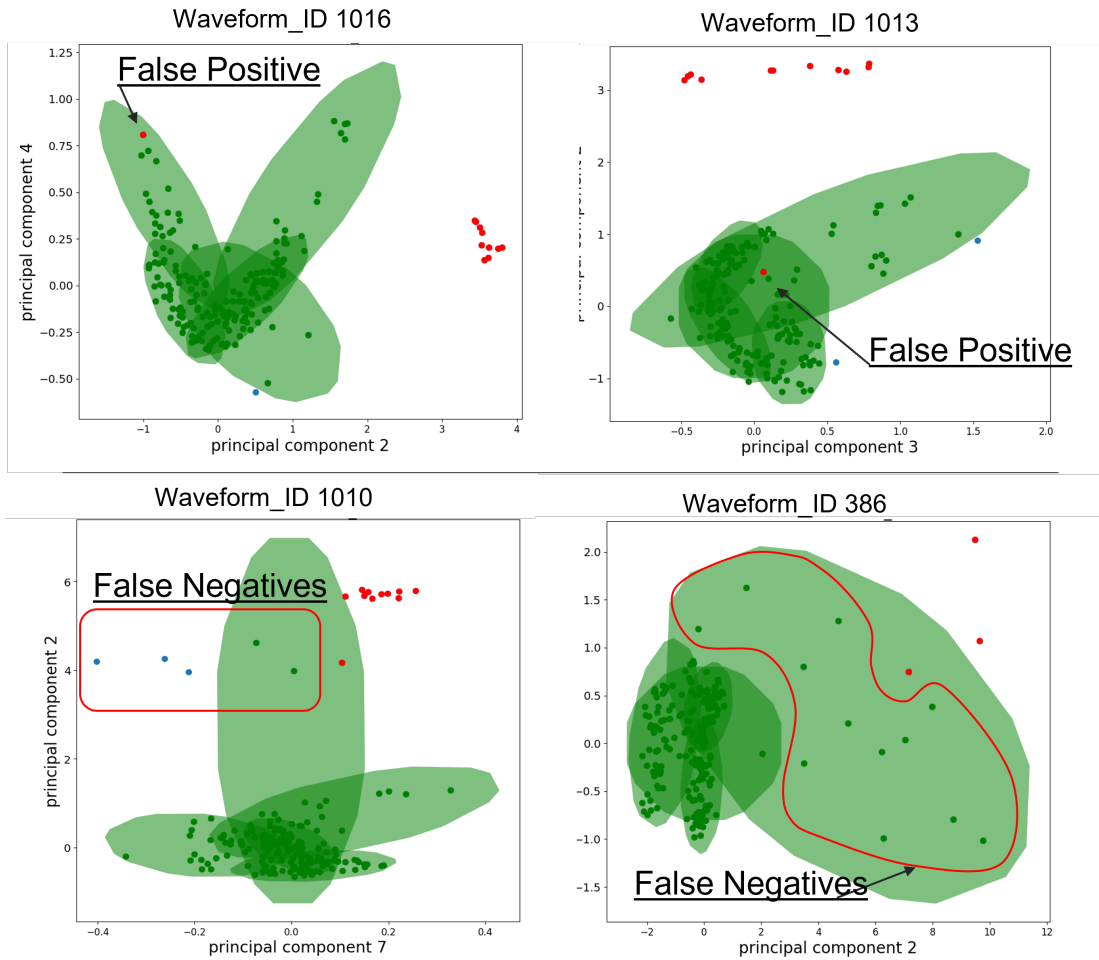


Figure 4.14: Four examples of waveforms included in the off-line trial. Any process data within the green boundary is classified as 'Normal' by the proposed PCA-GMM model. Green points indicate process data that are classified as 'Normal' by both the current PCA-DBSCAN model and the proposed PCA-GMM model. Blue points show process data that has been labelled as 'Normal' but the current PCA-DBSCAN model, but classed as a 'True Anomaly' by the proposed PCA-GMM model. Red points are classified as 'True Anomalies' by the current PCA-DBSCAN model. Because of the lack of validation, data that are suspected to be False Positives or False Negatives are highlighted.

red, while all other points are assumed to be normal but have not been validated. The GMM classifies any points that fall outside the anomaly threshold regions as an anomaly, while anything inside the threshold is classified as normal. Upon reviewing the data, multiple false positives and false negatives were identified using the visualisations, and the respective waveforms were sent to test engineers for review. These false positives and false negatives are highlighted in the example plots. Two conclusions can be drawn from these examples. Firstly, the current PCA-DBSCAN approach can result in a high false negative rate, as shown by Waveform ID 1010 and 386, however, the false negative rate is largely unknown. Secondly, sampling the training data from the 'normal' data is not an effective approach for the PCA-GMM method, as it is sensitive to any false negatives contaminating the training data, also shown by Waveform ID 1010 and 386.

#### 4.7.1 Opportunities for Future Work

At the time of writing, two more trials are being planned at Cologne Vehicle Assembly and Halewood Transmission Plant to explore further opportunities to implement nut runner anomaly detection using the PCA-GMM and LSTM methods. From the results presented in this chapter, there are multiple lessons learned that should be addressed ahead of these future trials. The following opportunities and recommendations for future work on nut runner anomaly detection:

- **Continued Analysis of Dagenham Trial Data:** Further work is required to continue the analysis of the datasets produced by the Dagenham trial and apply statistical approaches to select training data to train the model automatically. This will require input from test engineers to validate the labelled data.
- **Improving Methods of Data Labelling:** The lack of toolsets to efficiently label data to develop supervised and semi-supervised machine learning models was identified as a major gap at Ford motor company as part of this study. It is recommended that the data labelling tool developed as part of this study is used to support future data labelling to minimise the time requirements to perform this task. There are opportunities to improve this tool to make it easier to distribute to users who currently require admin rights to use the tool when on the Ford network. Further reserach may consider how to present the same waveform multiple times to multiple users to explore further how to deal with data for which test engineers disagree.

- **Tensorflow GPU Environment:** Input is required to support the setup of Tensorflow GPU environments to explore the use of LSTM for nut runner anomaly detection.
- Further research is required to explore how to reliably select normal training data from a sample of data contaminated with anomalies. It is recommended that for a fixed hyperparameter setup, random samples of 200 training data are selected  $n$  times to train  $n$  number of models. The performance of the models can be cross-validated and training data that lead to high prediction errors. However, this approach also requires high-quality testing and validation data.
- **Training Data Sampling Methods:** Initial results from the Dagenham off-line trial suggest that the PCA-GMM model works well on some waveforms but not others. With these new datasets, there are further opportunities to develop additional anomaly detection models and automatically select the optimal model to apply based on the features of the data.

## 4.8 Conclusions

This chapter proposes two solutions to detect anomalies in nut runner processes. Multiple reasons make nut runner data a challenging anomaly detection problem, including process staging, human-induced variability, and the subjectivity and ambiguity of the anomalous class. These process characteristics lead to an anomaly detection scenario where anomalies are not outliers, and the normal operating conditions are difficult to define. For these reasons, previous unsupervised attempts to automate nut runner anomaly detection have had limited success.

To develop a solution to address these challenges, two bespoke datasets are developed using data collected from two nut runner processes, one manual and one automated. In developing these datasets, multiple challenges of real-world data are faced. A major gap was identified in the company's machine learning development strategy, as no standard method was in place to support data labelling tasks. To address this, a simple user interface and labelling methodology are developed to minimise the human resource requirements to label large amounts of time series data. This dashboard has since been used to support additional projects at Ford Motor Company with further plans to develop this dashboard as an internal application.

In addition to the data labelling dashboard, a novel concept is introduced to label the training and testing data. When asked to label data, domain experts were given the opportunity to label data as 'Anomaly No Concern', in addition to the traditional labels of 'True Anomaly' and 'Normal'. Introducing this new term helped address knowledge gaps between data scientists and domain experts by highlighting conditions where some processing error had occurred but could be clearly explained as something that would not impact part quality or require any maintenance actions. The inclusion of the ANC class became a key consideration throughout the model development and testing to help clean data, build testing and training data, and address disagreement when labelling data. Furthermore, the ANC class provided further insights into model performance when analysing the results and can be used as further justification for the business case when estimating the solution's impact on quality metrics.

To overcome the challenges of nut runner anomaly detection, multiple solutions are presented that use the available normal data to train machine learning models. These semi-supervised approaches significantly outperform current methods at Ford Motor Company, increasing F-scores by a factor of ten in some cases. The methods presented fall under two main categories: semi-supervised clustering and anomaly detection using GMM, and semi-supervised LSTM forecasting.

For the GMM approach, three dimensionality reduction methods are compared: PCA, t-t-SNE, and UMAP. Of the three methods, t-SNE and UMAP were found to produce the best visualisations when developing the models, allowing data scientists and domain experts to identify mistakes when labelling data and support data cleaning and model development. However, the combination of PCA and GMM gave the best results when tested on two real-world datasets. A notable finding of this research was that because anomalies are not outliers, the principal components that gave the highest variance did not necessarily produce the optimal clusters for anomaly detection. To address this, a unique approach is used to apply PCA. Instead of selecting the selected principal components based on the highest variance, random combinations are included in the parameter search when training the combined PCA-GMM model.

For the time series forecasting approach, a LSTM model is trained using normal data, and the prediction error is used to identify anomalies. This study is the first attempt in the reviewed literature to identify anomalies in manual manufacturing processes. The LSTM achieves high levels of accuracy and demonstrates that accurate anomaly detection can be achieved on complex and highly variable manufacturing processes.

Following the success of the PCA-GMM and LSTM approaches, these methods were selected to be included in a trial at Dagenham Engine Plant to detect anomalies in an off-line setting. Several challenges were faced during the implementation. Human resource limitations and the complexity of setting up Tensorflow and GPU environments presented significant barriers to implementing the LSTM solution. These challenges are discussed in detail to be addressed ahead of future trials planned at Cologne and Halewood production sites. Due to these limitations, the PCA-GMM was used to compare performance with current anomaly detection methods. The results show that the current anomaly detection methods have high false negative rates and that there are further opportunities to deliver quality improvements in nut runner anomaly detection. The LSTM and PCA-GMM solutions are presented as the most promising areas in which to focus future work, and further input by test engineers to support data labelling to determine the success of these solutions. Future research should also explore improved methods of selecting training data for semi-supervised settings.



# Chapter 5

## Identifying Houses Suitable for EV Charging

### 5.1 Summary

The number of publicly available EV charging points in the UK is currently limited [268, 269]. According to UK government guidance, approximately 80% of all EV charging occurs at home [270]. Despite this, there have been limited efforts to survey residential areas to understand how many homes are suitable for home EV charging. This presents a challenge to local governments, automotive companies and other industries that require this data to understand future trends of EV uptake in order to deliver products, services, and infrastructures to support demand.

This chapter addresses these challenges and research gaps by presenting a novel method of surveying the built environment. Using geospatial and image data, 3 workflows to deliver automated surveys and identify houses suitable for EV charging are presented. Each method has its merits, developed with different end users in mind to address differences in the aims and scope of the different industries.

This work has led to a journal article published in Artificial Intelligence (AI), as well as 2 conference papers: one published in the proceedings of the 2021 IEEE International Geoscience and Remote Sensing Symposium (IGARSS) [271], and the second in IGARSS 2022 [272]. The works presented at IGARSS 2022 were done in collaboration with the Swansea University Geography department which assisted with the GIS analysis and Sentinel-2 analysis discussed in sections 5.5.3 and 5.5.3.

This chapter is structured as follows: Section 5.3 includes a literature review of prior research on machine learning applications using streetscape imagery. This section outlines various methods that have previously successfully identified building features from streetscape and overhead imagery using computer vision systems. This section also describes various datasets and data types used in previous studies, exploring how each of these data types is best suited to the different end users. Section 5.4 then introduces the image processing methodology used in all three of the final workflows. After first providing a definition of key terms and outlining some key considerations of the tools' end use, the network architectures used for image processing are described, as well as justification for their selection based on experimental results. A description of the geospatial and image datasets used is also provided. Section 5.5 introduces the three workflows, each of which is presented in individual sub-sections, along with three separate example surveys to highlight the benefits of each method and the opportunities for value creation for their targeted users. While each sub-section includes a brief description of the results, 5.6 includes a more detailed reflection on the three workflows, expanding on key findings and outlining suggestions for future work. Finally, the research conclusions are presented in section 5.7.

## 5.2 Introduction

In recent years, Plug-in Electric Vehicle (PEV) sales have grown considerably and are expected to represent 21% of the market by 2030 [273]. To manage this continued growth, companies and local governments must collaborate to ensure the appropriate infrastructures are in place to support the growing demands of charging electric vehicles. The number of EV charging points in the UK is currently limited [268, 269]. Significant research has been done to determine the optimal locations for public EV charging points [274, 275]. However, previous research is focused on delivering larger-scale aggregate data and does not consider the availability of home charging in local residential areas or at individual property levels. According to guidance by the UK government, approximately 80% of all EV charging occurs at home [270]. For a house to be classed as suitable for EV charging, according to the UK government's guidance for customers, a property must have some form of off-street parking, such as a driveway or a garage[270].

There is limited prior research that explores identifying houses with such features. Currently, no mapping agencies provide information about which properties are suitable for the installation of home EV charge points. This lack of data causes a challenge for local

governments as it makes it difficult to identify the best locations for community charge points. Such data is also required to plan and manage the infrastructure of the national grid as district network operators are required to manage the additional loads caused by home EV charging. Automotive companies could also use the data to better identify areas to target marketing, and better predict the future uptake of electric vehicles in different areas.

In the available literature, only one approach to this challenge was identified in a report published by Field Dynamics, a data analytics consultancy [276]. The approach uses GIS techniques to measure the distance between the edge of a residential building and the adjacent road. However, due to the commercial nature of the research limited information is available on the full methodology. No previous research has attempted to identify individual houses suitable for home EV charging, nor has there been research to identify residential off-street parking availability. It is critical that this research gap is addressed to provide the necessary data to support the rapid uptake of electric vehicles and the subsequent infrastructure requirements these products require.

To address this research gap, an image classification pipeline is proposed that uses a combination of geospatial and streetscape image data. Streetscape imagery is a type of remotely sensed RGB image data source taken from the perspective of a road user. Google Street View is perhaps the most well-known example of streetscape data with millions of freely available images available at 5m intervals taken from a drivers perspective. Each image in a streetscape dataset provides a vast amount of visual information about nearby properties that image recognition systems can leverage to identify a wide range of building attributes. The value of streetscape image data is already recognised by some UK local governments who already use Google Street View imagery to survey built-up areas. In these cases, surveyors manually inspect areas property by property in this virtual environment. This virtual approach improves on traditional surveying methods that would require an in-person inspection, however, this process is still labour-intensive. By combing these streetscape data with geospatial data it is shown that this process of auditing the built environment can be automated using Automated Programming Interfaces (API) to deliver a wide range of data outputs for various end users.

In this chapter, 3 workflows to deliver automated surveys and identify houses suitable for EV charging are presented. While each workflow differs slightly, they all rely on the same image processing steps using Google Street View images and pre-trained neural networks to classify properties as suitable or unsuitable for EV charging. It is the application of these

neural networks that enable large surveys of tens of thousands of households to be carried out in a matter of hours, at a very low cost. Each of the three workflows has its merits to address differences in aims and scope between the various end-users in mind, although the main focus of this research is to explore new data value chains for the automotive industry.

## 5.3 Related Research

There has been considerable research on classifying Land Use and Land Cover (LULC) in aerial and satellite remote sensing imagery with applications including environmental monitoring and natural hazard detection [277, 278], agriculture and vegetation mapping [279], and various urban planning applications [280, 281, 282, 283, 284, 285]. However, there has been significantly less research in the public domain that explores the use of streetscape imagery despite the widespread availability of tools such as Google Street View [286], Mapillary[287], and KartaView [288]. Streetscape imagery sources such as these capture a vast amount of high-resolution information on urban areas, much of which is impossible to capture using aerial and satellite imagery due to occlusion by roofs, large structures, trees, and the details lost in the low-resolution of these images.

Advancements in Industry 4.0 technologies such as image classification and Application Programming Interfaces (API) have enabled researchers to improve on traditional approaches to object classification by combining aerial, and streetscape imagery [280, 289, 290]. The success of these studies, and other works discussed in this section, highlight this emerging field of streetscape remote sensing as a possible route to develop a system to map the urban environment of residential properties suitable for home EV charging. This section includes an overview of the existing literature on remotely sensed streetscape imagery as well as recent advancements in machine learning tools used for image classification.

### 5.3.1 Machine Learning in Remote Sensing

Data-driven computational geographical approaches have been used since the 1980s when early LULC research relied on techniques to extract information at the level of individual pixels [291, 292]. This approach is often difficult as spectral information is often not sufficient and results in high miss classification [281, 282, 293]. For example, concrete rooftops, car parks, and road junctions may share similar spectral responses from an aerial perspective, making them difficult to distinguish using early handcrafted feature-based

methods such as Color Histograms [291] and Texture Descriptors [292]. Over the past decade, as distributed smart systems and the internet of things have become increasingly widespread, new technologies have emerged to process and analyse these vast amounts of data. The subsequent development of smart devices, artificial neural networks and computer vision systems has meant more recent approaches to LULC classification rely on machine learning techniques such as SVMs and CNNs to classify structures, and urban land use [280, 282, 283, 284].

Computer vision algorithms have proven to be especially powerful tools due to their versatility, scalability, and low-cost [294]. CNN's in particular have been widely used in remote sensing applications and are a well-established approach for pattern recognition, object detection, image classification, and image segmentation. CNN's have been used to distinguish between types of buildings in satellite and streetscape imagery, and have been shown to be very successful for some building types, such as religious buildings or entertainment buildings that typically have very clear distinguishing features [295, 280]. However, previous research has shown performance drops when identifying classes such as residential properties and garages where classes can be overlapped or have high variability [295, 280]. Inspired by the perception mechanisms of the visual cortex, a CNN architecture consists of multiple layers of artificial neurons in a stacked arrangement to perform three main operations: convolution, non-linearity, and pooling/subsampling [296, 297]. For image classification, a series of multispectral images are fed into the first layer in the form of 2-D arrays. The following sequential layers are then represented as input and an output feature maps calculated by alternatively stacking convolutional and pooling layers. The final layer is a fully connected layer in which classification is performed. For a further description of CNN, the reader is referred to the appendix, Section .1.2.

Most of these works use pre-trained networks such as AlexNet and GoogLeNet that have been pre-trained on the ImageNet database. ImageNet is a large image database containing 1000 images for 1000 different categories [298]. This open-source dataset leads to the introduction of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC). Engineers could now take advantage of this data, avoiding the otherwise expensive task of labelling data, allowing a platform to progress the field of image classification through competition. Other image databases similar exist for more specific purposes, such as Places: A 10 million image database for scene recognition [299].

Several of the networks presented in this subsection were developed as entries to the ILSVRC, and these pre-trained networks are also made open source. Re-purposing these

pre-trained networks using a fine-tune transfer learning approach means fewer images are required for training, as these networks retain information on low-level features that are transferable between tasks such as the detection of edges and object boundaries. The most common networks in the reviewed literature include AlexNet, VGG variants, GoogLeNet, and ResNet variants, all of which are pre-trained on ImageNet.




Figure 5.1: The AlexNet architecture is one of the most simple CNN architectures with only 8 layers and 60 million paramters [1].

## AlexNet

Over the years, various CNN models have been proposed with specific architectural traits and nuances to optimise performance. The eight-layer CNN was developed by Krizhevsky et al. in 2007 for the ImageNet ILSVRC-2010 contest is among the most popular due to its relative simplicity [1]. AlexNet was the first network to incorporate the ReLU activation function, which significantly improved the performance of feed-forward networks by making it possible to develop much deeper networks than was previously possible. This increased efficiency is because the gradient of the sigmoid activation function is between 0 and 1, which for many layers can result in gradients that are exponentially small. This means very small changes occur between each step of the gradient descent, which leads to slow convergence. For the ReLu activation function, the gradient is 1 if the input is positive, and 0 if negative. This results in more stable gradients for deeper networks, as well as faster computation as the derivatives, are quicker to compute.

Alexnet consists of 8 layers, with five convolutional layers and three max pooling layers. This architecture results in approximately 60 million parameters, making it the least complex network in this review. With only 8 layers, AlexNets simplicity makes it suitable for

situations where near real-time processing is required. This is demonstrated by Amato et al., who used AlexNet to detect parked cars in real-time using a Raspberry Pi as part of an occupancy detection system [300]. In recent years, other networks such as GoogLeNet, ResNet, and VGG variants have been shown to outperform Alexnet in a number of specific tasks. However, Alexnet remains the baseline approach for remote sensing applications, including LULC [301, 284], building instance classification [280, 282], and object detection [302].

## GoogLeNet

In 2015, Christian et al. proposed a new CNN building block called 'Inception Modules' as part of the GoogLeNet architecture. This innovation leads to GoogLeNet winning the 2015 ILSVRC14 competition. The inception module is characterized by multiple filters branching off in parallel from a single input layer [2]. The resulting outputs of the filters are then concatenated to form a single input to the next layer. This allows for more spacial information to be retained while using fewer parameters, making the network less sensitive to over-fitting as well as requiring less computational expense [2].

The most simple structure of an inception model is the 'naive version', consisting of a max pooling layer and three filters of sizes 1x1, 3x3, and 5x5, as shown in Figure 5.2 (a). In practice, the naive arrangement would lead to an increased number of outputs with each layer. To overcome this, 1x1 convolutions are introduced before the expensive 3x3 and 5x5 filters to reduce the dimensionality of the input and prevent computational blow up at deeper layers [2].

GoogLeNet has been shown to outperform AlexNet, and other networks in some aerial LULC applications [303, 304], and has been used to classify buildings in Google Street View images [305]. Previous work has shown that fine-tuning the network achieves better results compared to training the network from scratch or using feature vector approaches [303].

## VGG

Visual Geometry Group Networks (VGG), used very small convolutional filters to enable deeper networks ranging from 11 (VGG-11) to 19 (VGG-19) layers. The VGG networks were entered as part of the ILSVRC challenge in 2014, achieving very high accuracies [3]. The network achieved top 5 performance, outperforming AlexNet but losing to GoogLeNet

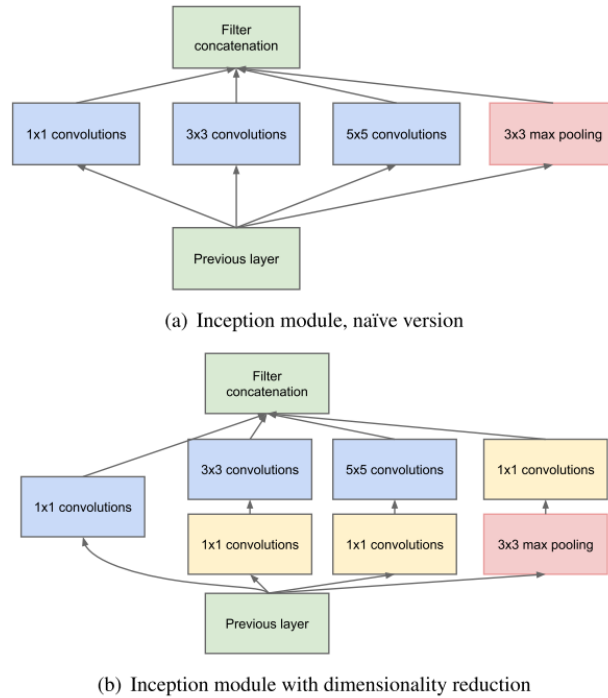


Figure 5.2: The two types of inception modules introduced in the GoogLeNet CNN architecture [2].

[3]. The VGG network is constructed with the smallest possible convolutional filters of size  $3 \times 3$  throughout all layers, although  $1 \times 1$  filters are also used to apply linear transformations using ReLU. Initial layers use 64 filters, with this number doubling with every increase in layer depth. This results in a deep network with a very large number of parameters, around 138 million. An example of the VGG-16 network is shown in Figure 5.3, consisting of 13  $3 \times 3$  convolutional layers and three fully connected layers.

Although GoogLeNet outperforms VGG in some applications, VGG is still widely used as it has been shown to generalise well to a wide range of tasks and datasets. This being said, the memory requirements of deeper VGG variants can make it challenging for some applications, requiring large GPUs [306]. VGG has been used in related works to classify urban scenery in streetscape images [307], as well as being successfully used to identify buildings in streetscape imagery [280].

## ResNet

In 2015, He et al. introduced the ResNet architecture which won the ILSVRC challenge through the novel application of feedforward 'shortcut' connections between layers [4]. This






Figure 5.3: The VGG-16 network contains 13 3x3 convolutional layers and three fully connected layers. [3].

architectural change was introduced to address the vanishing gradient problem faced by researchers when exploring very deep neural networks. For a mathematical description of the vanishing gradient problem, see Chapter 2 Section .1.2. He et al. propose a deep residual learning framework such that for an input  $F(x)$  and output  $H(x)$ , each residual block optimises the residual function  $H(x) = F(x) + x$  under the assumption that this is easier to optimise than the original unreferenced mapping  $H(x) = F(x)$ . The residual block is shown in Figure 5.4. The shortcut connections provide this residual mapping through each layer. This approach is shown to be successful, allowing for more layers to be stacked without increasing the number of parameters allowing for very deep architectures. ResNet variants can be of various depths, with 18-, 34-, 50-, and 101- depth networks most commonly used. ResNets have been shown to outperform other network architectures at mapping tree cover in Google Street View imagery [308, 289]. Li et al. also show that ResNet50 outperforms AlexNet at estimating building ages in Google Street View images [309]. There has been no published research that aims to identify individual houses suitable for home EV charging, nor has there been research to identify residential off-street parking availability. While no work has been done to identify driveways, some works have attempted to identify garages in Google Street View images. One related publication on building instance classification includes garages as a target class but had limited success, partly due to the overlap of classes in cases where many residential properties in the surveyed area contain integrated garages [280]. Zhao et al. improve on this approach




Figure 5.4: The ResNet architecture uses an identity connection to feed information forward between layers to address the vanishing gradient problem [4].

Table 5.1: A comparison of the most commonly used CNNs in related remote sensing works in the reviewed literature.

CNN	Layer Depth	Parameters (x10 <sup>6</sup> )	Defining Features
AlexNet [1]	8	60	First to use ReLU
GoogLeNet [2]	22	7	Inception Modules
ResNet18 [4]	18	11.4	Shortcut Connections
VGG16 [3]	16	138	Very small conv filters

by using a bounding box approach to identify types of buildings, including residential houses and garages [310]. By using object detection, this method allows for images to have multiple labels, avoiding the challenge of overlapping classes.

### 5.3.2 Streetscape Remote Sensing

Streetscape imagery is a type of remotely sensed RGB image data source taken from the perspective of a road user. Each image in a streetscape dataset provides a vast amount of visual information about nearby properties that image recognition systems can leverage to identify a wide range of building attributes. The use of streetscape imagery in remote sensing is a very recent research topic, made widely accessible in 2007 when Google Street View, the most popular source of such imagery, first released its API for users to automate the request of streetscape imagery. Soon after came the release of similar services, including KartaView in 2009 and Open Street View in 2013, although Google Street View has

remained the dominant source of such imagery for researchers in this time.

One of the first papers to explore this topic was on the use of SVMs to identify architectural characteristics from streetscape imagery [311]. Since then, following the advancements of image classification networks such as Alexnet and ResNet, more recent works have focused on fine-tuning CNN and DCNN variants for object detection [280, 312, 313]. Some previous works have explored using streetscape imagery and image recognition to identify external features and attributes of buildings [314, 309, 315, 316, 280, 282]. These approaches all use CNNs and Google Street View images. Li et al. demonstrated how CNNs can be used to extract characteristics of Google Street View images, such as the building materials and architectural styles of houses, to estimate the ages of buildings [309]. Other works using Google Street View imagery for geospatial analysis include: pedestrian mapping [268], quantifying outdoor space aesthetics [290, 269], and mapping green areas within cities [289, 317].

A major challenge when using streetscape imagery is the occlusion of the target by objects in the foreground, such as trees, vehicles, pedestrians, etc. Some researchers overcome these challenges by retrieving Google Street View images from multiple locations surrounding a target [280, 294]. Alternatively, a combination of multiple data types can be used to overcome occlusion challenges. There have been a limited number of studies where both streetscape and aerial imagery have been used to survey urban environments [280, 289, 290, 305]. In these cases, initial detections of the target domain are made using aerial imagery and streetscape imagery is then used as a high-resolution auxiliary dataset to either validate these assumptions or to gather further information on the target not visible from the aerial views. These works also focus on identifying high-level characteristics of buildings and scenes, such as whether they are residential, commercial or public. Many of these works focus on small-scale proof of concepts. However, Yin (2015) goes further to develop more detailed work pipeline to demonstrate how large-scale audits of the built environment can be automated [268]. Other than the bounding box approach by [310] and the building age estimation by [309], few works go beyond the high-level classification to identify more detailed features or characteristics of individual properties.

### 5.3.3 Geographic Data Sources and Digital Tools

To address the problem statement outlined in 5.2, any given solution to deliver surveys of houses unsuitable for EV charging should be developed with two main end users in

mind: Automotive companies and local governments. Automotive companies can use the derived data of this system to predict consumer trends, target locations to install privately operated charging stations, or sell the data directly to local authorities, District Network Operators (DNOs), or other interested parties. This data is extremely valuable to DNOs as no other mapping agencies, including Ordnance Survey (OS), currently provide information on off-street parking availability, making it difficult to predict the increase in energy requirements due to EVs in the medium- to -long-term. Similarly, local authorities need this data to understand the uptake potential of the EV market to help meet environmental targets, guide policies surrounding EV charging and plan future EV infrastructure projects to support the growing number of EVs.

The selection of any geographical data sources and digital tool sets used to develop the solution must consider these end users. This is to maximise the value of the final derived data and ensure the development is aligned with the Industry 4.0 vision outlined in Chapters 2 and 3. The availability of digital skill sets in the relevant industries should be considered to ensure the full potential of human resources can be leveraged to generate value. Similarly, the selection of any datasets should also be aligned with current industry standards to enable integration with existing datasets to create further value. Digital tools should also have the potential to be automated while ensuring transparency and simplicity.

A common source for geographic information is Open Street Maps [318], an open-source online mapping service used by many of the papers in the literature review to collect geographical information [295, 280, 290, 282, 319]. The main benefits of OSM are that it is fully open source and provides various attribute data in easy-to-use formats. Most GIS software support OSM data, and a wide range of open-source programming libraries are available to import and pre-process OSM data. A downside of OSM is that because it relies on Volunteered Geographical Information, the data quality can vary between locations, particularly in suburban and rural areas where fewer users are available to contribute data. Another disadvantage of OSM in the use case is that there is very limited attribute data for specific properties when compared to other sources such as Ordnance Survey.

The Ordnance Survey is the most extensive mapping agency in the UK and is the standard source of geographical information for local authorities and other government services. Local authorities in the UK typically have dedicated GIS officers who are familiar with OS data to support geospatial and geographical analysis works. The disadvantage of using OS data is that private companies' licences are extremely expensive. However, if a service is being provided to a local authority, access to OS data can be obtained through a Standard

Contractor Licence and derived data can be shared by both parties [320].

As part of the research, the authors collaborated with GIS Officers at various local authorities and private companies to understand how to maximise the value of the end data. It was found that a key consideration should be the exact definition of properties suitable and unsuitable for EV charging. While there are some general UK guidelines for what makes a property suitable for EV charging, this can vary between local authorities depending on the application. Any solution should be able to use a transfer learning approach to ensure flexibility for the different use cases of the end users.

A report by Field Dynamics was the only other research identified with similar objectives to identify houses suitable for EV charging. Field Dynamics is a data analytics consultancy that offers a service to map off-street parking availability for local governments using GIS tools. This method focuses on identifying houses with driveway access, using GIS techniques to measure the distance between the edge of a residential building and the adjacent road. However, limited information is available on the exact methodology and data processing steps to perform this analysis [276]. This method research was explored further by Swansea University expand on this method to identify suitable locations for installation of public electric EV charging points using GIS [5].

One local authority was working on a related project using Google Street View to survey a town to identify residential properties with various forms of off-street parking, including driveways and garages, that would be suitable for home EV charging. This survey was carried out by visually inspecting each Google Street View image in the town and collecting data manually in a spreadsheet. The government sector has been slow to adopt Industry 4.0 technologies such as Big Data analytics and the advanced programming skills required to automate workflows such as this [321].

## 5.4 Image Processing Methodology

This section discusses the methods used to develop the image processing portion of the 3 workflows. While each workflow differs slightly, they all rely on the same image processing steps using Google Street View images and pre-trained neural networks to classify properties as suitable or unsuitable for EV charging. The application of these neural networks enables large surveys of tens of thousands of households to be carried out in a matter of hours at a very low cost. This section describes the methods of data acquisition for the training and testing data, as well as details on how the models were selected and trained.

Figure 5.5: Locations from which training, testing, and case study data were retrieved.



#### 5.4.1 What makes a property EV suitable?

To identify residential properties suitable for home EV charging, a definition of an 'EV Suitable' and 'EV Unsuitable' property must be established. The current UK requirements are access to private off-street parking where a wall box can be installed [270], however, when developing a method, it is important to consider whether additional information on a property can be leveraged to produce added value for the end user. For example, the works by Field Dynamics classify a suitable property as one that has the potential for installing an EV charging point. This means they do not distinguish between houses with accessible driveways suitable for wall box installation and houses with frontal features blocking car access, such as access steps or a walled front garden. Here, the assumption is made that if sufficient space exists to park a car, the homeowner may decide to invest in adding a driveway to the property in the near future. Predictions made using this method are subject to some inaccuracy and speculation, as some homeowners may not be able to get planning permission or may not be able to afford the development of off-street parking. For these reasons, the approach is most appropriate for users considering longer-term projects >10 years.

In this study, a different approach is taken to that presented by Field Dynamics. This research focuses on providing the most accurate possible data on the location and distri-

bution of properties that are immediately ready for the installation of home charging wall boxes. The 'EV Suitable' property is therefore defined as a residential property for which there is a high probability that the property has access to at least one off-street parking space suitable for immediately installing a home-charging wall box. An 'EV Unsuitable' is defined as a residential property without clearly accessible off-street parking or where no residential property is visible. This may lead to some false negatives where homeowners may have access to some external parking access, such as a detached garage or other parking space away from the property. These instances are not considered in this study.

It is also important to note that the analysis is restricted to the property level, not the household. Due to limited attribute data, some UPRNs may be houses of multiple occupancy or converted into residential flats. However, insufficient data are available to distinguish between households. Future works should use additional data sources to identify and exclude these properties from the survey to reduce false positives.

## 5.4.2 Data Acquisition

Google Street View image data were collected from 9 different UK towns and cities, shown on the map in Figure 5.5. Six of these locations were used to collect training data, one was used to test the models and optimise hyperparameters, and the final two locations were used as case studies to analyse the accuracy of a scale full survey. These towns and cities were selected due to their varying geographical locations, local histories, and population density, all of which affect the housing style, age, and size. A diverse training dataset ensures high accuracy when classifying properties from new, previously unseen locations. The geographical data used to retrieve the Google Street View images at each training location was obtained from Open Street Maps. For each location, an .xml file was exported of all geographical information, and the road network coordinates were extracted at intervals of 20m using an open source OSM interface [322]. At each road coordinate, four images were requested from the Google Street View API at 0, 90, 180, 270-degree headings, while the pitch was kept constant at 0 degrees. During the training phase, coordinates from the entire road network are used to collect training data to ensure a wide sample that contains a diverse range of images, including images for which no houses are present.

For testing, data are gathered from four locations: Oswestry, Petersfield, Birmingham, and Gloucestershire. The Oswestry data is used in the early testing phase to explore different approaches, compare models, and optimise network parameters. Petersfield, Birmingham,

and Gloucestershire locations are used as case studies to perform surveys and evaluate the performance of each of the three workflows described further in section 5.5.

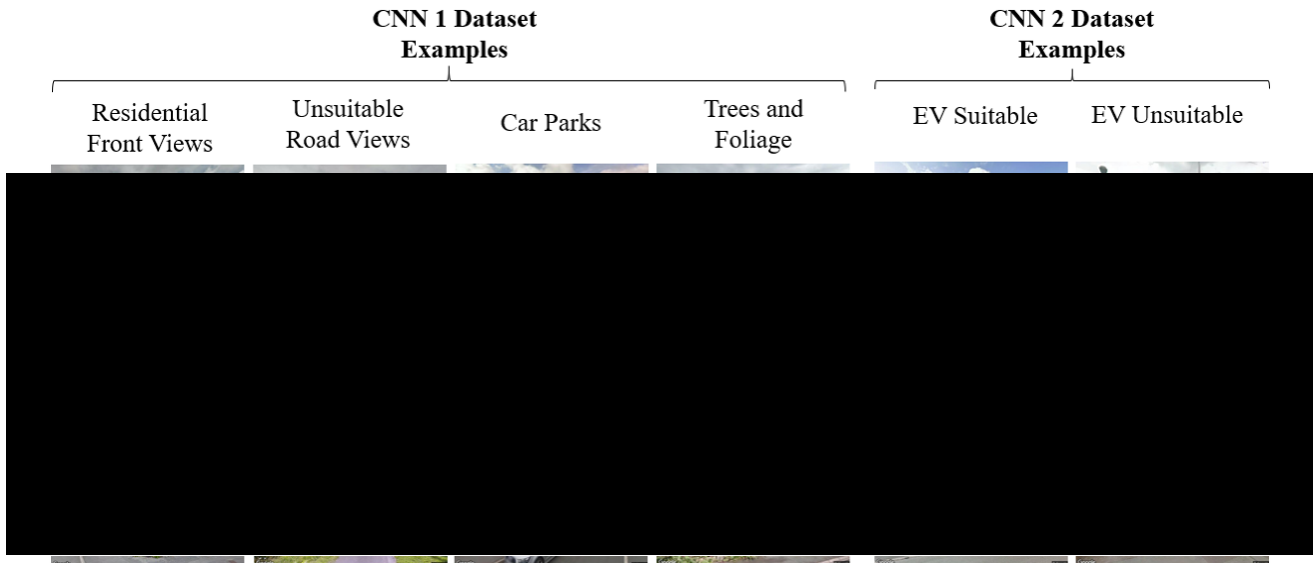


Figure 5.6: Example images from the 6 categories used to train the CNNs.

### 5.4.3 Image Classification

#### Definition of Classes

Initially, the image classification task was treated as a simple binary classification, using a single CNN to distinguish between the positive and negative classes. This approach was unsuccessful, and even when revisiting this approach using the optimised networks described in section 5.4.3, a maximum F-score of 25.4% was achieved. These early approaches likely failed due to the high diversity of the 'EV Unsuitable', leading to over-fitting. Increasing the number of positive output classes to reduce the diversity had limited success when new categories were introduced to identify photos of houses with driveways (contains driveway), houses with garages (contains garages), and all other unsuitable images (not suitable). Upon inspection of the 'not suitable' images, it was found images generally fall into 4 classes: buildings, car parks, trees/fields / other greenery, or images taken at angles facing directly down the road where buildings may be obscured by cars or trees. Examples of these images can be found in Figure 5.6.

To address the high diversity of the dataset, it was decided to introduce pre-processing step needed to filter out a large number of unsuitable images. The proposed solution is to



split the image classification task in two, using two separate CNNs in a combined workflow. The first CNN (CNN1), acts as a filter to remove unsuitable images. This is achieved by training the network to distinguish between images of 'Car Parks', 'Trees and Foliage', 'Unsuitable Road Views', and 'Residential Front Views'. Only images classified by CNN1 as 'Suitable Views' are passed on to CNN2 for further processing. CNN2 then classifies these images as either suitable or unsuitable for EV charging. This configuration of two CNNs in a series gave better results than a single CNN. Table 5.2 presents a summary of the class definitions.

While these class definitions were useful guidelines for most data, some cases presented

Table 5.2: The following definitions were used when labelling the training and testing data.

<b>Category</b>	<b>Definition</b>
Car Parks	Images of surface-level car parks, either occupied or unoccupied. This includes car parking facilities in commercial and residential areas but does not include roadside parking.
Trees and Foliage	Images of trees, fields, roadside greenery, parks, or other green spaces.
Unsuitable Road Views	Images taken at an angle facing down the road in the direction of travel or at 180 degrees to travel. For ambiguous cases where the angle is not exactly at 0 or 180 degrees, if the vanishing point of the road is visible, it should be classed as an 'Unsuitable Road View'.
Residential Front Views	Images of residential properties taken at an angle perpendicular to the road / directly facing the property front.
EV Suitable	Images showing at least one residential property with a driveway, integrated garage, or any other safe place to park and charge a vehicle that is clearly visible in the image. 'EV Suitable' properties must also meet the 'Residential Front View' criteria.
EV Unsuitable	Images of residential properties that have no off-street parking accessible by road shown in the image. 'EV Unsuitable' properties must also meet the 'Residential Front View' criteria.

challenges. For instance, some images would contain multiple instances of classes. To address this, when selecting training data for CNN1, it was decided that the positive class would always dominate. Any image containing a residential property taken at a suitable view should be labelled as 'Residential Front Views', even if trees and foliage are present and largely obscure the property. Two examples of this have been included in Figure 5.6.

Table 5.3: A breakdown of the image dataset used to retrain CNN 1. The Oswestry data is used for validation and hyper-parameter optimisation.

Category	B'pl	P'gh	S'sea	C'ran	Co'ter	Ex'th	Os'try	Total
Car Parks	349	424	308	405	90	117	200	1,893
Trees and Foliage	318	1,658	573	311	348	402	200	3,810
Road Views	354	388	434	334	337	375	200	2,422
Suitable View	1,251	937	947	1,019	972	1,033	200	6,359
Totals	2,272	3,407	2,262	2,069	1,747	1,927	800	14,484

Table 5.4: A breakdown of the image dataset used to retrain CNN 2. The Oswestry data is used for validation and hyper-parameter optimisation.

Category	B'pl	P'gh	S'sea	C'ran	Co'ter	Ex'th	Os'try	Total
EV Suitable	627	743	295	451	898	759	500	4,273
EV Unsuitable	542	379	411	94	322	415	500	3,232
Totals	1,169	1,122	706	545	1,220	1,174	1000	7,505

One example shows a house is only partially in the frame, and a fence partially obscures the driveway, this is classed as a 'Residential Front View' as the image is taken perpendicular to the road, facing directly to the house. This same image was also included in the 'EV Suitable' training dataset, as the driveway for a residential property is clearly in view. Another example shows where a property lies on the corner of a junction. This is a common case that had to be addressed and was somewhat subjective as this could be considered as either an 'Unsuitable Road View' or a 'Residential Front View'. It was decided that if the property took up more than half of the image, it would be classed as a 'Residential Front View', otherwise, it would be labelled as 'Unsuitable Road View'. However, more complex situations meant there was some subjectivity to unusual cases such as this.

CNN2 posed a more significant challenge. Initially, the definition of an EV-suitable property was based on the UK Governments' guidance for customers in its Electric Vehicle Home-charge Scheme documentation [270]. Here it says that a property needs access to off-street parking to be eligible for the scheme i.e. a garage or a driveway. However, this may include garages or parking spaces located away from the main building and out of view of the Google Street View images. For this study, EV-suitable properties are defined as residential properties with off-street parking adjacent to the house.

## Model Selection

To minimise the amount of training data required a fine-tune transfer learning approach is used. This allows us to use networks that have been pre-trained on large datasets and have already learnt to detect edges, shapes, intensity and other low-level features of images. By leveraging this knowledge of the pre-trained networks, less data is required when training new classes.

Based on similar works described in section 5.3, it was decided to compare three different




Figure 5.7: A box plot showing the range of F-scores resulting from the hyperparameter optimisation experiments for three different architectures. Googlenet performs the best for both CNN1 and CNN2.

pre-trained networks for CNN1 and CNN2: ResNet, GoogLeNet, and AlexNet. VGG variants were initially included, but the memory requirements were too high. The datasets used to compare the performance of each of the networks are shown in Tables 5.3 and 5.4, sourced from the Oswestry data. Prior to the comparison, a separate testing dataset from Oswestry was used to explore how different training parameters affected the overall accuracy, including the initial learn rates, batch size, learn rate drop factor, learn rate drop period, and L2 regularisation. For all three networks, the initial learn rate and batch size had the greatest effect on performance. A grid search approach was used to optimise these parameters for each network, the results of which are shown in Figure 5.7.

GoogLeNet achieved the highest overall accuracy for both CNN1 and CNN2. The confusion

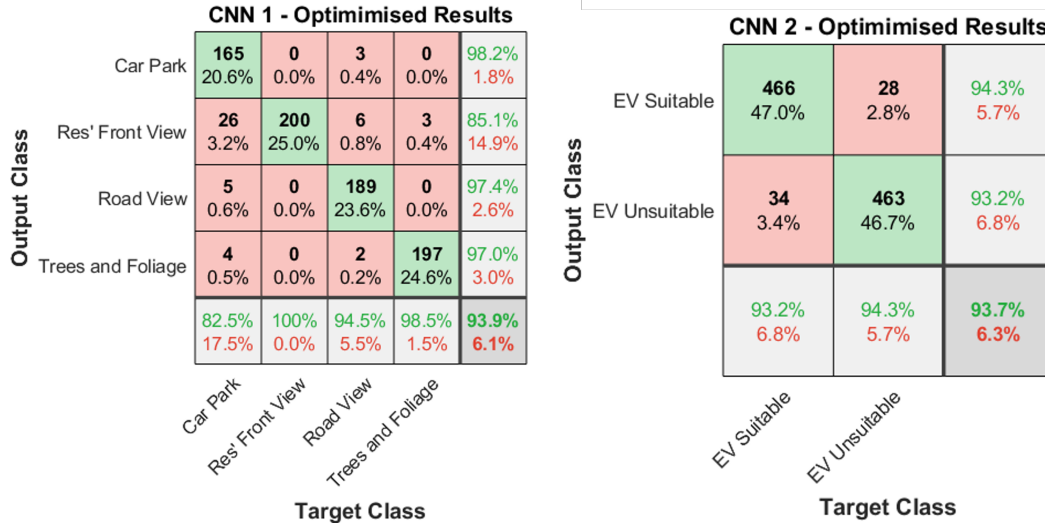


Figure 5.8: Confusion matrices of the optimal networks selected for CNN1 and CNN2

matrices from the best-performing runs are shown in Figure 5.8 For CNN1, a learning rate of  $1e-3$  and a batch size of 32 gave the best results of 93.9%. More importantly, GoogLeNet achieved a 100% recall for CNN1 at identifying the 'Residential Front Views' category. This means in this test dataset, all of the images would have been correctly passed onto CNN2. However, it is important to consider the range in results, particularly for CNN1. In real-world examples of this system, high-quality testing datasets may not always be available. Depending on the application, users may want to select AlexNet, which gave a more robust result at the expense of a lower maximum achievable accuracy. Given these results, it was decided to use GoogLeNet for CNN1, given this high recall rate. For CNN2, GoogLeNet was also selected, given the high F-score of 93.7%. The optimal model parameters are a learn rate of  $1e-3$  and a batch size of 32.

## 5.5 Proposed Workflows to Identify EV Suitable Houses

This section presents 3 workflows to deliver automated surveys of a specified area to identify residential properties suitable for home EV charging. Each of these workflows uses the image processing steps described in section 5.4 with varying pre-processing and post-processing steps to meet the requirements of different end-users. For each workflow, a survey is carried out on the testing locations to demonstrate the performance of each approach.




Figure 5.9: Flow diagram for Workflow 1.

### 5.5.1 Workflow 1

The first method uses the same image data acquisition as outlined in section 5.4.2, extracting the entire road network coordinates from OSM at intervals of 20m using an open source interface run in MATLAB [322]. This approach is advantageous when gathering training data as it ensures sufficient diversity and quantity of training data, however, it is less suitable for surveying use cases where cost is a limiting factor as many of the images downloaded will be removed by CNN1 and not used. While this approach is not cost-effective for surveying use cases, it is very useful during the development stage as all geospatial analysis, processing, and post-processing can be done by one engineer using a single software. This allows for rapid testing and the development of new ideas. For these reasons, this workflow is developed to be used by commercial entities such as automotive companies for applications such as targeted marketing and predicting consumer trends. The full workflow for the Birmingham case study is shown in Figure 5.9.

#### Birmingham Survey

The Birmingham test area was limited to roughly 4km<sup>2</sup> located east of the city centre due to the high density of houses and the limitations of the Google Street View API. The survey area is shown in Figure 5.10. A total of 28,728 Google Street View images were retrieved from this location, of which 13,962 were duplicate images and had to be removed prior to further processing. The remaining 14,766 images were manually labelled into the respective categories to produce the test dataset and evaluate the networks' performance.

The high number of duplicates is due to two main reasons. Firstly, road network coordinates are retrieved at 5m intervals, which was found to be a higher resolution than the Google Street View image locations. However, this resolution can vary between areas. When requesting an image from a given coordinate, the API returns the closest image

Figure 5.10: Survey area for the Birmingham case study.

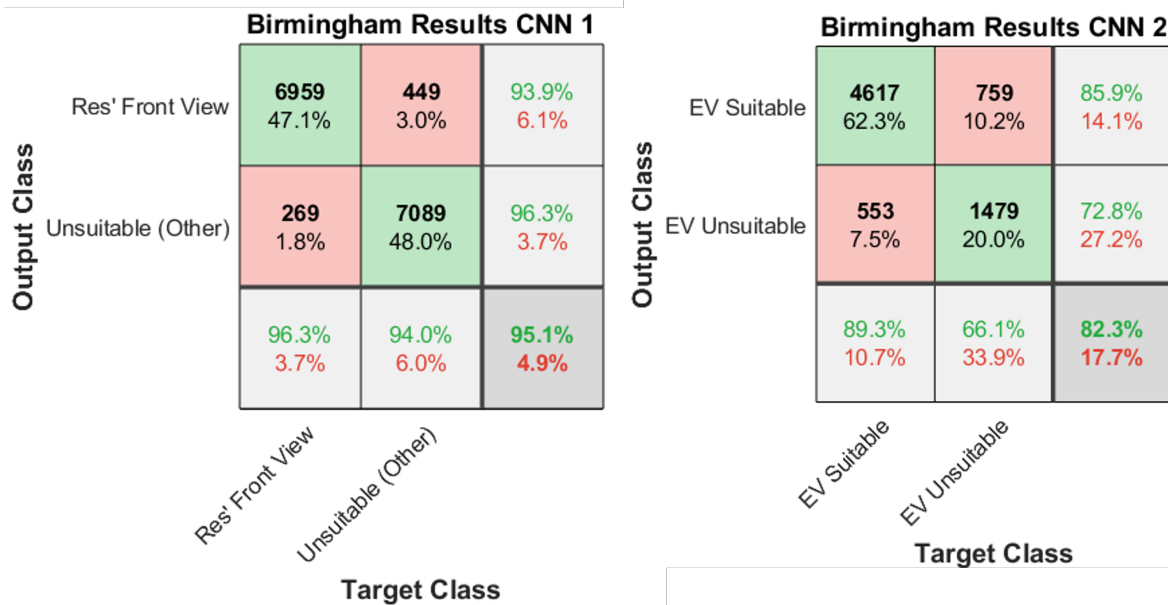
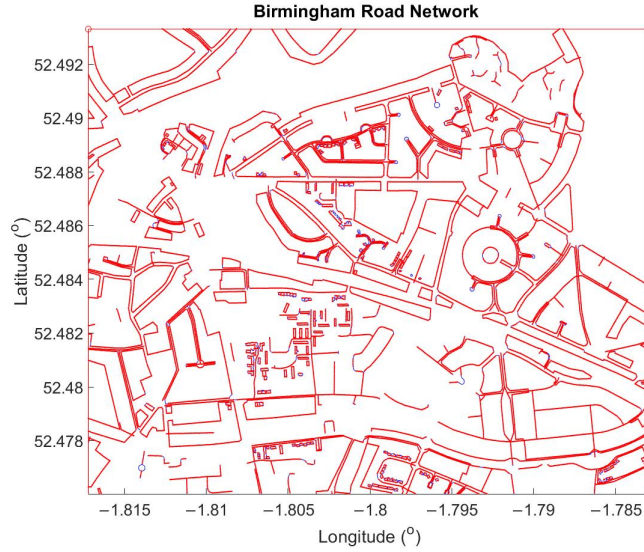


Figure 5.11: Confusion matrices showing the performance of CNN1 and CNN2 on the Birmingham dataset.

within a certain radius. There is a balance here between getting sufficient coverage and minimising duplicate downloads. Secondly, limited attributes are listed in the OSM documentation that enables the user to distinguish between types of A roads, B roads, one-way streets, no vehicle access roads, public footpaths, ect. This leads to some roads being

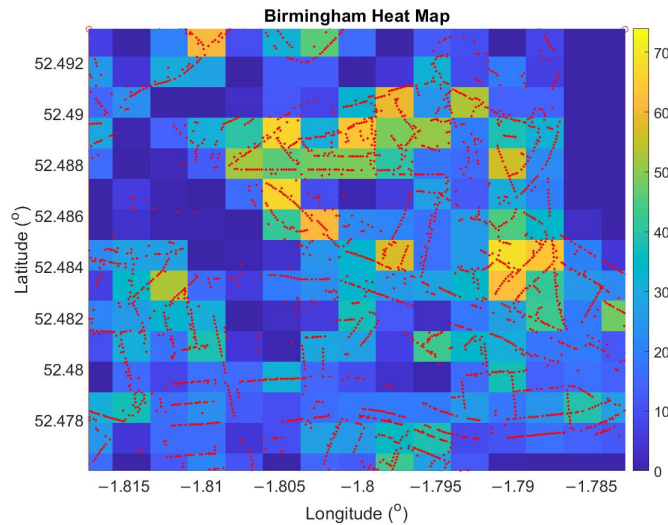


Figure 5.12: A heat map of EV-suitable properties in the Birmingham survey area.

overlayed on top of each other and also explains some of the gaps in the road network in figure 5.10.

Once the images were downloaded, they were passed through CNN 1 for pre-processing. For the Birmingham test area, CNN 1 achieves a recall of 96.3% when identifying images of Residential Front Views, as shown in Figure 5.5.1. Following the CNN 1 step, 7408 images were passed onto CNN 2 for processing, which then identified 5376 images as suitable for EV charging, giving a high recall of 89.3%, as shown in Figure 5.5.1. Overall, the system recognises 4617 of the total 5306 images suitable for EV charging. Figure 5.5.1 shows all EV-suitable properties plotted on a heat map to show which areas have the highest and lowest density of houses suitable for home EV charging point installation.

From these results, it is found that this data acquisition method does not lend itself to large-scale surveys due to the high number of duplicate images and the high number of images that do not include suitable views of houses. This is particularly true in rural areas. This method is more useful in the early development stages to gather training and testing datasets and develop custom post-processing tools such as the heat map in Figure 5.5.1. Because this method allows for the full workflow to be run in a single platform such as MATLAB or Python, this makes it more suited to integrating with existing systems and datasets.

### 5.5.2 Workflow 2

Following the findings from workflow 1, a second method was developed to address the data acquisition challenges and ensure the method’s scalability by exploring the use of alternative datasets and geoprocessing tools to retrieve location data for each property. With a proof of concept already established, it was also decided to narrow the scope of this research to focus specifically on developing a tool for local governments. As discussed in section 5.3.3, the advanced data analytics and coding skillsets required to use, develop and maintain workflow 1 are not widely available within local government. Therefore, to improve data acquisition, it was decided to use Ordnance Survey data rather than OSM data, as these data are already widely used by local authorities and provide more accurate and detailed attribute information of each property to enable more efficient retrieval of the image data.




Figure 5.13: Flow diagram for Workflow 2

### Ordnance Survey Data

OS data is readily available for all local governments in the UK and provides the most accurate attribute information on UK properties, each of which is identifiable using the Unique Property Identification Number (UPRN) associated with it. Workflow 2’s data acquisition step combines several Ordnance Survey datasets pre-processed using Geographical Information Software (GIS). Specifically, detailed information of buildings was attained from the AddressBase® dataset and property boundaries were obtained from Digimap OS Mastermap Topography Ordnance Survey datasets. The vector layers were pre-processed using QGIS 3.18 such that all non-residential buildings were removed. There are several advantages to using OS data and GIS in this way. Firstly, these data are widely used for geospatial analysis in various industries and make collaborating and sharing information



much easier. Secondly, OS data provides a range of attribute data for all properties, including geographical location and whether the property is residential or commercial. Therefore, this approach enables us to retrieve coordinate data at the individual property level, which significantly reduces data acquisition costs, making this approach highly scalable. Finally, being able to pre-process the data in GIS allows coordinate data to be generated based on specific attributes, making future works easier to adapt to other use cases.

### **Automation Considerations**

As part of this study, researchers met with GIS Officers at local authorities in England and Scotland who were interested in better understanding the local EV infrastructure. Discussions took place on how to maximise the value creation of the output data to create the greatest value for local governments in a surveying use case. It became clear that due to varying regulations, political motivations, and environmental targets between the local authorities, they each had different ideas of how the system might be used to deliver the greatest value. For example, one local authority had a longer-term view of EV suitability, defining EV-suitable houses as those at which a charging wall box could feasibly be installed immediately or in the near future if the property owner was to modify some external characteristics. Examples are cases where a front garden could be converted into a driveway, which may involve removing a fence or hedgerow. This highlights one of the major advantages of combining image classification with GIS in this way. When transferring this methodology to a new local authority with different but similar class definitions, the networks can be easily retrained using a new training sample developed collaboratively with local GIS officers to meet the demands of that specific end-user.

### **Dataset Validation**

To ensure the quality of the training and testing datasets, the authors collaborated with GIS officers to validate a sample of 100 images selected from both the 'EV Suitable' and 'EV Unsuitable' classes. Given that in this study only considers properties that are immediately suitable for home EV charging, the GIS officer the following definition for this labelling task:

Table 5.5: Geographical data sources

<b>Data</b>	<b>Scale</b>
OS MasterMap Topography	1:1 250
OS MasterMap Highways-Roads	1:2 500
OS Open UPRN	1:1 000
2011 Census MSOA Code: E02004708	MSOA

*Images showing at least one residential property with a driveway, integrated garage, or any other safe place to park and charge a vehicle that is clearly visible in the image.*

Following this guidance, the GIS officers' labels agreed with ours for all 200 images. With this result, the datasets were assumed to be of sufficiently high quality, and it was decided not to revise the training or testing datasets for either the Petersfield or Birmingham locations.

## **Petersfield Survey**

To test workflow 2, it was decided to explore a more rural area to see how the network performed compared to the more urban environment of the Birmingham study. For the Petersfield test dataset, the survey area is defined by the Middle Super Output Area code MSOA E02004708. This MSOA encompasses 4,317 buildings included in the survey giving a total of 17,268 images downloaded. A map of the survey area is shown in Figure 5.15. For this paper, QGIS 3.10 software was used to gather coordinate data for all properties within the survey area. The full workflow for the Petersfield case study is shown in Figure 5.9.

At the Petersfield location, of the total 17,268 images downloaded, 1,700 image requests failed to download and 8443 duplicate images were removed. Because Petersfield is a much more rural area, it is assumed that the failed downloads are either because no Google Street View data has been gathered in these specific locations, or these locations are instances where the centroid of the building polygon is too far away from the nearest public road for Google to recognise the request. This may be due to private road access to a property, or a very large building polygon. Of the remaining 7,125 images passed onto CNN1 for pre-processing, 3,295 images were identified as suitable views. The confusion matrix for CNN1

and CNN2 are shown in Figure 5.14, achieving F-scores of 93.6% and 88.7% respectively. Figure 5.15 shows all suitable and unsuitable houses plotted on a map of the survey area. Once again, CNN 1 is highly accurate at filtering out unsuitable images despite the change to a more rural environment. CNN 2 achieved a much lower false negative rate in this rural location while maintaining a similar false positive rate. However, given the different data acquisition methods for workflow 1 and 2, it is not possible to draw direct comparisons. This is addressed in the next case study presented in the following section.

The majority of the time taken to complete this survey lies in the GIS data acquisition and the Google Street View API requests. With the lessons learned from the Birmingham case study, this workflow allowed the author to survey all 4300 houses within a single working day. This excludes time spent labelling data and training models, which are significant tasks. However, this only needs to be completed once, after which multiple surveys can be quickly conducted at any location with similar architectural characteristics to the training dataset.

Figure 5.14: Confusion matrices showing the performance of CNN1 and CNN2 on the Petersfield dataset.


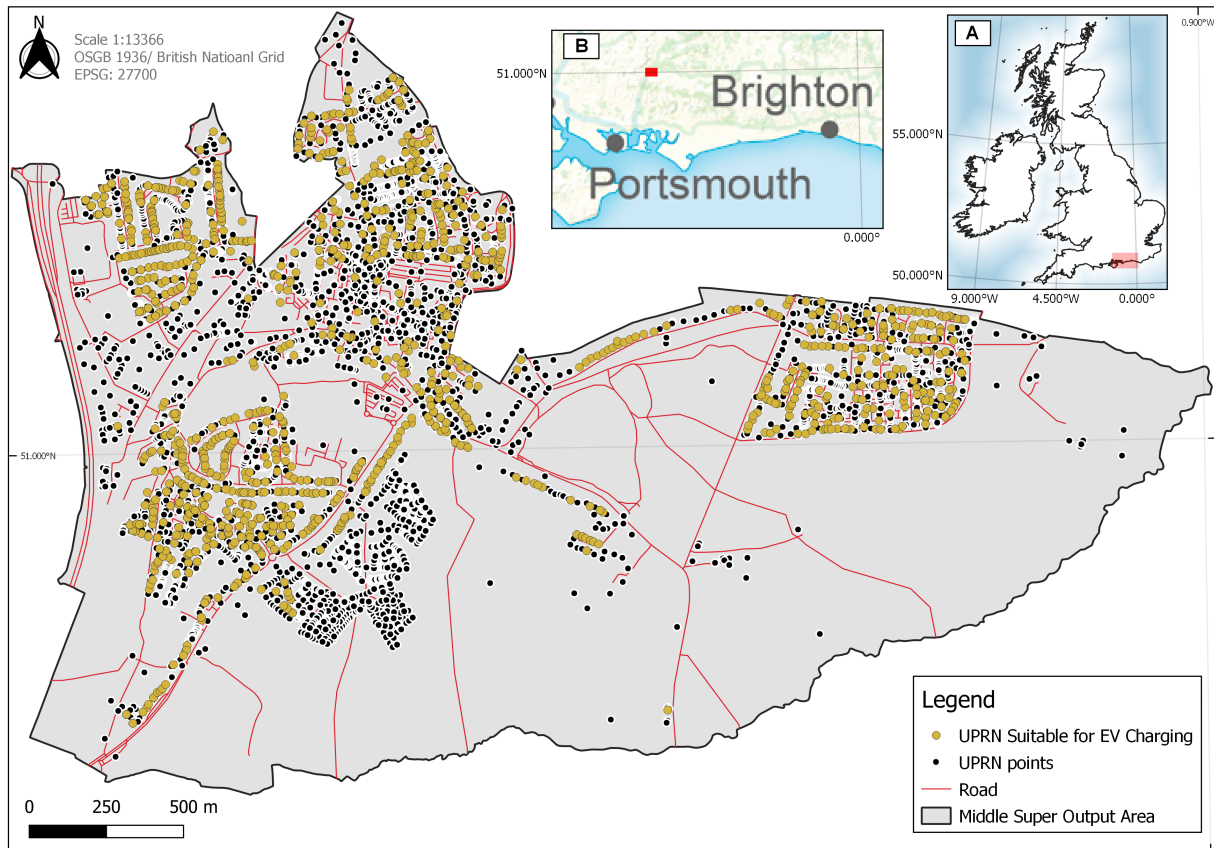


Figure 5.15: A map showing all locations of suitable and unsuitable properties in the Petersfield survey area.



### 5.5.3 Workflow 3

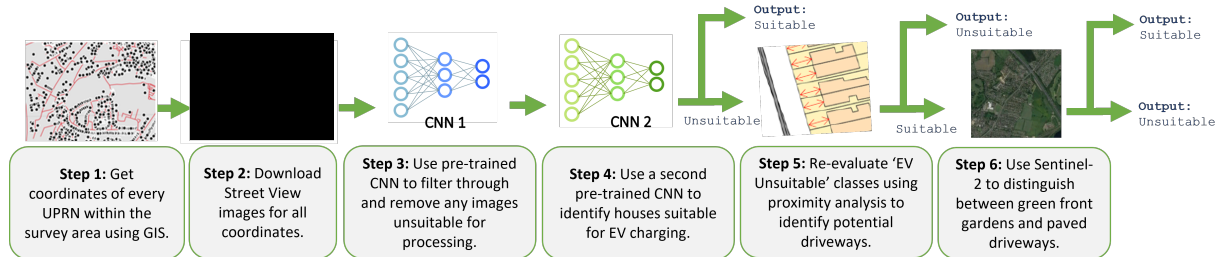
From the findings from the Birmingham and Petersfield case studies, further opportunities to improve the methodology are identified. Firstly, these findings suggest that by improving the methods of data acquisition, further reductions in the number of failed downloads can be achieved and improved accuracy. However, further testing is required to validate this, as the two test locations have contrasting population densities and architecture. Secondly, the property is immediately labelled as unsuitable at UPRN locations where no Google Street View data exists, resulting in a higher number of false negatives. To address this challenge, alternative datasets must be explored.

This section addresses these opportunities by introducing a third iteration of the methodology and performing two additional surveys at locations in Gloucestershire. Workflow 3 expands on Workflow 2, focusing on local authorities as an end users and aiming to reduce

the false negative rate by adding two additional processing steps that utilise additional data sources. Prior research shows that use satellite imagery can be combined with streetscape imagery to achieve high accuracy of building instance classification [280, 289, 290]. Drawing on lessons from this research, this section explores how extracting information on building footprints can be used to infer additional information on the architectural features of a building to inform the classification model. To achieve this Sentinel-2 optical satellite imagery is used, as well as building footprint information from the OS datasets.

A diagram of Workflow 3 is shown in Figure 5.16. By applying the 'double-check' on UPRNs classified as unsuitable by CNNs, the predictions can be re-evaluated to reduce false positives and provide a means of analysing UPRNs for which no Google Street View images are available.

Figure 5.16: A flow diagram for Workflow 3.



## GIS Proximity Analysis

In section 5.4.1 research by Field Dynamics is discussed, where GIS tools are used to identify properties with a high likelihood of off-street parking based on the space between the building footprint and the adjacent road. Using the average UK parking size of 2.4m x 4.8m, the proximity analysis method assumes that if insufficient space exists to park a car, then this property is classified as unsuitable for EV charging. If sufficient parking space does exist, the UPRN is passed onto an additional processing step to check the vegetation index of the frontal space using Sentinel-2 data to distinguish between a paved driveway and a green front garden. The size of the frontal areas is calculated using the OS Topography Layer processed in QGIS by measuring the distance between the edge of the building polygon to the adjacent road.

## Sentinel-2

Sentinel-2 is a multi-spectral satellite with 13 bands, providing high-resolution images that have been widely used for LULC applications such as distinguishing households plots [18], mapping urban areas [17], and analysing public urban green spaces and private gardens [19, 20].


If a UPRN has been found by the proximity analysis to have a suitable space for off-street parking, Sentinel-2 data are used to determine if this area is either a paved driveway or a green front garden. This is done by calculating the Normalized Difference Vegetation Index (NDVI) using bands 4 and 8, which capture the red and near-infrared at 10m resolution. NDVI is a common means of detecting vegetation in remotely sensed images. On a scale of -1 to 1, values between -1 and 0.19 indicate surfaces similar to rock or concrete, while values 0.19 and above indicate vegetation. For the selected UPRNs, the vector polygons from the OS MasterMap Topography Layer are obtained to extract the frontal area for calculating the NDVI value.

This combined approach of proximity analysis and Sentinel-2 assumes that the frontal area of any property that is large enough to park a car and has a low vegetation index is a driveway. This is not always the case, as it may be a paved front garden or other areas inaccessible from the road. However, these cases are assumed to be the exceptions, and while this will result in a slightly increased false positive rate, the overall accuracy should be improved overall due to the reduced false negative rates.

## Gloucestershire Survey

To test this methodology, additional test datasets are developed using images from two postcode locations in Gloucestershire, GL55 and GL53, as shown in Figure 5.17. Gloucestershire County Council declared a climate emergency and aims to reduce their carbon output by 80% by 2030 and become fully carbon neutral by 2050. At the time of this study, these locations have limited public charging points, making these areas an ideal study site to support the future uptake of EVs. The GL53 location is a rural area with low population density, making it a good location to test the new approach to identifying driveways that are likely to be occluded by greenery. The GL55 location is more urban, with parts of the area within Cheltenham, Gloucestershire's second-largest town with a population of 117,500. Surveying these two locations will not only allow us to directly compare the different methodologies but also allow us to see how each method performs in

Figure 5.17: Gloucestershire Map [5]



different environments.

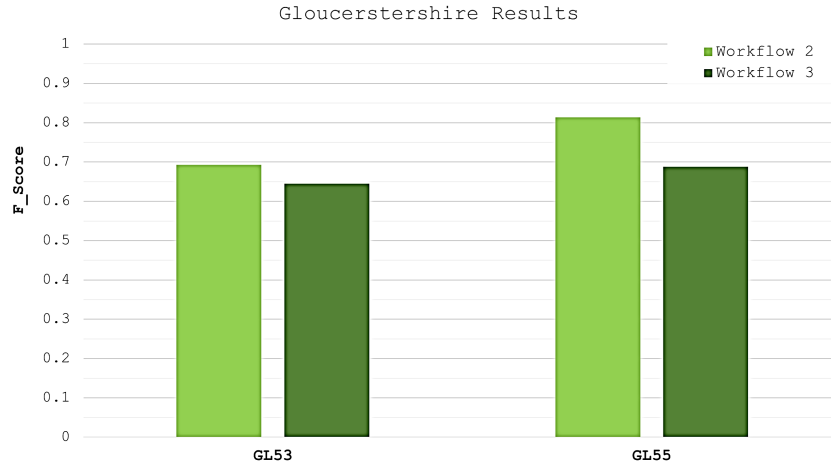
To evaluate this method at this new location, a random sample of 200 UPRNs was selected from each location and the respective Google Street View images were labelled by a licenced EV charger installer. This resulted in a dataset of 157 households from GL53, and 172 households from GL55 that were successfully downloaded to be used in the test datasets. The discrepancy between these numbers and the original sample of 200 is due to some Google Street View images failing to be downloaded.

In this comparison, workflow 1 was not included due to its different method of data acquisition, which prevents us from retrieving the same image data for an accurate comparison.

The results for the Gloucestershire locations are shown in Figure 5.19 with F-scores of 69.3% and 81.4% at GL53 and GL55, respectively. At both locations, the addition of the proximity and multi-spectral analysis steps resulted in a slight increase in precision due to the decreased false negative rate. However, there is a significant reduction in the recall due to a further increase in the false positive rate. As a result, the overall F-score of method 3 is lower than method 2 at both test locations.

Further investigation finds that one of the most likely reasons for this significant increase in false positive rate is the low resolution of Sentinel-2. Sentinel-2 has a resolution of 10m, while the minimum size of the parking space to be analysed was 2.4m x 4.8m. This low resolution means that nearby features will affect the NDVI value, such as roadside

Figure 5.18: Workflow 3 F-scores



trees, neighbours’ gardens, or other greenery, within the target property’s frontal space. Because the false positive rate is increasing, the pipeline incorrectly reclassifies unsuitable properties as suitable during step 6.

## 5.6 Discussion

Section 5.5 presents 3 methodologies to identify residential properties suitable for EV charging. Each of these methodologies is built around a pair of CNNs that perform image classification in series to first filter out images that are unsuitable for processing and then classify the remaining images as either suitable or unsuitable. It is demonstrated that using a CNN as a pre-processing step to identify and filter out unwanted images is highly effective, achieving recall rates of up to 100%. It is also shown at four separate locations that the three variations of this general method achieve high accuracy.

Figure 5.15 shows how a local authority or other commercial entity might use this workflow for a range of urban planning projects. For example, in the bottom right of the map there is a cluster of houses highlighted in black that are unsuitable for home EV charging. This highlights a potential location for installing community charge points to improve the accessibility of green transport options within the local area. Local authorities can use this data to support EV uptake as part of efforts to meet clean air targets outlined by the UK government and the Paris Climate agreement. Once models are trained and tested, being able to conduct surveys like this in a matter of hours represents a major advancement in surveying technology compared to current manual methods.



Figure 5.19: The confusion matrices of the Gloucestershire test locations for Workflows 2 and 3.



Using the results from the Petersfield location, a rough cost estimate can be made to deliver a survey of this type at the local authority level. Assuming the same number of images would be required to survey each of the 15 MSOAs in East Hampshire District Council, including the Petersfield MSOA, and given a cost of \$0.007 per image download from the Google Street View API, a total cost of \$1930 is estimated, excluding labour. Licensing costs can also be excluded as OS data can be accessed under a Standard Contractor Licence with UK local authorities. Using this same approach, to survey Birmingham, the largest English authority containing roughly 442,000 dwellings, the estimated total cost scales to approximately \$20,000 +- \$5000 when considering duplicate images. Furthermore, once image data is acquired from an area, these same images can be easily repurposed for additional image classification tasks for surveying use cases. Further work should explore adapting this methodology to identify other characteristics of residential properties to provide alternative surveying services. Possible research directions may include estimating the energy efficiency of homes or identifying properties suitable for solar PV installation.

The methods for data acquisition in methods 2 and 3 were found to be much more effective, resulting in fewer duplicate images and providing more useful output data for real-world applications when collaborating with local councils and other researchers who often required data on individual properties, identifiable by their UPRN and respective coordinates. The data acquisition approach used in Workflow 1 made it difficult to link images to these UPRNs, hence why this method was not used in the comparison in section 5.5.3.

While the improved method does result in a high level of accuracy, a significant number of duplicate images still needed to be removed in pre-processing. Further works should improve the GIS data acquisition to ensure the coordinates of each building polygon are retrieved as close to the road as possible to ensure the highest accuracy when requesting the respective image from the Google Street View API.

Compared to the Petersfield survey, the Gloucestershire results show a reduction in F-score of around 15%. In both survey areas, the false positive rate was found to have the greatest negative effect on the overall accuracy of the method. Upon inspecting the types of images that result in false positives, there are three main attributes that a large proportion of these images share in common. Firstly, are images of properties taken on the corner of a road intersection. Secondly, are properties with large front gardens with either low fences or no fences at all. Finally, are images of car parks that are adjacent to residential properties. While these features result in a large number of miss classified images, there

are also many more unique architectural quirks that also lead to false positives. Given the specificity and high variance of these features, it is impractical to address these issues by modifying the training datasets as this would add considerable time to what is already the most laborious process in developing and replicating this research. Instead, it is considered how prior research has addressed the challenges of using streetscape imagery by using aerial and satellite imagery as ancillary datasets [280, 289, 290]. By extracting information on building footprints, additional information can be obtained on the architectural style of a building that can be used to inform the classification model [280]. Such building footprint information is already available in the existing OS datasets.

## 5.7 Conclusion

In this chapter, three novel methods are presented to automate current approaches to survey the external characteristics of residential properties. Specifically, this research focuses on identifying houses with private off-street parking, a key requirement for installing home EV charging systems in the UK. The methods are developed with two main end users in mind: automotive companies and local governments. It is discussed how these industries can use these tools to perform large-scale surveys to support the future development of EV charging infrastructures in both the short- and long-term, as well as discuss the use of these tools for target marketing of EV products.

In developing these methods, a novel image processing pipeline is presented using two CNNs in series to address over-fitting in a classification task where labels are ambiguous and there is high diversity within classes. An initial CNN is used to filter out unwanted images, while the second CNN performs the binary classification of 'Suitable' or 'Unsuitable'. It is demonstrated that using a CNN as a pre-processing step to identify and filter out unwanted images achieves recall rates of up to 100%, leading to high overall accuracies of the full pipeline.

The methodologies presented in the chapter are shown to achieve high levels of accuracy in identifying residential properties suitable for EV charging. Three case studies are carried out at different locations in the UK to demonstrate the models, each of which is developed with different end users in mind. Each of the models achieves high accuracies in identifying EV-suitable properties. It is found that the methods that rely solely on Google Street View imagery have some limitations, achieving slightly lower accuracies in rural

areas where residential properties are obscured by roadside greenery such as hedgerows and trees. To overcome this, Workflow 3 also incorporates satellite imagery to improve performance. However, the low resolution of satellite imagery was found to be insufficient to improve the system's accuracy.

The methods developed in this chapter represent a major advancement towards fully automated remote surveying capability to audit the built environment. Further research is recommended to explore the use of this technology in identifying suitability for photovoltaic installation, external wall insulation, and other factors to help improve housing sustainability. In automotive applications, further research is also required to explore the issues relating to the intellectual property of the derived data and how to create value from these data insights. Possible directions may include a web-based application for local governments to access these data as a service to help improve EV charging infrastructure. With further advancements in autonomous driving, Ford should consider developing their own streetscape datasets to overcome the IP challenges and explore further use cases for data-as-a-service models such as these.

# Chapter 6

## Conclusions

The automotive industry is in the middle of a major change due to new environmental policies, changing consumer demands, and rapid advancements in data-driven manufacturing technologies. These factors have resulted in an industry-wide shift towards electric powertrain, which is widely expected to become the dominant mode of powertrain in the near future [9]. Ford Motor Company are still in the early stages of the EV revolution, which requires considerable investment to develop new manufacturing solutions for battery and rotor-stator manufacturing and new product lines.

In addition to the challenges of the transition towards electric powertrain, the automotive industry also finds itself amid a wider industrial revolution as emerging digitalization and automation technologies and business practices are redefining the manufacturing sector. This fourth industrial revolution, or Industry 4.0, presents additional opportunities for automotive manufacturers to digitalize and integrate all areas of business operations throughout the product life-cycle, delivering productivity improvements and reduced operating costs, as well as the opportunity to explore new data-driven business models to remain competitive in changing automotive markets.

Successfully managing the transition to electric powertrain and Industry 4.0 will require high levels of technological and organisational innovation that challenge the existing cultures of the automotive industry. It is widely agreed that data will become an organisation's greatest asset in delivering this change, as well as people with the advanced data analysis skills to create value from these data. Therefore, this thesis explores the sustainable application of machine learning solutions in the automotive industry.

In this thesis, a comprehensive review of machine learning applications in the automotive industry is presented. The review offers new perspectives on how to accelerate the

development of machine learning technologies to maximise the value creation of existing data value chains and how to mitigate the financial, strategic and cultural risks of these technologies. In addition, through multiple case studies, further reviews are conducted to critically analyse the most recent literature on machine learning in anomaly detection, image recognition, change management, and digitalisation strategy. From the findings of these reviews, a strategic framework is developed to support further organisational growth in the automotive industry. The framework provides senior management at the factory level with a stepwise roadmap to guide digitalisation and automation strategy to achieve Industry 4.0 objectives and identify innovation opportunities. In collaboration with industrial partners, this framework is used to perform a gap analysis and assess current Industry 4.0 maturity at three Ford UK manufacturing sites. The results of these assessments are presented, and multiple innovation projects are identified. Two of these innovation projects involving the applications of machine learning in the automotive industry are explored in-depth, providing further examples of how opportunities and barriers to machine learning adoption.

The innovations in Industry 4.0 change management and applied machine learning presented in this thesis are important steps in developing organisational knowledge for manufacturers to ensure a sustainable transition towards EV powertrain and Industry 4.0. The significance of these works is recognised by the industrial sponsor, and the proposed strategic framework has since been adopted by the European powertrain manufacturing teams to support ongoing digitisation efforts in Europe. The research presented in Chapter 4 has led to the development of a dashboard to deliver time savings at Valencia and Dagenham to label data for machine learning projects. This dashboard is now being developed by Ford's R&D teams as an internal application to support future machine learning projects in the company. Furthermore, a trial is ongoing to implement the machine learning solution presented in Chapter 4 at Dagenham Engine Plant, with further plans to expand this solution in vehicle operations at Valencia and Cologne.

To conclude, multiple advancements are presented that contribute to the academic literature as well as developing and implementing sustainable manufacturing solutions to deliver added value to current manufacturing operations. The proposed industrial solutions are supported by clear stepwise frameworks to develop, implement and replicate these solutions across the business. These contributions have been used to support ongoing digitalisation efforts, implement state-of-the-art anomaly detection solutions, and explore new data-as-a-service business models in one of the world's largest automotive companies.

## 6.1 Main Research Contributions

The work presented in this thesis contributes to the current research in applied machine learning, Industry 4.0 management, and remote sensing. Some of the research has been disseminated through journal articles and conference papers [323, 271, 324, 325]. A summary of the main contributions of the research concerning the original objectives are as follows.

**Objective 1:** *Identify the main machine learning technologies used in automotive manufacturing, identify the barriers and opportunities for further sustainable growth and value creation in this field, and understand the current barriers.*

A comprehensive systematic literature review of machine learning applications in the automotive industry is presented in Chapter 2. This review includes a low-level description of the most popular machine learning models used in manufacturing, as well as critically reviewing how these models are applied across various automotive manufacturing use cases. This review provides new insights into the opportunities of these technologies and the technological and cultural barriers preventing their further uptake in the industry.

The review finds that many papers lack of information on the methodology applied when conducting machine learning experiments, as well as the potential impact of the proposed solutions. Many studies are found to lack information on data collection, labelling, cleaning, and validation in existing research, as well as the absence of clear threshold limit-setting in anomaly detection applications. Few researchers also consider the economic impact of the proposed solution. Future research in machine learning applications must ensure sufficient information on these aspects is presented to support future research as well as the implementation of these solutions in industrial settings. These findings are carefully considered in our own research in chapters 4 and 5, where our own experimental methodologies are described in detail and resource requirements to implement these solutions are estimated. The literature review also explores cultural and organizational challenges associated with the adoption of machine learning technologies in the manufacturing industry. Common findings in the reviewed literature include a lack of data science skills and a reluctance to embrace emerging technologies due to complexity and uncertainty. The study also identifies limited knowledge of machine learning requirements at the managerial and process levels as a significant barrier to the development and implementation of these technologies. These findings highlight the need for human-centric change management approaches and strategies to drive innovation, and overcome cultural barriers in the automotive industry.

These findings are addressed in Chapter 3 where best practices in change management are incorporated in the Industry 4.0 assessment framework to support the future adoption of machine learning and Industry 4.0 technologies in automotive manufacturing.

In addition to these research findings in Chapter 2, further reviews are conducted in later chapters. Chapter 3 compares various maturity models to deliver Industry 4.0 in the automotive industry. In reviewing these models and exploring their use at Ford Motor Company, multiple shortcomings of these approaches are highlighted to explain why these assessments often fail to provide a clear roadmap to guide change in practice.

Chapter 4 reviews machine learning methods for anomaly detection time series focusing on semi-supervised approaches. Best practices to apply these methods using real-world data are presented, as well as the common challenges faced when acquiring these data. Chapter 4 includes low-level descriptions of dimensionality reduction approaches and machine learning models for anomaly detection. In Chapter 5, popular image classification algorithms and their use in auditing the built environment are reviewed.

**Objective 2:** *Development of a strategic framework to support the future uptake of machine learning in the automotive industry with a focus on sustainability.*

Chapter 3 addresses research gaps identified in the literature review by presenting a maturity model to assess Industry 4.0 readiness in automotive production facilities and develop a sustainable roadmap towards Industry 4.0. The assessment focuses on accelerating digitalisation efforts and identifying opportunities to advance data analytics capabilities. These aspects are key prerequisites to delivering machine learning solutions and maximise the value creation of existing data.

Previous research explains the slow adoption of Industry 4.0 in the automotive industry due to the limited availability of skills, poor change management, and a lack of organisational knowledge [39]. Furthermore, researchers suggest that this problem is more complex than a skills shortage and that the Industry 4.0 paradigm at its core needs to be better aligned with social sustainability goals [39]. The research in this thesis supports these claims, and therefore, the maturity model focuses on social sustainability and human-centric innovation approaches.

The proposed maturity model is applied across three of Ford Motor Companies' manufacturing plants in the UK, and the results of these assessments are critically reviewed. The industrial impact of the assessment is discussed in detail. Several projects are identified that aim to drive technological and organizational innovation, with tangible improvements already observed in simulation models, production efficiency, and upskilling initiatives.



This research also demonstrates the value of the proposed assessment framework in guiding strategic and organizational changes at a major automotive company. The questionnaire findings provide quantitative information and clear visualizations that support proposed changes, while the interview findings help identify innovation and replication opportunities. The debrief meeting was found to be a crucial aspect of the final methodology, as it facilitates discussion and action planning among senior management.

In addition to the industrial impact of this research, these findings also contribute to the academic literature by identifying and validating barriers related to data science skills, knowledge gaps among management, and the broader applicability of barriers to complex data-driven solutions in the context of Industry 4.0. It underscores the need for comprehensive data science skills training and talent acquisition in order to support future digitalisation and automation opportunities. The insights gained from this research can inform organisations and policymakers in addressing these barriers in order to further facilitate the successful implementation of Industry 4.0 technologies and practices for improved business outcomes.

**Objective 3:** *Using the proposed framework, develop machine learning solutions to create value from existing data sources in the automotive industry.*

Chapter 4 presents an industrial case study exploring anomaly detection in DC Nut Runner assembly processes. DC Nut Runner is a complex anomaly detection problem where anomalies are not outliers. In developing this solution, we introduce the novel concept of the 'Anomaly No Concern' class, in addition to the typical labels of 'Normal' and 'Anomaly'. Introducing this new term helped address knowledge gaps between data scientists and domain experts by highlighting conditions where some processing error had occurred but could be clearly explained as something that would not impact part quality or require any maintenance actions. The inclusion of the ANC class became a key consideration throughout the model development and testing to help clean data, build testing and training data, and address disagreement when labelling data. Furthermore, the ANC class provided further insights into model performance when analysing the results and can be used as further justification for the business case when estimating the solution's impact on quality metrics. A data labelling dashboard is also developed. The value of this proof of concept has been recognised by the company and is now being developed as an internal dashboard to support future data labelling efforts.

Further work carried out in Chapter 4 has contributed to successfully delivering the

ADAPT project, a machine learning strategy to address production anomalies and enhance quality in powertrain manufacturing. The author worked as part of a global cross-functional team to develop and implement a novel anomaly detection solution to detect outliers in nut rundown assembly processes. Project ADAPT has been successfully implemented in two trials at Ford's Dagenham engine plant and is estimated to deliver greater than £10m per annum savings per plant. Following the success of these trials, the anomaly detection method has demonstrated application readiness and is currently being rolled out globally within Ford Motor Companies' manufacturing operations. As a subject matter expert in machine learning within Ford's Power Train Manufacturing Engineering team, the author continues to work on project ADAPT to explore further opportunities for cost savings and quality improvements elsewhere within the company. Further applications of the authors' contribution to the ADAPT project have since been identified in vehicle assembly operations, with ongoing collaboration with the Cologne Vehicle Assembly plant to deliver further process optimization.

A key finding from the literature review is that few researchers consider the economic impact of proposed real-world manufacturing solutions. This research shows that without this information, organisations with low Industry 4.0 readiness, or those in the early stages of their digital transformation, find it difficult to justify investment in innovation projects. Without understanding the value proposition of the proposed solution, managers find it difficult to present a business case to invest in these technologies. Therefore, the case study in Chapter 4 focuses on the proposed solution's economic sustainability. Current investment strategies in the automotive industry are discussed as a barrier to innovation. This expands on the research presented in Chapter 2, providing a practical example of the challenges of estimating the ROI of innovation projects and how this should be considered when economically appraising projects of this type. In addition to the industrial impact of the proposed solution in Chapter 4, this research has also contributed to the academic literature. A paper has been submitted to the journal 'AI' and is currently under review [325].

A second case study is explored in Chapter 5 explores how open source data can be used to automatically survey the built environment and identifying houses with private off-street parking. Off-street parking is a key requirement for installing home EV charging systems in the UK, and understanding the current infrastructures to support the EV transition is valuable to predict consumer trends and target marketing. Furthermore, the value of these data for other industries is also discussed, including local governments and district network

operators. For these industries, this novel approach to surveying the built environment is a valuable tool to understand the future energy grid infrastructure requirements as energy demands change with widespread EV ownership. Local governments can also use these survey data plan and communicate grant schemes to promote EVs' uptake to help meet sustainable development goals. As Ford Motor Company continues to explore data-driven business models, the value of these data presents new opportunities for the company to promote its products, add value for customers, and deliver data-as-a-service business models. While some issues are identified with using Google Street View data for commercial purposes, as Ford Motor Company continues to explore self-driving technologies, these challenges will be overcome by using the company's own streetscape imagery instead. In addition to the business opportunities of this case study, the methods developed in this chapter represent a major advancement towards fully automated remote surveying capability to audit the built environment. This research resulted in 3 published papers in the field of AI and remote sensing [323, 271, 324].

## **6.2 Future Work**

The transition to Industry 4.0 and the adoption of machine learning technologies is a complex process, requiring organisational changes that challenge well-established business practices. Ford Motor Company is in the early stages of the Industry 4.0 transition and machine learning adoption in its manufacturing operations. The work presented in this thesis addresses some of the major strategic and technical barriers to further adoption. However, further work is required to continue this development in both academic and industrial settings.

### **6.2.1 Industry 4.0 Assessment Outcomes and Future Work**

The Industry 4.0 assessment presented in Chapter 3 identified multiple opportunities to develop, implement, and replicate machine learning applications and related technologies in Ford's UK manufacturing sites. Multiple projects were identified from these assessments to drive technological and organisational innovation at the respective sites. Two of these projects were developed and implemented by the author are discussed in detail.

The first of project explored how AGV vehicle monitoring data and error logs could be used to improve the company's simulation models as well as reduce the cycle times of AGV

routes on-site. This information has led to Ford Motor Company's simulation models being updated to reflect real-world AGV cycle times. This data has also been used by on-site material handling teams to support route cause analysis of AGV errors. The second project presents a new project management tool to address a gap in the company's strategy to support digital growth in sites with low Industry 4.0 readiness. The tool is designed to be used by department managers and supervisors to identify digitalisation and automation projects that can deliver quick wins using existing skills within teams. This tool has since been used by Halewood teams to support local digital growth, identify skills gaps and guide training and hiring strategies related to digitalisation and process automation. Within the scope of this thesis, a limited number of these opportunities could be explored in detail. Further work is ongoing within the sponsor company to continue the development and implementation the following projects that were identified as a direct result of the research presented in Chapter 3:

- **Cobot Voice Command Systems** Cobots will play an important role in future assembly operations. Opportunities to deliver process improvements and improved safety in Cobot systems were identified. These findings resulted in a collaborative research project between Ford and Swansea University to explore voice command systems for Cobots.
- **Decentralise IT Skills** The results of the Industry 4.0 assessment at 3 manufacturing sites and interviews with 121 employees across multiple organisations find the most significant barrier to digitalisation and Industry 4.0 at Ford Motor Company is the centralisation of IT skills at the plant level. Further work is required to decentralise these skills and ensure people on the shop floor have skills in IP addressing, networking, MQTT, IT standards, and IoT infrastructures. Multiple examples are identified where a lack of access to these skills has prevented innovation and IoT implementation in logistics, production, maintenance, and process teams. As Ford Motor Company continues to transition towards higher levels of data and systems integration, and adopts Cloud infrastructures, the demand for these skills will increase.
- **Cobot Implementation at Halewood** Multiple Cobot systems will be introduced at Halewood with the upcoming launch of a new EV transmission line. The Industry 4.0 assessment identifies a knowledge gap at the Halewood site related to these systems, particularly in process and maintenance teams. Opportunities are identified to install Cobots on existing lines fitted with computer vision for quality inspection.

These opportunities are of minimal risk to production and low cost. Installing these systems ahead of launch is important to proactively address skills gaps and deliver training on Cobots and computer vision.

- **Digitalising Kanban** Some manufacturing sites still use paper-based Kanban boards to manage production. Transitioning to digital project management in production areas presents significant risks. The proposed strategy to address these challenges is presented in Chapter 3. Further work is required to deliver this solution at the Halewood Transmission Plant.

Due to time constraints and COVID-19 restrictions, the long-term impact of proposed innovation projects could not be thoroughly examined. Work is ongoing to understand the economic impact and quantify the value of the assessment methodology. Developing this organisational knowledge will play an important role in communicating the business case to support future investment in Industry 4.0 and machine learning solutions and understand their economic sustainability. Given that these projects were identified using the Industry 4.0 assessment tool, quantifying their impact will help communicate the value of performing Industry 4.0 assessments to other manufacturers, providing further opportunities to support digital growth at a global scale.

In addition to the technical opportunities, this research also identifies further opportunities for strategic and organisational growth. Further work is required to define a benchmark for cultural aspects of Industry 4.0 across industries and organisations, as well as to understand how maturity models help support growth in these areas. While this was not possible within the scope of this research, within Ford Motor Company it is proposed that the Industry 4.0 questionnaire is sent out to sites that are known to have the highest Industry 4.0 readiness. These responses can then be used to define the internal benchmark for Industry 4.0 culture. Comparing responses of future assessments from sites with high maturity may give improved guidance to management on what actions need to be taken to make cultural changes to support ongoing digitalisation efforts. Furthermore, in addition to defining the Industry 4.0 benchmark, it is also recommended that following an assessment the same questionnaire is sent to all salaried employees every 6 months after the initial assessment. By comparing responses over time this would help measure progress towards Industry 4.0 objectives.

Future research is also required to explore how the findings presented in this study can be generalised to other industries or organisations. The research presented in Chapter

3 focuses specifically on the application of the proposed assessment methodology within Ford Motor Company’s UK manufacturing operations. While the findings provide valuable insights for this context, further research is needed to explore the application of this methodology in different contexts.

Finally, while the research mentions using the questionnaire to measure ongoing cultural changes at the factory level, this analysis was not possible due to the time constraints of this study.

### **6.2.2 Anomaly Detection Future Work**

This research presents a major advancement in anomaly detection research and highlights multiple opportunities for continued efforts in this space. This project presents the first attempt to perform anomaly detection in manual production processes. These data present new challenges yet to be explored by anomaly detection researchers. Further work is required to understand how to address data for which there is disagreement among domain experts on the category. This project overcame this challenge by having multiple engineers label the same data, however, this approach is time-consuming and labor-intensive. Further research should consider how to quantify the disagreement between labellers and how to use these contrasting data to optimally train an anomaly detection model.

Due to time constraints, the long-term impact of the anomaly detection solution presented in Chapter 4 has not been quantified in real-world production environments, although estimates are made based on available warranty data. This opportunity for future work is recognised by the industrial sponsor, with trials planned to implement the nut runner solution in engine assembly and vehicle operations as part of Project ADAPT. Further work is also required by Ford Motor Company to overcome the IT challenges of the LSTM implementation, specifically to enable Tensorflow GPU environments in existing machine learning pipeline architectures.

The lack of data was a challenge throughout the development of this solution. With newly available data from the recent trial at Dagenham, further research is also required to explore unsupervised approaches. The author proposes that future research explore the use of cross-validation to identify training data that impact final accuracy when randomly sampling training data. This requires high-quality testing datasets and therefore requires additional work from test engineers to support data labelling efforts.

This research shows that the proposed method performs better on automated nut runner

datasets than on manual datasets. As more data are available on different process variants, it should also be considered how different anomaly detection approaches perform on each type. Further research is required to explore how to automatically select the appropriate anomaly detection method based on the features of a specific process.

### **6.2.3 Mapping Homes Suitable for EV Charging Future Work**

The methods developed in this Chapter 5 represent a major advancement towards fully automated remote surveying capability to audit the built environment. However, in the proposed workflows to identify houses suitable for EV charging, the analysis is restricted to the property level, not the household. Due to limited attribute data, some UPRNs may be Houses of Multiple Occupancy (HMO) or converted into residential flats. Future works should explore the use of additional data sources to consider these attributes in the analysis.

Further research is recommended to explore the use of this technology for other use cases. Some examples that gained commercial and academic interest include, identifying suitability for photovoltaic installation, external wall insulation, double glazing, and other external characteristics that can be surveyed to help better understand the energy efficiency of residential housing.

In automotive applications, further research is also required to explore the issues relating to the intellectual property of the derived data and how to create value from these data insights. Possible directions may include a web-based application for local governments to access these data as a service to help improve EV charging infrastructure in their constituencies. As Ford continues to develop autonomous driving solutions, the company should consider developing their own streetscape datasets to overcome the IP challenges and explore further use cases for data-as-a-service models.

While the improved method does result in a high level of accuracy, a significant number of duplicate images still needed to be removed in pre-processing. Further works should further improve the GIS data acquisition to ensure the coordinates of each building polygon are retrieved as close to the road as possible to ensure the highest accuracy when requesting the respective image from the Google Street View API.

# References

- [1] A. Krizhevky, S. Ilya, and H. Geoffrey, “ImageNet Classification with Deep Convolutional Neural Networks,” *Handbook of Approximation Algorithms and Metaheuristics*, pp. 45–1–45–16, 2007.
- [2] S. Christian, L. Wei, and J. Yangqing, “Going Deeper with Convolutions,” 2015.
- [3] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” pp. 1–14, 2014.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2016-Decem, pp. 770–778, 2016.
- [5] E. Brealy, “Multi-criteria approach using neural networks, gis and remote sensing to identify households suitable for electric vehicle charging.” unpublished thesis, 2021.
- [6] V. R. Laguduva, S. Mahmud, S. N. Aakur, R. Karam, and S. Katkoori, “Dissecting convolutional neural networks for efficient implementation on constrained platforms,” *Proceedings - 33rd International Conference on VLSI Design, VLSID 2020 - Held concurrently with 19th International Conference on Embedded Systems*, pp. 149–154, 2020.
- [7] United Nations, “Adoption of the paris agreement,” 2015.
- [8] B. Zakeri, K. Paulavets, L. Barreto-Gomez, L. G. Echeverri, S. Pachauri, B. Boza-Kiss, C. Zimm, J. Rogelj, F. Creutzig, D. Üрге-Vorsatz, D. G. Victor, M. D. Bazilian, S. Fritz, D. Gielen, D. L. McCollum, L. Srivastava, J. D. Hunt, and S. Pouya, “Pandemic, War, and Global Energy Transitions,” *Energies*, vol. 15, no. 17, pp. 1–23, 2022.



- [9] McKinsey, “McKinsey Electric Vehicle Index: EV Market Trends & Sales — McKinsey,” *McKinsey Global Institute*, no. July, 2020.
- [10] W. W. H. Kagermann, W. D. Lukas, “Industrie 4.0: Mit dem Internet der Dinge auf dem Weg zur 4. industriellen Revolution.,” apr 2011.
- [11] W. Wahlster, J. Helbig, A. Hellinger, M. A. V. Stumpf, J. Blasco, H. Galloway, and H. Gestaltung, “Recommendations for implementing the strategic initiative INDUSTRIE 4.0,” Tech. Rep. April, Acatech, Frankfurt, 2013.
- [12] O. Agca and J. Gibson, “An Industry 4 readiness assessment tool,” tech. rep., Warwick University.
- [13] P. M. Deane and P. M. Deane, *The first industrial revolution*. Cambridge University Press, 1979.
- [14] R. H. S. Joel Mokyr, “The Second Industrial Revolution, 1870-1914,” 1998.
- [15] D. Romero, J. Stahre, T. Wuest, O. Noran, P. Bernus, Å. Fast-Berglund, and D. Gorecky, “Towards an Operator 4.0 Typology: A Human-Centric Perspective on the Fourth Industrial Revolution,” pp. 29–31, 2016.
- [16] P. Afonso, A. Santana, P. Afonso, A. Zanin, and R. Wernke, “Manufacturing in the fourth industrial revolution,” 2018.
- [17] M. Brettel, M. Keller, and M. Rosenberg, “How Virtualization, Decentralization and Network Building Change the Manufacturing Landscape: An Industry 4.0 Perspective,” *International Journal of Information and Communication Engineering*, vol. 8, no. 1, 2014.
- [18] G. Li, H. Yun, and W. U. Aizhi, “Fourth Industrial Revolution: Technological Drivers, Impacts and Coping Methods,” *Chin. Geogra. Sci*, vol. 27, no. 4, pp. 626–637, 2017.
- [19] T. D. Oesterreich and F. Teuteberg, “Understanding the implications of digitisation and automation in the context of Industry 4.0: A triangulation approach and elements of a research agenda for the construction industry,” *Computers in Industry*, vol. 83, pp. 121–139, dec 2016.

- [20] E. Congress, “ON THE ROAD TO THE SMART FACTORY,” *Information Unlimited The Copa-Data Magazine*, pp. 2015–2017, oct 2015.
- [21] P. PRISECARU, “CHALLENGES OF THE FOURTH INDUSTRIAL REVOLUTION Petre PRISECARU,” *Knowledge Horizons - Economics*, vol. 8, no. 1, pp. 57–62, 2016.
- [22] S. Mittal, M. A. Khan, D. Romero, and T. Wuest, “Smart manufacturing: Characteristics, technologies and enabling factors,” *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 233, no. 5, pp. 1342–1361, 2017.
- [23] J. Lee, H. A. Kao, and S. Yang, “Service innovation and smart analytics for Industry 4.0 and big data environment,” *Procedia CIRP*, vol. 16, pp. 3–8, 2014.
- [24] M. Babic, M. A. Farahani, and T. Wuest, “Image Based Quality Inspection in Smart Manufacturing Systems: A Literature Review,” *Procedia CIRP*, vol. 103, pp. 262–267, 2021.
- [25] M. Kauffman, M. N. Soares, and M. E. Kauffman, “Industry 4.0: The challenges to intellectual property in Manufacturing Industry 4.0: The Challenges to Intellectual Property In Manufacturing,” tech. rep., 2018.
- [26] G. Miragliotta, A. Sianesi, E. Convertini, and R. Distanto, “Data driven management in Industry 4.0: a method to measure Data Productivity,” *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 19–24, 2018.
- [27] M. Javaid, A. Haleem, R. P. Singh, and R. Suman, “Enabling flexible manufacturing system (FMS) through the applications of industry 4.0 technologies,” *Internet of Things and Cyber-Physical Systems*, vol. 2, no. April, pp. 49–62, 2022.
- [28] A. Luckow, M. Cook, N. Ashcraft, E. Weill, E. Djerekarov, and B. Vorster, “Deep learning in the automotive industry: Applications and tools,” *Proceedings - 2016 IEEE International Conference on Big Data, Big Data 2016*, pp. 3759–3768, 2016.
- [29] S. Huang, Y. Guo, D. Liu, S. Zha, and W. Fang, “A Two-Stage Transfer Learning-Based Deep Learning Approach for Production Progress Prediction in IoT-Enabled Manufacturing,” *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 10627–10638, 2019.

- [30] A. Theissler, J. Pérez-Velázquez, M. Kettelgerdes, and G. Elger, “Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry,” *Reliability Engineering and System Safety*, vol. 215, p. 107864, 2021.
- [31] R. J. Hsieh, J. Chou, and C. H. Ho, “Unsupervised online anomaly detection on multivariate sensing time series data for smart manufacturing,” *Proceedings - 2019 IEEE 12th Conference on Service-Oriented Computing and Applications, SOCA 2019*, pp. 90–97, 2019.
- [32] C. Llopis-Albert, F. Rubio, and F. Valero, “Impact of digital transformation on the automotive industry,” *Technological Forecasting and Social Change*, vol. 162, no. September 2020, p. 120343, 2021.
- [33] A. Camuffo and G. Volpato, “Dynamic capabilities and manufacturing automation: Organizational learning in the Italian automobile industry,” *Industrial and Corporate Change*, vol. 5, no. 3, pp. 813–838, 1996.
- [34] IHS Technology, “IoT platforms: enabling the Internet of Things,” *IHS Technology*, vol. Whitepaper, no. March, pp. 1–19, 2016.
- [35] R. Darnley, M. Diplacido, M. Kerns, and A. Kim, “INDUSTRY 4.0: DIGITIZATION IN DANISH INDUSTRY,” tech. rep., WORCESTER POLYTECHNIC INSTITUTE, 2018.
- [36] S. Fahle, C. Prinz, and B. Kuhlenkötter, “Systematic review on machine learning (ML) methods for manufacturing processes - Identifying artificial intelligence (AI) methods for field application,” *Procedia CIRP*, vol. 93, pp. 413–418, 2020.
- [37] V. Parida, D. Sjödin, and W. Reim, “Reviewing literature on digitalization, business model innovation, and sustainable industry: Past achievements and future promises,” *Sustainability (Switzerland)*, vol. 11, no. 2, 2019.
- [38] G. Martínez-Arellano, T. Nguyen, C. Hinton, and S. Ratchev, “A data analytics model for improving process control in flexible manufacturing cells,” *Decision Analytics Journal*, vol. 3, no. June, p. 100075, 2022.
- [39] A. Grybauskas, A. Stefanini, and M. Ghobakhloo, “Social sustainability in the age of digitalization: A systematic literature Review on the social implications of industry 4.0,” *Technology in Society*, vol. 70, no. April, p. 101997, 2022.

- [40] Engineering USA, “ROADMAP TO INDUSTRY 4.0 A Strategic Partnership with Engineering Digital Industry,” tech. rep., 2019.
- [41] G. Schuh, R. Anderl, R. Dumitrescu, A. Krüger, and M. Hompel, “Industrie 4.0 Maturity Index: Managing the Digital Transformation of Companies,” tech. rep., Acatech, 2020.
- [42] A. D. Carolis, M. Macchi, E. Negri, and S. Terzi, “A Maturity Model for Assessing the Digital Readiness of Manufacturing Companies A maturity model for assessing the digital readiness of manufacturing companies,” no. August, 2017.
- [43] A. Schumacher, S. Erol, and W. Sihn, “A maturity model for assessing Industry 4 . 0 readiness and maturity of manufacturing enterprises,” *Procedia CIRP*, vol. 52, pp. 161–166, 2016.
- [44] TDWI, “TDWI Analytics Maturity Model Assessment,” 2022.
- [45] S. Lemsä, “Framework to Build an Advanced Analytics Maturity Assessment Model : Questionnaire Design,” in *SOCIETY. TECHNOLOGY. SOLUTIONS.*, no. December 2021, pp. 11–12, 2022.
- [46] C. Leyh and T. Schäffer, “The Application of the Maturity Model SIMMI 4 . 0 in Selected Enterprises Full Paper Chair of Information Systems Chair of Information Systems Abstract Chair of Information Systems,” no. August, pp. 0–10, 2017.
- [47] G. Li, C. Yuan, S. Kamarthi, M. Moghaddam, and X. Jin, “Data science skills and domain knowledge requirements in the manufacturing industry: A gap analysis,” *Journal of Manufacturing Systems*, vol. 60, no. April, pp. 692–706, 2021.
- [48] S. Bag, J. H. C. Pretorius, S. Gupta, and Y. K. Dwivedi, “Role of institutional pressures and resources in the adoption of big data analytics powered artificial intelligence, sustainable manufacturing practices and circular economy capabilities,” *Technological Forecasting and Social Change*, vol. 163, no. November 2020, p. 120420, 2021.
- [49] F. J. Brandl, N. Roider, M. Hehl, and G. Reinhart, “Selecting practices in complex technical planning projects: A pathway for tailoring agile project management into the manufacturing industry,” *CIRP Journal of Manufacturing Science and Technology*, vol. 33, pp. 293–305, 2021.

- [50] D. Ibarra, J. Ganzarain, and J. I. Igartua, “Business model innovation through Industry 4.0: A review,” *Procedia Manufacturing*, vol. 22, pp. 4–10, 2018.
- [51] M. Antikainen, T. Uusitalo, and P. Kivikytö-Reponen, “Digitalisation as an Enabler of Circular Economy,” *Procedia CIRP*, vol. 73, pp. 45–49, 2018.
- [52] S. M. Borodo, S. M. Shamsuddin, and S. Hasan, “Big data platforms and techniques,” *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 1, no. 1, pp. 191–200, 2016.
- [53] R. F. Babiceanu and R. Seker, “Big Data and virtualization for manufacturing cyber-physical systems: A survey of the current status and future outlook,” *Computers in Industry*, vol. 81, pp. 128–137, sep 2016.
- [54] A. B. Lopes de Sousa Jabbour, C. J. C. Jabbour, M. Godinho Filho, and D. Roubaud, “Industry 4.0 and the circular economy: a proposed research agenda and original roadmap for sustainable operations,” *Annals of Operations Research*, vol. 270, pp. 273–286, nov 2018.
- [55] R. Y. Zhong, X. Xu, E. Klotz, and S. T. Newman, “Intelligent Manufacturing in the Context of Industry 4.0: A Review,” *Engineering*, vol. 3, no. 5, pp. 616–630, 2017.
- [56] J.-p. Skeete, “Technological Forecasting & Social Change Level 5 autonomy: The new face of disruption in road transport,” *Technological Forecasting & Social Change*, vol. 134, no. May 2018, pp. 22–34, 2020.
- [57] J. Lee, E. Lapira, B. Bagheri, and H.-a. Kao, “Recent advances and trends in predictive manufacturing systems in big data environment,” *Manufacturing Letters*, vol. 1, pp. 38–41, oct 2013.
- [58] M. Woodward, “Big Data and analytics in the automotive industry,” tech. rep., Deloitte, 2015.
- [59] M. Johanson, S. Belenki, J. Jalminger, M. Fant, and M. Gjertz, “Big Automotive Data: Leveraging large volumes of data for knowledge-driven product development,” in *2014 IEEE International Conference on Big Data (Big Data)*, pp. 736–741, IEEE, oct 2014.
- [60] “Ford mobility.” <https://www.ford.co.uk/shop/specialist-sales/motability>. Accessed: 2022-04-19.

- [61] J. Monios, R. Bergqvist, D. D. Luminy, and R. A. Bourdelle, “Technological Forecasting & Social Change Logistics and the networked society : A conceptual framework for smart network business models using electric autonomous vehicles ( EAVs ),” *Technological Forecasting & Social Change*, vol. 151, no. March 2019, p. 119824, 2020.
- [62] R. Want, B. N. Schilit, and S. Jenson, “Enabling the Internet of Things,” vol. 48, pp. 28–35, jan 2015.
- [63] J. Friederich, D. P. Francis, S. Lazarova-Molnar, and N. Mohamed, “A framework for data-driven digital twins for smart manufacturing,” *Computers in Industry*, vol. 136, p. 103586, 2022.
- [64] R. van Dinter, B. Tekinerdogan, and C. Catal, “Predictive maintenance using digital twins: A systematic literature review,” *Information and Software Technology*, vol. 151, no. February, p. 107008, 2022.
- [65] S. Rajput and S. P. Singh, “Connecting circular economy and industry 4.0,” *International Journal of Information Management*, vol. 49, no. November 2018, pp. 98–113, 2019.
- [66] M. Borg, C. Englund, K. Wnuk, B. Duran, C. Levandowski, S. Gao, Y. Tan, H. Kaiser, H. Lönn, and J. Törnqvist, “Safely Entering the Deep: A Review of Verification and Validation for Machine Learning and a Challenge Elicitation in the Automotive Industry,” *Journal of Automotive Software Engineering*, vol. 1, no. 1, p. 1, 2019.
- [67] E. E. Makarius, D. Mukherjee, J. D. Fox, and A. K. Fox, “Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the organization,” *Journal of Business Research*, vol. 120, no. July, pp. 262–273, 2020.
- [68] N. Boavida and M. Candeias, “Recent automation trends in portugal: Implications on industrial productivity and employment in automotive sector,” *Societies*, vol. 11, no. 3, pp. 1–17, 2021.
- [69] D. R. Sjödin, V. Parida, M. Leksell, and A. Petrovic, “Smart Factory Implementation and Process Innovation: A Preliminary Maturity Model for Leveraging Digitalization in Manufacturing Moving to smart factories presents specific challenges that can be addressed through a structured approach focused on people, p,” *Research Technology Management*, vol. 61, no. 5, pp. 22–31, 2018.

- [70] S. Albukhitan, “Developing Digital Transformation Strategy for Manufacturing,” *Procedia Computer Science*, vol. 170, pp. 664–671, 2020.
- [71] A. Kampker, H. Heimes, U. Bühner, C. Lienemann, and S. Krottil, “Enabling Data Analytics in Large Scale Manufacturing,” *Procedia Manufacturing*, vol. 24, pp. 120–127, 2018.
- [72] R. Sharma and B. Villányi, “Evaluation of corporate requirements for smart manufacturing systems using predictive analytics,” *Internet of Things (Netherlands)*, vol. 19, no. May, p. 100554, 2022.
- [73] A. Dacal-Nieto, J. J. Areal, V. Alonso-Ramos, and M. Lluch, “Integrating a data analytics system in automotive manufacturing: Background, methodology and learned lessons,” *Procedia Computer Science*, vol. 200, pp. 718–726, 2022.
- [74] S. Zhai, B. Gehring, and G. Reinhart, “Enabling predictive maintenance integrated production scheduling by operation-specific health prognostics with generative deep learning,” *Journal of Manufacturing Systems*, vol. 61, no. March, pp. 830–855, 2021.
- [75] P. J. Pereira, A. Pereira, P. Cortez, and A. Pilastrri, “A Comparison of Machine Learning Methods for Extremely Unbalanced Industrial Quality Data,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12981 LNAI, pp. 561–572, 2021.
- [76] W. Kirchgässner, O. Wallscheid, and J. Böcker, “Data-Driven Permanent Magnet Temperature Estimation in Synchronous Motors with Supervised Machine Learning: A Benchmark,” *IEEE Transactions on Energy Conversion*, vol. 36, no. 3, pp. 2059–2067, 2021.
- [77] R. Redondo, Á. Herrero, E. Corchado, and J. Sedano, “A decision-making tool based on exploratory visualization for the automotive industry,” *Applied Sciences (Switzerland)*, vol. 10, no. 12, 2020.
- [78] T. Grüner, F. Böllhoff, R. Meisetschläger, A. Vydrenko, M. Bator, A. Dicks, and A. Theissler, “Evaluation of machine learning for sensorless detection and classification of faults in electromechanical drive systems,” *Procedia Computer Science*, vol. 176, pp. 1586–1595, 2020.

- [79] A. Tufano, R. Accorsi, and R. Manzini, “A machine learning approach for predictive warehouse design,” *International Journal of Advanced Manufacturing Technology*, vol. 119, no. 3-4, pp. 2369–2392, 2022.
- [80] D. Knittel, H. Makich, and M. Nouari, “Milling diagnosis using artificial intelligence approaches,” *Mechanics and Industry*, vol. 20, no. 8, 2019.
- [81] R. Espinosa, H. Ponce, and S. Gutiérrez, “Click-event sound detection in automotive industry using machine/deep learning,” *Applied Soft Computing*, vol. 108, p. 107465, 2021.
- [82] R. P. Monteiro, M. Cerrada, D. R. Cabrera, R. V. Sánchez, and C. J. Bastos-Filho, “Using a Support Vector Machine Based Decision Stage to Improve the Fault Diagnosis on Gearboxes,” *Computational Intelligence and Neuroscience*, vol. 2019, 2019.
- [83] I. E. Hassani, C. E. Mazgualdi, and T. Masrour, “Artificial Intelligence and Machine Learning to Predict and Improve Efficiency in Manufacturing Industry,” 2019.
- [84] S. Hatanaka and H. Nishi, “Efficient GAN-Based Unsupervised Anomaly Sound Detection for Refrigeration Units,” *IEEE International Symposium on Industrial Electronics*, vol. 2021-June, 2021.
- [85] R. S. Peres, M. Azevedo, S. O. Araújo, M. Guedes, F. Miranda, and J. Barata, “Generative adversarial networks for data augmentation in structural adhesive inspection,” *Applied Sciences (Switzerland)*, vol. 11, no. 7, 2021.
- [86] R. S. Peres, M. Guedes, F. Miranda, and J. Barata, “Simulation-Based Data Augmentation for the Quality Inspection of Structural Adhesive with Deep Learning,” *IEEE Access*, vol. 9, pp. 76532–76541, 2021.
- [87] M. Mazzetto, M. Teixeira, É. O. Rodrigues, and D. Casanova, “Deep Learning Models for Visual Inspection on Automotive Assembling Line,” *International Journal of Advanced Engineering Research and Science*, vol. 7, no. 3, pp. 473–494, 2020.
- [88] I. Cerro, I. Latasa, C. Guerra, P. Pagola, B. Bujanda, and J. J. Astrain, “Smart system with artificial intelligence for sensory gloves,” *Sensors*, vol. 21, no. 5, pp. 1–18, 2021.



- [89] L. Malburg, M. P. Rieder, R. Seiger, P. Klein, and R. Bergmann, "Object detection for smart factory processes by machine learning," *Procedia Computer Science*, vol. 184, no. 2019, pp. 581–588, 2021.
- [90] P. Stavropoulos, A. Papacharalampopoulos, L. Athanasopoulou, K. Kampouris, and P. Lagios, "Designing a digitalized cell for remanufacturing of automotive frames," *Procedia CIRP*, vol. 109, pp. 513–519, 2022.
- [91] Z. J. Viharos and R. Jakab, "Reinforcement Learning for Statistical Process Control in Manufacturing," *Measurement: Journal of the International Measurement Confederation*, vol. 182, no. April, 2021.
- [92] V. Alcácer and V. Cruz-Machado, "Scanning the Industry 4.0: A Literature Review on Technologies for Manufacturing Systems," *Engineering Science and Technology, an International Journal*, vol. 22, no. 3, pp. 899–919, 2019.
- [93] D. D. Kho, S. Lee, and R. Y. Zhong, "Big Data Analytics for Processing Time Analysis in an IoT-enabled manufacturing Shop Floor," *Procedia Manufacturing*, vol. 26, pp. 1411–1420, 2018.
- [94] P. Stavropoulos, A. Papacharalampopoulos, and D. Petridis, "A vision-based system for real-time defect detection: A rubber compound part case study," *Procedia CIRP*, vol. 93, pp. 1230–1235, 2020.
- [95] X. Zhu, H. Manamasa, J. L. Jiménez Sánchez, A. Maki, and L. Hanson, "Automatic assembly quality inspection based on an unsupervised point cloud domain adaptation model," *Procedia CIRP*, vol. 104, pp. 1801–1806, 2021.
- [96] M. Schlüter, C. Niebuhr, J. Lehr, and J. Krüger, "Vision-based Identification Service for Remanufacturing Sorting," *Procedia Manufacturing*, vol. 21, no. 2017, pp. 384–391, 2018.
- [97] A. Papavasileiou, P. Aivaliotis, S. Aivaliotis, and S. Makris, "An optical system for identifying and classifying defects of metal parts," *International Journal of Computer Integrated Manufacturing*, vol. 35, no. 3, pp. 326–340, 2022.
- [98] A. I. M. Schwebig and R. Tutsch, "Compilation of training datasets for use of convolutional neural networks supporting automatic inspection processes in industry

- 4.0 based electronic manufacturing,” *Journal of Sensors and Sensor Systems*, vol. 9, no. 1, pp. 167–178, 2020.
- [99] I. G. Courville, Y. Bengio, and Aaron, “Detecting Teeth Defects on Automotive Gears Using Deep Learning,” *Sensors*, vol. 21, no. 8480, 2021.
- [100] R. Ferreira, J. Barroso, and V. Filipe, “Conformity Assessment of Informative Labels in Car Engine Compartment with Deep Learning Models,” *Journal of Physics: Conference Series*, vol. 2278, no. 1, 2022.
- [101] M. X. Bastidas-Rodriguez, F. A. Prieto-Ortiz, and L. F. Polania, “A textural deep neural network combined with handcrafted features for mechanical failure classification,” *Proceedings of the IEEE International Conference on Industrial Technology*, vol. 2019-Febru, pp. 847–852, 2019.
- [102] M. Kebisek, P. Tanuska, L. Spendla, J. Kotianova, and P. Strelec, “Artificial intelligence platform proposal for paint structure quality prediction within the industry 4.0 concept,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 11168–11174, 2020.
- [103] B. Lindemann, F. Fesenmayr, N. Jazdi, and M. Weyrich, “Anomaly detection in discrete manufacturing using self-learning approaches,” *Procedia CIRP*, vol. 79, pp. 313–318, 2019.
- [104] E. Wescoat, M. Krugh, A. Henderson, J. Goodnough, and L. Mears, “Vibration analysis utilizing unsupervised learning,” *Procedia Manufacturing*, vol. 34, pp. 876–884, 2019.
- [105] A. Saadallah, O. Abdulaaty, J. Büscher, T. Panusch, K. Morik, and J. Deuse, “Early Quality Prediction using Deep Learning on Time Series Sensor Data,” *Procedia CIRP*, vol. 107, pp. 611–616, 2022.
- [106] B. Einabadi, A. Baboli, and E. Rother, “A new methodology for estimation of dynamic Remaining Useful Life: A case study of conveyor chains in the automotive industry,” *Procedia Computer Science*, vol. 201, no. C, pp. 461–469, 2022.
- [107] M. Subramaniyan, A. Skoogh, A. S. Muhammad, J. Bokrantz, B. Johansson, and C. Roser, “A generic hierarchical clustering approach for detecting bottlenecks in manufacturing,” *Journal of Manufacturing Systems*, vol. 55, no. January, pp. 143–158, 2020.

- [108] P. Torres, J. Arents, H. Marques, and P. Marques, “Bin-Picking Solution for Randomly Placed Automotive Connectors Based on Machine Learning Techniques,” *Electronics (Switzerland)*, vol. 11, no. 3, 2022.
- [109] A. Rahimi, M. Anvaripour, and K. Hayat, “Object Detection using Deep Learning in a Manufacturing Plant to Improve Manual Inspection,” *2021 IEEE International Conference on Prognostics and Health Management, ICPHM 2021*, 2021.
- [110] D. Gankin, S. Mayer, J. Zinn, B. Vogel-Heuser, and C. Endisch, “Modular Production Control with Multi-Agent Deep Q-Learning,” *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA*, vol. 2021-Septe, 2021.
- [111] J. Dalzochio, R. Kunst, E. Pignaton, A. Binotto, S. Sanyal, J. Favilla, and J. Barbosa, “Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges,” *Computers in Industry*, vol. 123, p. 103298, 2020.
- [112] T. P. Carvalho, F. A. Soares, R. Vita, R. d. P. Francisco, J. P. Basto, and S. G. Alcalá, “A systematic literature review of machine learning methods applied to predictive maintenance,” *Computers and Industrial Engineering*, vol. 137, no. September, p. 106024, 2019.
- [113] Y. Fan, S. Nowaczyk, and T. Rögnvaldsson, “Transfer learning for remaining useful life prediction based on consensus self-organizing models,” *Reliability Engineering and System Safety*, vol. 203, 2020.
- [114] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [115] O. Chapelle, B. Scholkopf, and A. Zien, *Semi-supervised learning*. MIT Press, 2006.
- [116] S. Kulvisaechana, “The Role of Communication Strategies in Change Management Process,” tech. rep., 2001.
- [117] A. S. K. et al., A. Senthil Kumar et al., “An Industrial IOT in Engineering and Manufacturing Industries - Benefits and Challenges,” *International Journal of Mechanical and Production Engineering Research and Development*, vol. 9, no. 2, pp. 151–160, 2019.

- [118] H. Boyes, B. Hallaq, J. Cunningham, and T. Watson, “The industrial internet of things (IIoT): An analysis framework,” *Computers in Industry*, vol. 101, no. December 2017, pp. 1–12, 2018.
- [119] S. Tweneboah-Koduah, K. E. Skouby, and R. Tadayoni, “Cyber security threats to iot applications and service domains,” *Wireless Personal Communications*, vol. 95, no. 1, pp. 169–185, 2017.
- [120] Y. Lu and L. Da Xu, “Internet of Things (IoT) Cybersecurity Research: A Review of Current Research Topics,” *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2103–2115, 2018.
- [121] C. M. Roberts, “Radio frequency identification (RFID),” *Computers and Security*, vol. 25, no. 1, pp. 18–26, 2006.
- [122] F. X. Diebold, L. Brown, X. Cheng, F. Cunha, S. Diebold, D. Foster, M. Halperin, S. Lohr, J. Mashey, T. Nickolas, L. Olson, M. Pai, M. Pospiech, F. Schorfheide, M. Shin, and M. Steele, “A Personal Perspective on the Origin(s) and Development of ”Big Data”: The Phenomenon, the Term, and the Discipline \*,” tech. rep., 2012.
- [123] M. Chen, H. Jin, Y. Wen, and V. Leung, “Enabling technologies for future data center networking: A primer,” *IEEE Network*, vol. 27, no. 4, pp. 8–15, 2013.
- [124] A. Oussous, F.-Z. Benjelloun, A. Ait Lahcen, and S. Belfkih, “Big Data technologies: A survey,” *Journal of King Saud University - Computer and Information Sciences*, vol. 30, pp. 431–448, oct 2018.
- [125] A. Gandomi and M. Haider, “Beyond the hype: Big data concepts, methods, and analytics,” *International Journal of Information Management*, vol. 35, pp. 137–144, apr 2015.
- [126] R. Kune, P. K. Konugurthi, A. Agarwal, R. R. Chillarige, and R. Buyya, “The anatomy of big data computing,” *Software: Practice and Experience*, vol. 46, pp. 79–105, jan 2016.
- [127] A. McAfee and E. Brynjolfsson, “Big Data: The Management Revolution,” tech. rep., Harvard Business Review, 2012.

- [128] Z. D. Stephens, S. Y. Lee, F. Faghri, R. H. Campbell, C. Zhai, M. J. Efron, R. Iyer, M. C. Schatz, S. Sinha, and G. E. Robinson, “Big Data: Astronomical or Genomic?,” 2015.
- [129] W. Fan and A. Bifet, “Mining Big Data: Current Status, and Forecast to the Future,” *ACM SIGKDD Explorations Newsletter*, vol. 14, no. 2, pp. 1–5, 2012.
- [130] S. Kaisler, F. Armour, J. A. Espinosa, and W. Money, “Big data: Issues and challenges moving forward,” *Proceedings of the Annual Hawaii International Conference on System Sciences*, pp. 995–1004, 2013.
- [131] S. Shahrivari, Shahrivari, and Saeed, “Beyond Batch Processing: Towards Real-Time and Streaming Big Data,” *Computers*, vol. 3, pp. 117–129, oct 2014.
- [132] R. Casado and M. Younas, “Emerging trends and technologies in big data processing,” *Concurrency and Computation: Practice and Experience*, vol. 27, pp. 2078–2091, jun 2015.
- [133] S. Abiteboul, “Querying Semi-Structured Data,” in *International Conference on Database Theory*, 1997.
- [134] L. Wang and C. A. Alexander, “Big Data in Design and Manufacturing Engineering,” *American Journal of Engineering and Applied Sciences*, vol. 8, no. 2, pp. 223–232, 2015.
- [135] A. Katal, M. Wazid, and R. H. Goudar, “Big data: Issues, challenges, tools and Good practices,” *2013 6th International Conference on Contemporary Computing, IC3 2013*, pp. 404–409, 2013.
- [136] M. A.-u.-d. Khan, M. F. Uddin, and N. Gupta, “Seven V ’ s of Big Data Understanding Big Data to extract Value,” *Proceedings of the 2014 Zone 1 Conference of the American Society for Engineering Education*, pp. 1–5, 2014.
- [137] Puget Jean, “Optimization Is Ready For Big Data: Part 4, Veracity (IT Best Kept Secret Is Optimization),” 2015.
- [138] S. Athey, “Beyond prediction: Using big data for policy problems,”
- [139] A. Kankanhalli, J. Hahn, S. Tan, and G. Gao, “Big data and analytics in healthcare: Introduction to the special section,” *Information Systems Frontiers*, vol. 18, pp. 233–235, apr 2016.

- [140] Palfreyman John, “Big Data – Vexed by Veracity? - IBM Government Industry Blog,” 2013.
- [141] B. Saha and D. Srivastava, “Data quality: The other face of Big Data,” *Proceedings - International Conference on Data Engineering*, pp. 1294–1297, 2014.
- [142] N. Koseleva and G. Ropaite, “Big Data in Building Energy Efficiency: Understanding of Big Data and Main Challenges,” *Procedia Engineering*, vol. 172, pp. 544–549, jan 2017.
- [143] A. De Mauro, M. Greco, and M. Grimaldi, “A formal definition of Big Data based on its essential features,” *Library Review*, vol. 65, pp. 122–135, apr 2016.
- [144] P. Tabesh, E. Mousavidin, and S. Hasani, “Implementing big data strategies: A managerial perspective,” *Business Horizons*, mar 2019.
- [145] A. Mehmood, I. Natgunanathan, Y. Xiang, G. Hua, and S. Guo, “Protection of Big Data Privacy,” *IEEE Access*, vol. 4, pp. 1821–1834, 2016.
- [146] INFORMS, “IBM report: Most companies lack IT infrastructure strategy — Analytics Magazine,” 2018.
- [147] A. M. Wilson, “Understanding organisational culture and the implications for corporate marketing,” *European Journal of Marketing*, vol. 35, pp. 353–367, apr 2001.
- [148] L. S. Dalenogare, G. B. Benitez, N. F. Ayala, and A. G. Frank, “The expected contribution of Industry 4.0 technologies for industrial performance,” *International Journal of Production Economics*, vol. 204, no. July, pp. 383–394, 2018.
- [149] S. Wang, J. Wan, D. Li, and C. Zhang, “Implementing Smart Factory of Industrie 4.0: An Outlook,” 2016.
- [150] D. Mourtzis, J. Angelopoulos, and N. Panopoulos, “Design and Development of an Edge-Computing Platform Towards 5G Technology Adoption for Improving Equipment Predictive Maintenance,” *Procedia Computer Science*, vol. 200, no. 2019, pp. 611–619, 2022.
- [151] M. Schreiber, J. Klöber-Koch, J. Bömelburg-Zacharias, S. Braunreuther, and G. Reinhart, “Automated quality assurance as an intelligent cloud service using machine learning,” *Procedia CIRP*, vol. 86, pp. 185–191, 2020.

- [152] A. Carvalho, N. O' Mahony, L. Krpalkova, S. Campbell, J. Walsh, and P. Doody, "At the edge of industry 4.0," *Procedia Computer Science*, vol. 155, no. 2018, pp. 276–281, 2019.
- [153] Y. Chen, "Integrated and Intelligent Manufacturing: Perspectives and Enablers," *Engineering*, vol. 3, no. 5, pp. 588–595, 2017.
- [154] U. G. Schulte, "New business models for a radical change in resource efficiency," *Environmental Innovation and Societal Transitions*, vol. 9, pp. 43–47, 2013.
- [155] C. T. E Sutherland, G Patterson, C Dediccoat, S Holiday, "Towards Circular Economy," tech. rep., Ellen MacArthur Foundation, 2013.
- [156] EUROPEAN COMMISSION, "On the implementation of the Circular Economy Action Plan," tech. rep., EUROPEAN COMMISSION, Brussels,, 2019.
- [157] M.-L. Tseng, R. R. Tan, A. S. Chiu, C.-F. Chien, and T. C. Kuo, "Circular economy meets industry 4.0: Can big data drive industrial symbiosis?," *Resources, Conservation and Recycling*, vol. 131, pp. 146–147, apr 2018.
- [158] EM Foundation, "Towards a Circular Economy: Business Rationale for an Accelerated Transition," *Greener Management International*, p. 20, 2015.
- [159] A. Alcayaga, M. Wiener, and E. G. Hansen, "Towards a framework of smart-circular systems: An integrative literature review," *Journal of Cleaner Production*, vol. 221, pp. 622–634, 2019.
- [160] G. Bressanelli, M. Perona, and N. Saccani, "Reshaping the Washing Machine Industry through Circular Economy and Product-Service System Business Models," *Procedia CIRP*, vol. 64, pp. 43–48, 2017.
- [161] M. Suvarna, K. S. Yap, W. Yang, J. Li, Y. T. Ng, and X. Wang, "Cyber-Physical Production Systems for Data-Driven, Decentralized, and Secure Manufacturing—A Perspective," *Engineering*, vol. 7, no. 9, pp. 1212–1223, 2021.
- [162] E. Oztemel and S. Gursev, "Literature review of Industry 4.0 and related technologies," *Journal of Intelligent Manufacturing*, vol. 31, no. 1, pp. 127–182, 2020.
- [163] M. Javaid, A. Haleem, R. Pratap Singh, and R. Suman, "Significance of Quality 4.0 towards comprehensive enhancement in manufacturing sector," *Sensors International*, vol. 2, no. May, p. 100109, 2021.

- [164] C. Burchardt and B. Maisch, “Digitalization needs a cultural change – examples of applying Agility and Open Innovation to drive the digital transformation,” *Procedia CIRP*, vol. 84, pp. 112–117, 2019.
- [165] S. J. Arora, C. Ebbecke, M. Rabe, and J. Fisch, “Methodology for the assessment of potentials, selection, and design of Predictive Maintenance solutions,” *Procedia CIRP*, vol. 104, pp. 708–713, 2021.
- [166] A. Mayr, M. Weigelt, A. Kühn, S. Grimm, A. Erll, M. Potzel, and J. Franke, “Lean 4.0-A conceptual conjunction of lean management and Industry 4.0,” *Procedia CIRP*, vol. 72, pp. 622–628, 2018.
- [167] T. Wuest, A. Liu, S. C. Lu, and K. D. Thoben, “Application of the stage gate model in production supporting quality management,” *Procedia CIRP*, vol. 17, pp. 32–37, 2014.
- [168] G. Hahn, N. Doganaksoy, and R. Hoerl, “The evolution of six sigma,” *Quality Engineering*, vol. 12, pp. 317–326, 01 2000.
- [169] M. Meister, J. Beßle, A. Cviko, T. Böing, and J. Metternich, “Manufacturing analytics for problem-solving processes in production,” *Procedia CIRP*, vol. 81, pp. 1–6, 2019.
- [170] B. Ding, X. F. Hernández, and N. A. Jané, “Combining lean and agile manufacturing competitive advantages through industry 4.0 technologies: an integrative approach,” *Production Planning & Control*, vol. 0, no. 0, pp. 1–17, 2021.
- [171] M. Sony, “Industry 4.0 and lean management: a proposed integration model and research propositions,” *Production and Manufacturing Research*, vol. 6, no. 1, pp. 416–432, 2018.
- [172] E. Colangelo, C. Fries, T. F. Hinrichsen, Á. Szaller, and G. Nick, “Maturity Model for AI in Smart Production Planning and Control System,” *Procedia CIRP*, vol. 107, no. 2021, pp. 493–498, 2022.
- [173] R. D. Freeze, “Understanding the Main Phases of Developing a Maturity Assessment Model Understanding the Main Phases of Developing a Maturity Assessment Model,” no. January, 2005.



- [174] E. S. Rosa, R. Godina, E. M. Rodrigues, and J. C. Matias, “An Industry 4.0 Conceptual Model Proposal for Cable Harness Testing Equipment Industry,” *Procedia Computer Science*, vol. 200, pp. 1392–1401, 2022.
- [175] M. Menolotto, D. S. Komaris, S. Tedesco, B. O’flynn, and M. Walsh, “Motion capture technology in industrial applications: A systematic review,” *Sensors (Switzerland)*, vol. 20, no. 19, pp. 1–25, 2020.
- [176] P. Runeson, T. Olsson, and J. Linåker, “Open Data Ecosystems — An empirical investigation into an emerging industry collaboration concept,” *Journal of Systems and Software*, vol. 182, p. 111088, 2021.
- [177] D. Mourtzis, S. Fotia, N. Boli, and E. Vlachou, “Modelling and quantification of industry 4.0 manufacturing complexity based on information theory: a robotics case study,” *International Journal of Production Research*, vol. 57, no. 22, pp. 6908–6921, 2019.
- [178] M. Kotarba, “MEASURING DIGITALIZATION - KEY METRICS,” *Foundations of Management*, vol. 9, 2017.
- [179] N. Tuptuk and S. Hailes, “Security of smart manufacturing systems,” *Journal of Manufacturing Systems*, vol. 47, no. May, pp. 93–106, 2018.
- [180] A. Telukdarie, E. Buhulaiga, S. Bag, S. Gupta, and Z. Luo, “Industry 4.0 implementation for multinationals,” *Process Safety and Environmental Protection*, vol. 118, pp. 316–329, 2018.
- [181] J. Gomes, “Linking Benefits to Maturity Models,” in *15th International Academy of Management and Business Conference*, no. April, pp. Online Proceedings ISSN 1949 – 9108, 2013.
- [182] M. V. Steenbergen, R. Bos, O. U. Nederland, and S. Brinkkemper, “The Design of Focus Area Maturity Models,” no. May 2014, 2010.
- [183] B. Axmann and H. Harmoko, “Industry 4.0 Readiness Assessment,” *Tehnički glasnik*, vol. 14, no. 2, pp. 212–217, 2020.
- [184] A. Amaral, “A Framework for Assessing SMEs Industry 4.0 maturity,” no. June, pp. 283–312, 2019.

- [185] A. Koska, A. Professor, M. Banu ERDEM, and H. Fettahliođlu, “Measuring the Maturity of a Factory for Industry 4.0,” *International Journal of Academic Research in Business and Social Sciences*, vol. 7, no. 7, p. 52, 2017.
- [186] CMMI, “Cmmi institute,” 2022.
- [187] J. A. Krosnick and S. Presser, *Question and Questionnaire Design*. 2010.
- [188] C. G. V. Wangenheim, J. Carlo, R. Hauck, C. F. Salviano, and A. V. Wangenheim, “Systematic Literature Review of Software Process Capability / Maturity Models Systematic Literature Review of Software Process Capability / Maturity Models,” no. April, 2010.
- [189] A. FABRİKALAR, “Digital maturity test.” <https://akillifabrikalar.com.tr/dijital-olgunluk-testi/>. Accessed: 2022-07-10.
- [190] M. Gattullo, G. W. Scurati, M. Fiorentino, A. E. Uva, F. Ferrise, and M. Bordegoni, “Towards augmented reality manuals for industry 4.0: A methodology,” *Robotics and Computer-Integrated Manufacturing*, vol. 56, no. November 2018, pp. 276–286, 2019.
- [191] A. G. Frank, L. S. Dalenogare, and N. F. Ayala, “Industry 4.0 technologies: Implementation patterns in manufacturing companies,” *International Journal of Production Economics*, vol. 210, no. September 2018, pp. 15–26, 2019.
- [192] M. Schnappinger and J. Streit, “Efficient Platform Migration of a Mainframe Legacy System Using Custom Transpilation,” *Proceedings - 2021 IEEE International Conference on Software Maintenance and Evolution, ICSME 2021*, pp. 545–554, 2021.
- [193] Á. Bányai, B. Illés, E. Glistau, N. I. C. Machado, P. Tamás, F. Manzoor, and T. Bányai, “Smart Cyber-Physical Manufacturing: Extended and Real-Time Optimization of Logistics Resources in Matrix Production,” *Applied Sciences*, vol. 9, no. 7, p. 1287, 2019.
- [194] E. Gökalp, P. Erhan, and U. Şener, “Development of an Assessment Model for Industry 4.0: Industry 4.0-MM,” *Computer Standards and Interfaces*, vol. 54, no. September, pp. 117–118, 2017.

- [195] S. Hines, S. Hines, and N. Meissner, “Strategic Transformation of Ford Motor Company Strategic Transformation of Ford Motor Company A project submitted in partial fulfillment of the requirements for,” no. September, 2014.
- [196] F. Shrouf, J. Ordieres, and G. Miragliotta, “Smart factories in Industry 4.0: A review of the concept and of energy management approached in production based on the Internet of Things paradigm,” *IEEE International Conference on Industrial Engineering and Engineering Management*, vol. 2015-Janua, pp. 697–701, 2014.
- [197] A. Raj, G. Dwivedi, A. Sharma, A. B. Lopes de Sousa Jabbour, and S. Rajak, “Barriers to the adoption of industry 4.0 technologies in the manufacturing sector: An inter-country comparative perspective,” *International Journal of Production Economics*, vol. 224, no. November 2019, p. 107546, 2020.
- [198] T. Dietterich, “Anomaly detection: Algorithms, explanations, applications,” Mar 2018.
- [199] L. Salmina, R. Castello, J. Stoll, and J. L. Scartezzini, “Dimensionality reduction and clustering of time series for anomaly detection in a supermarket heating system,” *Journal of Physics: Conference Series*, vol. 2042, no. 1, 2021.
- [200] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, “A Review on Outlier/Anomaly Detection in Time Series Data,” *ACM Computing Surveys*, vol. 54, no. 3, 2021.
- [201] R. Chalapathy and S. Chawla, “Deep Learning for Anomaly Detection: A Survey,” pp. 1–50, 2019.
- [202] M. Ali, M. W. Jones, X. Xie, and M. Williams, “TimeCluster: dimension reduction applied to temporal data for visual analytics,” *Visual Computer*, vol. 35, no. 6-8, pp. 1013–1026, 2019.
- [203] P. Kamat and R. Sugandhi, “Anomaly detection for predictive maintenance in industry 4.0-A survey,” *E3S Web of Conferences*, vol. 170, pp. 1–8, 2020.
- [204] P. Ramos, J. M. Oliveira, and P. Silva, “Predictive maintenance of production equipment based on neural network autoregression and ARIMA,” *21st International EurOMA Conference - Operations Management in an Innovation Economy*, pp. 1–10, 2014.

- [205] K. Hundman, V. Constantinou, C. Laporte, I. Colwell, and T. Soderstrom, “Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding,” *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 387–395, 2018.
- [206] N. Kolokas, T. Vafeiadis, D. Ioannidis, and D. Tzovaras, “A generic fault prognostics algorithm for manufacturing industries using unsupervised machine learning classifiers,” *Simulation Modelling Practice and Theory*, vol. 103, no. January, p. 102109, 2020.
- [207] M. Braei and S. Wagner, “Anomaly Detection in Univariate Time-series: A Survey on the State-of-the-Art,” 2020.
- [208] E. Borghini, C. Giannetti, and J. Flynn, “Data-Driven Energy Storage Scheduling to Minimise Peak Demand on Distribution Systems with PV Generation,” pp. 1–22, 2021.
- [209] A. Essien and C. Giannetti, “A Deep Learning Model for Smart Manufacturing Using Convolutional LSTM Neural Network Autoencoders,” *IEEE Transactions on Industrial Informatics*, vol. 16, no. 9, pp. 6069–6078, 2020.
- [210] C. Giannetti and A. Essien, “Towards scalable and reusable predictive models for cyber twins in manufacturing systems,” *Journal of Intelligent Manufacturing*, 2021.
- [211] C. Point, “One-class Classification in the presence of Point, Collective, and Contextual Anomalies,” no. November 2018, 2019.
- [212] V. Ceronmani Sharmila, K. R. Kumar, R. Sundaram, D. Samyuktha, and R. Harish, “Credit Card Fraud Detection Using Anomaly Techniques,” *Proceedings of 1st International Conference on Innovations in Information and Communication Technology, ICICT 2019*, vol. 10, no. 1, pp. 7–12, 2019.
- [213] T. M. Mitchell, *Machine Learning*. McGraw-Hill, 1997.
- [214] V. L. Cao and J. Mcdermott, “Collective Anomaly Detection based on Long Short Term Memory Recurrent Neural Network,” *Agence nationale pour l’amélioration des conditions de travail*, 2017.
- [215] M. Ahmed and A. Mahmood, “Network traffic analysis based on collective anomaly detection,” pp. 1141–1146, 2014.

- [216] M. A. Salama, H. F. Eid, R. Ramadan, and A. Darwish, “Hybrid Intelligent Intrusion Detection Scheme,” no. January, 2011.
- [217] J. Cao, Z. Li, and J. Li, “Financial time series forecasting model based on CEEMDAN and LSTM,” *Physica A: Statistical Mechanics and its Applications*, vol. 519, pp. 127–139, 2019.
- [218] P. Malhotra, L. Vig, G. Shroff, and P. Agarwal, “Long Short Term Memory networks for anomaly detection in time series,” *23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2015 - Proceedings*, no. April, pp. 89–94, 2015.
- [219] P. Malhotra, A. Ramakrishnan, G. Anand, L. Vig, P. Agarwal, and G. Shroff, “LSTM-based Encoder-Decoder for Multi-sensor Anomaly Detection,” 2016.
- [220] S. Chauhan and L. Vig, “Anomaly detection in ECG time signals via deep long short-term memory networks,” *Proceedings of the 2015 IEEE International Conference on Data Science and Advanced Analytics, DSAA 2015*, 2015.
- [221] O. Fink, Q. Wang, M. Svensén, P. Dersin, W. J. Lee, and M. Ducoffe, “Potential, challenges and future directions for deep learning in prognostics and health management applications,” *Engineering Applications of Artificial Intelligence*, vol. 92, no. May, p. 103678, 2020.
- [222] M. Canizo, I. Triguero, A. Conde, and E. Onieva, “Multi-head CNN–RNN for multi-time series anomaly detection: An industrial case study,” *Neurocomputing*, vol. 363, pp. 246–260, 2019.
- [223] H. Qiao, T. Wang, P. Wang, S. Qiao, and L. Zhang, “A time-distributed spatiotemporal feature learning method for machine health monitoring with multi-sensor time series,” *Sensors (Switzerland)*, vol. 18, no. 9, 2018.
- [224] O. Janssens, V. Slavkovikj, B. Vervisch, K. Stockman, M. Loccufer, S. Verstockt, R. Van de Walle, and S. Van Hoecke, “Convolutional Neural Network Based Fault Detection for Rotating Machinery,” *Journal of Sound and Vibration*, vol. 377, pp. 331–345, 2016.
- [225] T. Zonta, C. André, R. Righi, M. José, D. Lima, E. Silveira, and G. Pyng, “Computers & Industrial Engineering Predictive maintenance in the Industry 4 . 0 : A

- systematic literature review,” *Computers & Industrial Engineering*, vol. 150, no. August, p. 106889, 2020.
- [226] E. Quatrini, F. Costantino, G. Di Gravio, and R. Patriarca, “Machine learning for anomaly detection and process phase classification to improve safety and maintenance activities,” *Journal of Manufacturing Systems*, vol. 56, no. May, pp. 117–132, 2020.
- [227] M. Carletti, C. Masiero, A. Beghi, and G. A. Susto, “A deep learning approach for anomaly detection with industrial time series data: A refrigerators manufacturing case study,” *Procedia Manufacturing*, vol. 38, no. 2019, pp. 233–240, 2019.
- [228] A. Diez-Olivan, J. Del Ser, D. Galar, and B. Sierra, “Data fusion and machine learning for industrial prognosis: Trends and perspectives towards Industry 4.0,” *Information Fusion*, vol. 50, no. July 2018, pp. 92–111, 2019.
- [229] D. Rumelhart, G. Hinton, and R. Williams, “Learning internal representations by error propagation,” Tech. Rep. ICS 8504, University of California: Institute for Cognitive Science, San Diego, 1985.
- [230] S. Hochreiter and J. Schmidhuber, “LONG SHORT-TERM MEMORY,” *Neural Computation*, vol. 9, no. 8, pp. 1–32, 1997.
- [231] E. Boser, N. Vapnik, I. M. Guyon, and T. B. Laboratories, “Training Algorithm Margin for Optimal Classifiers,” pp. 144–152.
- [232] S. Hawkins, H. He, G. Williams, and R. Baxter, “Outlier Detection Using Replicator Neural Networks,”
- [233] A. Chmielewski and S. T. Wierzcho, “V-Detector algorithm with tree-based structures,” no. November 2006, 2014.
- [234] A. Agata, “Demo files for predictive maintenance.” <https://www.mathworks.com/matlabcentral/fileexchange/63012-demo-files-for-predictive-maintenance>, 2022.
- [235] S. Maleki, S. Maleki, and N. R. Jennings, “Unsupervised anomaly detection with LSTM autoencoders using statistical data-filtering,” *Applied Soft Computing*, vol. 108, p. 107443, 2021.

- [236] S. Thudumu, P. Branch, J. Jin, and J. J. Singh, “A comprehensive survey of anomaly detection techniques for high dimensional big data,” *Journal of Big Data*, vol. 7, no. 1, 2020.
- [237] N. Amruthnath and T. Gupta, “A research study on unsupervised machine learning algorithms for early fault detection in predictive maintenance,” *2018 5th International Conference on Industrial Engineering and Applications, ICIEA 2018*, no. August 1993, pp. 355–361, 2018.
- [238] A. Singhal and D. E. Seborg, “Clustering multivariate time-series data,” no. January, pp. 427–438, 2006.
- [239] A. Diez-Olivan, J. A. Pagan, R. Sanz, and B. Sierra, “Data-driven prognostics using a combination of constrained K-means clustering, fuzzy modeling and LOF-based score,” *Neurocomputing*, vol. 241, pp. 97–107, 2017.
- [240] P. Zschech, K. Heinrich, R. Bink, and J. S. Neufeld, “Prognostic Model Development with Missing Labels: A Condition-Based Maintenance Approach Using Machine Learning,” *Business and Information Systems Engineering*, vol. 61, no. 3, pp. 327–343, 2019.
- [241] Z. Xue, Y. Shang, and A. Feng, “Semi-supervised outlier detection based on fuzzy rough C-means clustering,” *Mathematics and Computers in Simulation*, vol. 80, no. 9, pp. 1911–1921, 2010.
- [242] G. Manco, E. Ritacco, P. Rullo, L. Gallucci, W. Astill, D. Kimber, and M. Antonelli, “Fault detection and explanation through big data analysis on sensor streams,” *Expert Systems with Applications*, vol. 87, pp. 141–156, 2017.
- [243] E. Marchi, F. Vesperini, F. Weninger, F. Eyben, S. Squartini, and B. Schuller, “Non-linear prediction with LSTM recurrent neural networks for acoustic novelty detection,” *Proceedings of the International Joint Conference on Neural Networks*, vol. 2015-Septe, 2015.
- [244] D. A. Reynolds, “Gaussian mixture models.,” *Encyclopedia of biometrics*, vol. 741, no. 659-663, 2009.
- [245] A. Reddy, M. Ordway-west, M. Lee, M. Dugan, J. Whitney, R. Kahana, B. Ford, J. Muedsam, A. Henslee, and M. Rao, “Using Gaussian Mixture Models to Detect

- Outliers in Seasonal Univariate Network Traffic,” in *IEEE Symposium on Security and Privacy Workshops Using*, pp. 229–234, IEEE, 2017.
- [246] K. P. F.R.S., “Liii. on lines and planes of closest fit to systems of points in space,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [247] R. H. Lesch and D. Lowe, “Component Analysis in Financial Time Series,” in *Proceedings of the IEEE/IAFE 1999 Conference on Computational Intelligence for Financial Engineering (CIFEr)(IEEE Cat. No. 99TH8408)*, pp. 183–190, 1999.
- [248] D. Palac, J. Borr, and L. Thaise, “Multivariate Six Sigma: A Case Study in Industry 4.0,” pp. 1–20, 2020.
- [249] L. V. D. Maaten and G. Hinton, “Visualizing Data using t-SNE,” vol. 9, pp. 2579–2605, 2008.
- [250] L. Mcinnes, J. Healy, and J. Melville, “UMAP : Uniform Manifold Approximation and Projection for Dimension Reduction,” 2020.
- [251] G. Hinton and S. Roweis, “Stochastic Neighbor Embedding,”
- [252] S. Wang, J. Xiang, Y. Zhong, and Y. Zhou, “Convolutional neural network-based hidden Markov models for rolling element bearing fault identification,” *Knowledge-Based Systems*, vol. 144, pp. 65–76, 2018.
- [253] S.-k. S. Fan, D.-m. Tsai, C.-h. Jen, C.-y. Hsu, F. He, and L.-t. Juan, “Data Visualization of Anomaly Detection in Semiconductor Processing Tools,” vol. 35, no. 2, pp. 186–197, 2022.
- [254] M. Ali, R. Borgo, and M. W. Jones, “Concurrent time-series selections using deep learning and dimension reduction,” *Knowledge-Based Systems*, vol. 233, p. 107507, 2021.
- [255] A. Nanduri and L. Sherry, “ANOMALY DETECTION IN AIRCRAFT DATA USING RECURRENT NEURAL NETWORKS ( RNN ) Previous Research in Flight Data Anomaly Detection,” in *2016 Integrated Communications Navigation and Surveillance (ICNS) Conference*, pp. 1–8, IEEE, 2016.



- [256] C. Syms, “Principal Components Analysis,” in *Encyclopedia of Ecology, Five-Volume Set*, pp. 2940–2949, 2008.
- [257] K. R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf, “An introduction to kernel-based learning algorithms,” *Handbook of Neural Network Signal Processing*, vol. 12, no. 2, pp. 4–1–4–40, 2001.
- [258] T. M. Cover and P. E. Hart, “Nearest Neighbor Pattern Classification,” *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.
- [259] J. Tang, J. Liu, M. Zhang, and Q. Mei, “Visualizing Large-Scale and High Dimensional Data,” in *International World Wide Web Conference committee (IW3C2)*, pp. 287–297, 2016.
- [260] H. Shimodaira, *Informatics 2B*, vol. 2. University of Edinburgh: University of Edinburgh, 2014.
- [261] S. Akogul and M. Erisoglu, “A comparison of information criteria in clustering based on mixture of multivariate normal distributions,” *Mathematical and Computational Applications*, vol. 21, no. 3, 2016.
- [262] H. Elayan, M. Aloqaily, and M. Guizani, “Digital twin for intelligent context-aware iot healthcare systems,” *IEEE Internet of Things Journal*, 2021.
- [263] M. C., “Uniform manifold approximation and projection (umap).” <https://www.mathworks.com/matlabcentral/fileexchange/71902>, 2022.
- [264] H. Vinutha, B. Poornima, and B. Sagar, “Detection of outliers using interquartile range technique from intrusion dataset,” in *Information and decision sciences*, pp. 511–518, Springer, 2018.
- [265] C. L. Liu, W. H. Hsaio, and Y. C. Tu, “Time Series Classification with Multivariate Convolutional Neural Network,” *IEEE Transactions on Industrial Electronics*, vol. 66, no. 6, pp. 4788–4797, 2019.
- [266] A. Baldominos, A. Cervantes, Y. Saez, and P. Isasi, “A comparison of machine learning and deep learning techniques for activity recognition using mobile devices,” *Sensors (Switzerland)*, vol. 19, no. 3, 2019.

- [267] F. Lateef and Y. Ruichek, “Survey on semantic segmentation using deep learning techniques,” *Neurocomputing*, vol. 338, pp. 321–348, apr 2019.
- [268] L. Yin, Q. Cheng, Z. Wang, and Z. Shao, “‘Big data’ for pedestrian volume: Exploring the use of Google Street View images for pedestrian counts,” *Applied Geography*, vol. 63, pp. 337–345, 2015.
- [269] F. Y. Gong, Z. C. Zeng, F. Zhang, X. Li, E. Ng, and L. K. Norford, “Mapping sky, tree, and building view factors of street canyons in a high-density urban environment,” *Building and Environment*, vol. 134, no. March, pp. 155–167, 2018.
- [270] Office for Low Emission Vehicles, “Electric Vehicle Homecharge Scheme: Guidance for customers,” Tech. Rep. November, Government of the United Kingdom, London, 2021.
- [271] J. Flynn, E. Brealy, and C. Giannetti, “Making Green Transport a Reality: A Classification Based Data Analysis Method to Identify Properties Suitable for Electric Vehicle Charging Point Installation,” *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, no. 2018, pp. 6229–6232, 2021.
- [272] E. Brealy, J. Flynn, and A. Luckman, “Multi-criteria approach using neural networks, gis, and remote sensing to identify households suitable for electric vehicle charging,” in *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, pp. 283–286, 2022.
- [273] “Ev volumes datacenter.” <https://www.ev-volumes.com/datacenter/>. Accessed: 2022-04-19.
- [274] F. Ahmad, A. Iqbal, I. Ashraf, M. Marzband, and I. Khan, “Optimal location of electric vehicle charging station and its impact on distribution network: A review,” *Energy Reports*, vol. 8, pp. 2314–2333, 2022.
- [275] S. Hosseini and M. D. Sarder, “Development of a Bayesian network model for optimal site selection of electric vehicle charging station,” *International Journal of Electrical Power and Energy Systems*, vol. 105, no. April 2018, pp. 110–122, 2019.
- [276] Field Dynamics, “On-street households: the next EV challenge and opportunity,” p. 19, 2021.

- [277] K. S. Chen, M. M. Crawford, P. Gamba, and J. S. Smith, “Guest editorial introduction for the special issue on remote sensing for major disaster prevention, monitoring, and assessment,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 6, pp. 1515–1518, 2007.
- [278] T. R. Martha, N. Kerle, C. J. Van Westen, V. Jetten, and K. V. Kumar, “Segment optimization and data-driven thresholding for knowledge-based landslide detection by object-based image analysis,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 12 PART 1, pp. 4928–4943, 2011.
- [279] X. Li and G. Shao, “Object-based urban vegetation mapping with high-resolution aerial photography as a single data source,” *International Journal of Remote Sensing*, vol. 34, no. 3, pp. 771–789, 2013.
- [280] J. Kang, M. Körner, Y. Wang, H. Taubenböck, and X. X. Zhu, “Building instance classification using street view images,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 145, pp. 44–59, nov 2018.
- [281] S. Srivastava, J. E. Vargas Muñoz, S. Lobry, and D. Tuia, “Fine-grained landuse characterization using ground-based pictures: a deep learning solution based on globally available data,” *International Journal of Geographical Information Science*, pp. 1–20, 2018.
- [282] S. Srivastava, J. E. Vargas-Muñoz, and D. Tuia, “Understanding urban landuse from the above and ground perspectives: A deep learning, multimodal solution,” *Remote Sensing of Environment*, vol. 228, no. October 2018, pp. 129–143, 2019.
- [283] D. Leung and S. Newsam, “Exploring geotagged images for land-use classification,” *GeoMM 2012 - Proceedings of the 2012 ACM International Workshop on Geotagging and Its Applications in Multimedia, Co-located with ACM Multimedia 2012*, pp. 3–8, 2012.
- [284] B. Huang, B. Zhao, and Y. Song, “Urban land-use mapping using a deep convolutional neural network with high spatial resolution multispectral remote sensing imagery,” *Remote Sensing of Environment*, vol. 214, no. October 2017, pp. 73–86, 2018.

- [285] G. Cheng and J. Han, “A survey on object detection in optical remote sensing images,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 117, pp. 11–28, 2016.
- [286] “Google street view.” <https://www.google.com/streetview/>. Accessed: 2022-04-14.
- [287] “Mapillary.” <https://www.mapillary.com/>. Accessed: 2022-04-14.
- [288] “Kartaview.” <https://kartaview.org/landing>. Accessed: 2022-04-14.
- [289] E. Barbierato, I. Bernetti, I. Capecchi, and C. Saragosa, “Integrating remote sensing and street view images to quantify urban forest ecosystem services,” *Remote Sensing*, vol. 12, no. 2, pp. 1–22, 2020.
- [290] P. Stubbings, J. Peskett, F. Rowe, and D. Arribas-Bel, “A Hierarchical Urban Forest Index Using Street-Level Imagery and Deep Learning,” *Remote Sensing*, vol. 11, no. 12, p. 1395, 2019.
- [291] M. Swain and D. Ballard, “Color indexing,” *International Journal of Computer Vision*, vol. 7, no. 1, pp. 11–32, 1991.
- [292] R. M. Haralick, I. Dinstein, and K. Shanmugam, “Textural Features for Image Classification,” *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-3, no. 6, pp. 610–621, 1973.
- [293] N. D. Riggan Jr, R. C. Weih Jr, N. D. Jr, R. C. Jr, N. Riggan, and R. Weih, “Comparison of Pixel-based versus Object-based Land Use/Land Cover Classification Methodologies Recommended Citation A Comparison of Pixel-based versus Object-based Land Use/Land Cover Classification Methodologies,” tech. rep., 2009.
- [294] F. Enríquez, L. M. Soria, J. Antoni3 Alvarez-García, F. Velasco, and O. Déniz, “Existing Approaches to Smart Parking: An Overview,” in *Smart-CT 2017: Smart Cities*, pp. 63–74, 2017.
- [295] W. Zhao, Y. Bo, J. Chen, D. Tiede, B. Thomas, and W. J. Emery, “Exploring semantic elements for urban scene recognition: Deep integration of high-resolution imagery and OpenStreetMap (OSM),” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 151, no. March, pp. 237–250, 2019.

- [296] M. Kampffmeyer, A.-B. Salberg, and R. Jenssen, “Semantic Segmentation of Small Objects and Modeling of Uncertainty in Urban Remote Sensing Images Using Deep Convolutional Neural Networks,” tech. rep.
- [297] J. Schmidhuber, “Deep Learning in neural networks: An overview,” *Neural Networks*, vol. 61, pp. 85–117, 2015.
- [298] L.-J. Li, K. Li, F. F. Li, J. Deng, W. Dong, R. Socher, and L. Fei-Fei, “ImageNet: a Large-Scale Hierarchical Image Database,” *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- [299] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, “Places: A 10 Million Image Database for Scene Recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 6, pp. 1452–1464, 2018.
- [300] G. Amato, F. Carrara, F. Falchi, C. Gennaro, C. Meghini, and C. Vairo, “Deep learning for decentralized parking lot occupancy detection,” *Expert Systems with Applications*, vol. 72, pp. 327–334, 2017.
- [301] B. Zhao, B. Huang, and Y. Zhong, “Transfer Learning with Fully Pretrained Deep Convolution Networks for Land-Use Classification,” *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 9, pp. 1436–1440, 2017.
- [302] G. Cheng, P. Zhou, and J. Han, “Learning Rotation-Invariant Convolutional Neural Networks for Object Detection in VHR Optical Remote Sensing Images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 12, pp. 7405–7415, 2016.
- [303] M. Castelluccio, G. Poggi, C. Sansone, and L. Verdoliva, “Land Use Classification in Remote Sensing Images by Convolutional Neural Networks,” pp. 1–11, 2015.
- [304] L. Liu, H. Wang, and C. Wu, “A machine learning method for the large-scale evaluation of urban visual environment,” no. Harvey 2014, 2016.
- [305] E. J. Hoffmann, Y. Wang, M. Werner, and J. Kang, “Model Fusion for Building Type Classification from Aerial and Street View Images,” pp. 1–20, 2019.
- [306] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, B. Schiele, D. A. Rd, and T. U. Darmstadt, “The Cityscapes Dataset for Semantic Urban Scene Understanding,” tech. rep.

- [307] A. Dubey, N. Naik, D. Parikh, R. Raskar, and C. A. Hidalgo, “Deep Learning the City: Quantifying Urban Perception at a Global Scale,” in *Computer Vision – ECCV 2016*, pp. 196–212, Springer, Cham, 2016.
- [308] B. Y. Cai, X. Li, I. Seiferling, and C. Ratti, “Treepedia 2.0: Applying Deep Learning for Large-Scale Quantification of Urban Tree Cover,” *Proceedings - 2018 IEEE International Congress on Big Data, BigData Congress 2018 - Part of the 2018 IEEE World Congress on Services*, pp. 49–56, 2018.
- [309] Y. Li, Y. Chen, A. Rajabifard, K. Khoshelham, and M. Aleksandrov, “Estimating Building Age from Google Street View Images Using Deep Learning,” in *GIScience 2018* (M. S. Stephan Winter, Amy Griffin, ed.), pp. 40:1 – 40:7, LIPICS, 2018.
- [310] K. Zhao, Y. Liu, S. Hao, S. Lu, H. Liu, and L. Zhou, “Bounding Boxes Are All We Need : Street View Image Classification via Context Encoding of Detected Buildings,” vol. 60, 2022.
- [311] C. Doersch and S. Singh, “What Makes Paris Look like Paris ?,” 2011.
- [312] F. Alhasoun and M. González, “Streetify: Using Street View Imagery And Deep Learning For Urban Streets Development,” *arXiv*, pp. 2001–2006, 2019.
- [313] T. L. Hamilton and E. B. Johnson, “Using Machine Learning and Google Street View to Estimate Visual Amenity Values,” 2018.
- [314] X. Li, C. Zhang, W. Li, R. Ricard, Q. Meng, and W. Zhang, “Assessing street-level urban greenery using Google Street View and a modified green view index,” *Urban Forestry and Urban Greening*, vol. 14, no. 3, pp. 675–685, 2015.
- [315] S. Law, Y. Shen, and C. Seresinhe, “An application of convolutional neural network in street image classification,” *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery - GeoAI '17*, pp. 5–9, 2017.
- [316] J. Zhong, J. Gao, R. Chen, and J. Li, “Digital recognition of street view house numbers based on DCGAN,” no. x, pp. 19–22, 2019.
- [317] I. Seiferling, N. Naik, C. Ratti, and R. Proulx, “Green streets Quantifying and mapping urban trees with street-level imagery and computer vision,” *Landscape and Urban Planning*, vol. 165, no. July 2016, pp. 93–101, 2017.

- [318] “Open street map.” <https://www.openstreetmap.org/>. Accessed: 2022-04-15.
- [319] W. Li, C. He, J. Fang, J. Zheng, and H. Fu, “Semantic Segmentation-Based Building Footprint Extraction Using Very High-Resolution Satellite,” *Remote Sensing*, vol. 1, 2019.
- [320] “Ordnance survey licencing agreements.” <https://www.ordnancesurvey.co.uk/business-government/licensing-agreements/standard-form-contractor-licence>. Accessed: 2022-04-14.
- [321] World Economic Forum, “The Future of Jobs Report 2020 — World Economic Forum,” *The Future of Jobs Report*, no. October, p. 1163, 2020.
- [322] I. Filippidis, “Openstreetmap: Interface to openstreetmap.” <https://github.com/johnyf/openstreetmap>, 2013.
- [323] J. Flynn and C. Giannetti, “Using Deep Learning to Map Houses Suitable for Electric Vehicle Home Charging,” *AI*, pp. 1–14, 2020.
- [324] E. Breal, J. Flynn, and A. Luckman, “Multi-Criteria Approach Using Neural Networks, GIS, and Remote Sensing to Identify Households Suitable for Electric Vehicle Charging,” *International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 2022-July, pp. 283–286, 2022.
- [325] J. Flynn, H. V. Dijk, and C. Giannetti, “Anomaly Detection of DC Nut Runner Processes in Engine Assembly,” *AI (Submitted for Review)*, 2022.
- [326] T. Ouyang, H. Huang, Y. He, and Z. Tang, “Chaotic wind power time series prediction via switching data-driven modes,” *Renewable Energy*, vol. 145, pp. 270–281, 2020.
- [327] S. Hagemann, A. Sünnetcioglu, and R. Stark, “Hybrid artificial intelligence system for the design of highly-automated production systems,” *Procedia Manufacturing*, vol. 28, pp. 160–166, 2019.
- [328] W. S. McCulloch and W. Pitts, “A logical calculus of the ideas immanent in nervous activity,” *The bulletin of mathematical biophysics*, vol. 5, no. 4, pp. 115–133, 1943.
- [329] J. Mao, “Artificial neural networks: a tutorial,” *Computer*, vol. 29, no. 3, 1996.

- [330] P. Skalski, “Gentle dive into math behind convolutional neural networks.” <https://towardsdatascience.com/gentle-dive-into-math-behind-convolutional-neural-networks-79a07dd44cf9>, 2019.
- [331] Y. Bengio and P. Haffner, “Gradient-Based Learning Applied to Document Recognition,” vol. 86, no. 11, 1998.
- [332] G. A. Hembury, V. V. Borovkov, J. M. Lintuluoto, and Y. Inoue, “Deep Residual Learning for Image Recognition,” *Cvpr*, vol. 32, no. 5, pp. 428–429, 2016.
- [333] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [334] P. Radhika, “The mathematics behind support vector machine algorithm (svm).” [analyticsvidhya.com/blog/2020/10/the-mathematics-behind-svm/](https://analyticsvidhya.com/blog/2020/10/the-mathematics-behind-svm/), 2020.
- [335] T. Evgeniou and M. Pontil, “Support Vector Machines : Theory and Applications WORKSHOP ON SUPPORT VECTOR MACHINES : THEORY AND APPLICATIONS,” in *Machine Learning and Its Applications, Advanced Lectures*, no. January 2001, pp. 249–257, 2001.
- [336] L. BREIMAN, “Random Forests,” *Machine Learning*, vol. 45, pp. 5–32, 2001.
- [337] C. D. Sutton, “Classification and Regression Trees, Bagging, and Boosting,” *Handbook of Statistics*, vol. 24, no. 04, pp. 303–329, 2005.



# APPENDICES

## .1 Machine Learning Models and Descriptions

The following section provides details of the various machine learning scenarios, as well as the most common machine learning models presented in the literature reviewed in Chapter 2.

### Hyper-Parameter Optimisation and Evaluation Metrics

Many papers reviewed in Chapter 2 lack details on the specific methods used to select hyper-parameters, but those that do typically use simple methods such as trial and error [91], random search [78], or grid search [79, 79]. Some paper mentions using more advanced solutions such as Bayesian optimisation, which uses gradient descent to optimize multiple hyper-parameters [76, 74]. Bayesian optimisation demands high computational requirements that increase in proportion to the number of parameters selected to optimize, a challenge that some researchers overcome through the combination of both grid search and Bayesian optimisation [81].

The most common metric for evaluating machine learning methods for classification tasks is F-score [98, 80, 88, 81, 87, 101]. F-score is defined as:

$$F - score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (1)$$

where precision is the ratio of between true positives and all positives,

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

And recall is the measure of the method to identify true positives correctly

$$Recall = \frac{TP}{TP + FN}. \quad (3)$$

Some researchers only discuss precision and recall when evaluating quality inspection performance in real-world settings as it gives a more detailed insight into the true positive and false negative rates of the proposed inspection method than F-score alone [108, 109]. Only one researcher was found to use the area under the Receiver Operating Characteristics (ROC) curve, a metric commonly called AUC [75]. The ROC curve is produced by plotting the true positive rate over the false positive rate as the classification threshold varies, as shown in Figure 1. The author notes that despite not appearing in the searches, AUC-ROC is a common evaluation metric in machine learning, particularly in instances that involve binary classification as it gives a good quantitative and visual measure of how well a model distinguishes between two classes. However, Pereira et al. note in their research that this is an uncommon metric in the industrial quality detection literature [75]. For research on time series prediction, researchers often present multiple evaluation met-




Figure 1: The area under the ROC curve is a metric used to evaluate the performance of classifiers. (Source: [https://en.wikipedia.org/wiki/Receiver\\_operating\\_characteristic](https://en.wikipedia.org/wiki/Receiver_operating_characteristic))

rics to evaluate their findings, including Root Mean Squared Error (RMSE), Mean Square Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) [106, 103, 29, 83, 76]. In the reviewed literature, no researchers that explore time series forecasting justified the metrics used, nor any mathematical definitions of these metrics. To address this gap, the author explores additional external sources that include a good description and mathematical formulas for these metrics [326]. Ouyang et al. discuss that there are two main evaluation metrics in time series prediction: longitudinal and transverse [326]. Longitudinal errors are a measure of the deviation of the prediction amplitude from the actual values and include RMSE, MSE, MAE, and MAPE [326]. The mathematical

definitions of the metrics as mentioned above are defined as follows:

$$RMSE = \sqrt{\sum_{i=1}^n \frac{(y_i - \hat{y}_i)^2}{n}}$$

$$MSE = \sum_{i=1}^n \frac{(y_i - \hat{y}_i)^2}{n}$$

$$MAE = \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{n}$$

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100$$

where  $y_i$  are the actual values,  $\hat{y}_i$  are the predicted values. The error should be minimised in each of the above metrics, meaning the forecast perfectly matches the actual values. RMSE and MAE are in the same units as the target variable, while MAPE is given as a percentage. The main difference between RMSE and MAE is that because RMSE squares the errors before they are averaged, RMSE gives higher penalties for larger errors and therefore is more sensitive to outliers. For this reason, RMSE can be preferable in cases where the impact of unit increases in error become exponentially worse, i.e., being off by 10 is more than twice as bad as being off by 5. On the other hand, if a model is prone to making occasional large mistakes that do not impact overall performance, the MAE may be more appropriate. Note that the RMSE will always be larger or equal to the MAE. The sensitivity to outliers is further exaggerated by MSE, given the removal of the square root. By including multiple metrics for comparison, these factors can be inferred, and comparing the RMSE and MAE can indicate the model's tendency to give occasional errors. However, it is difficult to tell if inaccuracies are due to the prediction or anomalies in the actual data.

## .1.1 Machine Learning Approaches

The reviewed literature discusses five main types of machine learning: Supervised, Unsupervised, Semi-supervised, Reinforcement learning, and Transfer learning. Each of these learning scenarios is outlined below. A list of all machine learning methods compared in the reviewed literature is summarised in Table 1 along with the respective references.

Table 1: A list of all machine learning methods used in the reviewed literature and the respective references arranged from most common to least common.

Machine Learning Method	Abb.	No. Papers	References
Convolutional Neural Network	CNN	19	[96, 90, 105, 89, 108, 98, 100, 85, 175, 88, 81, 28, 87, 82, 101, 99, 76, 6, 109]
Support-Vector Machine	SVM	8	[103, 150, 78, 79, 80, 175, 81, 82]
Random Forest	RF	7	[78, 79, 175, 81, 75, 83, 76]
Artificial Neural Network	ANN	6	[94, 78, 327, 102, 79, 175]
Auto Encoder	AE	5	[74, 75, 75, 29, 84]
k- Nearest Neighbour	k-NN	5	[78, 80, 175, 81, 76]
Multilayer Perceptron	MLP	5	[79, 175, 81, 82, 76]
k-means	n/a	4	[103, 93, 104, 77]
Principal Component Analysis	PCA	4	[94, 102, 77, 29]
Generative Adversarial Network	GAN	3	[85, 86, 84]
Decision Tree	DT	2	[79, 80]
Support Vector Regression	SVR	2	[83, 76]
XGBoost	XGB	2	[78, 83]
Long Short Term Memory	LSTM	2	[106, 103]
Regression	n/a	1	[79]
Autoregressive Integrated Moving Average	ARIMA	1	[106]
Reinforcement Learning	n/a	1	[91]
Hierarchical Clustering	n/a	1	[107]
Agglomerative clustering	n/a	1	[77]
AutoGluon	n/a	1	[75]
Classical Multidimensional Scaling	CMDS	1	[77]
Deep Belief Network	DBN	1	[29]
Gradient Decent	n/a	1	[93]
Halt-Winter	n/a	1	[106]
Gaussian Mixture Model	GMM	1	[96]
Naïve Bayes	n/a	1	[79]
Transfer Learning	n/a	1	[91]
Maximum-Likelihood Hebbian Learning	MLHL	1	[77]
Cooperative Maximum-Likelihood Hebbian Learning	CMLHL	1	[77]
Automated Machine Learning	AutoML	1	[75]
Sammon Mapping	SM	1	[77]
Factor Analysis	FA	1	[77]

## Supervised Learning

Supervised learning tasks are when features from a labelled dataset consisting of pairs  $\mathbf{x}_i, \mathbf{y}_i$  are used to train a model to learn the mapping function between  $X$  and  $\hat{Y}$ , where  $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  is the set of targets and  $Y = (\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n)$  is the set of corresponding labels [115]. In the supervised learning setting, the problem is well-defined, and performance can be easily measured by making predictions on a labelled test dataset. When labels are continuous, the task is known as regression [83, 115].

## Unsupervised Learning

In contrast to supervised learning, unsupervised approaches train models without needing class labels. This is achieved by identifying trends and patterns in the structured data to extract the relevant features, which can then be used to group or segment the input data into distinct categories [104]. Some unsupervised methods assume that data are drawn independently and identically distributed from the distribution  $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  [115]. Therefore, the goal of the learning problem is to estimate the density likely to have generated  $X$  [115].

## Semi - Supervised Learning

As the name implies, semi-supervised methods lie in between supervised and unsupervised approaches. In an anomaly detection setting, because some amount of labelled data is required to evaluate unsupervised models and the scarcity of anomalies, most of the available labelled data is likely, normal data [30]. Semi-Supervised methods leverage this available data to provide algorithms with some, but not all, information on the targets. The "standard" semi-supervised scenario describes an instance where data contains two classes  $X_l = (\mathbf{x}_1, \dots, \mathbf{x}_l)$  and  $X_u = (\mathbf{x}_1 + \mathbf{1}, \dots, \mathbf{x}_1 + \mathbf{u})$ , where the labels  $Y_l = (\mathbf{y}_1, \dots, \mathbf{y}_l)$  are given, but no labels are given for  $X_u$  [115]. In the anomaly detection case presented in Chapter 4, the normal labels  $Y_l$  are known, while the anomaly labels are unknown  $X_u$ . This scenario where only normal data are used to train a semi-supervised model for anomaly detection is referred to as a "clean" approach or "clean-semi-supervised" approach [198]. While other forms of learning are possible using partial supervision, the clean-semi-supervised setting is the only example of semi-supervised learning in the reviewed literature and, therefore, the

only example discussed in this review. For further information on other semi-supervised settings, the reader is referred to Chapelle et al. [115].

## **Transfer Learning**

Other machine learning techniques include transfer and reinforcement learning; however, these learning scenarios are less frequent in the reviewed literature. Transfer learning is another machine learning technique used when limited labelled data are available. Transfer learning uses existing pre-trained models trained on a target domain and repurposes that model to re-train on a new domain, the idea being that the pre-trained network already has some learned features from the relevant data to the new target. The most common application of transfer learning is in supervised image classification settings, as many open-source image datasets such as ImageNet, MNIST, and CIFAR-10 can be used to train initial classifiers on a target domain [99]. When this pre-trained network is repurposed on a new target, the network retains information on low-level features that are transferable between tasks such as the detection of edges and object boundaries, meaning fewer images are required to learn the mapping between  $x$  and  $y$ . Luckow et al. state that the applicability of transfer learning largely depends on the application. For example, transfer learning has been successfully demonstrated in examples such as social media analytics and computer vision. However, this does not extend to enterprise use cases, the majority of which require custom datasets [28].

## **Reinforcement learning**

In reinforcement learning, an agent learns through trial and error what actions to perform in a specified setting where favourable actions are rewarded and unfavourable actions are penalised [83, 91, 30]. Reinforcement learning has been a key technology in the automotive industry over the past few decades as it has been an enabling factor in autonomous driving; however, applications of reinforcement learning in automotive manufacturing and production are limited [30]. Different types of reinforcement learning can be applied to perform classification or process control. Only temporal difference learning approaches for process control are discussed in the reviewed literature. Temporal difference learning is a learning approach where a numeric reward value is used to incentivize certain actions selected by a machine learning model. The reward is scaled inversely to the time between the initial

action and the respective reward [91]. For a mathematical description of a reward function for temporal difference learning case, see research by Viharos et al. [91].

## .1.2 Common Machine Learning Models

This subsection provides further mathematical details of some of the most commonly used machine learning methods in the reviewed literature, as shown in Table 1. The mathematics behind the LSTM, GMM, PCA, t-SNE, and UMAP are not introduced in this section, as they are discussed in Chapter 5.

### Artificial Neural Network

Artificial Neural Networks (ANNs) have been an area of significant scientific interest since the early 1940s in which McCulloch and Pitts first introduced the artificial neuron depicted in Figure 2 [328]. The neuron performs two computations. First, the weighted sum of some input  $X = \mathbf{x}_1, \dots, \mathbf{x}_n$  is calculated and bias is added. The result is then passed onto some activation function  $\sigma$ , such as tanh or sigmoid, to add some non-linearity to the output such that the output  $\hat{Y}$  is given by,

$$\hat{Y} = \sigma\left(\sum_{i=1}^n \mathbf{x}_i \cdot \mathbf{w}_i + b\right)$$

. This process is called forward propagation [329].

This architecture enables ANNs to perform a range of tasks including classification and




Figure 2: An artificial neuron used in neural networks. (Source: <https://lanstonchu.files.wordpress.com/2021/03/cell.jpeg?w=750>)

regression. However, given that the majority of cases in the reviewed literature explore ANNs for image classification, this subsection focuses on this application to describe the network architecture.

By combining many artificial neurons, a network of fully connected neurons can create a weight-directed graph that maps the input  $X$  onto an output  $\hat{Y}$  as shown in Figure 3. In a supervised setting, the output values  $\hat{y}_i$  can then be compared with the target labels  $y_i$  to evaluate the performance of the classification model. By varying network variables such as the number of layers, biases, weights, activation functions, and input variables, the user can manually tune an ANN to give the best-performing model. Popular types of ANNs include models such as Multilayer Perceptron, which contains 2-3 layers, and Deep Neural Network (DNN), which contains more than 3 layers [329].

In addition to forward propagation is backward propagation, the second learning mechanism used to train ANNs. Backpropagation is a much more complex process in which gradient descent is used to find optimal parameters that minimise the loss function and feed the updated weights and biases back between layers. For a detailed mathematical description of forward and backward propagation, the reader is referred Skalski's online article [330]. Networks that are trained only using forward and backward propagation are called feed-forward networks.




Figure 3: An example of a 3 layer ANN, sometimes called a Multilayer Perceptron. Each node in the network is an artificial neuron depicted in Figure 2 Adding at least one more layer to the network would make it a Deep Neural Network.



## Convolutional Neural Network

The Convolutional Neural Network (CNN) is a supervised learning architecture most commonly used for image recognition. The CNN architecture consists of multiple stacked layers of artificial neurons that take an input  $X$ , usually a 2D array, and performs consecutive transformations  $\phi(X) : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$  onto new reduced feature spaces  $\phi(X) = \mathbb{R}^{d'}$  [6, 105]. Over the multiple layers, the mapping function can be expressed as a nested transformation  $f(x) = f^{(l)}(f^{(l-1)}(\dots f^{(1)}(x)))$ . There are various types of layers throughout the stacked arrangement. The CNN begins with a convolution layer, followed by activation functions, pooling layers, and finally a fully connected layer. This architecture is shown in Figure 5. The convolutional layers each have filter matrices  $K$  that are passed over the input feature space  $X^l$  in a sliding window fashion to produce a new feature map in a reduced feature space  $X^{l-1}$ . For a mathematical description of this process, see [6]. The filter matrices  $K$ , also referred to as kernels, are updated throughout the learning process through forward propagation and backward propagation. Each convolutional layer has a tuneable bias value  $b$  which is used in forward propagation to pass learned information onto the following layer by adding bias to the resultant feature map. An activation function such as tanh, sigmoid, or Rectified Linear Unit (ReLU) is then applied to the output of the feature map. ReLU is the most commonly used activation function in convolutional layers, while sigmoid or softmax are typically used in fully connected layers [6].

Pooling layers are primarily used to reduce the feature space between convolution layers to




Figure 4: Plots showing the most common activation functions used in machine learning algorithms.

speed up the training process. The most common approach to pooling is the max pooling layer where the output matrix  $X^l$  is given by the maximum value of each corresponding subsample in the input space  $X^{l-1}$ , denoted by  $X_{pq}^l = \max(X_{ij}^{l-1})$  where  $p \in \mathbb{R}^{[i,j]}$  and  $q \in \mathbb{R}^{[i,j]}$ .

The final layers in the network architecture are the fully connected layer and a final activation layer, usually sigmoid or softmax. These layers are responsible for combining and computing the scores of the output classes. The class with the highest assigned score is used to determine the label.

Machine learning engineers can create different CNNs for specific tasks by varying input and output filter sizes, tuning weights and biases, and stacking specific combinations of layers. Public datasets such as ImageNet have been a key tool to advance computer vision Research. ImageNet is a large image database containing 1000 images for 1000 different categories [298]. This open source dataset lead to the introduction of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) from which multiple CNN networks have been demonstrated to have high accuracies, including: AlexNet [1], GoogleNet [2], LeNet [331], VGG [3], and ResNet [332]. Further details of these individual CNN variants are discussed further in Chapter 5.




Figure 5: An Example of the CNN architecture [6].

## Autoencoder

An autoencoder is a feed-forward neural network consisting of two parts: an encoder and a decoder. The idea of an encoder network layer dates back to the 1980s as a means of reducing the dimensionality of the input layer by compressing the input data  $X$  into a reduced feature space using an encoding function  $H = f(X)$  [333]. This encoding can occur over a single layer with a high compression ratio or over many layers with gradually decreasing feature space. Encoders are still used today as feature extraction input layers prior to applying some other machine learning techniques [68]. Sometime later in the late 1980s, researchers added a decoder layer that reconstructs the compressed layer  $H$  into a higher dimensional feature space equal to that of the encoder input  $X^l = g(f(X))$  [333]. This process of compression and decompression means the encoder-decoder network, also called an auto-encoder, is forced to learn the most important features to retain during the

compression process such that information on the target variable is retained upon decompression. Figure 6 shows an image showing the autoencoder architecture. The autoencoder uses the same principles of forward and backward propagation to train the weights and biases of the network to minimize some function  $L(X, g(f(X)))$  that penalizes  $g(f(X))$  from being different to  $X$ . In other terms, the network aims to minimize the reconstruction error  $X - g(f(X))$ . This reconstruction error can then be used to evaluate the output. This unique architecture has been applied to a range of use cases such as classification, de-noising and manifold approximation; however, the most common application in the reviewed literature is anomaly detection [333, 68]. In anomaly detection, a clean-semi-supervised approach is used to train a model on normal training data. The network learns to reconstruct normal data onto the output space. This can be done using either images or time series as input. When new unlabelled data is fed into the network, if the reconstruction error is over some threshold  $K$ , the observation is classified as an anomaly [75].

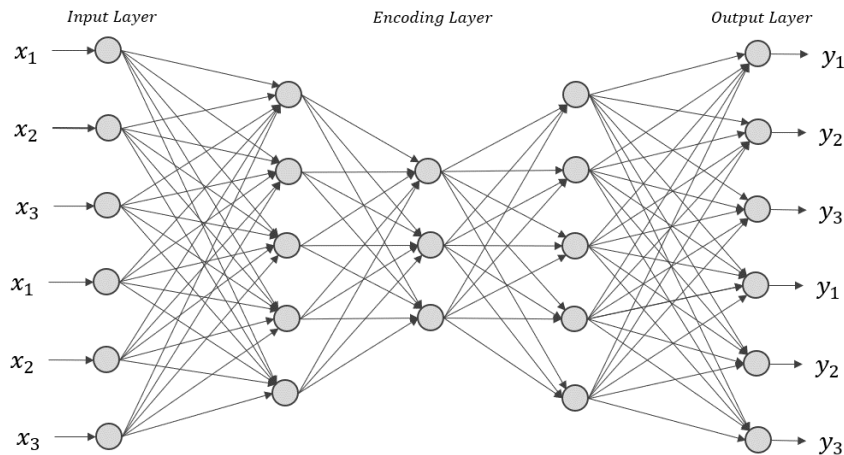


Figure 6: An example of an autoencoder architecture where each node of the network is an artificial neuron.

## Support Vector Machines

This subsection begins by introducing the simplest SVM, a linear SVM in one dimension, before moving on to the higher dimensional SVM problem. Consider a supervised learning scenario where some labelled dataset  $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  with labels  $Y = (\mathbf{y}_1, \dots, \mathbf{y}_n)$  has been mapped onto a 2D feature space  $X'$  shown in Figure 7. A 1D linear decision boundary

given by  $f(x) = ax + b$  can be introduced to separate the two classes by maximising the margin  $w$  between the support vectors of each class. For a detailed explanation of the convex optimisation problem and how this is solved to select the optimal hyperplane, the author refers the reader to the following references [334, 335]. The same problem can be solved in 3D space, where the hyperplane  $H$  is a 2D plane given by  $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x} + b$ . In reality, our dataset is unlikely to be perfectly separated by a linear function in the feature space  $X'$  as shown in Figure 7. In this instance, the so called 'kernel trick' can be used to learn non-linear decision functions by using a kernel  $K$  to transform input space  $X'$  into some higher dimensional feature space in which to apply the hyperplane [30]. The kernel trick is shown graphically in Figure 8.




Figure 7: example of a hyperplane generated by a linear SVM algorithm to separate classes in 2D feature space (Source: <https://www.analyticsvidhya.com/blog/2020/10/the-mathematics-behind-svm/>).




Figure 8: A graphic demonstrating the kernel trick in which the kernel function  $\phi$  being used to map points from a 2D input space into a 3D feature space to learn non-linear relations between the classes (Source: <https://medium.com/@KunduSourodip/finding-non-linear-decision-boundary-in-svm-a89a97a006d2>).

## Random Forest

Random Forest (RF) is a supervised Learning technique first proposed by Breiman in 2001 [336] and has since been applied to a wide range of classification and prediction problems. This section includes references to Breiman's original paper [336] as well as a book by Sutton 2005 on Classification Trees [337].

To understand the RF algorithm, we must first introduce decision tree classifiers, in which a dataset  $X$  is consecutively split into subsets  $B_1, B_2, \dots, B_j$  such that the predicted class of an observation  $\mathbf{x}_i$  is  $j$  if  $\mathbf{x}_i \in B_j$ . A decision tree classifier can be visualised as a tree diagram consisting of two nodes: a decision node and a 'leaf' node. Decision nodes take inputs  $X$  and  $Y$  and select a feature to split the data that results in the highest possible node purity in the child nodes i.e. one child node has a higher proportion of data belonging to one class, while the other child is given data mainly belonging to a different class. While this split will be sub-optimal at the beginning of the tree, through consecutive splits, the resulting child nodes become increasingly purer until a node only contains data from a single class, at which point the branch terminates. These terminating nodes are called leaf nodes. Figure 9 shows an example of a simple decision tree depicting decision nodes, leaf nodes, and their respective class labels. Once a decision tree is trained, new data can be fed in the input root node and passed through the decision nodes and is classified based on the final leaf node it terminates in.

At each decision node, various methods can be used to select the relevant feature to analyse and split the data. For example, the Classification And Regression Trees (CART) algorithm, uses the Gini Index to select the optimal feature where  $GiniIndex = 1 - [(P_-)^2 + (P_+)^2]$  where  $P_+$  is the probability of a positive class and  $P_-$  is the probability of a negative class [337]. Multiple features are tested and the one that gives the lowest Gini Index is selected as this results in the highest purity.

RF differs from the standard decision tree in that it uses an ensemble decision tree, each trained independently with randomly selected training samples and random feature selection at each decision node. While a single tree tends to overfit, by training multiple trees and averaging the result, any bias of the individual trees is reduced [336]. Different approaches can be used to choose the final results such as calculating the majority results of all trees, or calculating the Gini Index. The RF architecture is shown in Figure 10.

Random Forest uses a bagging approach at the input, sometimes known as bootstrapping or bootstrap aggregating. Given a training dataset  $X$ , bagging generates  $m$  new training samples  $X_i$  by sampling randomly and uniformly from  $X$ . In RF, this sampling process is

done with replacement, meaning some values  $x_i$  may appear multiple times in any given sample. Typically, 70% of the bagged samples are used for training (in-bag samples), while the other 30% are used for internal testing through cross-validation (out-of-bag samples) to give some indication to the user of the model's accuracy. A common challenge faced by earlier rule based decision tree approaches was overfitting, as the random selection of training samples means the result is less sensitive to any biases of the original training data [336]. This is overcome in random forests, as Breiman finds that applying this bagging approach increases the overall accuracy of the model while overcoming overfitting [336]. Similarly, the random selection of features at the decision node ensures further variability in the ensemble trees.




Figure 9: A simple decision tree with class labels of each node depicted by coloured dots. Notice the increase in class purity through the tree with nodes terminating when only 1 class label exists in the subset.

## GAN

GAN is a popular approach to creating new synthetic image data for image data. GANs consist of two separate neural networks, a generator and a discriminator. Through an iterative training process, the generator creates new samples of data which are then presented to the discriminator. The discriminator then has to identify if the image has been created synthetically by the generator, or if it was part of the original labelled training set. A simplified diagram of the GAN architecture is shown in Figure 11.

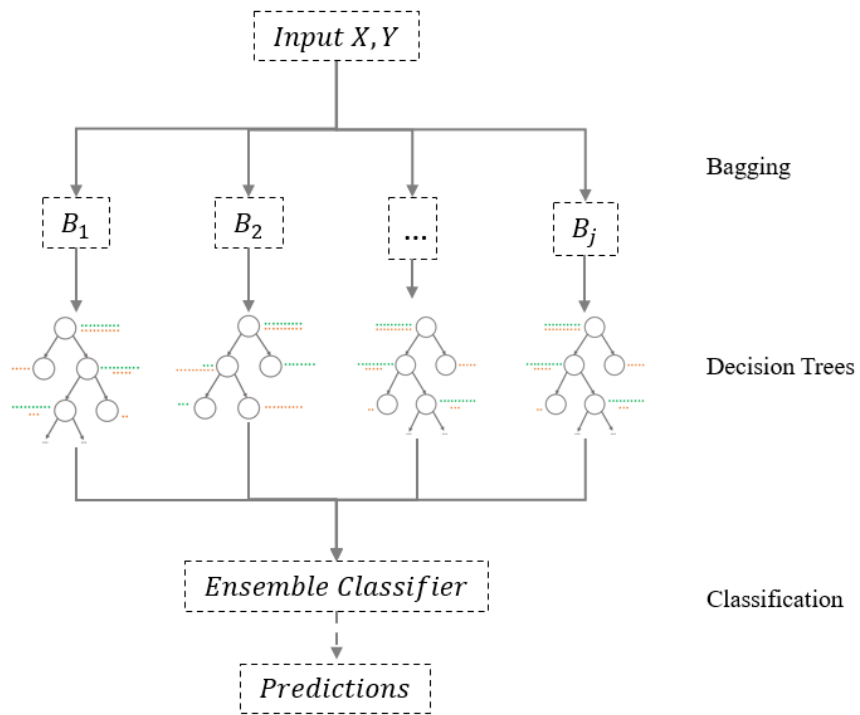


Figure 10: The random forest architecture showing bagging and ensemble tree classification.

The generator's input is simply random noise in the form of a 1D vector sampled from a fixed distribution of the latent space, which is then mapped onto the 2D output space using a series of connected weights [87]. The discriminator is usually some form of basic CNN classifier initially trained on the available labelled data. Both networks are trained using backpropagation of the discriminator's loss resulting in an adversarial setting where the generator is tasked with maximising the error rate of the discriminator, while the discriminator aims to minimise classification error [84]. Initially, the generated images are poor representations of the final image. However, over many iterations, the loss function of both networks will converge, and as the discriminator can no longer distinguish between real and synthetic images. For more details on the mathematics of the GAN architecture, the reader is referred to the original paper by Goodfellow et al. [114].

While GAN is discussed as one of the best ways to generate new image data to expand existing datasets, a large amount of training data (thousands of images) are required to train the discriminator models similar to any other CNN [86]. While GAN are popular for generating image data, they are less suited to create synthetic data for complex multivariate time series data due to high computational requirements during training [75].

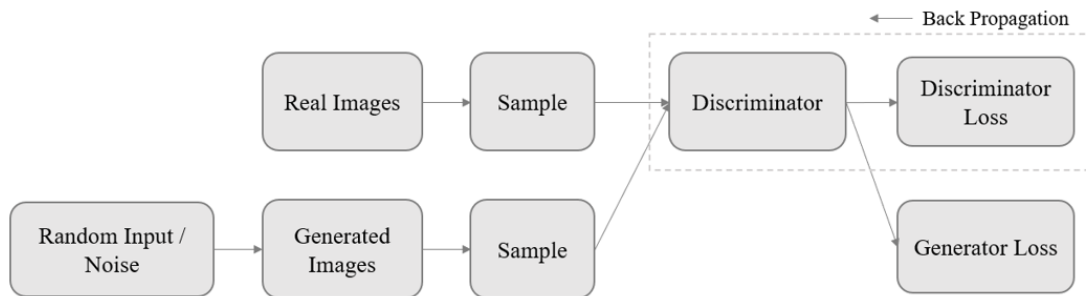


Figure 11: Architecture of a Generative Adversarial Networks (GAN)

## .2 Applying the Assessment tool: A Case Study at Ford Motor Company

A bar on access has been placed on this section due to the sensitive nature of the information presented in this section. Access is restricted to Swansea University students and researchers. This information will be made public on June 1st 2025.



## .3 Industry 4.0 Assessment Questionnaire

Figure 12: Industry 4.0 Assessment Questionnaire page 1

10/12/2022, 07:49

I4 Questionnaire

### I4 Questionnaire

The following questionnaire is estimated to take 10 minutes to complete.

#### Data protection and fair use statement:

This questionnaire is part of an ongoing assessment to better understand the organisations technological and strategic readiness for digitalisation and automation. These findings will be used as part of a wider assessment process to provide feedback to senior management and guide future organisational changes. Results from the aggregate data gathered in this survey will be fully anonymised prior to being used in any reports.

Participation in this assessment is optional. Microsoft Forms does gather information on your name and email should you decided to partake in this assessment. Results are stored on a dedicated SharePoint with the strictly controlled security permissions to ensure that only those who are collating results can view responses and personal information. Based on your responses, we may use this information to get in touch to ask further questions related to your answers to gain further insights into ongoing technological innovation, R&D, and other organisational practices. This information will only be accessible to the primary assessors, Jay Flynn and Michael Wallace. No personally identifiable information or details on responses be shared with others or included in any reports. Personal information will be deleted after 2 months.

\* Required

1. I have read the above data protection and fair use statement and understand that my name, email, and responses to questions are collected as part of this assessment and will be stored for up to 2 months. \*

Yes

## Figure 13: Industry 4.0 Assessment Questionnaire page 2

10/12/2022, 07:49

H Questionnaire

### Defining Industry 4.0

2. In your own words, please describe your understanding of the term 'Industry 4.0'. \*

If you have not heard of Industry 4.0, please enter 'n/a'

## Figure 14: Industry 4.0 Assessment Questionnaire page 3

10/12/2022, 07:49

H Questionnaire

### Automation

3. I am concerned about how automation may effect my future job security.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

4. Which of the following is the most common cause of non-value added administration time in your job role?

- Populating excel spreadsheets
- Weekly report writing
- Mandatory Training
- Email
- Face-to-face meetings
- Non-value added Webex meetings
- Other

5. Real time production data helps support decision making in my department / team / business function.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

Figure 15: Industry 4.0 Assessment Questionnaire page 4

10/12/2022, 07:40

14 Questionnaire

6. Reporting is largely automated in my department / team / business function

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

7. Outdated technologies / legacy systems present challenges in my department / team / business function

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

8. Team members can access structured and unstructured data through an easy to use platform when necessary.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

9. Is your department / team / business function using service-oriented cloud based solutions?

- Yes, cloud solutions are widely used.
- Somewhat. Cloud solutions only used in selected areas.
- No, I understand what cloud services are but I'm not aware of their use in my department.
- I don't know what cloud services are.
- Not applicable

Figure 16: Industry 4.0 Assessment Questionnaire page 5

10/12/2022, 07:49

H Questionnaire

10. Is your department / team / business function using edge computing solutions?

- Yes, edge solutions are widely used.
- Somewhat. Edge solutions are only used in selected areas.
- No. I understand what Edge solutions are but I'm not aware of their use in my department.
- I don't know what Edge computing / Edge solutions are.
- Not applicable

11. Have you had any cloud migration experience in your department / team / business function?

- I don't know what cloud services are.
- Yes, we are actively migrating with external support.
- Yes, internal team members are supporting cloud migration.
- No, we would like to but do not have the required support.
- No ongoing cloud migration.
- Unsure

Figure 17: Industry 4.0 Assessment Questionnaire page 6

10/12/2022, 07:49

H Questionnaire

12. To the best of your knowledge, which of the follow technology does your department / team / business function use for data management?

- Only use spreadsheets and reports
- Various data warehouses that are not well integrated
- Data warehouse, data lake, cloud and other supporting architectures
- Cloud / GCP
- Other
- Unsure

## Figure 18: Industry 4.0 Assessment Questionnaire page 7

10/12/2022, 07:49

I4 Questionnaire

### Main Drivers of Automation

Automation is not limited to hardware but also extends to automating software, data collection, business decisions, and other backroom manufacturing processes

Figure 19: Industry 4.0 Assessment Questionnaire page 8

10/12/2022, 07:49

H Questionnaire

13. How important is the role of automation to deliver improvements in the following areas:

	Not Important at all	Low Importance	Somewhat Important	Very Important	Extremely Important
Increasing production productivity	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Improving workplace experience	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reducing operating costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reducing material costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reducing labour costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Improving traceability through data insights	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reducing lead times	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reducing non-value added admin time	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Improving health and safety	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Improving quality	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>



Figure 20: Industry 4.0 Assessment Questionnaire page 9

10/12/2022, 07:49

H Questionnaire

### Data Analytics

14. Data analytics is important in my department / team / business function.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

15. Which of the follow best describes the capabilities of data analytics in your department / team / business function?

- Descriptive** – the condition, environment, and operation of equipment or processes can be tracked and reported.
- Diagnostic** – Root causes of failures or poor equipment performance can be easily identified using data.
- Predictive** – Data driven models can predict future events
- Prescriptive** – actions are automatically identified and scheduled using data driven models.

16. Based on the descriptions above, have you been involved in any projects that deliver **predictive** or **prescriptive** analytics?

- Yes
- No

17. My department / team / business function is good at creating value from existing data.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

Figure 21: Industry 4.0 Assessment Questionnaire page 10

10/12/2022, 07:49

14 Questionnaire

18. My department / team / business function is good at identifying instances where data are not well utilised.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

19. There is a culture of innovation in my department / team / business function that extends to data analytics

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

20. I can easily get access to data when I need it.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

21. Clear data management practices ensure data are well managed in my department / team / business function.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

Figure 22: Industry 4.0 Assessment Questionnaire page 11

10/12/2022, 07:49

I4 Questionnaire

### Human Resources and Digital Skills

22. Training opportunities are available for those interested in expanding their data analytics knowledge.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

23. I have a good understanding of my departments KPIs

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

24. Business strategy is well communicated throughout the organisation

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

25. A team of analytics experts are available to support me with data analytics when necessary

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

Figure 23: Industry 4.0 Assessment Questionnaire page 12

10/12/2022, 07:49

H Questionnaire

26. How often do you discuss your professional development with your manager / supervisor / team leader / equivalent.

- Weekly
- Monthly
- Quarterly
- Yearly
- Never

Figure 24: Industry 4.0 Assessment Questionnaire page 13

10/12/2022, 07:49

I4 Questionnaire

## Innovation

27. In order for something to be considered an innovation, it must be successful.

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

28. In your own opinion, how would you rank innovation culture in your department / team / business function compared to other industries / competitors

- Industry leading
- High above the industry standard
- Aligned with the industry standard
- Below the industry standard
- Far below the industry standard
- Unsure

29. How would you rank innovation culture in your department compared to other organisations within Ford Motor Company?

- Best in Ford
- High above the organisational standard
- Aligned with the organisational standard
- Below the organisational standard
- Far below the organisational standard
- Unsure

Figure 25: Industry 4.0 Assessment Questionnaire page 14

10/12/2022, 07:49

I4 Questionnaire

30. Innovation plays an important role in my department / team / business function

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

31. Management and senior management actively encourage innovation in my department / team / business function

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

32. I often find time in my work week for self-directed learning and develop my skills

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

33. I would benefit from more opportunities and allocated time for self-directed learning to develop my skills

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---

Figure 26: Industry 4.0 Assessment Questionnaire page 15

10/12/2022, 07:49

H Questionnaire

34. Which of the following best describes your feelings towards innovation in your department / team / business function?

- Innovation often exceeds expectations
- Innovation often meets expectations / fairly successful
- Too early to tell
- Innovation sometimes fails to meet expectations
- Innovation often fails to meet expectations and aborted / planning to abort

35. My department is open to organisational innovation.

(Organisational innovation is the implementation of a new organizational method in the firm's business practices, workplace organisation, or external relations.)

Strongly Disagree

Strongly Agree

1	2	3	4	5	6	7
---	---	---	---	---	---	---