

An Intelligent Inspection System
for
Non-Ferromagnetic Pipes

by
Liam D. Mackenzie

A thesis presented in
fulfilment of the requirement
for the degree of Doctor of Philosophy.

Centre for Ultrasonic Engineering
Department of Electrical and Electronic Engineering
University of Strathclyde

2010

Copyright

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

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

Acknowledgements

I would like to thank Professor Gordon Hayward and Dr Gareth Pierce of the Centre for Ultrasonic Engineering (CUE), for the opportunity to conduct this research and for their support throughout. This research was supported by the Engineering and Physical Sciences Research Council (EPSRC) and the Research Centre for Non-Destructive Evaluation (RCNDE). Thanks to Nexia Solutions Ltd for their support in producing a specification for the inspection system. A big thank you to Gordon Dobie for always being about to bounce ideas off and for generally keeping me amused for three and a half years with his absent-mindedness! Thanks also, to Walter for being on hand to answer my circuit design problems and to Tommy and Alex for the great job they did in fabricating all the mechanical parts for the robot. Finally, thanks to my mother, father and step-father for their support (moral, financial and academic!)

Abstract

The demand for remote inspection of pipework in the nuclear, oil and gas industries provides significant challenges of access, navigation, inspection technique and data communication. Nuclear processing cells, in particular, typically contain several kilometers of densely packed pipework whose actual physical layout may be poorly documented. Access to these pipes is typically available through the radiation shield, via a small removable concrete plug, which may be several meters from the actual inspection site. The high levels of radiation within such cells means that manual inspection is often impossible and, typically, bespoke solutions are utilised on a case specific basis. This thesis describes the development of an intelligent inspection system for non-destructive evaluation on non-ferromagnetic pipes. This wireless robotic platform has a range of kinematic abilities suited to the infrastructural demands of typical industrial facilities. Inspection is conducted using an integrated camera coupled with statistical pattern classification which complements a swept-frequency, eddy current array for rapid imaging of surface breaking defects. Extensive evaluation of the performance of the classification system is provided along with the establishment of a combined theoretical and finite element framework for the eddy current array. Data from each sensor is fused to provide the final inspection result, yielding enhanced detection capabilities.

Table of Contents

Copyright	i
Acknowledgements	ii
Abstract	iii
1 Introduction	1
1.1 Background.....	1
1.1.1 Non-Destructive Evaluation and the Power Industry.....	1
1.1.2 The Move Towards Remote Inspection	3
1.1.3 Typical Inspection Scenario	4
1.2 Aims of this Thesis.....	5
1.3 Contributions to the Field.....	6
1.4 Thesis Content	7
1.5 Publications to Date Arising from this Thesis.....	9
2 Review of Material, Systems and Techniques	10
2.1 Introduction.....	10
2.2 Overview of Austenitic Stainless Steel	10
2.3 Pipe Inspection Robots	13
2.3.1 Traction Methods	13
2.3.2 Pipe Crawling Robots	16
2.4 Influence of Material on Test Method	18
2.4.1 Ultrasonic Testing.....	18
2.4.2 Eddy Current Testing	20
2.5 Review of Eddy Current Testing.....	20
2.5.1 Probe Types.....	20
2.5.2 Applications of Eddy Current Inspection	21
2.5.3 Probe Configuration for Mobile Robotic Applications.....	22
2.5.4 Multi-frequency Eddy Current Techniques.....	23

2.5.5	Commercial Instrumentation.....	24
2.6	Visual Pattern Recognition Techniques	26
2.6.1	Motivation for Incorporating Pattern Recognition.....	26
2.6.2	Review of Applications.....	26
2.6.2.1	Facial Recognition	26
2.6.2.2	Medical Applications.....	27
2.6.2.3	Astronomical Applications.....	30
2.6.2.4	Non-Destructive Evaluation.....	30
2.7	Concluding Remark.....	32
3	Specification of Inspection System	33
3.1	Introduction.....	33
3.2	Robotic Platform	33
3.2.1	Environmental Constraints	33
3.2.2	Performance Constraints.....	35
3.2.3	Levels of Autonomy.....	36
3.2.4	System Specification.....	37
3.3	Pattern Recognition System.....	38
3.3.1	Discussion.....	38
3.3.2	System Specification.....	41
3.4	Eddy Current System	41
3.4.1	Discussion.....	41
3.4.2	System Specification.....	43
4	Robotic Pipe Crawling Platform	44
4.1	Introduction.....	44
4.2	Locomotion and Climbing Mechanisms	44
4.2.1	Traction System	44
4.2.2	Chassis Design	45
4.2.3	Design of Gripper for Mechanical Adhesion.....	49
4.2.4	Cornering Mechanism.....	50
4.3	Mark I - Linear Crawler	51
4.3.1	Overview	51
4.3.2	Selection of Motors	53
4.3.2.1	Gripping System	53
4.3.2.2	Axial Motion.....	53
4.3.2.3	Selection of Gripper Material to Maximise Traction	54
4.3.3	Finite Element Analysis.....	54
4.3.3.1	Introduction	54
4.3.3.2	Simulation Loads and Constraints	55
4.3.3.3	Assessment Criteria and Assumptions	55
4.3.3.4	Results	56
4.3.4	Overview of Mark I Design.....	58

4.4	Mark II - Pipe Crawling System	61
4.4.1	Evaluation of Mark I platform	61
4.4.1.1	System for Increasing Gripping Force	61
4.4.1.2	Weight Reduction Measures	63
4.4.2	Optimisation of Collar Design	65
4.4.3	Finite Element Analysis	70
4.4.3.1	Simulation Loads and Constraints	70
4.4.3.2	Results	70
4.5	Electronic Systems.....	73
4.5.1	Electronic and Processing System	73
4.5.2	Signal Interface Board (SIB)	73
4.5.3	Motor Driver Board	75
4.5.4	Processor Board.....	78
4.5.5	Payload board.....	78
4.6	Communications.....	78
4.6.1	Wireless Protocols	78
4.6.1.1	Wireless RS232.....	78
4.6.1.2	Bluetooth (Class 1 and Class 2).....	79
4.6.1.3	IEEE 802.11g.....	79
4.6.1.4	Certified Wireless USB	79
4.6.1.5	Discussion.....	79
4.6.2	Robot to Host	80
4.6.3	Internal	80
4.7	Power.....	81
4.8	Performance	83
4.8.1	Vertical Ascent/Descent	83
4.8.2	Translation Around Pipe Bend.....	83
4.9	Future Mechanical Development.....	86
4.10	Graphical User Interface	87

5 Pattern Recognition for NDE Images 89

5.1	Introduction.....	89
5.2	Pattern Recognition Methods	89
5.2.1	Introduction.....	89
5.2.2	Feature Selection and Extraction	90
5.2.3	Types of Classifiers.....	91
5.2.4	Overview of Principal Component Analysis	93
5.2.4.1	Theory	93
5.3	Development of Visual Inspection System	96
5.3.1	Image Pre-processing.....	96
5.3.1.1	Illumination	96
5.3.1.2	Grayscale Image.....	97
5.3.1.3	Spatial Frequency Spectrum.....	98

5.3.1.4	Frequency Selection	99
5.3.1.5	Definition of Classes	100
5.3.2	Classification Metrics	102
5.3.3	Statistical Classification: a Bayesian Approach.....	104
5.3.4	Review of Proposed System.....	106
5.3.5	Effect of Pre-Processing and Dimensionality Reduction.....	107
5.4	Binary Classification Method.....	110
5.4.1	Introduction.....	110
5.4.2	Training Process	110
5.4.3	Validation/Test Process	112
5.4.4	Optimising Classification Parameters.....	112
5.4.4.1	Influence of Image Bandwidth.....	112
5.4.4.2	Sensitivity to Volume of Training Data	115
5.4.4.3	Generalisation	116
5.4.5	Performance of Individual Binary Classifiers.....	123
5.4.6	Combining Multiple Classifiers	124
5.4.6.1	Introduction	124
5.4.6.2	Combination Algorithms.....	125
5.4.6.3	Combination Theory.....	126
5.4.6.4	Binary Fusion Results	129
5.5	Multi-class Sensors to Enhance Probability of Detection.....	132
5.5.1	Overview.....	132
5.5.2	Performance of Sensors Approach.....	133
5.5.2.1	Individual and Fused Classifiers.....	133
5.5.2.2	Explanation for Misclassifications	139
5.6	K-Nearest Neighbour Classification.....	140
5.6.1	Overview	140
5.6.2	Results	141
5.7	Performance Comparison Using Real Defects	142
5.7.1	Creating Samples and Image Set	142
5.7.2	Training and Optimisation.....	143
5.7.3	Classification Results.....	147
5.7.3.1	Individual Binary Classifiers.....	147
5.7.3.2	Binary Fusion	147
5.7.3.3	Multi-Class Sensors	148
5.7.3.4	K-Nearest Neighbour	151
5.8	Discussion.....	152
6	Eddy Current Array	155
6.1	Introduction.....	155
6.2	Background.....	156
6.2.1	Eddy Current Theory.....	156
6.2.2	Probe Design.....	158

6.2.2.1	Coil Configuration.....	158
6.2.2.2	Design of Coil.....	159
6.3	Analysis of Resistive, Inductive and Capacitive Circuit.....	160
6.3.1	Fundamentals of an RLC Inspection Circuit.....	160
6.3.2	Effect of Series Resistance.....	163
6.3.3	Optimum Frequency Selection.....	166
6.3.4	Suitable Defect Metrics.....	167
6.4	Finite Element Modelling.....	171
6.4.1	Introduction.....	171
6.4.2	Basic Model.....	171
6.4.2.1	Overview.....	171
6.4.2.2	Mesh Sensitivity Analysis.....	174
6.4.3	Lift-off.....	176
6.4.4	Sensitivity to Lateral Position and Radius.....	176
6.4.5	Sub-Surface Defect.....	177
6.4.6	Depth Detection Using Frequency Sweep.....	182
6.5	System Overview.....	185
6.5.1	Introduction.....	185
6.5.2	Signal Generation.....	185
6.5.3	Drive Signal Conditioning.....	186
6.5.4	Coil Selection.....	187
6.5.5	Received Signal Conditioning.....	187
6.5.6	Digital Signal Processing (DSP).....	190
6.5.7	Phase Measurement Error.....	193
6.6	Transducers.....	195
6.6.1	Linear Array.....	195
6.6.2	Circular Array.....	195
6.6.3	Pipe Inspection Array (Fusion Probe).....	195
6.7	Results.....	199
6.7.1	Coil Performance.....	199
6.7.2	Eddy Current Array Imaging.....	202
6.7.2.1	Imaging Algorithm.....	202
6.7.2.2	Imaging Simulation.....	205
6.7.3	C-Scan Results.....	208
6.7.3.1	Comparison with Commercial System.....	208
6.7.3.2	Using Multi-Frequency Data to Reduce Noise.....	210
6.8	Discussion.....	212
7	Multi-Sensor Data Fusion	216
7.1	Introduction.....	216
7.2	Development of a Data Fusion Framework.....	217
7.2.1	Scope.....	217
7.2.2	Overview of Procedure for Artificial Corrosion Pits.....	219

7.2.2.1	Geometric Active Contours.....	219
7.2.3	Hough Transform.....	222
7.3	Data Fusion on Real Data.....	225
7.4	Summary	235
8	Conclusions and Future Work	236
8.1	Conclusions	236
8.1.1	Robotic Platform	236
8.1.2	Visual NDE Pattern Classification	238
8.1.3	Eddy Current Array.....	241
8.1.4	Data Fusion and the Integrated Inspection System	243
8.2	Future Work	245
8.2.1	Robot Design	245
8.2.2	NDE Payloads	248
8.3	Technology Transfer.....	249
Appendix A	Eddy Current Array	251
A.1	Analysis of Resistive, Inductive and Capacitive (RLC) Circuit.....	251
A.1.1	RLC Impedance.....	251

List of Figures

1.1	Model of a typical inspection scenario in a nuclear processing cell	4
2.1	Gecko-inspired foot consisting of micro and nano-scale adhering fibers...14	
2.2	Suction cup array.....	15
2.3	Stewart-Gough platform for climbing square cross-section tube [29].....	17
3.1	Pipe support clips.....	34
4.1	Profiles for pipe gripper	49
4.2	Cornering Mechanism	52
4.4	Loads and constraints.....	55
4.5	Von Mises stress analysis of Mark I collar	57
4.6	Displacement analysis of Mark I collar design	57
4.7	Mark I robotic pipe crawler (RPC) on 40mm diameter pipe (Axial view)..	59
4.8	Image of collar 1 on 40mm diameter pipe	59
4.9	Profile of Mark I design on 40mm diameter pipe.....	60
4.10	Navigating past pipe stands	60
4.11	System for increasing gripping force through the use of springs.	62
4.12	Weight distribution of Mark I design.....	64
4.13	Weight breakdown of Mark II robot	64
4.14	Collar diameter minimisation.	67
4.15	Mark II collar 1 minimised weight design.	68
4.16	Mark II collar 2 minimised weight design.	68
4.17	Cornering mechanism showing modified power transmission system.....	69
4.18	Mark II RPC platform featuring optimised collar design.	69
4.19	Von Mises stress analysis of Mark II collar 1.	71
4.20	Displacement analysis of Mark II collar 1.....	72
4.21	Von Mises stress around the circumference of hole.	72
4.22	Main processing stack including payload board.	74
4.23	Block diagram of the Mark II platform	74
4.24	Battery level indication circuit	76
4.25	Coil winding diagram.....	76
4.26	Gripper motor drive circuit.	77
4.27	Stepper motor drive circuit.....	77
4.28	RPC power consumption breakdown in stand-by mode.....	82

4.29	Total power consumption during standard tasks.....	82
4.30	Mark II RPC ascending a vertical pipe.	84
4.31	Mark II RPC travelling around a vertical bend.....	85
4.32	Mark II(r) crawler manufactured using rapid prototyping facilities	86
4.33	Robot control interface with visual inspection support.....	87
4.34	Eddy current array control interface	88
5.1	Principal component analysis	93
5.2	Effect of raw images on dimensionality	97
5.3	2D FFT of illustrative defect images.....	98
5.4	Normalised frequency response of image bandpass filter.	100
5.5	Aluminium pipe (40mm diameter) containing synthetic defects	101
5.6	Classification tree	102
5.7	Illustration of classification metrics.....	103
5.8	Probability distribution functions (illustrative).....	105
5.9	Functional block diagram of the pattern recognition system.....	106
5.10	Mean 2D FFT magnitude of the filtered training images	108
5.11	Pre-processing and PCA stages.....	109
5.12	Bandwidth sensitivity.....	114
5.13	Correct classification rate vs variance	117
5.14	Mean classification rate vs number of training images.....	119
5.15	Generalisation.....	121
5.16	Scatter plot showing the weights of training image images	123
5.17	Binary system: unclassifiable images	130
5.18	Probability density functions using two principal components	135
5.19	Scatter plot showing the weights of all images	139
5.20	K-nearest neighbour results on validation and test sets.....	142
5.21	Corrosion defect images.....	144
5.22	Corrosion defects: training volume sensitivity.....	145
5.23	Corrosion defects: generalisation.....	145
5.24	Scatter plot of corrosion images	147
5.25	KNN applied to corrosion defects for varying value of K.....	151
6.1	Eddy current inspection illustration.....	157
6.2	Theoretical variation of skin depth with frequency	157
6.3	Resistance and reactance of coils measured in-air	161
6.4	Circuit diagram of resistor-inductor-capacitor circuit	162
6.5	Effect of series resistance on magnitude and phase of the coil voltage ...	164
6.6	Dependence of sensitivity and dynamic range on series resistance.....	165
6.7	Effect of change in inductance for RLC circuit	166
6.8	Optimum frequency as a function of L	168
6.9	Phase response of two coils for varying frequency.....	170
6.10	Comparison of PI and single frequency defect metrics	170
6.11	Coil FEM model.....	173

6.12	Lift-off characteristics (COMSOL)	178
6.13	Lift-off (qualitative distribution of eddy currents)	178
6.14	Effect of position of a slot from the centre of the probe.....	179
6.15	Defect position (qualitative distribution of eddy currents).....	179
6.16	Effect of the hole radius	180
6.17	Hole radius (qualitative distribution of eddy currents).....	180
6.18	Effect of the depth of a sub-surface cavity	181
6.19	Sub-surface cavity (qualitative distribution of eddy currents)	181
6.20	Phase vs frequency plots and PI vs hole depth for varying diameter	183
6.21	PI vs hole depth for varying diameter in 304 austenitic stainless steel.....	184
6.22	System overview separated into the functional units.....	186
6.23	Coil drive circuit	188
6.24	Coil selection circuit with the connections for channel 1 illustrated.	189
6.25	Received signal conditioning circuit	189
6.26	Depiction of upsampling process	191
6.27	Signal captured by ADC before and after upsampling.....	191
6.28	Active bandpass filter frequency response	194
6.29	Phase measurement error for five frequency intervals	194
6.30	Linear array.....	196
6.31	Circular array	197
6.32	Pipe inspection array (fusion probe)	198
6.33	Resistance and inductance vs frequency of a coil compared with FEM..	200
6.34	Resistance and inductance vs lift-off of a coil compared with FEM.....	201
6.35	Induced current distribution showing quantisation.....	202
6.36	Schematic of circular probe and imaging algorithm	204
6.37	Imaging model.....	206
6.38	Imaging simulation - coils 0 and 1 positioned symmetrically.....	207
6.39	Imaging simulation - centre of defect 1 is shifted by -2 mm.....	207
6.40	Aluminium test plate.....	208
6.41	C-scan images: comparison with commercial system.....	209
6.42	Best-fit noise reduction	211
6.43	Multiple solutions of imaging algorithm.....	215
6.44	Single solution imaging method.....	215
7.1	Functional block diagram of data fusion and shape extraction	218
7.2	Evolution of the level-set function.....	221
7.3	Illustration of the Hough transform.	223
7.4	Image registration using the fusion probe.....	224
7.5	304L austenitic stainless steel pipe containing two defects.	226
7.6	Edge processing stage.....	226
7.7	The Hough space representation of the edge detected image	227
7.8	Accumulator values that exceed the threshold	227
7.9	Active contour algorithm on defect 1	229

7.10	Active contour algorithm on defect 2.....	229
7.11	Defect mask extracted from camera image.....	230
7.12	2D representation of ECA data from fusion probe	230
7.13	Eddy current data wrapped to pipe circumference (no fusion)	231
7.14	Fusion of camera image and eddy current array data	231
7.15	Effect of registration mismatch on the fusion result.....	232
7.16	304L austenitic stainless steel pipe containing slot	233
7.17	Slot: Defect mask generated by manually specifying an initial contour....	233
7.18	Slot: ECA data (no fusion)	234
7.19	Slot: Fused image and ECA data	234
8.1	Mark IIr RPC inspecting a 304L pipe using the fusion probe.	247
8.2	Integration of the inspection system.....	247
A.1	Circuit diagram of resistor-inductor-capacitor circuit (RLC)	253

List of Tables

2.1	Iron states with associated temperature range and lattice structure.....	11
3.1	Level of autonomy.....	36
4.1	Chassis collar options.....	46
4.2	Interconnection mechanisms.....	47
4.3	PVC-CAW material properties.....	54
4.4	Stepper motor firing sequence.....	75
5.1	Binary Fusion: Confusion matrix using validation set.....	131
5.2	Binary Fusion: Confusion matrix using test set.....	131
5.3	Summary of classification rates for multi-class sensors.....	134
5.4	Confusion matrix for ND sensor (Validation Set).....	137
5.5	Confusion matrix for HC sensor (Validation Set).....	137
5.6	Confusion matrix for CP sensor (Validation Set).....	137
5.7	Confusion matrix for DC sensor (Validation Set).....	137
5.8	Multi-Class Fusion: Confusion matrix for max rule on validation set.....	138
5.9	Multi-Class Fusion: Confusion matrix for max rule on test set.....	138
5.10	Multi-Class Fusion: Confusion matrix for product rule on validation set.....	138
5.11	Multi-Class Fusion: Confusion matrix for product rule on test set.....	138
5.12	Summary of classification rates for multi-class sensors on real images.....	148
5.13	Binary Fusion: Confusion matrix using validation set.....	149
5.14	Binary Fusion: Confusion matrix using test set.....	149
5.15	Confusion matrix for ND sensor (Validation Set).....	149
5.16	Confusion matrix for SL sensor (Validation Set).....	149
5.17	Confusion matrix for CP sensor (Validation Set).....	149
5.19	Multi-class Fusion: Confusion matrix for max rule combination on test set.....	150
5.20	Multi-class Fusion: Confusion matrix for product rule on validation set.....	150
5.21	Multi-class Fusion: Confusion matrix for product rule on test set.....	150
5.18	Multi-class Fusion: Confusion matrix for max rule on validation set.....	150
5.22	Synthetic Data: Summary of classification rates for multi-class sensors.....	154
5.23	Real Data: Summary of classification rates for multi-class sensors.....	154
6.1	Optimum frequency sweep range.....	169
6.2	Properties of model sub-domains.....	174
6.3	Mean and standard deviation of error due to max. element size.....	175
6.4	Mean and standard deviation of error due to growth rate.....	175
6.5	Group delay for each multi-rate filter.....	193

List of Abbreviations

AC	Alternating Current
ADC	Analogue to Digital Converter
ANN	Artificial Neural Network
BCC	Body Centred Cubic
BWR	Boiling Water Reactor
CAD	Computer Aided Diagnosis
CP	Corrosion Pit
CUE	Centre for Ultrasonic Engineering
DC	Diagonal Crack (also Direct Current)
DFT	Discrete Fourier Transform
DMEM	Design, Manufacture and Engineering Management (dept. of)
DOF	Degrees of Freedom
DSP	Digital Signal Processing
DSP/FG	Digital Signal Processing/Function Generator
ECT	Eddy Current Testing
ECA	Eddy Current Array
FCC	Face Centred Cubic
FDA	Food and Drug Administration (USA)
FE	Finite Element
FEA	Finite Element Analysis
FEM	Finite Element Modelling
FFT	Fast Fourier Transform
FG	Function Generator
FIR	Finite Impulse Response
FN	False Negative
FP	False Positive
GAC	Geometric Active Contours

GPIO	General Purpose Input-Output
GUI	Graphical User Interface
HAZ	Heat Affected Zone
HC	Horizontal Crack
I2C	Inter-Integrated Circuit
ID	Inside Diameter
IP	Internet Protocol
KNN	K-Nearest Neighbour
LAN	Local Area Network
LED	Light Emitting Diode
LOS	Line-of-Sight
MDB	Motor Driver Board
MDOF	Multi-Degree of Freedom
MOSFET	Metal Oxide Semiconductor Field-Effect Transistor
MSE	Mean Squared Error
MUX	Multiplexer
ND	No Defect
NDE	Non-Destructive Evaluation
NDT	Non-Destructive Testing
OD	Outside Diameter
OVA	One vs All
PC	Principal Component
PCA	Principal Component Analysis
PDE	Partial Differential Equation
PDF	Probability Density Function
PEC	Pulsed Eddy Current
PHWR	Pressurised Heavy Water Reactor
PI	Phase Integral
POD	Probability of Detection
PSB	Primary Signal Bus
PVC	Polyvinyl Chloride
PWM	Pulse Width Modulation
PWR	Pressurised Water Reactor
RAM	Random Access Memory
RF	Radio Frequency
RPC	Robotic Pipe Crawler
SCC	Stress Corrosion Cracking
SCP	Secure Copy

SFEC	Swept Frequency Eddy Current
S-G	Stewart Gough (platform)
SIB	Signal Interface Board
SL	Slot
SNR	Signal to Noise Ratio
SSB	Secondary Signal Bus
SSDA	Sequential Similarity Detection Algorithm
SSH	Secure Shell
SSID	Service Set Identification
SVM	Support Vector Machine
THORP	Thermal Oxide Reprocessing Plant
TN	True Negative
TP	True Positive
USB	Universal Serial Bus
UT	Ultrasonic Testing

Chapter 1

Introduction

1.1 Background

1.1.1 Non-Destructive Evaluation and the Power Industry

The majority of nuclear infrastructure in the UK was constructed during the post-war era as the race to harvest the commercial power of the atom commenced. The working life of these assets was expected to be no greater than around 25 years and, such was the confidence in the quality of design, little thought was given to the need for inspection or decommissioning. Due to a number of high profile incidents, nuclear power lost favour with both the public and the policy-makers and, as a result, investment in new infrastructure was reduced and the working life of existing plant extended well beyond its initial scope. Calder Hall power station is an example of such plant. Originally commissioned for an operational life of 20 years, starting in 1956, generation at this facility was repeatedly extended until 2003! Consequently, the role of non-destructive evaluation (NDE) has become increasingly important. A 2009 report [1] showed that, despite the global financial crisis reducing the rate of market growth, the overall growth rate in the NDE equipment market was positive. This growth is, in part, due to the ever-present need to ensure the safety of personnel and the maximisation of plant operation. In 2008, the world market was estimated to be in the region of \$1.1billion [1] with forecasts predicting this to increase to \$1.3billion by 2013. The commitment of successive UK governments to the principle of new nuclear facilities, in the face of much negative publicity, follows the inevitable realisation that, with the steady

demise of fossil fuels, current renewable technology is unable to realistically meet the country's soaring energy demands. In order to ensure price stability and security of energy supply, coupled with increasing demands for CO₂ reduction, nuclear power is entering a global 'renaissance' period [2], expanding the potential for non-destructive evaluation. This trend towards lifetime extension and increasing usage of NDE is evident in all power generation industries. Precise figures are difficult to obtain however, it is estimated that the annual inspection of pipework in the oil and gas industry costs many tens of millions and, given that a typical gas refinery may contain thousands of kilometers of pipework, this presents considerable logistic challenges where access is limited to within a scheduled time window.

The ubiquity of certain defects varies across different industries. For much of the carbon steel pipework, corrosion is the major problem, however, corrosion pits in isolation are not necessarily of tremendous concern. In fact, reductions of 50% of the wall thickness can be tolerated on occasions before a section must be replaced. However, multiple instances of corrosion pitting within close proximity is of considerable concern, due to the prime conditions this creates for inter-pit crack generation. In the nuclear industry, due to the predominance of austenitic stainless steel usage, corrosion pitting is not the most significant problem. Fatigue cracking, however, is a particular concern and was instrumental in the 2005 THORP incident, where a failure in a nozzle resulted in approximately 83 m³ of highly radioactive uranium and plutonium dissolver product liquor leaking into the processing cell [3]. The total cost of this incident is estimated to be in the region of £300m [4] and resulted in the operator of the plant being fined £500,000 [3]. Furthermore, the site was forced to shut down entirely for a period of 22 months. Most plant is typically subjected to prolonged cyclic loading of some form or other and, for the majority of ageing plant in the nuclear, oil and gas industries, fatigue cracks are inevitable. The existence of such cracks in itself is not an immediate cause for concern, in many instances, however it is important to be able to assess the extent of the crack such that an estimation of the remaining useful life can be obtained. Such defects in ageing plant are likely to be in the tertiary (or most rapid) stage of creep, which ultimately precedes failure. Furthermore, as plant degrades there will be increasing numbers of defects, thus compounding the need for quick and efficient inspection. Given the relatively modern nature of materials such as austenitic stainless steel, the long-term performance is unknown and due to the volume of such material used in critical plant, the need for NDE is likely to be ever-growing.

There are a wealth of NDE inspection systems and techniques commercially available, however the equipment market is dominated by the four main modalities: eddy current, visual, ultrasonic and radiographic. In 2005, the eddy current and visual inspection test equipment markets accounted for 20.6% and 24.3% of the total NDE equipment market respectively with ultrasonic claiming 26.6% and radiographic 25.9% [5].

1.1.2 The Move Towards Remote Inspection

The motivation for increased research into remote inspection is, broadly speaking, to address three concerns: NDE in areas of restricted access, inspection efficiency (i.e. volume of inspection in a given time frame) and accuracy (resolution of defect characterisation/location). It can be shown that the latter two are linked, with more accurate inspection allowing more efficient deployment of NDE resources.

Nuclear processing cells process spent fuel and separate the re-usable by-products of fission, such as uranium and plutonium, from the radioactive waste which must be treated and stored. The high levels of radiation within such cells means that inspection is at best difficult and, in some cases, impossible by manual means. Furthermore, some areas may be physically too small to allow access and often, particularly in the nuclear industry, bespoke solutions are required. Such innovative and unique systems inevitably come at a premium and therefore a strong case must be made for this investment. A multi-use robotic system for pipe inspection offers a cost-effective means of protecting personnel safety while enabling inspection of all necessary plant. It is therefore a very attractive solution for the industrial end-user.

Inspection and maintenance is required by law, however the competitive nature of the markets means that there is a drive to be more efficient and minimise downtime and hence costs [6]. As intervals between inspection can be in the region of several years and there is a limit to the amount of inspection that can be achieved within a given time window, priority must be assigned to those areas of greatest concern. The preparation required to facilitate a given NDE inspection (scaffolding, insulation removal, grit blasting etc) is often as much as four or five times more expensive than the cost of the actual inspection itself. Furthermore, in spite of increasing expenditure on NDE training, the fallibility of human nature means that errors will inevitably occur. As many oil and gas pipelines operate at temperatures of several hundred degrees, marking of pipes is not possible and thus, in densely populated regions, mistakes may be as simple (and costly) as inspection

of the wrong pipe. Originally, NDE was a largely qualitative process aimed at simply identifying if a defect was present or not. As technology progressed, and driven by economic factors, there has been a move to a more quantitative approach where greater emphasis is placed on determining the residual life of an asset [6]. In doing so, it may be feasible to keep a part in operation until an outage is scheduled rather than unnecessarily interrupting operation. By gaining a more precise quantitative description of a particular structure, unnecessary interruptions can be minimised. Intelligent robotic systems offer the potential to make the most efficient use of plant down-time by maximising the volume of inspections that can be achieved while minimising the extraneous costs. Such a system would offer a significant reduction in overall maintenance costs and, consequently, increase profitability while offering the capability to expand the scope of NDE and provide quantitative data to a predictive maintenance regime.

1.1.3 Typical Inspection Scenario

At this point it is pertinent to provide an example of a specific inspection problem in order to ‘set the scene’ for this research. In the scenario shown in Figure 1.1, a weld requires inspection at the point where the red pipe connects to a vessel. Processing cells typically have small inspection ports through which access to the




Figure 1.1: Model of a typical inspection scenario in a nuclear processing cell. A weld on a pipe (shown in red) connecting into a vessel requires inspection for fatigue, however this is a considerable distance from the nearest inspection port.

pipe may be gained. The size of the cell can be quite large and, as there are only a limited number of such access points, it is quite likely that, as in this instance, the artefact of interest is some distance from the nearest inspection port. The level of radioactivity within the cell is high, making manual inspection dangerous and costly. The use of a multi-degree-of-freedom (MDOF) snake-arm robot to deploy a payload to the site has the problem that the arm must be long enough to span the distance required. Furthermore such systems are typically expensive and can be difficult to control. A more convenient solution would be to clamp a mobile robotic vehicle onto the pipe leading to the vessel where it passes an inspection port and then drive the system to the weldment of interest, navigating past obstacles en-route. Deployed on this 'pipe crawler' would be a range of payloads suited to the particular inspection. The results of the inspection would then either be transmitted wirelessly back to a receiver located at the inspection port or stored on-board for later post-processing. A platform such as this, rather ironically, has a significant advantage in terms of flexibility over a MDOF system since its range is not limited and, consequently, it is a more cost-effective solution.

1.2 Aims of this Thesis

One of the major difficulties in the field of nuclear facility inspection is the non-trivial matter of getting the NDE equipment to the area of interest. In many instances, the same is true of the oil, gas and petrochemical industries, where either there is significant expense in terms of creating a supporting infrastructure suitable to provide a safe platform for personnel to perform a manual inspection, or space is so restrictive that actually gaining access itself is highly problematic. The fundamental aim of this thesis is to create a robotic pipe inspection system which is capable of addressing this issue. A particular emphasis will be placed on the application of nuclear processing cell inspection, while offering broader applicability to the other relevant industries. The principal application is somewhat of an 'acid test' in that the restrictions and stipulations imposed by the operating environment are much stricter than in the other areas mentioned. Thus, a system that is designed to operate in such harsh environments could also, almost certainly, operate in the majority of cases in these other industries. Thus, by prioritising the application in this way, a platform can be developed that offers a high level of application flexibility and thus broader industrial appeal. Critical to the efficacy of such a system, is a set of NDE inspection tools that are integrated with the main platform. Off-the-shelf inspection equipment could quite easily be adapted so that they could be mounted on a robotic platform, however only their complete integration within the system offers the high levels of control that a truly flexible

inspection solution demands. Furthermore, bespoke payloads for specific inspection problems can be developed that offer advantages in terms of sensitivity and probability of detection over an adapted, commercially available solution.

Given the preceding argument, the specific aims of the research can be summarised as follows:

- i. Create and evaluate a robotic platform for the inspection of non-ferromagnetic pipes, particularly within nuclear processing cells.
- ii. Design and evaluate a visual inspection system incorporating both image capture and pattern recognition for defect classification.
- iii. Create and evaluate a sensor for integration with the main pipe inspection platform.
- iv. Fuse image and eddy current sensor data sets to enhance clarity and probability of detection (POD).

1.3 Contributions to the Field

Existing pipe inspection systems are typically either very bulky, external devices, requiring an umbilical to deliver power and communications, or are internal pipe inspecting ‘pigs’¹. Due to the dense nature of the internal infrastructure and the aggressive nature of the pipe contents, such systems are unsuitable for the inspection of pipework within nuclear processing cells. CUE has a long-term interest in remote inspection however, prior to this research, the focus has been on wheeled robots using magnetic traction [7]. The unsuitability of this platform for non-magnetic material and also for small diameter pipework required that a completely new and novel platform be developed for this task. The development of such a system for external pipe inspection, capable of operating wirelessly and of circum-navigating typical obstacles such as pipe hangers, flanges and pipe bends demonstrates a considerable contribution to the field and to the robotic inspection portfolio of the group. The creation of a stable and robust inspection platform necessarily precedes the development of payloads since without such stability and kinematic ability, the inspection sensors cannot be delivered to the location accurately. Surmounting the key issue surrounding traction on such materials, while providing these capabilities, is a significant achievement.

1. A ‘pig’ refers to a device that travels through a pipeline and does not require the flow of the pipeline product to be stopped.

The use of pattern recognition techniques to extract features in data is one where the results have often been overstated, particularly in the field of neural networks as will be outlined in Chapter 5. However the implementation of an automated visual recognition system for classifying defects, based on the *statistical* information within images captured using the on-board visual inspection system, is crucially different from these approaches and is novel. The application of the eigenface facial recognition technique to NDE images has not been attempted previously and therefore is a contribution of this thesis. A good quantitative evaluation of the performance of the system on both synthetic and real defect images has been provided.

The integration of a low power, swept-frequency, differential eddy current array with a robotic pipe inspection platform is a contribution and forms an essential non-contact payload. The establishment of an equation to calculate the optimum operating frequency for a differential phase eddy current system and a metric that makes use of the swept excitation method which is less sensitive to unwanted factors than the conventional single-frequency metric are also both contributions. A theoretical relationship between the defect geometry and the new defect metric is provided and has been shown to relate well to defect size allowing C-scan images with the relative depth of defects to be obtained.

Finally, a framework for combining data sets from the visual and eddy current array payloads has been presented with some rudimentary examples to satisfactorily prove the overall concept of an intelligent pipe inspection system. A significant body of further work could be focussed on this area and this is acknowledged.

1.4 Thesis Content

Chapter 2 provides an overview of the nature of austenitic materials and the influence this has on the suitability of certain inspection modalities. A review of state-of-the-art pipe crawling robots follows a focussed examination of suitable traction options for use on non-ferromagnetic materials. A literature review of eddy current techniques and visual pattern recognition systems is also provided.

Chapter 3 reflects upon the literature outlined in Chapter 2 and defines a specification for the intelligent inspection system.

Chapter 4 contains an in-depth study of the design of the robotic platform, indicating the reasoning behind the choice of traction system and chassis design

from the many alternatives. A finite element model (FEM) examines the structural integrity of the proposed design and an exhaustive analysis of the kinematic performance is provided which feeds a second, optimised design iteration. Finally, there is an overview of the graphical user interface (GUI) for control of the robotic platform and integrated control of the on-board payloads.

In Chapter 5, the automated visual recognition system is discussed. An overview of pattern recognition and types of classifier is provided together with a detailed examination of the use of principal component analysis (PCA) for image processing. Following a system overview and establishment of a Bayesian framework for the classifier, a detailed examination of the influence of image bandwidth, volume of training data and generalisation is given. An overview of classifier combination theory outlines the mathematical framework for combining the outputs from multiple binary and multi-class classifiers followed by an examination of the relative performance, benchmarked against a K-nearest neighbour (KNN) methodology. Finally, the system is tested using genuine corrosion defects and the performance of the various classifier ensembles compared.

Chapter 6 discusses the fundamentals of eddy current inspection and the design of the sensor. A detailed derivation of the optimum operating frequency and defect metric is provided and finite element modelling (FEM) is used to relate this to the underlying defect geometry. The effect of lift-off, defect position and radius as well as the performance on sub-surface defects are considered in terms of the coil impedance. A detailed description of the functional design and operation of the eddy current array system is provided along with an overview of the associated digital signal processing (DSP) that is required. An imaging algorithm is presented which combines the data from multiple probe positions and this is compared with a benchmark, commercially available device. The results of the imaging process are validated against results from an algorithm simulator.

Chapter 7 combines the three key threads of the research, namely the robotic platform, the visual inspection system and the eddy current array, through a process of data fusion. The data sets from the two sensor payloads, resulting from a 304L austenitic stainless steel pipe inspection, are combined into a single image and in so doing, it is demonstrated that the clarity of the defect image, and hence probability of detection (POD), is significantly improved. The effect of a registration mismatch is identified along with the effect of inadequate shape extraction performance.

In Chapter 8, a review of the complete developed system is provided including a discussion of the performance and suitability to the intended inspection tasks. Additionally, a programme for future work is outlined.

1.5 Publications to Date Arising from this Thesis

- L. Mackenzie, S.G.Pierce, G. Hayward, “Robotic Pipe Crawling Platform for NDE”, *Proc. of 46th Annual British Conference on NDT*, Glasgow, 18-20th September 2007
- L. Mackenzie, S.G.Pierce, G. Hayward, “Robotic Inspection System for Non-destructive Evaluation (NDE) of Pipes”, *AIP Conference Proceedings*, 2009, 1096, 1687
- S.G. Pierce, G. Dobie, R. Summan, L. Mackenzie, J. Hensman, K Worden, G. Hayward, “Positioning Challenges in Reconfigurable Semi-Autonomous Robotic NDE Inspection”, *Proceedings of SPIE*, 2010, 7650, 76501C
- L. Mackenzie, S.G.Pierce, G. Hayward, “NDE Image Pattern Classifier for Intelligent Pipe Inspection”, *Pattern Analysis and Applications*, 2011 (In progress)
- L. Mackenzie, S.G.Pierce, G. Hayward, “Robotic Pipe Inspection with NDE Image and Eddy Current Data Fusion”, *Insight*, 2011 (In progress)

Chapter 2

Review of Material, Systems and Techniques

2.1 Introduction

The purpose of this Chapter is to examine areas relevant to the project and present a summary of the important research that has been conducted with a view to establishing a formal specification in Chapter 3. Further details will be provided, on each topic, in the relevant chapters however here the aim is to provide a broad overview and the reader will be directed to the relevant section for a more detailed explanation where required. The nature of austenitic stainless steel will first be outlined, leading into a discussion of suitable traction methods for a robot operating on this material and a summary of various pipe robots from research papers. The influence of the pipe material on the suitability of different payloads will be discussed and it will be shown that a visual payload coupled with a pattern recognition algorithm in addition to an eddy current sensor offers a useful combination. A review of eddy current techniques and commercial equipment is provided followed by a discussion of how pattern recognition techniques have been applied to images in various fields including NDE.

2.2 Overview of Austenitic Stainless Steel

Austenitic stainless steel is widely used in industry for its relatively high tensile strength, malleability and excellent corrosion resistance. Typical applications within

the nuclear power industry include piping in pressurised water reactors (PWR) and boiling water reactors (BWR). The following section gives an insight into the factors that provide its unique properties.

Pure iron has three, temperature-dependent, solid states each having one of two crystalline configurations. Body centred cubic (BCC) refers to a lattice structure containing one lattice point at each corner of the cubic structure and one in its centre. A face centred cubic (FCC) structure similarly has lattice points at each corner, however additional points are found on each face of the cube rather than at its centre. This leads to a greater number of shared atoms in adjacent lattices, meaning the atoms can pack more closely together resulting in greater ductility. Unlike the more brittle BCC structures, such as carbon steel, unexpected changes in operating conditions (e.g. temperature, loading) are much less likely to compromise the performance of austenitic stainless steel. The three solid states of iron along with their associated lattice structure and stable temperature range are given in Table 2.1.

Table 2.1: Iron states with associated temperature range and lattice structure

State	Stable Temperature Range ($^{\circ}C$)	Crystalline Structure
Alpha Ferrite	<910	BCC
Gamma Ferrite (Austenite)	910 - 1400	FCC
Delta Ferrite	1400 - 1535	BCC

When cooling from its liquid phase, an iron alloy initially solidifies as delta ferrite, transforming to gamma ferrite (austenite) and then alpha ferrite with further cooling. The addition of nickel (Ni) increases the temperature range over which the FCC structure, and hence the austenite state, prevails. However, ferrite crystals can remain in the casting which will lead to a reduction in corrosion resistance. The corrosion resistance property of the material is achieved through a process called passivation. To prevent the formation of iron oxide (i.e. corrosion) on the surface of the steel upon exposure to air, chromium (Cr) can be added to the iron which preferentially forms a protective film of chromium oxide on the surface [8]. This oxide layer reforms if the surface is damaged, leading to such steels being termed 'stainless'. The 304 grade austenitic stainless steel, used a great deal in the nuclear industry, has a 18% Cr and 8% Ni content.

The rate at which the alloy is cooled can have a significant effect on the properties of the resulting material. During the cooling process, interstitial accumulations of impurities, such as sulphur and phosphorous, can occur, leading to an increased risk of cracking - the situation is exacerbated if the material is cooled too rapidly. If the alloy is cooled too slowly however, corrosion of the grains at the grain boundary can occur due to a chromium depletion in this area. At a temperature of around 650°C the relatively small carbon atoms diffuse easily through the lattice structure to the grain boundary. However, the larger chromium atoms cannot diffuse, leading to the formation of chromium carbides [8, 9]. The lack of available chromium atoms now to form protective oxides results in a reduced corrosion resistance. This process, known as sensitisation, can also occur during welding when there is a sudden increase in temperature. Low carbon grades of steel are available (e.g. AISI 304L, 316L) which reduce the available carbon for carbide formation and hence provide better resistance to corrosion.

Weldments in austenitic steel present a further microstructural complexity. In both ferritic and austenitic metals the initial stages of grain formation are the same in that, during the solidification process following welding, a columnar grain structure is produced. Grain growth is along the direction of the maximum thermal gradients and is faster along a $\langle 100 \rangle$ crystallographic axis leading to the creation of columnar grains [8, 9, 10, 11]. For a V-preparation weld, in the proximity of the fusion faces, the direction of heat flow will be normal to the fusion face, towards the cooler parent material, where-as in the centre of the weld this will be normal to the weld surface [10, 12]. The difference arises after the initial pass of the weld torch. In the case of a ferritic weld, reheating and then subsequent cooling of the weld bead causes an austenite to ferrite phase transformation resulting in the disappearance of the columnar grain structure. In the case of an austenitic weld there is no such transition and, provided that the penetration depth is not too great, epitaxial growth on existing grains may occur leading to a long columnar grain structure. In fact, the reheating process allows the crystallites to cross the grain boundaries, to an extent, resulting in epitaxial grain growth (i.e. having the same crystalline orientation) leading to the characteristically large grain structure [9, 10, 11, 12]. The net result of this is a highly anisotropic structure.

In addition to fatigue cracks, stress corrosion cracking (SCC) can occur within the heat affected zone (HAZ) of a weld, adjacent to the weld bead. Consisting of intergranular cracks, SCC is caused by the metal experiencing a certain level of tensile stress coupled with either the presence of chlorides in the process fluid or sensitised metal [8, 9]. SCC produces a significant reduction in mechanical strength

with little metal loss, complicating inspection and increasing the risk of rapid fracture.

2.3 Pipe Inspection Robots

The following sections introduce the role of robots in pipe inspection, starting with an overview of suitable traction schemes and finishing with a summary of existing robotic pipe inspection technology.

2.3.1 Traction Methods

The most common methods of traction for robotic systems are: magnetic, chemical, friction, biomimetic, vacuum and electroadhesion. Although magnetic adhesion is not suitable for austenitic steel, due to the non-ferritic nature of the material, it is included here for completeness.

Magnetic traction makes use of the attraction force between the persistent magnetic field generated by a permanent magnet and any ferromagnetic surface. In general, the size of the magnet dictates the magnitude of this force however, by definition, this is unsuitable for any material that is not ferromagnetic thus limiting its applicability.

Chemical adhesion, typically involves the use of an adhesive tape of some form, which sticks the robot to the surface. Since such systems offer no means of controlling the chemical bond, they can be described as ‘always on’. Consequently, there is a trade-off between the magnitude of the adhesion force and the size of the actuators required to break the bond and allow locomotion. The nature of the material is such that it attracts dust deposits which ultimately impairs the system’s ability to adhere in subsequent instances.

A friction based traction method is a particularly simple technique that utilises the inherent friction force that exists between any two surfaces in contact. All materials have a particular friction coefficient which dictates how large the resulting traction force will be. The coefficient of kinetic friction is the ratio of the lateral friction force to the normal component of the applied force when two objects are moving relative to each other. The coefficient of static friction is the dominating factor in determining the initial force required in order to get an object moving from stationary and is usually higher than the coefficient of kinetic friction. Neither coefficient is necessarily constant for a given material and are both affected by

factors such as the material surface roughness, temperature and humidity of the operating environment.

In recent years, significant work has been conducted in the field of biomimetic (or biologically inspired) traction systems, with particular interest in recreating the techniques employed by geckos; a member of the lizard family. A gecko can climb on rough or smooth surfaces at any orientation by utilising the molecular forces (primarily Van der Waal's force) between the nano-scale beta-keratin structures present on its foot and the contact surface [13,14] as shown in Figure 2.1.





Figure 2.1: Gecko-inspired foot consisting of micro and nano-scale adhering fibers [14]

Current experimental traction systems based on this principle appear to be approaching the performance achieved in nature by geckos [13] and studies have found the adhesion force to be as high as 10 N cm^{-2} [14]. Furthermore, since the adhesion is based on molecular forces, it can be maintained indefinitely with no power consumption. This is of considerable interest in robotic applications, particularly where there is a need to use battery power. Gecko inspired adhesion systems offer the possibility of very fast locomotion due to the ease with which attachment and detachment can be achieved. Furthermore, the hydrophobic nature of the synthetic fibres under development means that they are self-cleaning [14]; this is attractive to the inspection of many industrial facilities where a high level of particulate surface contamination is common. Similar technology using micro-spines to exploit the small asperities on surfaces such as concrete, sandstone and brick (rather than the molecular forces) has been developed [15]. This technology was shown to climb reliably on surfaces with asperities of the order of 25 microns, however was not suitable for surfaces with asperities smaller than this. While showing considerable promise, both technologies are relatively immature and not yet available as an off-the-shelf solution.

Vacuum adhesion works by creating a controlled region of low (or vacuum) pressure within a region between the robot and the surface it wishes to adhere to. The pressure differential between this and the surrounding atmosphere generates

a force that pushes the robot towards the surface. Providing this force is great enough to overcome the force due to gravity, traction will be achieved.

A number of wall climbing robots designed for non-ferritic surface applications (such as window cleaning) use a single suction cup design which results in a single pressure source. An issue with such a system lies in the intrinsic relationship between the traction force and the application of power. In environments where even momentary loss of traction cannot be tolerated or where the use of a safety line (for retrieval in this event) is impractical, single-cup, vacuum adhesion does not offer an attractive solution. Furthermore,



single-cup systems are generally sensitive to cracks in the supporting surface and can have difficulty in overcoming significant obstacles due to the loss of vacuum pressure that arises in this process [16, 17]. This sensitivity could potentially be overcome through the use of a suction cup array, shown in Figure 2.2. The concept of this system is that a conformable pad consisting of a number of miniature suction cups would generate a distributed vacuum force that would collectively support the load of the robot. A certain level of redundancy must be built into the system so that, in the event that a cracked surface is encountered, provided that enough suction cups make a sound coupling to the surface, the weight of the robot can be supported and traction therefore maintained. The conformable nature of the system would allow wrapping to the contours of a pipe and possibly more complex structures. In general, the performance of all vacuum adhesion technologies suffer from exposure to dust and damage to the suction cup material and is therefore only really suitable for clean and smooth surfaces.

Figure 2.2: Suction cup array consisting of multiple pneumatically isolated suction cups providing a distributed adhesion force.

Electroadhesion is a relatively new technology that works by inducing electrostatic charges on a surface through a conformable series of electrodes. This battery powered system is capable of traction on a range of both conductive and non-conductive surfaces [18]. Unlike passive systems, such as biomimetics, the adhesion force can be switched off, thereby minimising the power consumption required in climbing applications. The clamping force has been found to be in the range of 0.5

to 1.5 N cm^{-2} [19]. As with the gecko-inspired systems, while offering an attractive solution (particular in terms of controllability), there is currently no commercially available product on offer.

2.3.2 Pipe Crawling Robots

Pipe crawling robots are a particular subset of climbing robots that focus specifically on applications involving tubular structures. Although there are a diverse range of systems available, they can generally be grouped for either internal or external operation. Internal pipe inspection robots are commonly used in inspection applications in the oil and gas industry [20, 21]. The design of such systems varies significantly but for small bore operation the design is generally notable for its narrow, elongated shape. The propulsion mechanism varies from systems that use the kinetic energy of the fluid in the pipe [22] to self-propelled systems using robotic legs [23], wheeled systems [24, 25] and even snaking drive mechanisms [26]. In-line pipe robots require either access to the pipe ends or the presence of ‘pig traps’ to insert and remove the system. Furthermore, they cannot be used in situations where the native fluid is highly corrosive such as in nuclear processing cells.

External pipe inspection systems are less common but their designs are no less diverse. Their relatively limited use is most probably due to the large volume of pipework that contains outer insulation and the need to preserve this if at all possible while conducting the inspection. For uninsulated pipes, where insulation removal can be tolerated or where the pipe contents precludes in-line inspection, external pipe robots offer a viable solution. One such system, designed to deliver NDT sensors on 6 inch bore pipes having approximately 50 mm of insulation and cladding, consists of a single metal collar containing omni-directional wheels to provide both axial and circumferential locomotion [27]. A gap in the collar exists so that the robot can be placed on the pipe as well as to allow obstacles, such as pipe hangers and branches, to be overcome. Although the system can negotiate bends in the region of 1 to 3.5 times the standard diameter, the single collar design means that it is not suitable for pipe work that contains flanges. The overall weight of the design is 8 kg and is capable of carrying a 5 kg payload. However, it requires an umbilical to deliver power and control. Another design, intended for outlet feeder pipe inspection in pressurized heavy water reactors (PHWR), consists of two collars powered by pneumatic actuators [28]. As with the omni-wheel robot, a gap exists in the collar for inserting the device onto the pipe and an umbilical is required for both power and control. Locomotion is achieved through an inch-worm

mechanism which allows both axial and rotational movement. Miniature robotic vehicles using magnetic wheels have been used elsewhere [7] for NDE on large diameter carbon steel pipes as well as flat, steel plates. However, they are not suitable for either small diameter or non-ferromagnetic pipework.

In recent years, parallel robots based on the Stewart-Gough (S-G) platform, such as that shown in Figure 2.3, have been developed for the external climbing of tubular structures [29, 30, 31, 32, 33, 34]. Parallel robots are characterised by an end-effector connected to its base by two or more independent kinematic chains. Such systems have greater rigidity than their serial counterparts since the load is distributed rather than cumulatively added between each link, thereby offering greater load carrying capabilities. In terms of pipe inspection




Figure 2.3: Stewart-Gough platform for climbing square cross-section tube [29]

robots, this translates to a system whose forward and rear ends are connected by one or more independent actuators. The S-G platform is one such example, consisting of two rigid bodies connected by six identical linear actuators providing a highly rigid design with a high payload/weight ratio [34]. A parallel S-G robot design was presented in [32] for climbing palm trees and consists of two hinged hexagonal rings through which the tree trunk passes with four clamping devices on each allowing the robot to grip to its surface. The system has a movable arm fitted with an end-effector and the main chassis contains ultrasonic sensors to detect the centrality of the robot to the trunk. Although this design was found to be suitable for its intended application, it is limited to use on structures where the radius of curvature is low and where no support structures such as pipe hangers are used. A similar design was proposed for climbing metallic structures that differs only in that it incorporates arms extending from each ring to facilitate gripping. In this system, therefore, the supporting structure does not need to pass through the centre of the robot thereby allowing full use of the flexibility of the S-G platform to be made [30]. While mechanically less demanding compared with serial robots, in general parallel S-G robotic platforms suffer from the complexity of their direct kinematics with the existence of multiple solutions for a given set of actuator extensions [32, 33, 34]. Consequently, the control system is rather complex. In addition, the

application of the S-G platform to a parallel robot requires the design of special universal joints to accommodate the large rotations that are possible. A four degree-of freedom (DOF) hybrid serial-parallel design has been developed that uses three electrical cylinders and three DC motors to climb a pole for an arbitrary light bulb changing operation [31]. Using the two translational and two rotational DOF, the system was shown to be capable of negotiating bends in pipes and has a high load capacity of 800 N (or one very large light bulb). Although relatively heavy at 16 kg, the use of electric motors rather than pneumatic cylinders ensures that an umbilical is not needed thereby minimising the impact of adjacent infrastructure on its operational ability. Furthermore, the design uses an inter-meshing gripper to clamp on to the pole, thus allowing inline variations in diameter to be tolerated.

2.4 Influence of Material on Test Method

The anisotropic nature of austenitic stainless steel material has an inevitable influence on the most appropriate choice of inspection technique. Clearly, visual inspection takes no account of such factors, performing equally well on carbon and austenitic steels and is a valuable addition to any inspection modality. Magnetic flux leakage is unsuitable due to its reliance on the material being ferromagnetic. Radiographic inspection is technically suitable, however the hazardous nature of the inspection materials, and subsequent safety implications, means that it is not ideally suited to a cross-application, pipe inspection solution. The effect of the material properties will, therefore, be restricted to the other principal modalities: ultrasonic testing and eddy current testing.

2.4.1 Ultrasonic Testing

A general issue with the use of contact ultrasonic testing (UT) in active plant, regardless of the material, is in the need for a couplant to transmit the transducer energy to the metal. Aside from the obvious issue with carrying and dispensing the couplant, further problems arise where there is dirt, oil or oxide layers on the surface of the material. Additionally, high temperature surfaces can reduce the effectiveness of the couplant or, in extreme cases, cause it to evaporate entirely.

Two primary forms of attenuation arise in UT, of anisotropic materials, due to the material structure: beam skewing and scattering. The orientation of the austenitic steel fibres has a strong influence on the propagation of ultrasound, resulting in beam skewing occurring and leading to potentially misleading results in terms of the position and even existence of defects. It has been found that an ultrasonic

beam can be bifurcated (split in two) resulting in double images of a defect being detected [10]. Further skewing can occur through refraction at the weld fusion faces resulting, in some cases, to mode conversion of the incident wave. Since the extent to which the beam is skewed is dependent on the incident angle, the width of the beam will also change. In tests conducted on cylindrical specimens of 316 austenitic weld metal it was reported that the attenuation varied from a maximum of 0.4 dB mm^{-1} , with the direction of propagation at 0° or 90° to the grain axis, to a minimum value of 0.05 dB mm^{-1} at approximately 45° [10]. Consequently, for weld inspection in austenitic stainless steel, the optimum incident angle is at 45° to the grain axis [10, 11, 35, 36]. The level of scattering is dependent on the insonified volume of the metal which, in turn, is related to the level of beam spreading [10]. Beam spreading will tend to increase the insonified volume and hence increase the number of scatterers contributing to the overall coherent noise floor.

The signal-to-noise ratio (SNR) has an inverse-root relationship with the pulse length and so shorter pulses will result in lower noise levels [10]. However, as the energy input by a shorter pulse is also lower, the signal-to-electrical noise ratio may be low [9]. For ferritic welds, shear wave probes are commonly used due to the 'oriented' nature of the defects acting as good reflectors for such incident waves. In general, compression wave probes are preferable for austenitic welds, due to the reduced beam skewing and increased SNR that this mode offers in such media. The latter arises from the fact that the majority of scattered ultrasound is in the shear mode regardless of the incident wave and thus, compression probes can be designed such that they are unaffected by this mode [9, 35].

In general, where-as with ferritic welds, it can be assumed that, for the most part, significant and even small defects can be detected with a high level of success, the same cannot be said for austenitic welds where low sensitivity means that fracturable defects may not be detectable [10]. The problems associated with UT in austenitic stainless steel were considered too great for implementation at this time within a mobile, miniature robotic vehicle when compared with eddy current testing and it was therefore not pursued. Its use, however, will be discussed further in conjunction with a programme of future work.

2.4.2 Eddy Current Testing

Eddy current testing (ECT) [37, 38] is founded on the principles of electromagnetism as described by James Clerk Maxwell. This subject will be discussed in detail in Chapter 6 however for the purposes of an overview the following introduction is sufficient. An alternating current (AC) in a conducting coil produces a changing magnetic field that induces a current in any conducting material that intersects it. This, in turn, establishes an opposing magnetic field which reduces the coil impedance. Defects, modify this field and consequently affect the impedance of the coil. It is these changes, measured in terms of the magnitude and/or phase, that are measured in ECT.

Eddy current systems do not suffer the same level of microstructural complication as UT, however material properties can still influence the inspection results. The process of cold working (deformation of steel below its recrystallisation temperature) can cause a partial transformation from austenite to a ferromagnetic phase. The presence of this material causes a local variation in the relative permeability and hence increased density of magnetic flux from the coil. As eddy current testing is based on the interaction of the magnetic field with the material under test, this can lead to spurious results. A positive use of the effect of ferrite content on eddy current systems has been proposed to control the phase fraction in steel production processes [39]. Further detail on the theory of eddy current testing can be found in Chapter 6.

2.5 Review of Eddy Current Testing

2.5.1 Probe Types

Typically, coils are configured as either absolute, differential or reflection probes. An absolute probe consists of a single coil with the drive signal doubling up as the sensing signal. Such probes are particularly sensitive to lift-off, probe wobble and thermal, permeability and conductivity variations which can lead to spurious indications of defects. Differential probes consist of a pair of coils which usually each form one half of a balanced Wheatstone bridge. The detection signal, in this case, is the out-of-balance voltage from the bridge circuit. In this configuration, the unwanted sensitivities are cancelled out provided that each coil experiences them to the same degree. The differential probe is fundamentally a relative system and simply detects a difference in the conditions under either of its coils. A limitation naturally arises from this in cases where the defect length exceeds the coil spacing

since only the leading and trailing edges of the crack will be detectable. Reflection probes consist of electrically separate coils: one to induce the eddy currents and one or more sensing coils to detect variations in the induced field. Such probes allow flexibility in terms of optimisation of the driver/pick-up coils thereby offering the potential for high sensitivity to very small defects and their performance can be equivalent to that of the preceding configurations [40]. In all cases, the sensitivity of each coil is inversely proportional to the distance from the conducting surface.

2.5.2 Applications of Eddy Current Inspection

Eddy current applications can be broadly grouped into one of three categories: defect detection, thickness measurement and metal sorting. The latter is not relevant to automated NDE and therefore the following section will concentrate on the former two applications.

Absolute pencil or pancake probes are typically used for detecting surface breaking cracks operating in the frequency range of 100 kHz to several MHz and, provided the surface is not overly rough, are capable of detecting defects down to approximately 0.1 mm or less [41]. For inspection close to metal-air edges or regions where two dissimilar metals are joined (e.g. ferrous fasteners in aluminium) shielded probes can be used to provide a more focussed field. Differential probes are less common outside automated applications. Sub-surface cracks, of particular interest in the aerospace industry, can be detected by using a suitably low operating frequency (100 Hz to 10 kHz). As a rule of thumb, for a given probe diameter, the maximum depth of penetration at which defects can be reliably detected will be 33-50% of this. For surface-breaking cracks, the capacity to detect defects decreases as their size becomes less than the diameter of the coil [42]. Thickness measurement of a non-conductive coating (e.g. insulation) is possible using eddy current techniques and essentially equates to a measurement of lift-off. For non-ferrous materials, high frequencies (>100 kHz) are desirable and it should be ensured that the depth of penetration is small enough that the thickness of the underlying material is not important. A rule of thumb is that the operating frequency should be 10 times that required for a penetration depth equal to the material thickness. To obtain quantitative results, calibration blocks of varying thickness must be used to generate a calibration curve and absolute results derived from this.

2.5.3 Probe Configuration for Mobile Robotic Applications

There are three factors that make absolute probes unsuitable, or at the very least, non-ideal for robotic inspection applications. Firstly, temperature variations cannot easily be corrected without retrieving the vehicle and re-calibrating. This in itself, is likely to produce poor results since the temperature at the point of access may be different from that at the inspection location. Secondly, local conductivity and permeability variations would be difficult to identify remotely and could result in an incorrect designation of a section as being defective. Both of these mean that interpretation of the results is less straightforward. Thirdly, absolute systems are generally highly sensitive to probe lift-off. Automatic lift-off suppression, making use of the relatively slow change in signal due to lift-off compared with that of a defect, is possible through judicious high-pass filtering of the raw signal to eliminate this low frequency oscillation. Such a technique is useful for inline processing environments where the material moves at a constant speed relative to the probe, but not for NDE applications. Each of these issues is corrected implicitly by a differential system. Temperature compensation is inherent since each coil will be situated together and therefore, as they share the same drive circuit, at the same temperature. Material variations are unlikely to be confined to an area smaller than the coil spacing unless the coils are very large or situated very far apart. Finally, lift-off is automatically suppressed provided that both coils experience it to the same degree. The addition of an extra coil/winding results in a larger probe assembly however this can be minimised through the use of either small coils or a single core, differentially wound coil. Although such a design introduces limitations in terms of the detectability of large defects, this can be more easily overcome remotely than the issues of absolute probes, with potential solutions being either a mechanical scanning assembly or an array of differential coils. Furthermore, for pipe inspection, such a system is required in any case for a practical system capable of inspecting the entire circumference of a pipe. Although a successful inspection can be achieved using an absolute probe in an environment containing all of the aforementioned issues, the success of the inspection is largely a result of the experience of the NDT technician to eliminate the unwanted artefacts and interpret the data and is therefore not suited to a remote, automated inspection. Furthermore, while very sensitive results may be obtained by a skilled technician, the time required to perform the inspection is greater and a quicker inspection is possible using a differential setup. While suitably configured reflection probes can produce results equivalent to that of either absolute or differential probes, they are typically more complex and thus expensive to produce. For mobile robotic

applications the trade-off between the simplicity and sensitivity of the differential setup make it a good solution.

Array instruments are suited to the same inspection applications as conventional systems however offer the advantage of lower inspection times through increased speed. Parallel inspection may be possible with multiple coils employed at the same time. The resolution in the scanning direction is generally very good and is dependent on the increment interval of the probe head. The resolution in the transverse direction, however, is generally quite coarse since this is limited by the element size and spacing. The use of an array is an attractive solution in many automated applications for three primary reasons:

- i. Increases the speed of inspection
- ii. Eliminates the need for mechanical scanning of probe
- iii. Provides greater data accuracy

The latter is true because the fixed reference system of the array provides more accurate information on the position of the coils.

2.5.4 Multi-frequency Eddy Current Techniques

Multi-frequency eddy current techniques are used in many instances to increase the sensitivity of an inspection. The term can relate to systems using small numbers of discrete frequencies added together and fed into the drive coil to swept frequency eddy current (SFEC) systems and pulsed eddy current (PEC) systems. Due to the skin effect, different operating frequencies result in sensitivities to different artefacts. A common use of multi-frequency techniques is in the inspection of heat exchanger pipes where a number of additional infrastructural supports add noise to the eddy current signal. By subtracting a frequency that is very sensitive to the support structure from a signal that covers the full thickness of the tube, the effect of the supports can be largely eliminated, thus increasing sensitivity to defects.

SFEC differs from this approach in that it is essentially the repeated application of a single frequency over a wide range and is often used for thickness measurements in rolling mills. Although such applications use a differentially connected probe, one coil is in air and the other on the process line and thus only temperature compensation is achieved by the second coil. Furthermore, in this configuration there will always be a differential signal and so a major issue is achieving sensitivity to small changes superimposed on a much larger nominal signal. In [43], Yin et al investigated the use of the change in complex inductance (ΔL) for detection of

plate thickness in a hot steel rolling mill. It was shown that the $\text{Im}(\Delta L)$ has a frequency response containing a maximum and that the frequency of this is related to the thickness of the plate. Furthermore, it was found that the effect of lift-off was confined to a change in the magnitude at which this peak occurred. Consequently, the error due to lift-off could be eliminated by considering only the frequency, which was found to be inversely proportional to the plate thickness. For NDE, it is generally cracks that must be identified and thus lift-off can be much more simply eliminated through hardware using a differential probe with both coils oriented towards the test specimen. One advantage of a multi-frequency approach for NDE, however, is that data is gathered over a range of frequencies and can be averaged to potentially yield higher sensitivity.

Pulsed eddy current differs from both previous approaches in that it uses a step function rather than a sinusoidal excitation. The advantage of using this frequency rich signal is that all defects can be identified from one signal and very quickly using the same probe and setup, which is attractive for fast, practical inspection. The PEC signal is split into three parts for interpretation. The initial part shows near surface artefacts such as lift-off and surface breaking cracks, the middle part identifies internal and interlayer corrosion while the 'tail' signal shows any far surface cracks/corrosion and can be used to detect the overall thickness. Post-processing techniques can be used to compensate for the effects of lift-off and edges. However, where swept-frequency excitation allows noise from other frequencies to be filtered out through, for example, lock-in detection, this is not possible with PEC leading to SNR issues. In general however, the advantages associated with the inspection speed make PEC very attractive to the industrial end-user.

2.5.5 Commercial Instrumentation

Two dominant categories of eddy current system exist presently: handheld, graphical instruments for single/dual probe use or array instruments [44].

Handheld instruments have gained popularity due to the improvements in processing capabilities and battery technology and typically allow operation over a wide frequency range. The precise features included vary greatly depending on the model. General Electric Measurement & Control offer the Vector 22 [45] and Phasec series [46] systems, both of which do not require connection to a PC. The Vector 22 system is a dual channel instrument intended for automated testing lines or laboratory use and provides an impedance plane output. The Phasec 3 series

systems are single channel, handheld devices allowing visualisation on a TFT screen and weigh 1 kg including a 6 hr battery pack. The data output can be displayed as either one of x/y impedance plane, Y/t or bar graph. Multiple probes are interchangeable and the system has built-in lift-off and probe orientation compensation. There are also tunable filters available to remove unwanted frequency components and data can be stored on board up to 200 traces. Both systems offer a wide operating frequency range from tens of Hz up to 10 MHz and allow dual frequency mixing. The UniWest US454 EddyView and the US454-A Multi-Frequency EddyView [47] are similar in size and weight to the Phasec 3 and operate over a similarly large frequency range. The multi-frequency model allows up to four frequencies to be mixed together (two more than the GE system) to form the drive signal and can operate at up to 15 MHz. The MIZ-21B [48] from Zetec is the smallest form-factor, portable device and can display results as x/y impedance plane, bar graph or C-scan. Consistent with other offerings, dual-frequency mixing is possible however the upper frequency range is slightly more limited at 8 MHz. Common to all of the listed systems, is support for a variety of different probe heads to suit the desired inspection application.

There are two main commercial array products available. The Olympus NDT Omniscan MX ECA allows multiple probe heads to be used depending on the particular inspection application. The drive signal can contain up to 8 frequencies at the same time and nominally supports 32 coils (64 with the inclusion of an external multiplexer). Measurement of the amplitude, phase and probe position are available in a range of graphical representations. Built-in filtering options allow further customisation to a particular inspection need. Performance on sub-surface defects is good with a 10% lap-splice thickness loss detectable at a depth of 5mm [49]. Zetec offers flexible eddy current array probes for weld inspection that interface with their MS5800 [50] or TC7700 [51] remote data acquisition systems. In contrast to the Olympus array, this system is designed so that processing and visualisation of the data is done on a PC using their custom ECVision software [52] rather than providing a fully portable device. Consequently the total weight of this system is 12 kg (or 7.4 kg greater than the Olympus system) and is less suited to on-site inspection. The system performance is quoted in terms of the minimum detectable pit size which varies from 0.75 mm diameter by 0.75 mm deep to 2 mm diameter by 0.75 mm deep for coil diameters of 2 mm and 3 mm respectively in 300 series stainless steel [53].

2.6 Visual Pattern Recognition Techniques

2.6.1 Motivation for Incorporating Pattern Recognition

Despite the human eye being relatively poor in ocular terms, compared with most modern cameras, the capacity of humans to recognise complex patterns and shapes from noisy data is considerable. Such intrinsic processing abilities are due to the vast computational capabilities of the brain. Such performance cannot be mimicked using current technology however the principal of pattern recognition and characterisation has been applied widely, with fields ranging from speech, handwriting [54], eye [55] and face recognition systems [56, 57] through to medical [58, 59], process control [60], weather forecasting and information retrieval [61]. The potential has also been recognised for automated NDE processes [62, 63, 64, 65], although to a much lesser extent, where manual inspection of many images is highly time consuming and may lead to performance degradation due to mental fatigue. The application to the nuclear pipe inspection problem offers the opportunity for a low cost, automated inspection systems that can operate with limited supervision to assess or flag suspect areas for further scrutiny.

The purpose of the following section is to draw together a variety of different approaches to image processing and pattern recognition and to explore the potential for their application to NDE images. Existing applications in facial recognition, medical, astronomical and radiographic NDE will be considered, highlighting the key features of the input data and reviewing the associated techniques for them as outlined in literature.

2.6.2 Review of Applications

2.6.2.1 Facial Recognition

Facial recognition research has been fuelled to a large extent by the increasing need for personal identification for security and law enforcement reasons [66]. Depending on the specific application, input images may contain single or multiple faces at varying pose with partial occlusion and/or under varying lighting conditions. The role of low frequency components has been shown to be important in recognition task with studies showing that gender classification can be achieved using only the low frequency while specific identification of the individual requires the inclusion of the higher spatial frequencies [66]. Osuna et al [67] presented a support vector machine (SVM) implementation which achieved a

97.1% correct classification rate on an image set consisting of 313 high quality images and 74.2% on a set containing 23 mixed quality images. In both cases the images contained instances of multiple faces at a variety of pose. Their goal was to identify areas of each image that contained a face and highlight it. Each image was first segmented into 19x19 pixel overlapping sub-images and a best-fit illumination plane subtracted from each to correct for heavy shadows. Histogram equalisation was used, which scales the image intensities so that they occupy the full spectrum of values, to compensate for differences in illumination brightness and camera response. This process is repeated at multiple scales of the original image to account for the possibility of different sizes of face within the image. A SVM was used to determine if the sub-image contains a face. Eigenfaces [56] is an information theory approach which decomposes an image containing a face into a series of feature images and projects the original images onto this feature space. For each known face, a boundary is defined for the maximum allowable distance from the face class. A new image is classified by first projecting it into the feature space and then comparing it with the projected positions of known images within this feature subspace. If the image falls within the decision boundary for one of the face classes then it is assigned to that class otherwise it is labelled 'unknown'. This approach, adopted by Turk and Pentland, used principal component analysis (PCA) to extract the eigenfaces and, operating on images containing faces on a varying background, was shown to yield an average classification rate of 96% on images exhibiting lighting variations, 85% with face orientation variations and 64% with face size variations

2.6.2.2 Medical Applications

In medical applications, particularly oncology, pattern recognition takes the form of computer aided diagnosis (CAD). In contrast with facial recognition, the aim of CAD is to support the experienced physician by acting as a 'second reader' to highlight areas of potential interest (in a process known as 'prompting') and is not intended to replace him. In this system, the medical expert is required to view all images first and consult the CAD system only after recording an initial diagnosis. The input data varies depending on the specific medical application. Early detection of melanoma is of critical importance and, in automated systems, three classes of lesions may be of interest: benign, dysplastic (precursor for melanoma) or malignant [68]. Noise in skin images may come from the presence of hair follicles surrounding a lesion and may be removed through a combination of morphological operators to identify the hair location and substitution of the corresponding pixel values with adjacent skin pixel values [68,69]. The image is then segmented into lesions and healthy skin by either boundary reconstruction

using edge detection/linking procedures or colour segmentation. The latter may involve simple thresholding techniques (particularly where the histogram is bimodal), region growing or clustering. In some cases, the process of classification can be achieved by simply thresholding offering a computationally inexpensive method however practically speaking, the performance is likely to be poor due to image noise or the artefact of interest being small in comparison to the image frame. In general, the amount of diagnostic information contained within an image is low in relation to the total image content and so only a small subset of the total features necessary for the full image content are required to represent the information required for diagnostic purposes; this is similar to NDE images. Thus, having identified the region of interest, a set of features are determined which are typically either pigment colour (normalised for skin type) or shape (particularly asymmetry) related. A classification is made based on the *a priori* knowledge (garnered from a training procedure) of the likelihood of such features appearing in the defined classes. Statistical classifiers are the most common [70] although KNN methods have been employed. In [68] a correct classification rate of 73% for a malignant class was achieved using a KNN classifier (K=24) however for the three class problem (benign, dysplastic or malignant) this reduced to 61%. The suppression of false negatives is a high priority in melanoma detection due to the relative risk of this misdiagnosis. This has parallels with NDE in that the incorrect identification of a defective region has more serious implications than a non-defective region. Zang et al [70] presented a system to detect the borders of skin tumours using a radial search technique. This may be regarded as a semi-automated system since no interpretation of the results is provided.

Mammography has been a popular area for the use of CAD, where labels are placed on the mammogram indicating areas believed to have a high probability of malignancy and some commercial systems have gained FDA approval [71]. Due to the relatively low contrast of mammograms, false negatives (cancerous tissue deemed to be benign) can occur due to the highly demanding nature of manually performing a detailed search for small indications of such entities. Secondary reading procedures using an independent physician can significantly reduce the false negative rate however this can be costly and the resources may not always be available. Microcalcifications in the breast, when present in groups, can indicate a precursor to cancer. Noise within mammograms generally come from regions of tissue with varying radiographic density and connective tissue. The detection and classification of microcalcifications is relatively well suited to CAD because the radiographic signature (size, shape, brightness) of these artefacts is distinct compared with other mammographic features [72]. Brem et al [73] showed that a

98% classification rate could be achieved in this case although no statistically significant improvement was achieved over using a group of expert mammographers. There is, however, the potential to free up resources and/or assist less experienced clinicians with diagnosis. Yu et al [74] developed a system that first segments suspected microcalcification pixels using two wavelet and two grayscale statistical features (median contrast and normalised gray level) and groups them using spatial connectivity. These features then form the inputs to a multi-layer feed-forward ANN to generate a likelihood map of microcalcifications. In order to reduce the number of false positives due to blood vessels, radiographically dense tissue etc, a second step involving a 31 element feature set (including grayscale, shape and contrast measures) classifies these as true or false microcalcifications. A 90% true positive (TP) rate is achieved at the expense of 0.5 false positives per image however the test and training was performed using the same image database. An ANN approach was also adopted by Meinel et al [75] but this time to breast MRI images. An interactive region growing algorithm was used to segment the image by specifying an initial grayscale value and represents a non-automated step for each image. Thirteen features are selected relating to either shape (area, radial length etc), texture (grayscale intensity) and time/intensity curves and each of the 80 breast MRI images are defined in terms of these features. These features are then fed into a 13-3-1 back-propagating neural network, trained using the leave-one-out method (see Section 5.4.4.3) which yields an output ranging from +1 (definitely malignant) to -1 (definitely benign). This system showed improved performance for the CAD system supporting diagnosis over human readers alone. A statistical approach to segmenting microcalcifications from digitised mammograms using an iterative Bayesian update was presented by Karssemeijer [76]. In this approach each image is expressed by local contrast and shape images and each pixel label is updated based on the relative pixel value and the likelihood that the pixel is a dot or a line. The shape parameter is necessary in order to differentiate faint microcalcifications from the x-ray signatures of the connective tissue. The approach was shown to perform much better than standard local thresholding techniques however carried additional computational complexity and required that the parameters for the prior distribution be determined experimentally. The resulting three class image highlights either microcalcifications, connective tissue or background.

Pattern recognition techniques have been applied to the problem of identifying abnormalities in the heart wall movement from echocardiographical sequences [77, 78]. In this application, a diagnosis requires analysis of the inner and outer wall movement through a full cardiac cycle. In [78], the difference image between the

heart position at the contracted and relaxed phase is used and the principal components of the training set (where the heart condition has been determined by medical experts) form a set of basis vectors in the same way as is obtained in the eigenfaces approach. An image is first projected into the feature space and the Euclidean, city-block and Cosine distance (see Section 5.3.2) of each are calculated and a combined KNN classifier used. Assessment of the system showed a performance of ~83% correct classification as either normal or abnormal.

2.6.2.3 Astronomical Applications

A number of features appear in solar images that relate to a particular phenomenon. Sunspot identification and characterisation (particularly sizing) form part of quantitative study of the solar cycle and of forecasting of solar flare activity. Quality of images varies particularly where archive records have been digitised. Initial techniques used thresholding to separate sunspots from the background intensity of the sun using *a priori* estimations for the intensity threshold [70]. The size of the spot was characterised by the number of pixels that fell below the threshold. Histogram techniques removed the need for threshold estimation and for high quality images a more accurate approach involves using edge detection and boundary gradient intensity techniques. Bayesian methods are typically more computationally expensive (and hence slower) than these however are more insensitive to noise and background variations. Turmon et al presented a system employing a Bayesian framework to identify quiet sun, faculae, umbra and penumbra [79]. Each image pixel was assigned a class label based on its intensity and the system makes use of the assumption that physical processes will have a spatial coherence such that labels tend to form clusters in the image plane. Prior knowledge of the intensity distribution for each class label of interest is achieved through expert-input as a training stage.

2.6.2.4 Non-Destructive Evaluation

Pattern recognition has been applied to NDE radiographic images of structures and in particular weld seams for the detection of incomplete penetration, cracks, porosity and voids. Noise in such images is usually manifest as randomly distributed pixels having significant intensity differences from neighbouring pixels. In such an application [62], the raw radiographical images were first low pass filtered to remove the high frequencies due to noise pixels, followed by a histogram equalisation stage and background subtraction. The background was defined as the low spatial frequencies due to the slow varying intensity changes of the pixels defining the weld bead where-as the defect indications could be identified by high

spatial frequencies. Twelve features relating to shape, size, location and intensity are then extracted from the resulting image and the classification performance using a fuzzy KNN and multilayer perceptron neural network evaluated. Both classifiers yielded approximately a 92% classification rate however how this result is distributed between each of the six classes is not explained. Gayer et al [63] outlined a two step process designed to mimic that of a human expert examining a radiographic film. The first step attempts to locate defective regions in the weld seam with coarse resolution through a fast-search process while the second step examines these flagged areas in more detail for more accurate identification and location. The fast search algorithm is based on the assumption that a defect results in irregular gray levels within the image. Two fast search techniques were investigated (both on 32x32 pixel sub-images): a high frequency contribution algorithm and a gradient approach. The high frequency content was shown to be indicative of a defect with the derivative approach even more effective. Defect classification of the flagged areas was done using the sequential similarity detection algorithm (SSDA) which attempts to match the 1D row or column grayscale variation to a template variation for known defects quantified by an error function. The performance of the SSDA approach suffers from defect scaling and contrast variations. A neural network approach was implemented in [64] with the intention of improving on the performance achieved by linear classifiers. Low-pass filtering is used, again, to remove noise and histogram equalisation is employed. Four features were evaluated for each image: position, aspect ratio, ratio of the small axis to the area and the roundness. Classification of each image, expressed in terms of these features, was attempted to one of the following classes: undercut, lack of penetration, porosity and slag inclusions. In order to ensure equal numbers of class images, duplicates were chosen randomly to obtain 25 images per class which appears difficult to justify and the system is only tested using the data presented for training. Consequently a suspiciously high classification rate of 100% is achieved for all classes other than the slag inclusions (where 1 out of 50 is unclassifiable) through modifying the number of intermediate neurons with a further conclusion that the probability of correct classification is high for new samples. An SVM approach was adopted in [65] to detect linear defects in low contrast images. The weld area is first segmented from the rest of the image and then the resulting area is processed using a sliding window. In each position the grayscale value of the pixels within the block are ordered and successively higher elements are used as thresholds resulting in a series of feature vectors defined by the original block thresholded at each given level. The SVM is then trained to define a decision boundary between defect blocks and non defect blocks. Once a defect block in a new image has been identified by the SVM, the Hough transform is applied and an

algorithm invoked to predict the rectangular region that contains the defect. This approach assumes only one defect per image frame and that all defects are essentially straight lines.

2.7 Concluding Remark

This Chapter has provided an overview of a broad range of topics relating to the problem of austenitic stainless steel pipe inspection. Chapter 3 will now formalise a system specification based on this review.

Chapter 3

Specification of Inspection System

3.1 Introduction

The aim of this Chapter is to discuss the specific elements of the project which influence the design of the system. The nature of the nuclear industry is such that it is difficult to obtain site visits and/or specific details regarding the facilities therein. The specific details presented here were obtained through an informal research collaboration with Nexia Solutions Ltd (latterly National Nuclear Laboratory). As much of the detail does not exist in the public domain, in any formal sense, it has not been possible to provide references as would be desired. In fact, in many cases, as was pointed out in Chapter 1, the precise contents of facilities are in fact unknown to the operators themselves. The details used in this specification are therefore the result of numerous email exchanges, telephone conversations and meetings and represent the best estimates of the inspection environment. The specification for the mechanical design, visual payload and eddy current payload will be discussed in turn and the specification for each stated at the end of each section.

3.2 Robotic Platform

3.2.1 Environmental Constraints

In Chapter 1 it was observed that, within nuclear processing cells, the material of choice for pipes is typically austenitic stainless steel due to its structural strength

and resistance to corrosion. Limited ferromagnetic properties may exist due to factors associated with material fabrication however these are grossly insufficient to allow a magnetic form of traction to be employed. Therefore the search for an adhesion mechanism must be restricted to one of the remaining five technologies discussed in Chapter 2. Several decades' worth of ageing will inevitably result in the outer surface of the pipes harbouring significant levels of particulate deposit in addition to any cracks, corrosion pits or other artefacts that may be present. Such factors present issues both in terms of the method of traction and the inspection modality. The diameters of pipes within processing cells are typically in the range of 20mm to 220mm outside diameter (OD) with the exact size depending on function. Particularly in older plant, there is generally a higher proportion of 50mm OD pipe. An essential feature of any useful inspection robot must be its ability to be fitted on to a pipe without requiring the disassembly of any of its component parts. Such a requirement is likely to influence the size of the robot.

Due to the nature of the material flowing through the pipes, processing cells tend to rely on welded joints rather than flanges, however some may exist. Pipe support clips are used to support the weight of the pipe with hangers or stands used depending on whether the pipe is tethered from above or below respectively. Some typical support clips are shown in Figure 3.1; these wrap around the circumference



Figure 3.1: Pipe support clips

of the pipe with securing bolts at either one or both sides. For a 50mm pipe such a clip would be expected to add approximately 5mm to either side of the pipe and thus 10mm to the overall diameter. In addition to welds, flanges and pipe hangers, there are a number of other features including pipe bends, reducers, and expanders which impose kinematic constraints on the system. Consequently, a practical inspection robot should be capable of operating in the presence of all such fixtures and on any orientation of the pipe. Typically, in order to ensure laminar flow, pipe bends have a radius of curvature of no less than four times the outside diameter of the pipe.

The finite volume of the processing cells, coupled with the lack of a design-for-inspection philosophy in their construction and significant retrofitting of additional piping in the years following their commissioning, has resulted in a relatively densely populated environment with no standard clearance between adjacent pipework. Several consequences arise from this:

- i. The robotic inspection system must be physically small enough to ensure that potentially damaging collisions with, and obstructions due to, adjacent infrastructure do not occur. The dimensions of the access point to the processing cell (through which the robot must be deployed) defines the maximum allowable size of the robot to be 300mm, however, in the absence of a standard clearance policy, the design intent must be to ensure the robot is as small as possible.
- ii. The complexity of the internal infrastructure means that the design must be entirely wireless to avoid the risk of an umbilical becoming entangled and impairing the inspection. Furthermore, no tethered line may be used as a fail-safe to retrieve the robot in the event of failure and hence ideally the robot should be 100% reliable.
- iii. The likelihood that there will be no clear line-of-sight (LOS) path between the access point and the robot largely precludes the use of a visual or laser range finding positioning system and simple LOS communication technologies. Also, the potential for strong multi-path effects from such a densely populated structure further limits the options for a communications system to technologies with multi-path suppression capabilities.

3.2.2 Performance Constraints

The performance constraints can be broken down into three principal categories: traction, range and positional accuracy. The hazardous nature of the operating environment necessitates a performance requirement that traction should be maintained even in the event of a total power failure. Failure of the adhesion system at any point could result in damage being incurred to the robot or (less likely) the surrounding environment and invoke a costly retrieval operation. This task is made more complex by the incongruity of magnetic traction. The distribution of the access points to the cells is such that it is estimated that the distance from any desired inspection location to the nearest inspection port will be approximately 10m. Therefore the battery life must be such that the robot can travel to the desired location, perform the inspection and return to the starting point without requiring a recharge. Furthermore, the communications system must be capable of operating over this range. The positional accuracy of defect location desired by the end-user is ± 3 mm and to minimise inspection costs and disruption, any associated positioning system equipment should be deployable through a single inspection port.

3.2.3 Levels of Autonomy

Level of Autonomy	Description
Fully autonomous	User interaction is not required beyond the provision of the required inspection task(s). Having obtained this, the robot can then perform the inspection, dealing with any incidents and obstacles that may occur, intelligently and return to the entry point with the inspection results.
Semi-autonomous or automatic	The robot can guide itself to the inspection location while scanning its environment for obstacles or hazards and perform the specified task requested of it. However, if an obstacle is encountered which the robot does not recognise, it cannot intelligently adapt to the scenario for which it has not been trained. In this instance, the user will be alerted and manual control may be enabled to overcome the situation. The robot can run self-diagnostic checks (such as battery level measurement) and communication between the base station and the robot is automatic.
Drone	The user controls the robot to guide it to the inspection location and avoids any obstacles that may be encountered. The communications architecture is a master-slave system, with the robot acting as the slave and only responding to direct commands from the master. The user has remote control of the sensor during the inspection process.

Table 3.1: Level of autonomy

The term 'autonomous' is one that is often used incorrectly and is frequently confused with 'automatic', meaning 'self acting, under the conditions fixed for it'. Strictly speaking, autonomous is used to define something that is capable of operating entirely independently and has the ability of self governance. The assumption of this level of intelligence is subtly different from automatic. Three basic levels of autonomy have been identified in Table 3.1. In reality, such discrete levels are more appropriately replaced with a continuum of levels of robot

intelligence however, for the purposes of robot design, it is helpful to categorise them as shown.

A fully autonomous robot is the 'holy grail' of robotics and represents the peak of artificial intelligence. In its final form, the robotic pipe inspection system aims to slot into the semi-autonomous category where manual supervision will be on-hand if necessary.

3.2.4 System Specification

The aim is to create a robotic vehicle which can operate on non-ferromagnetic pipes however the wide range of pipe diameters initially specified is a significant engineering task. The operating environment plays a central role in the design of any robotic system and in this instance places certain constraints on the variety of suitable designs. As with all climbing robots, overall weight is a critical parameter and an optimum system will be the minimum weight solution that best satisfies the specification. Given that NDE is likely to be most urgently required on older plant, to simplify the mechanical complexity and cost, the RPC should be designed for use on pipes with a principal diameter of 50 mm with a tolerance of at least ± 10 mm to allow for inline changes in diameter and the width of the pipe support clip. By following such a development strategy, a proof-of-concept design can be created which can be scaled to greater diameters as necessary, while greatly simplifying the engineering problem and maintaining focus on the fundamental research element of the project. Similarly, the kinematics of the system should be limited to translation along any orientation of straight and curved sections of pipe. Circumferential locomotion, while ultimately necessary to allow the RPC to orient itself in any position on the pipe and essential for a practical, industrially deployed system, introduces a further degree of freedom which presents an additional engineering obstacle. This omission does not detract from the proof-of-concept, and therefore is regarded as a sound decision. Given the limited resources and expertise in robot development, creating a platform which can be guaranteed to be infallible is impractical for a prototype system. It is essential however that the system is wireless and therefore a fail-safe tether is not an option. Instead of providing a guarantee that once the system has been deployed in the processing cell then it can be guaranteed to make it out again under its own power the goal will be that, given a system power failure, the robot will maintain its grip on the pipe and will not fall off. A limitation of this modification is that, in practise, loss of power would require a separate robot-retrieval process to be initiated.

The desired accuracy of defect location is ± 3 mm which, given the 10 m range of the robot, requires a positioning system accurate to 0.6% of the expected range. Line of sight limitations preclude the use of laser range finding systems and ultrasonic positioning technology, such as the commercially available 'cricket' system [80], are limited to positional accuracies of the order of centimeters. Additionally, the desire to make the robot deployable through a single inspection port presents a significant challenge in terms of the usage of any positioning system relying on trilateration. It was concluded that a more achievable goal would be to design a robot which could be positioned to a rough global accuracy while the subsequent inspection could be conducted against the local reference of the robot to a much greater accuracy. This compromise ensures that the data measured by the sensors will be very accurate relative to the fixed position of the robot, however there will be an increased level of uncertainty in the global position of the defect within the processing cell.

Finally, given the limited resources and knowledge-base within CUE regarding radiation hardening of components, it was concluded that a proof of concept system should have the luxury of assuming that the environment is non-radioactive. The process of ruggedising the robotic hardware and electronics of the final system is an important engineering process in order to produce a commercially viable product. However, this constitutes a distraction from the more fundamental question of the viability of such an inspection system. The assumption taken here allows the principle of robotic inspection, miniaturisation and collection and fusion of different data to be the primary focus while acknowledging the ultimate need to radiation-harden the final system.

3.3 Pattern Recognition System

3.3.1 Discussion

For the given application it is not possible/practical to allow an NDE expert to perform the inspection manually. Thus, where in many NDE applications the use of an experienced technician employing visual inspection techniques is straightforward, albeit subjective, the automation of this process poses a significant challenge. The choice is therefore to either automate the process or have an expert continually monitor the visual data. Common to all of the reviewed applications, the motivation for the research has been primarily to reduce the effect of human limitations by increasing objectivity, accuracy and repeatability. In pipeline inspection applications, as with oncological screening, the vast majority of the

images processed would be expected to contain no artefacts of interest and automation would potentially increase performance where previously expert opinion was degraded by the mental fatigue associated with monotonous processes. Furthermore, the correct assessment of an image is highly dependent on the experience of the technician/physician and their ability to recognise minute details within noisy images. In oncology applications, the consequences of missing a cancerous instance are much higher than those associated with an increased workload to the physician. Therefore it is more important to be able to detect that something is present than it is to accurately quantify what that is since this can then be handled by the oncologist. Similarly, in NDE it is of greater value to be able to reliably detect that a defect exists than it is to be able to accurately quantify what the nature of that defect is. Thus, a good result (and the primary objective for an automated classifier) would be to be able to reliably distinguish images containing a defect from those that do not. A secondary objective would be to be able to characterise its nature provided this was not at the expense of an increase in false negatives in the primary objective.

Fatigue cracks are of significant interest due to their potential to cause failure however they are visually very difficult to see and penetrant techniques are often used to aid identification. As no depth information is obtained within the 2D image, it is unrealistic to presume that such defects could be differentiated from surface scratches by an automated system without similar visual enhancement techniques. Corrosion pits are often a precursor to the initiation of fatigue cracks and so early detection of these would be valuable. Corrosion pits in steel are characterised by their irregular but roughly circular shape and may penetrate up to the wall thickness of the pipe. Noise in images of fatigue cracks and corrosion pits may arise from surface scratches, dirt, associated support structures or lighting variations. The problem of hair follicles in melanoma detection applications has parallels with the presence of surface scratches in NDE images. Where-as a combination of morphology and pixel substitution was used in [68, 69], it is by no means certain that there would be sufficient contrast between scratches and the general surface to allow such techniques to be employed for NDE images. However, such techniques may be suitable for removing artefacts such as support structures where contrast will be high and there is little variation in physical shape. Low frequencies in facial recognition were identified as being suitable for determining gender (characterised by facial structure) and it seems possible that this could be applied to NDE where the low spatial frequencies could identify the general type of defect (e.g. significant cracks, corrosion pitting etc) based on its shape and in particular to distinguish between images that contain defects from

those that do not. The use of high frequencies (as was shown for radiographic NDE) may be more appropriate for fatigue cracks however will result in images containing higher levels of noise due to their similarity to surface scratches. Through the use of penetrant techniques it may be possible to use the high frequency content to identify fatigue cracks while removing the noise associated with scratches. The opposing frequency requirements for the description of corrosion pits and fatigue cracks leads to the conclusion that separate pre-processes may be required for each. As fatigue cracks are often generated through the structural weakness introduced by corrosion pits, for a first attempt at an automated system it is logical to limit the scope to consider these and thus the low frequency content.

In terms of yielding a classification result, the techniques used are often either template matching, neural network, SVM or pixel-level processes. The initial process usually involves some segmentation of the image to identify the artefact of interest from the image and then calculates some parameters (colour, shape, texture etc) that define its features. In some cases the image is divided into blocks or sub-images and the features calculated for each. In statistical pixel-level processes, the image may be segmented by iteratively evaluating the class label applied to each pixel based on its own value and the surrounding pixel labels. The eigenface approach has been successful in facial recognition and has been applied to some medical applications however as yet this has not been attempted on NDE data. In contrast with the previous methods of feature extraction, this technique derives the features automatically through maximisation of the variance and so only the most statistically relevant features are used. The tolerance of such an approach to noise is better than for more computationally simple algorithms with the added advantage that it does not adopt a black box approach and thus the performance may be traced directly back to the original training data set. It may be possible to define a set of classes to describe a structure in a similar manner to the eigenface approach and to then employ the statistical methodology that follows from this. Thus different conditions of the pipe would be represented by a combination of these features. In facial recognition systems, the subject of interest (i.e. the face) is the consistent parameter (in a template sense) upon a varying background. Considering the primary goal, it is realistic that the undamaged pipe can be assumed to be the consistent parameter and it should therefore be possible to represent this condition as a 'face'. The variance of the image set for this class will be small and therefore it should be possible to differentiate such images from ones that contain other artefacts (e.g. defects, support structure). Thus in the first instance, anything that is not classifiable to this is of interest to the NDE technician. If the grayscale

image is fed into such a system then it is likely that it would be possible to discriminate between pipes that contain no defect and those that do. However due to the positional uncertainty of defects within the image frame it would not be possible to determine the nature of the defect within the same system framework; this would require a separate algorithm to identify defect borders and classify according to perhaps its shape or some other relevant parameter. Furthermore, the variance associated with different classes of defect will also be greater due to the organic nature of their creation. However, if the data is pre-processed to map the original image to a space such that the signature for each defect is broadly similar then the technique could again be applied with the advantage of the inherent tolerance of the statistical approach.

3.3.2 System Specification

The aim is to design an automated pattern recognition system capable of detecting images that contain defects from those that do not. The eigenface approach should be adopted to achieve this in order to demonstrate whether or not this technique is appropriate for this type of NDE. A secondary goal will be to develop this system further so as to attempt to sub-classify an image to a particular type of defect. A defect/no-defect classification rate of 100% is desired however given the performance of systems produced by dedicated research teams in this field this is regarded as a goal that may not be achievable. A more appropriate measure, given the novelty of the technique to this application, is to benchmark the performance against an established technique such as the KNN algorithm. The system will be developed for finding defects consisting of slots and drill holes since these can be readily fabricated and will yield a signature composed mainly of low frequency content (and therefore will be more readily distinguishable from surface scratches). Artificial corrosion pits will be created to test the performance on a more noisy, and more representative, set of defects. The application of the system for fatigue cracks will not be considered and will instead be left for future work. Finally, the system will be developed in MATLAB to make use of the plethora of image processing toolboxes available.

3.4 Eddy Current System

3.4.1 Discussion

The argument for using a differential eddy current system in mobile robotics applications was made clear in Chapter 2. The need to inspect the full

circumference of the pipe makes an array particularly attractive. Fundamentally, there are two options that would allow full circumferential coverage during an inspection: a mechanical scanning system to translate the eddy current probe or a multi-coil array wrapping around the circumference. The advantage of the former is two-fold in that it requires a very simple electronic drive system and that the transverse scanning resolution can be made as high as desired by reducing the transverse scanning interval. However, although lift-off suppression is inherent for the differential probe type, the inverse relationship of sensitivity to coil lift-off means that the coils should be positioned as close to the pipe surface as possible. Achieving a consistent distance within a moving mechanism would require a high degree of precision and it is likely that a practical system would exhibit some degree of lift-off (and hence sensitivity) variation proportional to either the position of the coil on the pipe and/or the rate at which it is moved. Furthermore, such a design would result in slower inspection speeds. The opposite is true of a multi-coil array with mechanical positioning of each coil to achieve constant lift-off very simple while having the added electronic complexity of either including multiple drive circuits or using a single circuit in conjunction with a multiplexer. The latter solution is preferable for a measurement system where repeatability and compactness are desired qualities. As producing high precision mechanical components is of greater difficulty and expense than producing the array electronics, the array approach is the most appropriate solution. One drawback of the array within the RPC platform is that due to the constraints associated with obstacle avoidance the array cannot extend around the full circumference of the pipe and hence in this instance only half of the circumference will be inspectable. Since the addition of circumferential motion to the robot is an ultimate goal of future work, this is seen as an appropriate compromise.

The advantage of a multi-frequency eddy current array is that data is gathered over a range of frequencies and can therefore be averaged to potentially yield higher sensitivity. Single frequency inspection can achieve very sensitive results however relies on a feedback loop (usually in the form of an expert NDE technician) to adjust the frequency so that the best SNR is achieved for the defect. Swept frequency eddy current removes this need by exciting the coils across appropriate range of frequencies. Pulsed eddy current is undoubtedly a very popular and promising technique however, given that the research element of the eddy current array payload lies in its integration with the overall inspection system, the added simplicity of the SFEC system is more attractive.

Measurement of a signal's phase is generally regarded as a more reliable method of detecting defects than amplitude measurement. Phase measurement may be done in the time domain or the frequency domain. Time-domain measurement has a more simple algorithmic implementation however requires either a very fast analogue-to-digital converter (ADC) or a slower ADC and the use of upsampling (see Chapter 6). Frequency-domain implementations involve calculating the Fast Fourier Transform (FFT) which requires less data (and hence little or no upsampling) however is more complex to implement efficiently. In the first instance, to minimise the design complexity, the time-domain approach should be used.

The diameter of detectable defect is related to the probe diameter, as outlined in Chapter 2. In general, the performance is scalable and thus a system designed for one size of defect could be made to operate on smaller defects by simply changing the coil size; this principle explains the multiple probe offerings from commercial ECT equipment suppliers. Discussions with industrial end-users yielded the conclusion that an automated system capable of detecting defects in the region of 0.1 mm deep would be considered as satisfactory.

3.4.2 System Specification

The aim is to develop a differential, eddy current array using phase measurement to detect the presence of defects. A swept frequency system will be used to attempt to gain a high sensitivity while maintaining simple payload design and interpretation of the results. The coil diameter will be 12 mm (equivalent to that of a previous differential system developed within CUE [81]). To make the problem a little simpler, and given the scalability of the coils, the diameter of the defects that should be detected will be 10 mm, which is slightly smaller than that of the coils. This is quite a large defect however will allow the sensitivity to depth to be assessed while ensuring that the defects are of suitable size so as also to be detected visually by the visual payload. Given that the novelty of the research is in the synthesis of all the component parts into a single inspection system this is deemed to be justified however system scaling should be addressed in future work. The minimum depth of defect that should be detectable is 0.1 mm (i.e. 10mm x 0.1mm).

Chapter 4

Robotic Pipe Crawling Platform

4.1 Introduction

It has been established that robotic systems offer the attractive solution of conducting inspections remotely, thereby ensuring that there is reduced health and safety risks to personnel while providing significant savings in terms of cost and time. The design and development of such a system for the inspection of pipes is the subject of this Chapter.

4.2 Locomotion and Climbing Mechanisms

4.2.1 Traction System

Following the discussion of possible traction methods in Chapter 2, an analysis can be made as to the most appropriate solution for the pipe inspection system in light of the specification outlined in Chapter 3.

The substantial performance reduction that chemical adhesion suffers, when exposed to relatively mild levels of dirt and grime, coupled with the inherent lack of control over when the adhesion is applied, makes it unsuitable for the given inspection task. While showing considerable promise, the relative immaturity of the biomimetic and electroadhesion technologies coupled with the lack of commercial, off-the-shelf availability make it currently unsuitable for use in this application. Vacuum adhesion systems can provide good adhesion capabilities on flat, smooth

surfaces. In instances of rough or contoured surfaces, where containment of the vacuum may be compromised, such a system would provide an unreliable method of traction. A further concern with implementing this technology as the sole traction system arises from the reliance of the adhesion on the application of constant power, since the operating environment is such that a loss of traction cannot be tolerated. A hybrid traction system utilising this technology as the primary adhesion mechanism but with the incorporation of a fail-safe adhesion mechanism is conceivable but introduces a level of complexity into the design with no obvious benefit over simply using the fail-safe mechanism alone. The infrastructure required to generate and maintain the vacuum is significant in terms of weight and furthermore, the power requirements of such a system do not lend themselves particularly well to a battery powered system. A vacuum array would offer a solution that is lower power in comparison with the more traditional single cup design however an actuated force must still be applied to the array in order to extend/retract the array for contact with the pipe. This, then, can be seen as simply an extension to a mechanical adhesion system and therefore, unless the grip possible by friction is insufficient to support the robot, it is unnecessary and the additional traction provided by the array is thus redundant. The mechanical method of traction is both simple and well understood. Furthermore, through the appropriate choice of actuators and/or gearing, a robot could be made to maintain its grip on the pipe in the event of a power failure. Thus, mechanical adhesion is the most fail-safe system and provides the desired traction with minimum infrastructure and no functional redundancy. Two principal methods of mechanical adhesion present themselves. Firstly, a wheeled system where the traction force is generated between the wheel and the surface of the pipe and secondly, a gripper-pair system which, when engaged, forms a clamp around the outside of the pipe. Both methods are discussed in conjunction with the design of the chassis which forms the subject of the following section.

4.2.2 Chassis Design

Having elucidated the advantages of a mechanical traction system, the design of the chassis can be restricted accordingly to those suitable to incorporate this adhesion mechanism. Numerous morphologies of chassis and interconnect were identified and considered to allow selection of the optimum combination of partial solutions. Table 4.1 shows the possible partial solutions for the adhesion support structures and Table 4.2 shows the options for structures allowing locomotion along a pipe. By decoupling the design in this way, the best combination of partial solutions from

Table 4.1 and Table 4.2 can be selected to provide the optimum design to meet the given specification.

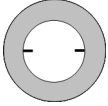
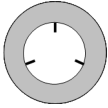
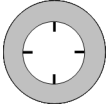




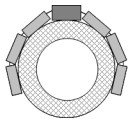
Collar Type	Illustration	Description
C1		Circular collar with two wheels
C2		Circular collar with three wheels
C3		Circular collar with four wheels
C4		Semi-circular collar with two wheels
C5		Semi-circular collar with three wheels
C6		Semi-circular collar with two grippers
C7		Semi-circular collar with two grippers fitted with wheels
C8		Dockable modules fitted wheels collectively gripping pipe

Table 4.1: Chassis collar options.

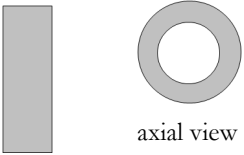
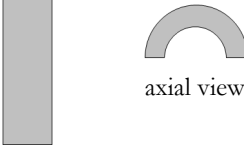
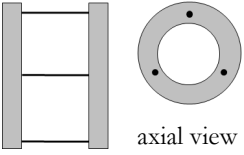
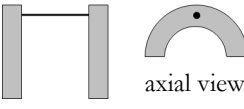


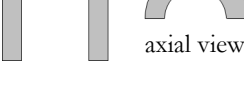
Interconnect Mechanism	Illustration	Description
IC1		Single circular collar
IC2		Single semi-circular collar
IC3		Two circular collars connected by rigid links on linear bearing at one end
IC4		Two semi-circular collars connected by one rigid link on linear bearing at one end
IC5		Two semi-circular collars connected by multiple rigid links on linear bearing at one end
IC6		Two semi-circular collars connected by one rigid link with hinge at one end. Linear bearing at one end
IC7		Two semi-circular collars connected by one rigid link with hinge at both ends. Linear bearing at one end

Table 4.2: Interconnection mechanisms.

In terms of simply obtaining the optimum chassis for support of a mechanical adhesion system, in this case, a circular collar of the type shown in C1, C2 and C3 provides the best solution and allows for motion around pipe bends. Critically however, this design does not have the ability to overcome pipe hangers and therefore this, coupled with its associated interconnect IC1, must be eliminated as design options. The modular design shown in C8 consists of multiple, small actuated units interfaced together to surround the pipe and scaling the design to larger diameters would be the trivial matter of adding additional modules. Both the semi-circular and modular morphologies are suitable for passing pipe hangers. The latter, however, has significantly more mechanical, electronic and control system complexity and it can therefore be concluded that, in this case, the most practical base structure is the semi-circular collar. Having identified this, further simplification can be made to the connection mechanisms through a simple stability analysis. A single collar design requires a wheeled or inchworm solution for locomotion however, if the collar is too thin, the system will be highly unstable and will require significant control to overcome the 'inverted pendulum' effect. A four wheeled, wide wheel-base system could be used, however the width of the frame required for stability would not represent a minimum weight solution and IC5 would be superior. Thus, in this instance, a dual collar design, of the form shown in IC4-IC7, is the most suitable one.

Having identified a dual semi-circular collar design as the most appropriate base structure for the chassis, the relative merits of a wheel and gripper system can be analysed in this context. The simplest system in terms of achieving traction is clearly the gripper system since adhesion is achieved through the application of a motor of some description to cause the gripper to clamp to the pipe. In contrast, a wheeled system would require a bi-directional gearing mechanism to allow the wheel to lock when stationary, so that grip could be maintained without either backward driving the motor or having to exert constant torque by the motor. This issue is particularly significant for vertical sections. Furthermore, in order to ensure that a constant contact force is generated and to allow for an element of abnormality in the contour of the surface, a clamping mechanism would be required for each wheel. In its most basic form, this would involve a non-actuated spring-loaded mechanism or, at the expense of added complexity and additional motors, a hybrid system such as that shown in C7 could be used to actively control the force applied by each wheel. An advantage of a gripper system, in comparison with all of the wheeled designs, is that the gripping mechanism is decoupled from the locomotion system. Thus, if the latter fails the results are unlikely to be catastrophic. Such a decoupled system would ultimately be slower due to the need

to apply and remove the gripper prior to moving forward. However, since robot safety is a higher priority than operational speed, it can be concluded that a gripper pair solution is the most suitable for this NDE robot.

All interconnect options now fall into two basic categories: either a rigid link (IC4 and IC5) or some form of flexible link such as a hinge (IC6 and IC7). The ability to negotiate corners was identified as a key kinematic requirement and any interconnect system must allow for this. A hinged mechanism allows a rotation about an axis perpendicular to the plane of the page, however in order for a robot to travel around a bend, each collar would need to be hinged at both ends. If this was not the case then when the forward collar attempted to move around the radius, the rigid connection at the rear collar would act to inhibit it. In order to make the mechanical system robust, each hinge would need to be actuated, thereby introducing a significant overhead in terms of the control system requirements. A flexible link fabricated from a conformable material could be used, however both this and the dual hinge option would lead to a similarly 'limp' design which would lack the durability for industrial use. A better alternative would be to connect each collar with rigid links whilst allowing the grippers to rotate. Such a system would require one collar to be mounted on the interconnecting links using linear bearings in order to allow translation. In this way the chassis can be made highly robust without sacrificing the overall kinematic goal. A single rigid link system as shown in IC4 while satisfying the minimum weight criteria, would allow torsion of the chassis which could result in damage to the robot components. The optimal combination of partial solutions is therefore determined to be a combination of collar configuration C6 and the multiple rigid links interconnect configuration IC5.

4.2.3 Design of Gripper for Mechanical Adhesion




Figure 4.1: Profiles for pipe gripper: (a) curved gripper, (b) semi-square gripper and (c) semi-triangular degree gripper

Having concluded the merits of a gripper pair traction system, the specifics of the design can now be considered. To eliminate the need for additional actuation, the use of a pincer gripper, whose action is similar to the opening and closing of a pair

of scissors, was eliminated and therefore the gripping force should be generated entirely by the force used to engage/retract the gripper unit. Three basic gripper profiles were identified from which other variants may be derived and are shown in Figure 4.1. The curved gripper provides a solution which, for the particular pipe diameter it is designed, gives the largest contact area and hence best traction. However, it is unsuitable for use on pipes of whose diameter is greater than that for which it was designed and its performance reduces for pipes of smaller diameter as the contact area is reduced. The semi-square profiled gripper is similarly unsuitable in these conditions and furthermore, due to the flat vertical surface, performs less well than the curved gripper for any diameter less than that for which it is designed. The semi-triangular gripper provides four points of contact at tangents to the pipe circumference. However in contrast to the previously discussed profiles, the grip performance is constant for varying pipe diameters providing the diameter is not too great that contact is made only by the tip of the gripper or too small that the tips of the gripper pair collide. The latter issue can be addressed through use of an intermeshing design [31] which allows the tips of the gripper to pass through each other in much the same way as a person clasps their hands. In accordance with the varying inline diameter stipulation in the specification, a semi-triangular, intermeshing gripper was selected.

In order to ensure a lightweight design with high structural strength, corrosion resistance and tolerance to the elevated temperatures that may be encountered at the pipe surface, the gripper should be fabricated from aluminium. Since the friction force between metal-metal surfaces is not typically very high, additional rubber pads can be added to the contacting face of each 'finger' of the gripper. These pads may be replaced due to friction and/or thermal wear as necessary

4.2.4 Cornering Mechanism

As discussed in Section 4.2.2, in order to allow the robot to travel around a bend, the gripper must be capable of rotation. This may be either passively or actively controlled. Active control allows the gripper to be precisely positioning however the control/sensing system required is complex since the gripper must orient itself to the correct angle prior to its engagement. If the orientation is incorrect then the gripper will not contact the surface in the optimum manner and sufficient grip may not be achieved to support the weight of the robot. Further control could be added to allow a successive approximation approach which would overcome this, however such a system would introduce greater complexity still. A passive system, in contrast, allows the grippers to freely rotate and hence, as each gripper is

engaged, the reaction force from the point of contact with the pipe causes the gripper to automatically orient itself to the correct angle. However, due to the dual-collar chassis design, in a passive system, if both gripper pairs are initially engaged and one subsequently releases, due to the rotational freedom the chassis will rotate about the axis of the other gripper pair. Thus during locomotion, the disengaged gripper pair will either drag or be pushed along the surface of the pipe. While both options have their negative attributes, the passive system offers the simplest and lowest power solution and therefore is the most suitable. In order to limit the extent of the 'drag' effect, each gripper can be constrained to provide only a small rotation ($\pm \frac{\theta_{max}}{2}$) which, while still allowing the pipe crawler to adjust its orientation with each application of the gripper, limits the extent to which the chassis can rotate back when one gripper pair is disengaged. This process may be understood more clearly through reference to Figure 4.2 which shows the robot in a position midway around a bend with the gripper for collar 1 at an angle of θ_{C1} and that of collar 2 at an angle of θ_{C2} to the axis of the chassis, and with each gripper perpendicular to the tangent of the bend, TC1 and TC2, respectively. As the platform progresses through the bend it will naturally attempt to follow the tangent to the radius of the curve however the step size must be limited to ensure that the rotation required by the grippers to correctly orient themselves does not exceed the limits of $\pm \frac{\theta_{max}}{2}$ from the axis of the chassis.

4.3 Mark I - Linear Crawler

4.3.1 Overview

In order to assess the performance of the semi-circular collar structure and refine the locomotion system, it was concluded that the Mark I robot should be a linear crawler, omitting the cornering capabilities discussed in Section 4.2.4. This two-stage design process ensured that the knowledge base derived from Stage 1 could be brought forward to refine the final system. In particular, the performance of the mechanical adhesion mechanism for vertical locomotion could be analysed in a simplified kinematic form.




Figure 4.2: Cornering Mechanism. Each gripper is free to rotate within the constraint of $\pm \frac{\theta_{max}}{2}$ from the axis of the chassis. As the robot travels around the bend, each gripper attempts to orient itself to the tangent to the radius of the bend.




Figure 4.3: Mark I collar design highlighting bearing block mounting points and shaft bearing mounting holes.

A modular approach to the design was taken, with the aim of creating standard mechanical units which could be connected together to form the working system. Such a design strategy minimised the impact of damage or manufacturing tolerances on the operability of the system. As discussed in Section 4.2.2, multiple shafts are required to connect the two collar units in order to prevent torsion of the chassis and, in this instance, three shafts were used as shown in Figure 4.3. Linear bushings were mounted in bearing block modules which could be attached to the main collar using screws at the locations shown. The motor to control axial separation of the collars was mounted on the top of the collar above the central interconnecting shaft. In order to ensure that all the manufactured components were lightweight and furthermore to facilitate easy fabrication, all chassis components were fabricated from grey PVC. Simona(R) PVC-CAW [82] was selected for its high rigidity, relatively low density and easy availability.

4.3.2 Selection of Motors

4.3.2.1 Gripping System

The gripping force is generated by four miniature linear actuators [83] each consuming 250 mA of current at peak efficiency and geared to provide 18 N of force and 6 mm s⁻¹ maximum speed. At no load, each actuator has a maximum speed of 12 mm s⁻¹ which results in the minimum time to reach full extension being 1.67 s. Frictional forces and slow-down effects at the instant that contact is made with the pipe result in a grip time that is slightly greater than this and dependent on the pipe diameter. When grip has been established, power to the motors is disconnected to maximise battery life and traction is maintained, providing the actuators' 40 N backward drive force is not exceeded. Each actuator comprises a built-in potentiometer to provide positional feedback with a rated non-linearity of no more than 1%.

4.3.2.2 Axial Motion

Axial motion is generated by a 12 V, two-phase, bi-polar stepper motor [84] with a 120 mm stroke connected to collar 1 and provides 25 N of thrust while consuming 230 mA/phase. The motor housing is 'non-captive', meaning the screw rail can pass through the centre in either direction allowing flexibility in the positioning of the motor on the robot. A stepper motor was used in this instance because it allowed a customisable stroke length through the use of longer screw rails, thereby offering scalability while providing the best trade-off between thrust, power consumption, speed and cost compared to alternatives such as the linear actuators.

4.3.2.3 Selection of Gripper Material to Maximise Traction

The following factors were identified as critical parameters for the gripper material:

- High friction coefficient
- Resistance to wear
- Suitable for use at rated temperature
- Machinable - ability to be cut to size and texture put on surface
- Resistance to water, oil and other chemicals
- Available in low order quantities
- Low cost

The surface temperature of the pipes on which the RPC will operate is a critical parameter and it must be ensured that the breakdown temperature of the material is greater than maximum temperature expected at the pipe surface. For the Mark I design, rubber gasket material was used which typically has operating temperatures in the range of -40 to +120 °C [85].

4.3.3 Finite Element Analysis

4.3.3.1 Introduction

In order to test the ability of the PVC collars to withstand the forces generated by the actuators a 3D, small displacement, linear-elastic, finite-element (FE) stress analysis was conducted using PTC ProEngineer/Mechanica Wildfire 2.0. The properties of PVC-CAW [82] are given in Table 4.3.

Material Property	Standard	Value
Density	DIN 53479	1.42 g cm ⁻³
Bending Modulus	DIN 53457	3000 N mm ⁻²
Yield Stress	DIN 53455	58 N mm ⁻²
Vicat Softening Temperature	DIN 53460	78 °C

Table 4.3: PVC-CAW material properties

4.3.3.2 Simulation Loads and Constraints

Observation of Figure 4.3 shows that the collar exhibits symmetry in the X-Y and Y-Z plane. Utilising this symmetry, a simplified collar model can be simulated, with a smaller number of mesh elements, thereby reducing the simulation time. The centre point on the inner circumference of the collar is fixed in the z-direction providing the datum point and stabilising the model. A bearing load of magnitude 10 N is applied, in the -x direction, to each of the two bearing block mounting holes as shown in Figure 4.4.




Figure 4.4: Loads and constraints. A 10 N bearing load is applied to both bearing holes, totalling the 20 N maximum reaction force exerted by the actuators. Two symmetry constraints are applied in the X-Y and Y-Z directions and the model is fully fixed in the z direction to stabilise the model.

4.3.3.3 Assessment Criteria and Assumptions

If the action of the grippers causes the collar to displace to a significant extent then the direction of the applied force from each gripper ceases to be coaxial, meaning that the gripper performance will be reduced, coupled with a marginal increase in the centre of gravity. The maximum displacement criterion was set to 1mm, meaning that the maximum allowable displacement of any part of the collar is 1

mm. Von Mises stress is a commonly used failure criterion for ductile materials and, if greater than the appropriate yield stress, indicates the failure of a design [86]. Flexural yield strength is a measure of a material's ability to resist deformation under load and, from Table 4.3, for PVC-CAW is 58 N mm^{-2} .

Any FE analysis requires a set of assumptions to be defined which dictate the boundaries over which the simulation may be regarded as valid. In the case of the collar model, the first assumption is that Young's modulus can be considered to be constant. In reality there will be a time dependent deformation of the material due to creep under load which will lead to a non-constant value of the modulus. During locomotion, the load is applied to each collar for only the time it takes to drive the opposite collar to its destination and apply its grippers. The longest period that each collar will be subject to the load will be during the inspection process when both collars are gripping. It is expected that any single inspection will not take longer than ten minutes to complete and so it is assumed that creep effects can be neglected.

The second assumption is that the RPC will be operated at $20 \text{ }^\circ\text{C}$ corresponding to the testing standards for which the data in Table 4.3 is valid. The Vicat B test [87] is used to define the softening point of materials such as PVC that have no defined melting point. The Vicat B softening temperature is that at which a 1mm^2 flat-ended needle penetrates to 1mm at a load of 49.05N and for the grade of PVC used in this simulation it is $78 \text{ }^\circ\text{C}$.

4.3.3.4 Results

Figures 4.5 and 4.6 shows the Von Mises stress distribution and the displacement of the model respectively. The maximum stress of 1.467 N mm^{-2} is well below the yield stress and occurs at the inner surface of the collar about the Y-Z plane. Further regions of high stress occur at the bearing block mounting holes which are to be expected since these are the points at which the force is transmitted from the bearing block to the collar. The maximum displacement to the applied 20 N net load can be seen to be 0.1928 mm which is well below the maximum displacement criterion and thus satisfies the assessment criteria outlined in Section 4.3.3.3. The mass of the collar predicted by the simulation is 212.8 g which is in reasonably good agreement with the measured value of 232.2 g .




Figure 4.5: Von Mises stress analysis of Mark I collar design using PTC ProEngineer/Mechanica. The maximum stress is 1.467 N mm^{-2} .




Figure 4.6: Displacement analysis of Mark I collar design using PTC ProEngineer/Mechanica. The maximum displacement is 0.1928 mm.

4.3.4 Overview of Mark I Design

The Mark I platform is shown in Figures 4.7, 4.8 and 4.9. As discussed in Section 4.3.4, the vehicle consists of two rigid, semi-circular collar units oriented axially on the pipe and connected by three support shafts. Collar 2 has three linear bearings and slides along these support shafts relative to collar 1, powered by the stepper motor identified in Section 4.3.2.2. This motor is mounted in an aluminium housing which is fitted to collar 1. Both collars 1 and 2 are fitted with a pair of bearing blocks. A set of two retractable grippers are mounted on linear bearings housed within these blocks and allow the robot to clamp on to the outside of the pipe. The multi-fingered, intermeshing gripper design provides flexibility for operation on pipes that have a varying in-line diameter while still providing a high degree of surface contact area. Grip performance is maximised through the addition of high friction coefficient, rubber pads to each finger. The power delivered by the linear actuators outlined in Section 4.3.2.1 is transmitted to the shaft via a motor-shaft coupling. Each linear actuator sits on a motor cradle to provide mechanical support and to minimise the lateral stress applied to the motor shaft. Safety caps are fitted to the ends of the support shafts as a fail-safe mechanism to prevent the system over-extending during axial motion. The electronic hardware is mounted on rear of collar 1 and acts as a partial counterweight to collar 2. The total weight of the vehicle is 1.375 Kg.

Starting in the fully retracted position and with both collars 1 and 2 gripping the pipe, forward motion is generated by repeating the following routine. The grippers of collar 2 are retracted and then the collar is extended along the pipe (away from collar 1). Once in the fully extended position, collar 2 then re-establishes its grip on the pipe. Collar 1 then releases its grip and is driven along the pipe (this time towards collar 2) before finally re-gripping when it is in the fully contracted position. The current prototype is configured to allow both forward and reverse motion. As outlined earlier, a common feature in any pipeline is the presence of flanges for interconnection and pipe hangers for support. It is critical to the practical use of a robotic pipe inspection system that such obstacles can be overcome with minimal overhead. Because of the mechanical design of the RPC, the actual locomotion algorithm does not change when such obstacles are encountered; the grippers are fully retracted and the RPC simply 'steps over' the obstacle as though it were not there, as shown in Figure 4.10.




Figure 4.7: Mark I robotic pipe crawler (RPC) on 40mm diameter pipe (Axial view).




Figure 4.8: Image of collar 1 on 40mm diameter pipe. Bearing blocks mounted on the collar house a linear bearing on which the gripper slides. The intermeshing gripper accommodates varying diameters of pipe.




Figure 4.9: Profile of Mark I design on 40mm diameter pipe. A stepper motor with a 120mm screw rail facilitates axial motion. Power transmission to the grippers is achieved through a motor-shaft coupling. Safety caps are included to pre-




Figure 4.10: Navigating past pipe stands. Both collars start in the gripping state. Collar 2 then opens and extends past the pipe stand before re-gripping. Collar 1 releases its grip and is then contracted before re-establishing its grip.

4.4 Mark II - Pipe Crawling System

4.4.1 Evaluation of Mark I platform

The Mark I platform is capable of travelling along horizontal pipes of diameter 50mm and below, however vertical locomotion presents a challenge. During tests on vertical pipes, the RPC could maintain traction when both sets of grippers were engaged, however it was not possible using only one gripper pair. Since the locomotion mechanism requires the release of one pair of grippers prior to the initiation of the axial translation mechanism, this issue meant that reliable vertical ascent/descent was not possible. In order to address this problem for the Mark II platform, two potential solutions were identified:

- Increase the gripping force
- Reduce weight of platform

A discussion of a strategy to increase the gripping force is provided in Section 4.4.1.1 followed by a detailed examination of the weight reduction measures in Section 4.4.1.2.

4.4.1.1 System for Increasing Gripping Force

A drawback of using motors with a greater output force to overcome the grip deficiency is that this is coupled with an increase in power consumption. In keeping with the low power, battery operated philosophy, a passive system utilising springs was designed to increase the gripping force generated by each of the existing linear actuators.

Springs come in two main varieties: compression and extension. Compression springs, as their name suggests, exert a reaction force as they are compressed while extension springs do the opposite, exerting their reaction force as the spring is extended. Due to space constraints between the gripper and the inside radial surface of the collar, the use of compression springs is problematic and thus extension springs, mounted as shown in Figure 4.11, is the preferred solution. Increasing the gripper force through the use of springs has the potential to reduce the working stroke of the actuator depending on the spring constant and the spring extension at zero actuator extension, since the force required to extend the spring increases with distance. The extent to which the stroke is reduced is determined by the spring length at which the force equals the 20 N maximum force of the actuator and hence the spring stiffness, k . In order to ensure no loss in stroke, the former




Figure 4.11: System for increasing gripping force through the use of springs.

should be greater than the 20 mm maximum stroke of the actuator. The optimum spring, however, will provide a maximum force of 20 N when the actuator extension is 20 mm whilst also maximising the force exerted at 0mm actuator extension. The total gripping force, F_t , generated by a spring-assisted actuator system is given by Equation (4.1)

$$F_t = kx + F_m \quad (4.1)$$

where x is the extension (m) and F_m is the force generated by the linear actuator. A suitable spring was sourced with a spring constant of 0.59 N mm^{-1} and nominal length 24.2 mm. As stated, the maximum extension possible of the spring corresponds to the extension at which the spring force is 20 N which, by rearranging the equation $F = kx$ for x and substituting the value for the spring constant and maximum force, is 33.9 mm. Setting this as the limit in order to maximise the contribution of the spring, the spring extension at zero actuator extension, and therefore the pre-load value x_{0p} must be $33.9 - 20 = 13.9$ mm. Thus for a 50 mm pipe diameter, corresponding to an actuator extension of 10 mm, an additional 14.1 N will be supplied by the spring. Similarly for a 40 mm and 60 mm pipe the additional force will be 8.2 N and 20 N respectively. As a consequence however, the power consumption of the actuators will increase during the grip release phase due to the additional energy required to extend the springs. Furthermore, the speed with which the grippers can be released will be reduced due to the greater load the actuator must now push against. Due to these issues, and the effectiveness of the weight reduction measures outlined in the next section, this system was not implemented.

4.4.1.2 Weight Reduction Measures

Figure 4.12 shows the composition, by mass, of the Mark I robotic pipe crawler. The main contributions arise from the collars, the bearing blocks and the grippers which total 1007.6 g or 69.1% of the total RPC weight. Of the remaining components, the stainless steel bushings form the most significant proportion followed by the payload bracket and then screws and miscellaneous fittings. Based on this data, a significant weight reduction strategy was implemented.

The diameter of the Mark I collar can be reduced slightly while still providing enough space for the stroke of the grippers and by removing material strategically from the body of the collar, a significant weight reduction can be achieved with little or no reduction in structural strength. By integrating the stepper motor into the body of the collar the weight reductions are two-fold through the creation of a 25 mm diameter cavity to accommodate the motor and the elimination of the need for a separate motor housing. The design of a new collar incorporating these changes is described in more detail in Section 4.4.2.

The desire to have a modular design was driven in part by the desire to have reusable building blocks thereby allowing an iterative approach to the initial proof-of-concept design. This strategy led to the adoption of separate bearing blocks to house the gripper bushings. The benefit of several design iterations was to provide valuable data on performance of various configurations however, such flexibility is no longer required for a Mark II system. Thus, it is now possible to incorporate the gripper bushing into the main collar thereby allowing the bearing blocks to be removed entirely at a saving of 18.5% of the total Mark I weight. Substitution of PVC in place of aluminium used in the Mark I gripper would allow a weight reduction of the order of 50% of the weight of the aluminium gripper.

Several smaller measures were introduced including the use of nylon rather than stainless steel bushings, minimising the number of screws/fittings and reducing the size of the motor-shaft coupling with a collective saving of approximately 100 g. The mass breakdown of the Mark II robot platform is shown in Figure 4.13 and totals 1043.5 g.

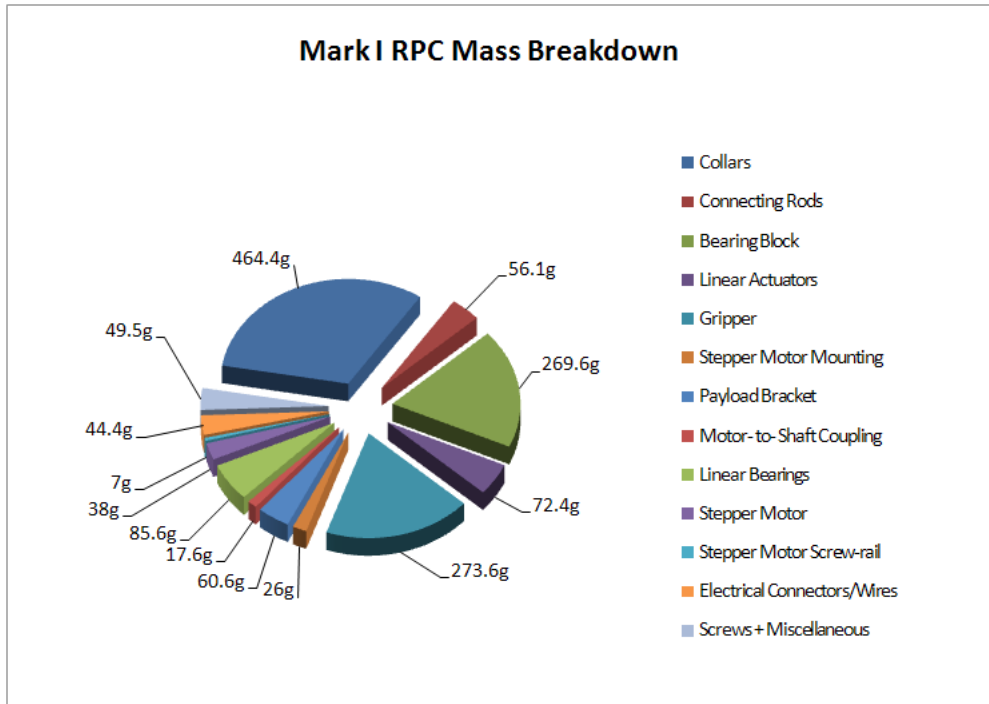


Figure 4.12: Weight distribution of Mark I design

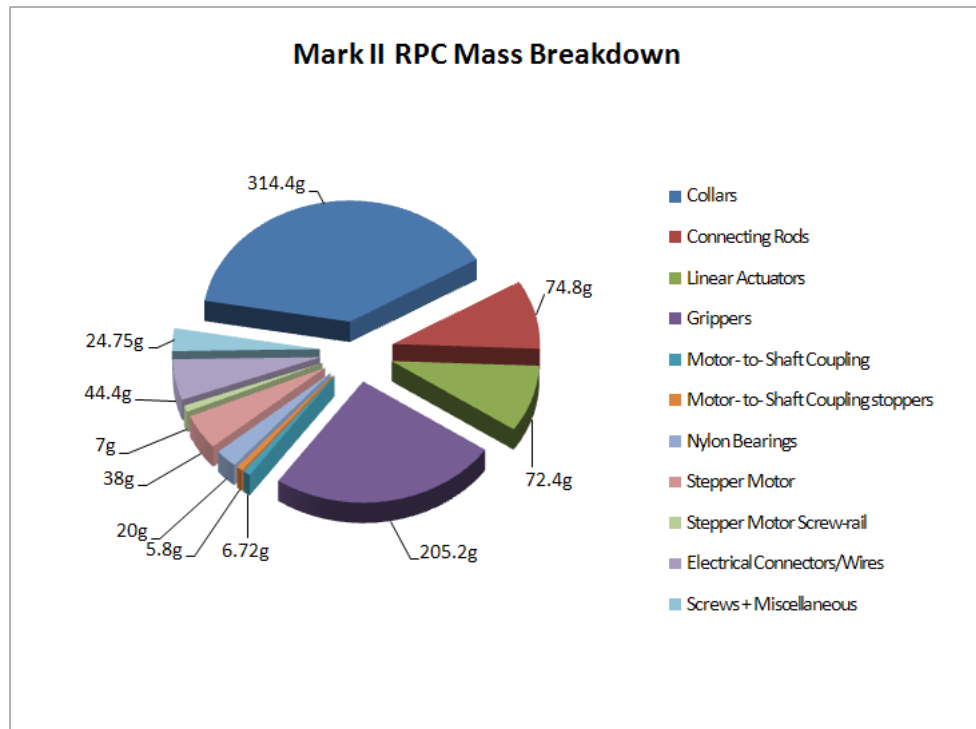


Figure 4.13: Weight breakdown of Mark II robot following implementation of the weight reduction measures.

4.4.2 Optimisation of Collar Design

In order to minimise size and weight, a minimum diameter design is required. With reference to Figure 4.14 (a), if the gripper is designed such that the angle between the two fingers is 90° then for all pipe diameters, $d = 2r_p$, the pipe will contact the gripper at a distance of $r = \frac{d}{2} = r_p$. Using Pythagoras' theorem, the distance, m , from the centre of the pipe to the vertex of the gripper is $m = \sqrt{r_p^2 + r_p^2}$ which reduces to $m = r_p\sqrt{2}$. Considering the arrangement depicted in Figure 4.14 (b), the clearance, c , between the pipe surface and the gripper vertex is given by $c = m - r_p$ and therefore:

$$c = r_p(\sqrt{2} - 1). \quad (4.2)$$

Given that the outside diameter of the collar, D , is defined by Equation (4.3).

$$D = 2(r_p + c + s + t) \quad (4.3)$$

a minimum diameter design is obtained by substituting Equation (4.2) into Equation (4.3) yielding:

$$D = 2(\sqrt{2}r_p + s + t) \quad (4.4)$$


where, s is the actuator stroke and t is the thickness of the collar. For a pipe radius of 25 mm, stroke of 20 mm and collar thickness 30 mm, from Equation (4.4), the minimum diameter is 170 mm, 30 mm less than the Mark I design. A 5 mm allowance was allocated to allow the shaft to be connected to the gripper thus raising the minimum realisable diameter to 180 mm.

Having established the optimum diameter for the collar design, further weight reductions measures must involve modifying the design of the collar itself. Due to limitations in the fabrication facilities within CUE, complex geometry parts are difficult to achieve and time consuming to manufacture. However, it is well understood that a structural member can still exhibit significant load bearing abilities even when a large portion of the structural material has been strategically removed. This is undoubtedly best performed in nature, with the honeycomb structure of a bee-hive or a spider's web being prime examples of optimisation of strength-weight ratio. Analysis of the Mark I collar design led to the conclusion that a large proportion of the collar material could be removed with an unnoticeable change in strength. As was outlined in Section 4.4.1.2, a cavity could be created in collar 1 in which the stepper motor could be housed resulting in a significant removal of material from its centre. However, in order to ensure correct balance


between collar 1 and 2, the design of each must be slightly different. Common to both collars are 24, 6 mm radial and 20, 4 mm axial holes as shown in Figure 4.15 and Figure 4.16 for collars 1 and 2 respectively. However, collar 2 has an additional six 4 mm axial holes to compensate for the stepper cavity. Furthermore, since the stepper motor is now housed at the centre of collar 1, the central supporting rod must be moved adjacently. A fourth supporting rod was added to the design, at the expense of an additional 18.7 g, to ensure symmetry (and hence balance) about the axis of the pipe and to enhance the rigidity of the design. A bearing is located at the centre of collar 2 to accommodate the screw rail of the stepper motor and a locking screw provides the link between the collar and motor. As the gripper bearing block has been eliminated, the bearings are now inserted directly into the collar and the linear actuator attached to the side face. In total the collar optimisation process achieved a Mark II design that contained 31% less material by mass than the Mark I design.

The final modification to the collars involved changing the location of the system electronics. Having the electronic mounting at the rear of the robot generated a moment force about the collar. In order to distribute this load more evenly across the structure, the mounting point was moved to the top of collar 1 so that it effectively straddles the two collars.

For simplicity, the Mark I design omitted the cornering mechanism outlined in Section 4.2.4 however, having established the concept of a linear crawler, this could be realised in the Mark II design. Due to the linear constraint of the Mark I crawler, the shaft could be attached directly to the motor for power transmission. However, since the shaft can now rotate, this is not possible as it would generate an unacceptable rotational force on the actuator shaft. The motor-shaft coupling was modified such that the gripper shaft could freely rotate and two transmission discs were added to either side of it which were fixed to the shaft. The mechanical system is shown in Figure 4.17. A 1 mm nylon disc bearing is fitted between the coupling and the transmission discs to provide a smooth rotation. When the linear actuator contracts or extends, the motor-shaft coupling pushes against the transmission discs and, since they are firmly attached, thereby moves the gripper shaft. If the gripper needs to rotate, it is free to do so within the limits dictated by the diameter of the cornering lugs and the rotational limit cavity. The Mark II platform is shown in Figure 4.18.



(a)



(b)

Figure 4.14: Collar diameter minimisation. (a) Showing pipe with right-angled gripper and (b) showing the gripper mounting on the collar.




Figure 4.15: Mark II collar 1 minimised weight design.




Figure 4.16: Mark II collar 2 minimised weight design.




Figure 4.17: Cornering mechanism showing modified power transmission system.




Figure 4.18: Mark II RPC platform featuring optimised collar design.

4.4.3 Finite Element Analysis


The effect of the optimisation strategy on the structural strength of the Mark II collar was assessed through a second finite-element analysis. This was once again performed using the integrated Mechanics package within PTC ProEngineer Wildfire 2.0. The assessment criteria and assumptions for the previous analysis given in Section 4.3.3.3 are still valid and remain the same for this design study.

4.4.3.1 Simulation Loads and Constraints


As per the simulation in Section 4.3.3, the collar exhibits symmetry in the X-Z and Y-Z planes which allows a simplified model to be simulated. The centre point on the inner circumference of the collar is, once again, fixed in the z-direction and acts as the datum point. However in this case, as the bearing block is no longer present, a bearing load vector of 20 N was applied to the motor mounting hole on the side of the collar.

4.4.3.2 Results

The results of the stress and displacement analyses for collar 1 are shown in Figures 4.19 and 4.20 respectively. The maximum stress can be seen to be 4.173 N mm^{-2} which is well below the yield stress given in Table 4.3. It may be observed that at two points the stress reaches this maximum value: at the motor mounting hole and at the opening of the first radial hole on the inner diameter of the collar, near the outer face. A marginally reduced stress of 3.9 N mm^{-2} occurs at the inside of this cavity. The stress profile around the circumference of this radial hole is given in Figure 4.21. The two maxima can be observed corresponding to the points of minimum wall thickness at the location between the hole and the axial surface and the hole and the plane of symmetry. The maximum displacement of 0.221 mm occurs at the outer edge of the collar along the X-Y plane and, although marginally greater than the 0.1928 mm of the mark I design, is still well within the maximum displacement criterion of 1mm. Thus the Mark II collar design achieves a weight reduction of 31% while easily satisfying the stress/displacement criteria.



(a)



(b)

Figure 4.19: Von Mises stress analysis of Mark II collar 1.




Figure 4.20: Displacement analysis of Mark II collar 1.




Figure 4.21: Von Mises stress around the circumference of hole.

4.5 Electronic Systems

4.5.1 Electronic and Processing System

The electronic hardware was designed to be highly modular, with the main control electronics centred on a 40-pin signal bus allowing peripheral circuit boards to be stacked while providing each layer with access to the same signals. Three layers are required for the basic operation, referred to collectively as the ‘main processing stack’, and are shown in Figure 4.22:

- Signal interface board (SIB)
- Motor driver board (MDB)
- Processor Board

An additional payload board is needed to add an inspection capability and is sensor specific. A block diagram of the vehicle architecture is shown in Figure 4.23. At the heart of the system is the main processing stack which controls all robot movement and sensor measurement. Each linear actuator has a positional feedback loop so that the shaft extension can be measured. The camera payload connects directly to a serial port on the processing board, however the eddy current array sensor requires a dedicated payload board as shown; the details of this will be discussed further in Chapter 6. The following sections will now examine the individual roles of each element of the main processing stack.

4.5.2 Signal Interface Board (SIB)

All input/output (I/O) signals associated with the robot operation and power supply come through the signal interface board (SIB). Each motor is connected to the SIB through a connector allowing each to be isolated if necessary for diagnostic purposes. The 11.1 V battery output is regulated to 5V with a 1 A maximum output current and fed into the Primary Signal Bus (PSB) for use by the upper layers of the stack. The battery voltage is passed through a voltage divider circuit and fed into the PSB thereby allowing higher levels to sample this on an available analogue-to-digital converter (ADC) if required. A power on and hardware battery low level indicator are also provided for convenience. This circuit is shown in Figure 4.24 and uses a transistor, which in normal operation is conducting and hence bypasses the battery ‘low’ LED. When the voltage drops below the threshold set by the ratio of resistors R2 and R3, it switches off resulting in the low LED illuminating. The threshold is set to 9.5 V and is determined by the minimum cell voltage that will allow recharge (9V) plus 0.5 V.

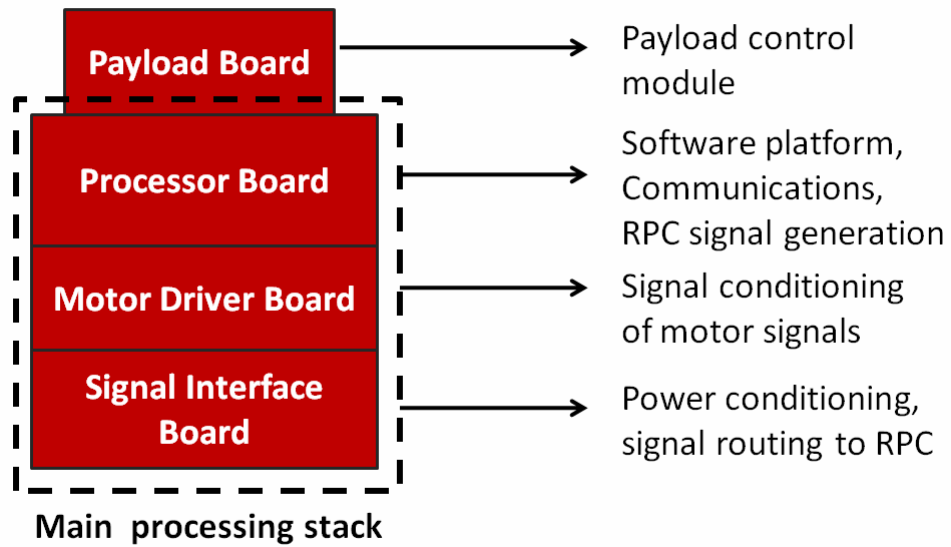


Figure 4.22: Main processing stack including payload board.

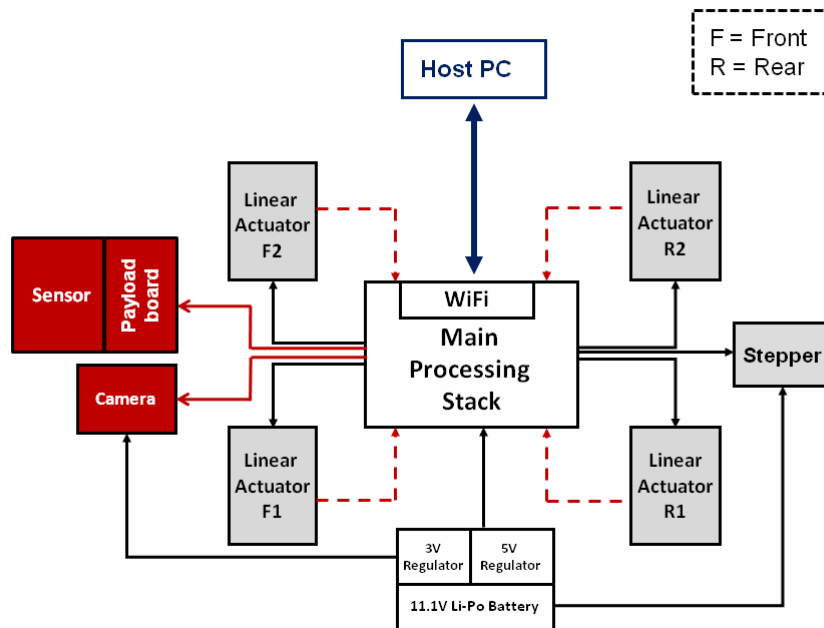


Figure 4.23: Block diagram of the Mark II platform with eddy current array and camera payloads.

The battery itself is connected to the SIB with a connector present to break the 12V supply, providing a convenient means of current measurement. In normal operation, this connector is shorted. A secondary signal bus (SSB) is present on the SIB and carries the motor drive signals as well as several free channels to cater for future expansion. The signals from this bus are connected directly to the output connectors.

4.5.3 Motor Driver Board

The motor driver board provides the drive signals for all the motors on the robot. Input pulse width modulated (PWM) signals from the Primary Signal Bus (PSB) are fed into two quadruple half H-bridge drivers [88] which provide the drive signals for the linear actuators. Each linear motor combines two half-bridges to form a full bridge providing bi-directional control. Motor direction is controlled through digital direction control signals from the PSB. Each collar has one full bridge dedicated to its pair of motors allowing the same direction control lines to be utilised thereby cutting the channel requirements by a factor of two. The output from each bridge is fed into the secondary signal bus for distribution to the motors.

The stepper motor is controlled using an alternative driver containing internal clamping diodes for inductive transient suppression. Four digital channels are fed into the driver from the microcontroller via the PSB and the four outputs are connected to either side of the stepper motor coils. The coils shown in the winding diagram in Figure 4.25 are pulsed in the sequence described in Table 4.4. Motor speed is controlled by the time duration between energising the coils and the direction of the screw rail can be changed by reversing the coil pulse sequence. The schematics for the gripper and stepper motor drive circuits are shown in Figures 4.26 and 4.27.

Sequence	Coil 1+	Coil 1-	Coil 2+	Coil 2-
1	ON	OFF	ON	OFF
2	OFF	ON	ON	OFF
3	OFF	ON	OFF	ON
4	ON	OFF	OFF	ON

Table 4.4: Stepper motor firing sequence.




Figure 4.24: Battery level indication circuit. In normal operation, the transistor, T1, is conducting thus bypassing the battery 'low' LED. When the voltage drops below the threshold set by the ratio of resistors R2 and R3, it switches off resulting in the low LED illuminating.




Figure 4.25: Coil winding diagram.

4.5.4 Processor Board

Processing is performed on a 400 MHz embedded computer [89] running the Linux operating system and having 64 MB RAM. Low level operations including actuator control and sensor monitoring are performed using an 8-bit microcontroller [90] interfaced with the main processor via an I2C bus. Four of the six PWM channels available on the microcontroller are used to generate the drive signals for the linear actuators, however all six are fed into the PSB leaving two channels available for future use. All nine ADC channels are connected to the PSB for sensor monitoring and 14 General Purpose Input-Output (GPIO) channels are available, of which eight are presently involved in motor control.

While the majority of the robot operation could be performed using a more basic microcontroller with an RF connection to the host, the use of a powerful processor allows flexibility and upgradeability which is invaluable in the development stage of a relatively complex robotic system and is therefore an appropriate choice for such an undertaking. Furthermore, the processing capabilities facilitate the future use of on-board image processing algorithms as discussed in Chapter 5.

4.5.5 Payload board

The payload board is an interchangeable layer with dedicated electronics designed specifically for each sensor and integrated with the main RPC electronics through the PSB. This will be discussed in Chapter 6.

4.6 Communications

4.6.1 Wireless Protocols

As stated in Chapter 3, the use of an umbilical is impractical and therefore a wireless communications solution is required. A brief discussion of some of the more common wireless protocols operating within unlicensed bands is provided in the following sections.

4.6.1.1 Wireless RS232

Wireless RS232 is a simple, low cost, low data-rate, bi-directional communications system. Long range communications is possible with basic off-the-shelf systems [91] capable of 250 m transmission with line-of-sight (LOS) and more advanced systems offering ranges of several miles. This technology does not support

networking and is intended as a point-to-point or point-to-multipoint communications system.

4.6.1.2 Bluetooth (Class 1 and Class 2)

Bluetooth operates in the 2.4 GHz band and is primarily designed for short-distance, peer-peer communications [92]. Class 1 and Class 2 Bluetooth can both achieve data rates of 1Mbps with maximum power outputs of 100 mW and 2.5 mW respectively. Class 2 Bluetooth is only capable of communications within a 10 m range however Class 1 offers a range of up to 100 m. Both Bluetooth classes are significantly lower power than their IEEE 802.11g counterpart, described in Section 4.6.1.3, however this is at the expense of data throughput.

4.6.1.3 IEEE 802.11g

The IEEE 802.11g standard permits data rates up to 54 Mbps in the unlicensed 2.4 GHz band and is commonly used for wireless local area networks (LANs). Typical power consumption is approximately 1 W which is ten times greater than Bluetooth Class 1 however as a factor of achieved data-rate, 802.11g is more power efficient [93].

4.6.1.4 Certified Wireless USB

Certified Wireless USB is a point to point communications system that is currently under development and aims to provide short range data-rates equivalent to wired USB 2.0 with 480 Mbps possible within a 3m range deteriorating to 110 Mbps at 10 m [94]. Currently, no commercial product is yet available with this technology.

4.6.1.5 Discussion

The nature of the inspection problem dictates that an on-board camera would be useful for navigational assistance inside the processing cell and hence a direct streaming video link may be required at some stage for a practical solution. The limited bandwidth of wireless RS232 places a constraint on the data throughput from the on board sensors to the base station and hence was eliminated as a design choice. The limited 10 m range supported by Class 2 Bluetooth precluded its use in this instance since this needlessly constrained the range of the device. In order to provide a development system which can be built upon without undue constraint and to provide the option for flexible networking of teams of robots in the future, it was elected that the communications system should be built around the 802.11g standard. A significant advantage of using this standard was the availability of an integrated 802.11g solution for the family of processor. While

providing a substantial excess in available bandwidth, the benefit in networking capability and infrastructure for future upgrades was felt to outweigh the increased power consumption that the technology demands compared with Class 1 Bluetooth. As typical industrial facilities contain significant numbers electromagnetically reflecting surfaces, multi-path is inevitable. For the purposes of this research, the impact of multi-path on the communications system will not be considered however it is recommended as an area for future work.

4.6.2 Robot to Host

The RPC communicates with a host computer over an IEEE 802.11g wireless connection supporting data rates up to 54 Mbps. Data is transferred using Secure Shell (SSH), a network protocol that facilitates the transfer of data between two networked devices. The SSID is the identifier given to a particular wireless network and in order to establish a network connection the two connecting devices must use the same SSID. The RPC is assigned a unique internet protocol (IP) address and when it is powered on the wireless communications automatically initiates the registration of the IP on the wireless network and prepares the device for incoming connections. An SSH client with a command line interface is used to control the low level functionality required for the wireless communications and is embedded in the graphical user interface (GUI) which is discussed in Section 4.10. Secure Copy (SCP) is a file transfer protocol associated with SSH. It is used exclusively to copy files between the RPC and the host and is the means by which images captured by the visual payload are transferred to the GUI.

4.6.3 Internal

All internal communications between the main embedded computer and the microcontroller are conducted through the I2C bus. I2C is a common two-wire (data and clock), bi-directional communications standard allowing 400 kb s^{-1} data rates with the current on-board hardware between the master (embedded computer) and any slave devices (microcontroller). Faster implementations are available through the 'high speed' mode although this is not supported by the current family of processors and in any case the current data-rate is ample for the purposes required.

4.7 Power

The system is powered from a 11.1 V Lithium-Polymer (Li-Po) battery regulated to 5 V and 3 V. The breakdown of the power consumption for the RPC in stand-by mode is shown in Figure 4.28. These are the statistics when no demand is placed on either the motors or payloads and constitutes the base power level. Communications accounts for the largest proportion of the overall budget, totalling more than the main embedded computer and microcontroller collectively. The SIB constitutes the lowest requirement which is to be expected since this board only supports signal routing and power regulation. The power consumption during the key tasks associated with locomotion and inspection is provided in Figure 4.29. The total stand-by power level is 3.6 W and peaks at 8.4 W during the application of either pair of grippers. Although both this and the 8.04 W for the stepper motor operation are quite high, the duration is typically only 2-3 s before the power is reduced to the standby level and thus the total drain on the battery is not too severe. The use of both the camera and ECA payloads constitutes a relatively small increase from the stand-by level to 4.44 W and 4.32 W respectively.

When performing vertical locomotion, in either direction, the effect of gravity requires that the set of grippers in contact with the pipe be energised at all times until the movement is completed. At this point, the robot can engage both grippers generating a friction force great enough to support the RPC weight and the power can then be removed, returning it to stand-by power levels. In order to meet the peak current demand of 1.07 A that occurs during such a task, a second voltage regulator is mounted on the motor driver board to contribute to the 1 A output of the primary regulator.

The total current drawn during stand-by mode is 300 mA (at 12 V) and, given the 910 mA h capacity of the battery, this provides a total stand-by time of a little over three hours on a single charge. For continual vertical locomotion, this reduces to around 50 minutes.




Figure 4.28: RPC power consumption breakdown in stand-by mode.




Figure 4.29: Total power consumption during standard tasks.

4.8 Performance

4.8.1 Vertical Ascent/Descent

The sequence of images in Figure 4.30 shows the robot climbing a 50mm diameter vertical section of stainless steel pipe. The overall linear motion performance of the system is very good with the new lightweight design solving the grip issues encountered in the Mark I platform. Due to the slight backlash in the linear actuators, when power is removed the gripping force reduces slightly. On horizontal or minimally inclined sections, this does not present a problem however on vertical or steeply oriented pipes, the reduction in grip is critical and causes the robot to lose traction. This is easily solved through the application of constant power to the gripping collar's actuators when the platform is engaged in either vertical ascent or descent. A 170 mm vertical distance was covered in a time of 36 seconds, equating to a speed of 4.7 mm s^{-1} .

4.8.2 Translation Around Pipe Bend

The sequence of images in Figure 4.31 shows the Mark II RPC crawling around the inside of a bend in a 50 mm diameter section of aluminium pipe. The radius of the bend was four times that of the pipe diameter (200 mm) and the cornering mechanism, outlined in Section 4.2.4, was found to work very well. As discussed earlier, it is critical that, during such locomotion, the step-size of the RPC is limited so that the difference in gripper angle between subsequent positions of the collar does not exceed those permitted by the rotational limits. Consequently, the speed of motion throughout such a manoeuvre is necessarily slower, at 1.1 mm s^{-1} , than along straight sections of pipe.

Referring back to Figure 4.17, it can be seen that the upper and lower rotational limits will act similar to a cantilever beam upon loading by the cornering lugs. This resulted in a small degree of deflection. Since the cornering lugs are involved when the grippers are engaged, on smaller diameter pipes this does not present an issue as the lugs will contact the rotational limits near the collar and the cantilever effect is minimal. However, for operation on larger diameters, the lugs will contact near the end of the limiting unit and hence there may be considerable stress induced at the collar interconnect. This may require an additional supporting member to be added to the end of the cantilever to eliminate any significant deflection.




Figure 4.30: Sequence of images showing Mark II RPC ascending a vertical pipe.




Figure 4.31: Sequence of images showing Mark II RPC travelling around a vertical bend.

4.9 Future Mechanical Development

As outlined in Section 4.4.2, the manufacturing facilities within CUE prohibited the fabrication of light-weight, complex parts. A collaboration with the Department of Design, Manufacturing and Engineering Management (DMEM) offered a unique opportunity to utilise their recently acquired rapid prototyping facilities to produce precise parts for the RPC. Due to advantages inherent in such computer aided fabrication, a high level of alignment precision can be achieved and resulted in the Mark II(r) prototype, shown in Figure 4.32 with highly smoothed axial translation. Furthermore, it was possible to manufacture each collar as a single unit meaning that the resulting design was both structurally strong and lightweight. This was particularly beneficial in stabilising the flexibility of the cornering control mechanism through the addition of a cradle both to reinforce the resistance to deflection and also to support and stabilise the linear actuators. Contoured internal cavities were generated to house and protect the wiring from external elements and entanglement while a network of triangular cavities within the collar ensured a lightweight but strong structure. A mounting point for the payloads was included allowing each sensor to be easily docked to the main platform.




Figure 4.32: Mark II(r) crawler manufactured using rapid prototyping facilities. Techniques such as this allow complex geometrical shapes to be manufactured as a single part increasing the robustness of the overall product.

4.10 Graphical User Interface

In order to provide a convenient and centralised way to control both the robot movement and inspection payloads, a graphical user interface (GUI) was developed for Windows. All robot control takes place within the main control tab, shown in Figure 4.33. Each linear actuator has a progress bar which reflects the measurement of the actuators built-in potentiometer allowing gripper extent to be determined. The message centre displays any communications from the RPC system received through the wireless link. The camera payload is also controlled here with a simple, single-click image capture process initiating communications with the camera and resulting in the captured image being displayed in the panel indicated. The eddy current array has a dedicated tab, shown in Figure 4.34, wherein two operational modes can be controlled. The primary mode involves the control of the pipe inspection array while the second offers flexibility for use with a two axis, linear stage for sensor development and accurate imaging of flat plates.

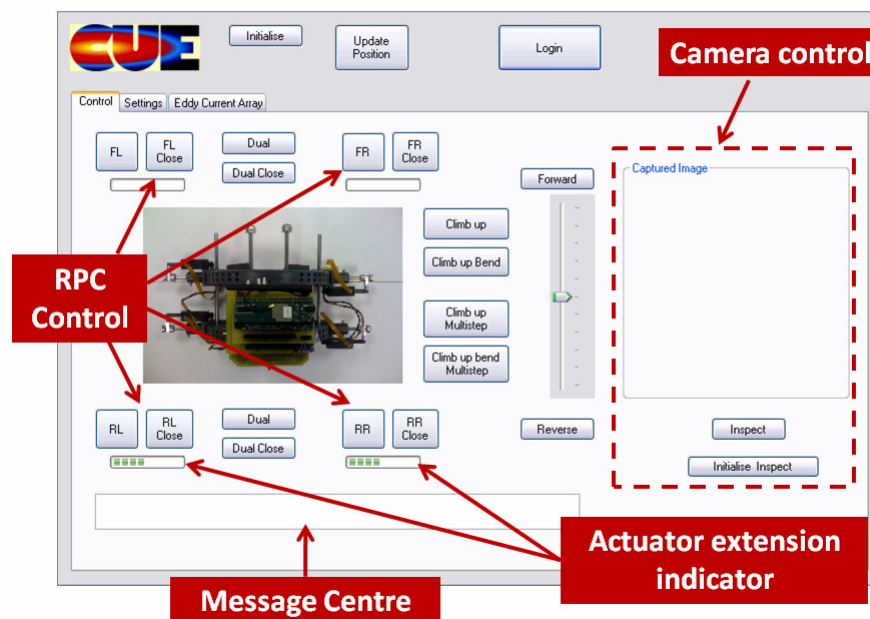


Figure 4.33: Robot control interface with visual inspection support.

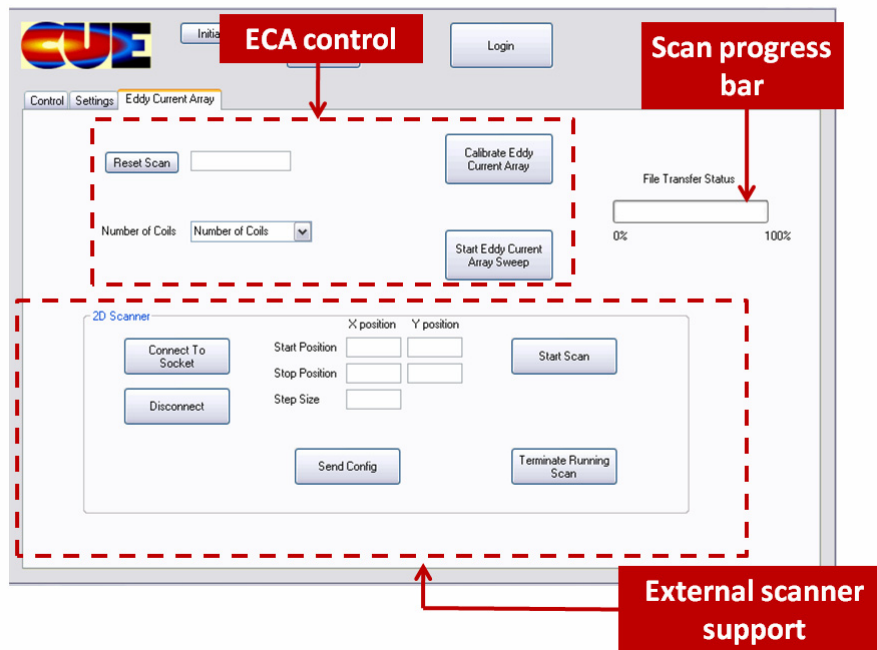


Figure 4.34: Eddy current array control interface featuring support for precise ECA control using an external, two axis linear stage.

Chapter 5

Pattern Recognition for NDE Images

5.1 Introduction

The aim of this Chapter is to determine if the techniques underlying the Eigenfaces [56] algorithm can be applied to create a framework for detecting defects in NDE images. The novelty of this work is primarily in its application to NDE images. An explanation of the key stages involved in pattern recognition are first discussed followed by an explanation of the theory behind principal component analysis; the latter being fundamental to the eigenface algorithm. A detailed explanation of the proposed system is then outlined followed by an optimisation of the system parameters. The performance of the optimised system is then benchmarked against a KNN implementation for synthetic and real corrosion defects.

5.2 Pattern Recognition Methods

5.2.1 Introduction

The complexity of pattern recognition systems has increased with the increasing processing power of microprocessors, due to their ability to perform very complex calculations quickly. Research in this area spans several decades [95]. Put simply, pattern recognition is the process of identifying sequences within a data set. The

three stages of pattern recognition design defined in [96] by Jain *et al* have been extended here to elaborate on the processes involved. That is:

- Data acquisition
- Pre-processing to convert data to a separable form
- Training of classifier
- Decision making
- Data fusion

The first two steps are rather self-explanatory, however it is perhaps worthwhile elaborating a little on the rest. Classifier training is a critical step and is required in order to teach the chosen classifier what patterns are to be identified and/or to expose the classifier to a representative sample of data. The decision making stage is the point at which the classification is made based on the input data provided. Several metrics may be used in this process however, in general, all will require the comparison of the data representation to a defined decision threshold. Data fusion, in this context, is the process of combining two or more pattern recognition systems together to enhance the probability of detection of the pattern(s) of interest and is common in situations where no single classifier provides sufficiently good results. This subject and its application will be discussed in detail in Section 5.4.6. Data fusion can also refer to the combination of two or more data sets and will be considered in this way throughout Chapter 7.

5.2.2 Feature Selection and Extraction

A feature is a variable and element of a larger feature set (or ‘vector’) to which an input pattern can be mapped. Using this feature set, a pattern of interest can be expressed as a weighted combination of the individual features. The term ‘feature’ in image processing is used to describe a wide range of entities ranging from individual pixels, to particular shapes or collection of pixels, to statistical parameters all of which can be used to describe the input image. The motivation for defining a pattern in this way is primarily to reduce the computational requirements arising from the high number of variables (or *dimensionality*) of the input data. A popular method for feature extraction is principal component analysis which will be discussed in detail in Section 5.2.4. It is clear that as the feature set is made simpler and smaller there comes a point at which the approximation of the original data by the feature set is too great and the classification ability of the subsequent algorithm will suffer as a consequence. It is perhaps less intuitive that the same may be said of a feature set which is too large. The ominously termed ‘*curse*

of dimensionality', coined by the mathematician Richard Bellman, is used to describe the problem of the exponential relationship between the number of variables in the input data and the volume of data required to train a system to classify such a pattern. This can be illustrated by first considering a single variable which can have ten discrete values. If a second variable is added which is again able to have ten values, there are now a total of $10 \times 10 = 100$ possible values that the two-variable system can have compared with only 10 for the single variable case. Thus for an increase in dimensionality from one to two, the number of samples required to cover 100% of the data points increases from 10 to 10^2 .

5.2.3 Types of Classifiers

The four main families of classification systems can be described as template matching, statistical classification, structural classification and neural networks. Template matching involves using a known template and attempting to find this in an input pattern through correlation. Such a process suffers from poor performance for instances where distortion of the pattern has occurred or where there is a difference in perspective. Statistical classifiers represent each pattern by a series of features and classification is determined by the probability of membership of the classes. Structural classification is a technique which is useful for patterns which can be formed from smaller building blocks or 'primitives'. Handwriting recognition is an area where such classifiers are useful [97]. An artificial neural network (ANN) is a classification system comprising of an interconnection of adaptive elements which seeks to mimic the structure and functionality of biological neurons. The weights of the interconnections of each of the artificial 'neurons' are modified based on a set of input training data in order to achieve a desired output, while minimising a particular cost function. Following the training process, the predictive 'power' of the neural network can then be used to classify new test data. The misuse of ANNs has had the unfortunate consequence of creating a rather negative view of pattern recognition as a whole. The main issue with all pattern recognition problems is in ensuring that the performance that is being observed is the result of an underlying pattern and not just a response to noise. Statistical classifiers will be used in this research because they allow their performance to be traced back to the underlying statistics of the data set. It is important, from the outset, to underline the very clear distinction between a neural network approach to classification problems and the statistical classification approach that will be outlined in the following sections.

The reasons for the criticism are well founded and are largely the result of overstated claims originating from an over-eagerness to confirm personal bias' towards this technology. In their (justifiably) scathing paper on the misuse of neural networks for the prognostic and diagnostic classification of tumours [98], Schwarzer *et al* provide an excellent summary of the common mistakes made. It concludes that often there is a fundamental misunderstanding of the mathematics of ANNs and of basic statistical theory. The principal problem with neural networks lies in the fact that they have a fundamentally non-linear transfer function. Consequently, they are *very* sensitive to initial conditions. Consider two identical networks trained using two identical sets of training data and with two identical desired outputs. Now, if the two networks are given different initial weights for the interconnections (prior to training), it is *highly* likely that the resulting network (following training) will perform differently on an independent test set. Furthermore, their 'black box' nature [99] means that there is no clear link between the weights of the interconnections and the actual classification performance. Thus, the ANN approach boils down to adjusting a set of weights to achieve the known correct classification on a given training set. The difference between this and the statistical approach outlined in the remainder of this Chapter, is that, in the latter, there is no adjustment of some variable in order to obtain a desired result. Quite simply, if there is no statistical separation *in the data itself*, then there can be no system designed to classify them correctly. This is crucially different to an ANN which *insists* that, given a large enough network, there must be a way to separate them. By applying enough neurons, the weights of the connections can be adjusted to *ensure* that the data is classified, however this will only work on the specific training set used. Introduce a new piece of data to the network and ask it to perform a classification and it will, in all likelihood, fail completely. This is entirely to be expected since what has actually happened is that the network has adapted specifically to the training data and not to an underlying differentiable variable. In the classifier presented here, by inspecting the data itself, it is possible to check whether there is actually an underlying statistical difference between the classes and, from this, derive confidence in the resulting performance.

A classifier may be defined as supervised or unsupervised depending on whether an input pattern is classified to one of a known set of predefined classes or whether the classes are learned by the system based on similarities with other patterns in a given data set respectively. The use of unsupervised learning, or clustering, is an important technique which has found widespread use in fields such as data mining [100] however for the purposes of NDE, as the defects of interest are known in advance, supervised classifiers are more suitable.

Classifier outputs can be described as either hard or soft depending on whether they output class labels or probabilities of class membership respectively. Xu *et al* [54] categorised the information output by individual classifiers to one of three groups:

- i. The *abstract* level: the classifier outputs a class label based on an assessment of the input pattern.
- ii. The *rank* level: the classifier outputs a list of the available class labels, ordered in terms of their relative likelihood.
- iii. The *measurement* level: the classifier outputs a confidence level for each of the available classes.

The abstract and rank levels are examples of a hard outputs while the measurement level can be regarded as a soft decision. The system outlined in this Chapter will be of the measurement level.

5.2.4 Overview of Principal Component Analysis

5.2.4.1 Theory

The following section outlines the theory of principal component analysis (PCA) that underpins the eigenface approach to pattern recognition. PCA is a statistical feature extraction technique which can be used to reduce the dimensionality of a multivariate data set. Such dimensionality reduction is achieved first through a transformation of the original data set to a new set of uncorrelated variables which are ordered according to their contribution to the total variance of the data set. Subsequently, it is possible to select only the components which contribute most




Figure 5.1: (a) x-y scatter plot of two variable (x_1 and x_2) data set. (b) axis x_1' selected to coincide with the vector of maximum variance of the original data in (a) and axis x_2' defined as perpendicular to x_1' . If the variance in the direction of x_2' is relatively small compared with x_1' then it can be assumed to be negligible thus reducing the dimensionality of the data to one. In a 2D case such as in this example, this reduces to a simple best-fit line as shown in (c).

significantly to the total variance and redefine the original data in terms of this lower dimensional basis set. As with any dimensionality reduction process there is an inevitable loss of information however PCA allows the loss of information to be quantified and controlled.

As an example, consider Figure 5.1(a) which shows an x-y scatter plot of a data set consisting of two variables x_1 and x_2 . By inspection it can be seen that there is a high level of correlation between the two variables, however the main variation in the data is along a vector roughly 45° from x_1 . If this vector is defined as x_1' and a second vector, x_2' , is defined as being perpendicular to it, the data can be viewed relative to these new axes and it can be seen, from Figure 5.1(b), that the level of variation in the x_2' direction is much smaller than in the x_1' direction. The dimensionality of the data can therefore be reduced by assuming that the variance in the x_2' direction is zero and thus entirely in the x_1' direction. For the simple 2D case, it can be seen from (c) that this reduces to a best-fit line approach, however the fundamentals can be derived using PCA.

The fundamental theory of PCA is relatively straightforward and can be regarded as a simple translation/rotation of the co-ordinate system of the original data set. More formally it can be expressed mathematically as an optimisation problem to find a set of orthonormal vectors that maximise the variance of a data set, x , of length m . The principal components pc_k where $k = 1, \dots, m$ of x can be defined as the functions $a_k^T x$ (for $k = 1, \dots, m$) which maximise the variance of x and which are uncorrelated with $a_j^T x$ for $j \neq k$. The latter corresponds to stating that the covariance of x , $\Sigma_{jk} = 0$ for $j \neq k$. The variance Var_k of x , in the direction k , is given by Equation (5.1)

$$Var_k = a_k^T \Sigma a_k. \quad (5.1)$$

The goal of PCA is first to maximise Equation (5.1) however, by inspection, it is clear that Var_k may be increased by scaling a_k therefore a normalising constraint must be imposed which, for simplicity, is typically the unit vector constraint $a_k^T a_k = 1$. Using the Lagrange multiplier method, the problem can be expressed in the form $\Lambda(p, \lambda) = f(p) - \lambda(g(p) - c)$ where $f(p)$ is the function to be optimised, $g(p) - c = 0$ is the constraint and λ is the Lagrange multiplier. This is demonstrated in Equation (5.2) for the first principal component.

$$\Lambda(a_1, \lambda) = a_1^T \Sigma a_1 - \lambda(a_1^T a_1 - 1) \quad (5.2)$$

Differentiating Equation (5.2) with respect to a_1 and equating to zero yields

$$\begin{aligned} 2\Sigma a_1 - 2\lambda a_1 &= 0 \\ (\Sigma - \lambda I)a_1 &= 0 \end{aligned}$$

which is the standard eigenvalue equation. For a non-trivial solution, λ must be an eigenvalue of Σ with a_1 its associated eigenvector. a_1 is selected as the eigenvector which has the largest corresponding eigenvalue λ_1 and hence variance. The remaining principal components can be calculated using the same method but with the additional constraint that the covariance between each new component and all previous components is zero. In so doing, a set of principal components $pc_k = \{a_1^T x, a_2^T x \dots a_m^T x\}$ with corresponding variances $\{\lambda_1, \lambda_2 \dots \lambda_k\}$ is produced. The eigenvectors are then sorted in descending order of variance such that the a_1 is the greatest and a_m the lowest. At this stage, no dimensionality reduction has been achieved. However, if only r of the m principal components are retained (where $r < m$) and the data is re-expressed according to this reduced feature set, then the dimension of the subsequent data will also be r . The compressed data is now expressed using the r most significant principal components by variance. In image processing, the term ‘weight vector’ is used to define the relative proportions of each of these components that are required to represent each image

For trivial examples, such as that in Figure 5.1, reducing the dimensionality by selecting the best-fit line for the data is obvious and does not require the problem to be considered in terms of maximising the variance. Where PCA offers an important contribution is when the data extends beyond the 2D case to higher orders of dimensionality. A typical image from a miniature camera may contain 640 x 480 pixels which means there are 307,200 variables. With the exception of an image containing white noise, the majority of the data is interrelated in some way and therefore PCA makes it possible to extract the relevant data and re-express these images using a set of features defined by a subset of the principal components.

In order to perform PCA on a set of images, each array of $M \times N$ pixels is first reshaped into a $MN \times 1$ vector Γ_n . Given the mean of the image set, ψ , the covariance matrix of the D images in the data set can be derived from [56] to be:

$$\begin{aligned} \Sigma &= \frac{1}{D} \sum_{n=1}^D (\Gamma_n - \psi)(\Gamma_n - \psi)^T \\ &= AA^T \end{aligned} \tag{5.3}$$

where $A = [\Phi_1, \Phi_2 \dots \Phi_D]$ and $\Phi_n = \Gamma_n - \psi$. However, due to the length of vector Γ_n , the dimensions of Σ are $MN \times MN$, which is computationally intractable for typical image sizes. It was shown in [56] that if the number of images in the data set, D , is less than the dimension of the image space ($D < MN$) then there will exist only $D - 1$ rather than MN useful eigenvectors (and hence principal components) the rest of which will have corresponding eigenvalues of zero and hence no contribution to the variance. The more computationally efficient method outlined there-in, involves first calculating the eigenvectors of the $D \times D$ matrix $L = A^T A$ which is shown to be related to the eigenvectors of AA^T by:

$$\begin{aligned} A^T A v_i &= \lambda_i v_i \\ (AA^T) A v_i &= \lambda_i (A v_i) \end{aligned} \quad (5.4)$$

Thus by pre-multiplying both sides of the eigenvector equation for L by A , it can be seen that the eigenvectors, u_i , of the original covariance matrix Σ are given by $A v_i$ and hence by Equation (5.5).

$$u_i = \sum_{k=1}^D v_{ik} \Phi_k \quad (5.5)$$

This computationally efficient method for calculating the covariance opens the power of PCA for use with typical image sizes.

5.3 Development of Visual Inspection System

5.3.1 Image Pre-processing

5.3.1.1 Illumination

Due to the curvature of the pipe, the reflected light detected by the camera will appear stronger down one column of the image, with the exact pixels involved dependant on the position and nature of the light source. Such an irregular illumination effect will generate a region of contrast which may yield similar intensity gradients between pixels as those due to defects. It should be assumed that there will be no ambient light within an enclosed processing cell and thus on-board illumination should be incorporated. Light emitting diodes (LEDs) are a simple and common analogue light source with a variety of wavelengths commercially available. The illumination may be made more uniform through either multiple lights sources, appropriately placed mirrored surfaces or an optically diffuse housing. Multiple light sources generate multiple instances of the light effect outlined earlier and so is not suitable without further intervention. The flat surface

of a mirror is not ideal since this will simply lead to more discrete regions of high intensity, similar to multiple light sources, rather than producing an even illumination. The nature of a diffuse housing is such that light will be reflected at a multitude of angles making it ideal for the purpose required. Thus, the camera used for image acquisition was placed within a custom housing manufactured using crushed aluminium foil (providing the diffuse surface) and illuminated using three white LEDs. It was observed that this setup yielded the most even illumination of the methods outlined.

5.3.1.2 Grayscale Image

The use of PCA for dimensionality reduction relies on the fact that there are underlying statistical similarities in the data set. The use of raw or filtered grayscale images as inputs to the algorithm suffers from the problem that each image is relatively unique in that, while several may contain a visually similar defect, the pixels that it occupies may be different. Consider, for instance, the images shown in Figure 5.2 (a), (b) and (c). The black line is intended to illustrate a long, thin crack in an otherwise sound piece of material with the pixel grid superimposed on top. It can be seen that the defect is 1x9 pixels in length in all three images however its position within the image is not consistent. Although these three images all belong to the same class of defect (i.e the ‘long, thin crack’ class) none of the defect pixels in the raw image represent defect pixels in either of the other two images. As was discussed in Section 5.2.2, a system using such an input would require an intractable volume of training data to cover all possible defect locations and sizes.




Figure 5.2: Illustration of how the use of the raw images leads to problems with reducing the dimensionality due to lack of similarity in the underlying data. Each image shows a simplified representation of a long, thin crack however at different locations in (a), (b) and (c). It can be seen that in each of the 3 images, none of the pixels shown in red, corresponding to the previous ‘defects’ are the same as in the associated defect image.

Furthermore, such a system would still lack the uniqueness of pixel values for each defect type to allow successful recognition. In other words, the variance within each class of images would be so great that it would be impossible to distinguish them from another class. Such sensitivity to the position of the artefacts requires the implementation of a pre-processing stage to re-express the original data in a more suitable form. There are a number of techniques available that would allow frequency selection as desired for the defect characterisation outlined in Chapter 3 and one such technique is the Fourier Transform.

5.3.1.3 Spatial Frequency Spectrum

Fourier theory can be applied to images to represent them in terms of their frequency and phase components. The magnitude spectrum carries information which defines the change in intensities of objects within the image while the phase contains information about where the objects are located in the image. The 2D Discrete Fourier Transform (DFT) is given by Equation (5.6)

$$F(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x, y) e^{-j2\pi\left(\frac{ux}{M} + \frac{vy}{N}\right)} \quad (5.6)$$

where u and v are spatial frequencies in the x and y directions respectively and $F(u, v)$ is the 2D spectrum of the image defined by $f(x, y)$. Expressed in polar form this becomes:

$$F(u, v) = |F(u, v)| e^{j\Phi(u, v)} \quad (5.7)$$




Figure 5.3: (a) and (b) show the magnitude (rad/pixel) of the FFT of the images in Figure 5.2 (a) and (b) respectively. It can be seen that the magnitude spectrum for each image is identical. (c) shows the effect of rotating the defect in Figure 5.2 (a) prior to calculating its FFT. It can be seen that there is a corresponding rotation in the magnitude spectrum.

The DFT of an image, $f(x - x_0, y - y_0)$, which is the same as the original image but translated by x_0 and y_0 in the x and y directions respectively, can be calculated as follows:

$$F_2(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x - x_0, y - y_0) e^{-j2\pi\left(\frac{ux}{M} + \frac{vy}{N}\right)} \quad (5.8)$$

Using the substitutions $j = x - x_0$ and $k = y - y_0$, Equation (5.8) yields:

$$\begin{aligned} F_2(u, v) &= \sum_{j=-x_0}^{M-1-x_0} \sum_{k=-y_0}^{N-1-y_0} f(j, k) e^{-j2\pi\left(\frac{u(j+x_0)}{M} + \frac{v(k+y_0)}{N}\right)} \\ &= \sum_{j=0}^{M-1} \sum_{k=0}^{N-1} f(j, k) e^{-j2\pi\left(\frac{uj}{M} + \frac{vk}{N}\right)} e^{-j2\pi\left(\frac{ux_0}{M} + \frac{vy_0}{N}\right)} \end{aligned}$$

which simplifies to:

$$F_2(u, v) = F(u, v) e^{-j2\pi\left(\frac{ux_0}{M} + \frac{vy_0}{N}\right)} \quad (5.9)$$

It can be seen from Equation (5.9) that $|F_2(u, v)| = |F(u, v)|$ and the translation has therefore appeared solely as a change in phase. This translation invariant property of the DFT is important since, unlike the raw image data, it allows the position of a defect to be eliminated as a variable in the input data to the PCA algorithm. Consequently a particular class of defect can be defined without concern of its position within the images. The Fourier transform is not rotation invariant and therefore defects of the same dimensions but different orientation result in unique magnitude spectra. The effect of translation and rotation is illustrated in Figure 5.3. In Figure 5.3 (c) the additional harmonics that can be seen are due to aliasing of the image edges but the principle can still be understood.

5.3.1.4 Frequency Selection

The $F(0,0)$ frequency of an image is usually referred to as the DC component due to the parallels with the Fourier transform of electrical signals and defines the average intensity of the image. In most images, this component is significantly greater than the higher frequencies and consequently tends to dominate the magnitude spectrum. Given a set of images, the variation in the DC component will be the most significant even though it carries no information regarding the detail of the image including any defects that may be present. Since PCA calculates the vectors of maximum variance, the presence of the DC component in the image set dictates that the percentage variance due to image detail is very small and




Figure 5.4: Normalised frequency response of image bandpass filter.

consequently the algorithm will perform poorly. Therefore, it should be removed and should not adversely affect the ability to correctly classify an image.

Based on the preceding arguments and the discussion of the frequencies of interest outlined in Chapter 3, a two-step pre-processing stage was designed. The first stage comprises of a finite impulse response (FIR) bandpass filter to allow the image frequencies to be selected. The second stage involves the calculation of the Fourier coefficients of the filtered image and the subsequent vectorisation of the image array to create a $1 \times MN$ vector. Through a trial and error approach it was found that a suitable frequency response was obtained for a filter order of 20 as shown in Figure 5.4. Following this stage, the image size is reduced by the filter order minus one in both dimensions.

5.3.1.5 Definition of Classes

The principal distinction required by the inspection algorithm is to differentiate between those images that contain defects and those that do not. As the size, shape and orientation of defects may vary, defining a single ‘defect’ class is not practical and consequently it is more suitable to define a ‘no-defect’ class since it can be assumed that the surface of the pipe will look reasonably similar in the absence of a defect. The main defects expected are cracks and corrosion pits. While it is




Figure 5.5: Aluminium pipe (40mm diameter) containing synthetic defects. (a) HC class, (b) DC class and (c) CP class.

principally the case that any cracks will be in the circumferential direction and hence the horizontal direction of the plane of the image, other orientations may be encountered which, as was discussed in Section 5.3.1, would provide unique magnitude spectra. It is impractical to have a class defined for every orientation of crack, however it is possible to define a discrete number of crack classes corresponding to, for instance, circumferential and 45° trans-circumferential cracks. For convenience, these will be referred to as ‘horizontal’ and ‘diagonal’ cracks respectively, in recognition of their appearance within the image frame. Thus, four classes are defined: *no defect* (ND), *corrosion pit* (CP), *horizontal crack* (HC) and *diagonal crack* (DC).

As discussed in Section 3.4.2, in order to test the feasibility of the algorithm, a set of synthetic defects was created in an aluminium pipe, as shown in Figure 5.5, comprising of 12 mm drill impressions and slots oriented at 45° and 90° to the axis of the pipe. A data set of 180 images was then created with 45 images captured for each class. As there are only a limited number of synthetic defects, the same defect was used multiple times, however in different positions within the image, under

different lighting conditions and in some cases with the defect only partially in-frame in order to maintain a level of uniqueness to every image. Each class was defined by the mean and covariance of the training image set. The classification tree shown in Figure 5.6 illustrates the potential classification results. The binary classification system requires that an *unclassifiable* class be defined however this is explained further in Section 5.4.6.4.

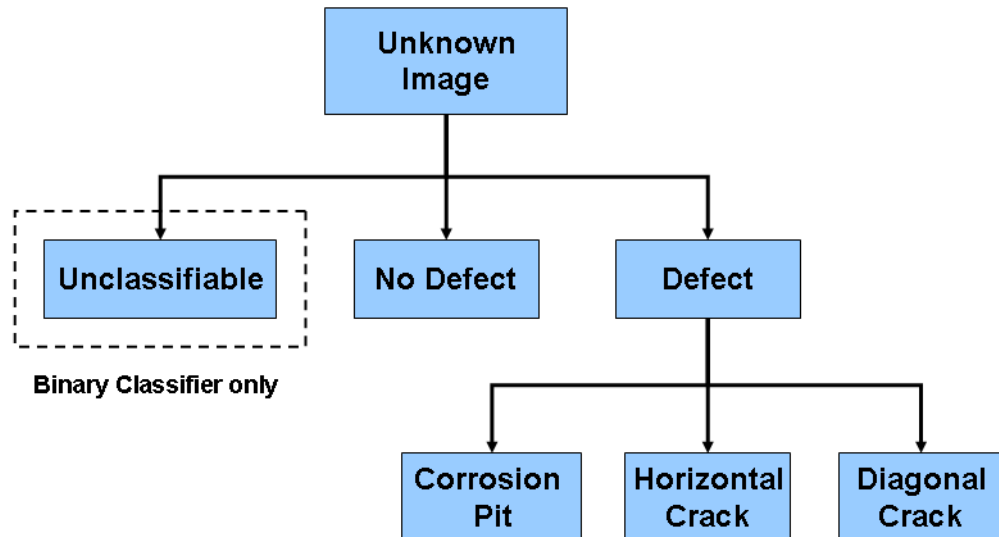


Figure 5.6: Classification tree showing the possible classification outcomes. An additional ‘unclassifiable’ class is defined for the binary classification system.

5.3.2 Classification Metrics

A classification metric is any measure from which the classification of a particular image can be determined. Three common metrics are the city block distance, Euclidean distance and Mahalanobis distance. Any of these classification measures can be fed into a classifier to determine to which class the corresponding image belongs.

The city block, or Manhattan, distance is so called because of the parallels between a two-dimensional problem to measure the distance that would be covered by travelling along the grid-like roads of a city such as New York. More specifically it can be thought of as the sum of the absolute distance between two co-ordinates. The Euclidean distance is a measure of the distance between two points in space. In the current context this specifically relates to the distance from the set of weights produced by the projection of the new image onto the feature space and the




Figure 5.7: Two distributions comprising of 2D representations of a number of images. The Euclidean distance, $d_1 < d_2$, indicates that the new image, I , is closer to the centroid of class c_1 however observation of the covariance of the two class sets shows that the image is much more likely to belong to class c_2 . The Mahalanobis distance accounts for such changes and is therefore better for distributions which are ellipsoidal in nature.

weights corresponding to a particular class and is essentially the Pythagorean extension for N -dimensional space. The Euclidean distance between a vector A_i and B_i where $i = 1, \dots, N$ is given by Equation (5.10).

$$d_{Euclidean}(A, B) = \sqrt{\left(\sum_{i=1}^N (A_i - B_i)^2 \right)} \quad (5.10)$$

In terms of N -dimensional space, the Euclidean distance defines the radius of a hypersphere. The Mahalanobis distance, commonly used in statistics, is a metric which considers both the mean and covariance of the data set and is given by Equation (5.11), where x is an image vector, μ is the mean vector of a known class and Σ is the covariance matrix of the known class.

$$d_{Mahalanobis}(x) = \sqrt{[(x - \mu)^T \Sigma^{-1} (x - \mu)]} \quad (5.11)$$

Because it considers the ‘shape’ of the distribution of known class images through the covariance term, when calculating the distance of a new image to that class, the Mahalanobis distance better reflects the actual distance in statistical terms. This point is demonstrated in Figure 5.7 where two distributions of 2D image weights exist with strongly elliptical patterns. If a new image, I , is to be classified to one of the two classes then, using the Euclidean distance metric to the centroid of the classes, it appears that the correct assignment would be to classify I as a member of class c_1 . However by taking account of the statistics of the two class data sets,

by using the Mahalanobis distance, it is clear that the image is actually more likely to belong to class c_2 .

The calculation of the Mahalanobis distance can be problematic since the required inversion of the covariance matrix can become intractable for high dimensions and if insufficient data is present to ensure it is well conditioned, a singular result may be obtained. As the dimensionality of the 2D FFT data is reduced through PCA, the use of Mahalanobis distance as a metric is both possible and preferable.

5.3.3 Statistical Classification: a Bayesian Approach

A Bayesian classifier is a statistical classification system which makes use of Bayes' rule given in Equation (5.12):

$$P(c|x) = \frac{p(x|c)P(c)}{p(x)} \quad (5.12)$$

where $P(c|x)$ is the *a posteriori* probability of the observed data, x , belonging to class c , $p(x|c)$ is the class-conditional probability density function (PDF) or *likelihood* for class c , $P(c)$ is the *a priori* probability of observing an image from class c and $p(x)$ is the probability density function for x . Bayes' rule changes a prior belief (i.e. before any data has been observed) into a posterior belief (i.e. after some data has been observed). Likelihood is a term that is often encountered in a Bayesian context. It is subtly different from *evidence* in that it defines how *likely* an observation of x is to occur, given a particular class, rather than the reverse [101].

The PDFs for a hypothetical two-class system with Gaussian distributions are shown in Figure 5.8(a). In any classification system there must be a threshold or decision rule by which an image can be classified to a particular class to the exclusion of the rest. In the context of a statistical system, it is somewhat intuitive that the classification should be based on classifying an image to the class which has the greatest posterior probability such that:

If $P(c_1|x) > P(c_2|x)$, then the image is classified as belonging to class c_1 .

If $P(c_1|x) < P(c_2|x)$, then the image is classified as belonging to class c_2 .

This leads to the definition of the optimum (or Bayes') decision threshold being defined by the intersection of the two PDFs for this two-class system. It is clear from the overlapping distributions that, even for this optimum boundary, there exists a finite probability that an image belonging to class c_2 could be classified as belonging to c_1 and vice-versa.

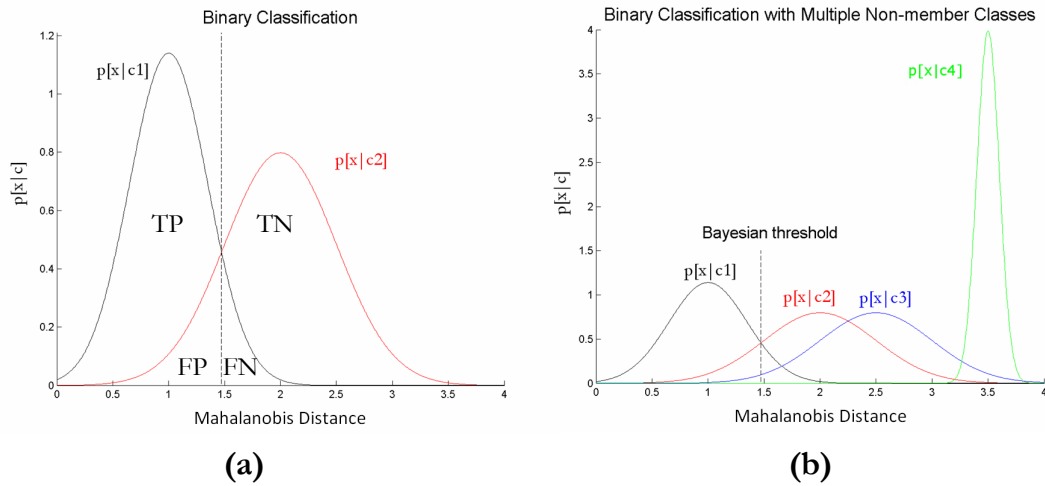


Figure 5.8: Probability distribution functions of (a) a two class system (member/non-member) showing the TP, FP, FN and TN regions separated by the Bayesian (optimum) threshold and (b) showing the binary classification principal in a four class system where the non-member class comprises of three classes (in this case c_2 , c_3 and c_4).

It is possible to sub-divide the possible classifications/misclassifications into four categories which, from the perspective of class c_1 , are as follows:

- True positive (TP) : Correctly classified as belonging to class c_1
- True negative (TN) : Correctly classified as belonging to class c_2
- False positive (FP) : Incorrectly classified as belonging to class c_1
- False negative (FN) : Incorrectly classified as belonging to class c_2

It can be shown that the Bayesian threshold is optimal in terms of minimising the probability of error and hence minimising the FP and FN regions [95].

When more than two classes exist, as in Figure 5.8(b), the definition of the classification categories becomes more complicated since the aforementioned regions can now be defined for each pair of density functions. However by considering the three non-member classes (in this case c_2 , c_3 and c_4) to be a subdivision of a single non-member class whose probability density function is a mixture of Gaussians the problem, again, reduces to one involving only two classes: *member* (c_1) and *non-member* (the rest). In the system proposed, the class with the maximum probability will be found and if this is any one of the non-member classes then it will be classified as a non-member.

5.3.4 Review of Proposed System

A functional block diagram of the proposed system is provided in Figure 5.9 to give a high-level overview of how the system operates. The training procedure is indicated by the flow of red, dashed arrows and normal operation (following training) is indicated by the solid blue arrows. Prior to first use, the system must be trained. Each member of the training image database is fed through the two pre-processing stages to yield their filtered, 2D FFT representations. The principal components of this training set are then extracted, a subset of these retained to form the feature vector and the training images then mapped to the feature space using this subset. The Mahalanobis distances for each image is then calculated and a set of distributions formed for each class. The mapping and distributions formation processes are explained in detail in Section 5.4.2. The system is then ready to be used for new images. Both the training and normal operation processes involve the same pre-processing stage and so each newly captured image is also fed through the two pre-processing stages. The image is then mapped to the feature space using the feature vector derived in training. The Mahalanobis distances are then calculated and the probability of class membership is obtained using Bayes' rule in conjunction with the set of known distributions obtained through training. For development and testing purposes, a validation and test set are used to optimise and independently quantify the performance of the system respectively.

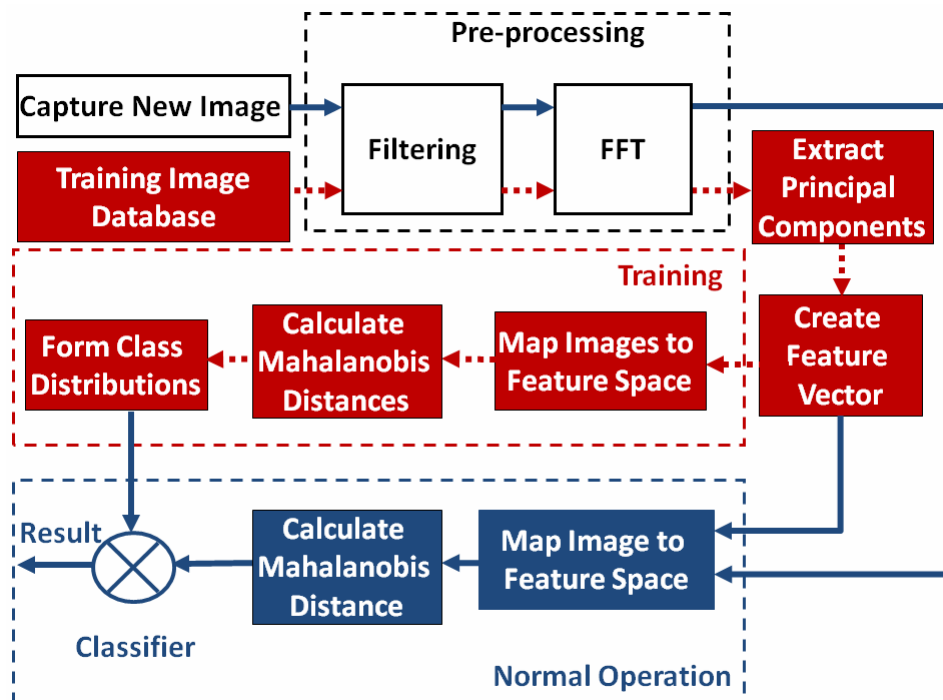


Figure 5.9: Functional block diagram of the pattern recognition system. Dashed, red and solid, blue lines represent training and normal operation respectively

Both of these sets follow the ‘normal operation’ route of Figure 5.9 and, unlike the training set, have fixed sizes of 10 and 5 respectively throughout.

5.3.5 Effect of Pre-Processing and Dimensionality Reduction

Figure 5.10 shows the mean 2D FFT magnitude calculated for each class, comprising of a set of 15 filtered training images, with no dimensionality reduction. There is good distinction between each of the four classes which satisfies the principal requirement for the design of an algorithm to successfully separate them. The effect of the pre-processing filter response can be seen, particularly on the *corrosion pit* class where the high frequencies at approximately 45° , 135° , 225° and 315° to the horizontal axis have been attenuated. The strongest inter-class similarities exist between the *no defect* and *horizontal crack* classes, where the chief difference is the relative magnitudes of the frequency components.

Figure 5.11 shows the effect of the pre-processing and dimensionality reduction stages through two step-by-step examples: one for an image with no defect and one with a drill hole (corrosion pit class). It can be seen from the comparison of (a) with (c) and (b) with (d) that the pre-processing filter helps to further reduce the effect of lighting variations across the image, from the on-board source, over and above that of the hardware method outlined earlier. In the *corrosion pit* image, it can be seen that the regions where there are some surface discontinuities, or significant changes in light intensity, have been retained while for the *no defect* image there remains some uncorrelated low frequency content with the central strip of vertical light still slightly visible. Comparison of the high dimensional 2D FFT data in (e) and (f) shows significant difference as would be expected from the results in Figure 5.10. The effect of reducing the dimensionality of the FFT magnitude data using two principal components is shown in (g) and (h) with certain frequency components of the original now eliminated, particularly in the *no defect* image. It should be noted that the magnitude data is now bipolar due to the mean adjustment of the PCA algorithm. For completeness, although it is not part of the classification process, both images have been regenerated from the dimensionally reduced frequency data and are shown in (i) and (j). There is no significant difference between the filtered originals of either image and the reconstructions with the defect in the corrosion pit image still clearly visible.





Figure 5.10: Mean 2D FFT magnitude (uncompressed) of the filtered training images belonging to (a) *no defect* class, (b) *horizontal crack* class, (c) *corrosion pit* class and (d) *diagonal crack* class.




(a)




(b)




(c)




(d)




(e)




(f)



(g)



(h)



(i)



(j)

Figure 5.11: Pre-processing and PCA stages. Respectively (a, b) show original images from the *no defect* and *corrosion pit* classes, (c, d) show the filtered images, (e, f) show the 2D FFT of the filtered images, (g, h) show the 2D FFT magnitudes using 2 principal components and (i, j) show the reconstructed images.

5.4 Binary Classification Method

5.4.1 Introduction

A binary classification system is one which categorises data as either belonging to a particular class or not: ‘The car *is red* or the car is *not red*’. There are two principal classification methods by which such classifiers can be used to determine membership in multi-class problems: *one-vs.-all* or *one-vs.-one*. The *one-vs.-all* (OVA) approach involves comparing each class against all the remaining classes, each time determining class membership or non-membership. Using the example of coloured cars once again, this may be illustrated for a three class problem as: ‘The car is *not red*, the car is *not yellow*. The car *is blue*’. It can be seen that this fits with the system outlined in the previous section since it equates to one class (c_1) against the three remaining classes. An alternative to OVA is the *one-vs.-one* approach where, instead of each class being compared against all the remaining classes, a pair-wise system is adopted and each classifier is trained to differentiate one class from a single other class. Such a system requires mutual exclusivity between the defined classes. Although there is no theoretical basis to suggest which approach is better in terms of maximising the classification rate, it is clear that the training and computation requirements are significantly greater for the latter due to the greater number of individual classifiers. Thus, in this instance, the OVA approach is favourable.

Recalling the specification outlined in Section 3.4.2 and the class definitions in Section 5.3.1.5, it can be seen that the ND classifier represents the defect/no-defect classifier. If an image is not classified as being ‘no-defect’, then it is assumed to have one. Therefore, the primary target is to achieve 100% classification with this classifier. The remaining binary classifiers (HC, CP and DC) can be combined later to try and yield a more specific result. The subject of classifier combination is covered in detail, a little later, in Section 5.4.6.

5.4.2 Training Process

The training data comprises of K images for each of the four defined classes, with K being a tunable system parameter. Following the pre-processing stage, the 2D FFT data for each of the $4K$, M -by- N pixel images are vectorised into $MN \times 1$ vectors and subsequently combined to form the $MN \times 4K$, ‘A’ matrix of Equation (5.3). The resolution of the camera used for the image capture is adjustable however 320 by 240 pixels was used in order to provide a good compromise between high image clarity and low computational expense. The

principal components are then calculated using the PCA algorithm outlined in Section 5.2.4 for this training set and a feature vector ($MN \times r$) is generated using the first r of these. Each training image is then re-expressed as a set of r weights into a ‘weight vector’ ($r \times 1$) using Equation (5.13):

$$WeightVector = FeatureVector^T \times FFTdata \quad (5.13)$$

This is essentially a process of mapping the 2D FFT data onto the feature space resulting in an output vector which defines the original data as a weighting of the feature vector. As the feature vector is composed of the principal components (which are unit vectors), the units of the WeightVector are the same as that of the original data (i.e. rad/pixel). Having constructed a representation of all images in terms of the relative weights of the feature vector, the *a priori* knowledge of each image’s parent class can be exploited. The Mahalanobis distance, x_i , for each of the training images, i , is calculated to a particular class, c_i , using the covariance and mean data relating to that class in Equation (5.11). Using the *a priori* knowledge of class membership, this yields a set of four probability density functions which describes the distribution of Mahalanobis distances for the four class types (ND, HC, CP and DC) when measured against class c_i . This process is repeated for each of the remaining three classes providing, in each case, a further four PDFs. The critical assumption made here is that images belonging to a particular class will tend to have similar Mahalanobis distances to a particular class and that the distribution of these distances will be Gaussian. Assuming a Gaussian distribution, the class-conditional PDF of the distances of images I belonging to the set of images c_i to the class c_j , $x_{I \in c_i, c_j}$, given class c_i is obtained through Equation (5.14):

$$p(x_{I \in c_i, c_j} | c_i) = \frac{1}{\sqrt{2\pi\sigma_{c_i, c_j}^2}} e^{-\left(\frac{(x_{I \in c_i, c_j} - \mu_{c_i, c_j})^2}{2\sigma_{c_i, c_j}^2}\right)} \quad (5.14)$$

where $i = 1 \dots 4$, $j = 1 \dots 4$, μ_{c_i, c_j} is the mean of the Mahalanobis distances of the images, known to be from class c_i to class c_j and σ_{c_i, c_j}^2 is the variance. The definition of these density functions along with the creation of the feature vector are the goals of the training process.

5.4.3 Validation/Test Process

As with the training stage, each newly captured image is first pre-processed to obtain the filtered, 2D frequency magnitude data. Following this, the feature vector defined in the training process is used in conjunction with Equation (5.13) to obtain a best approximation of the validation/test data defined in terms of weightings of this feature vector. The Mahalanobis distance to each of the four known classes is then calculated using the mean and covariance data derived from the training data. Using the PDFs obtained from the training process, the posterior probability, $P(c_i | x)$, of the new image belonging to the class c_i given Mahalanobis distance, x , can then be estimated. A [1/0] decision rule is applied to convert the soft outputs of the classifier to hard outputs, as shown in Equation (5.15), however the probabilities are retained as a confidence measure.

$$h(x) = \begin{cases} 1 & P(c_1|x) > P(c_2|x) \\ 0 & \text{else} \end{cases} \quad (5.15)$$

This can be paraphrased by stating that if the probability of the image belonging to class c_1 is greater than for c_2 then a logic 1 classification result will be returned, otherwise logic 0 will be returned. For the two-class system, classes c_1 and c_2 are redefined as *member* and *non-member* respectively. In total, four binary classifiers are used in the current system and the class which returns a logic 1 is the one to which the image should be assigned. Since each classification is decoupled from the rest, clearly there is the potential for conflicting classification results where an image is classified to more than one class (i.e. a ‘tie’). A thorough discussion of classifier combination techniques is given in Section 5.4.6 which addresses this issue.

5.4.4 Optimising Classification Parameters

The optimum operating point will be that which maximises the probability of detecting the particular defects of interest. As stated, this can be estimated for a multi-class problem by using the optimum Bayesian threshold and thus by minimising the error regions (FP and FN) or maximising the separation of the means. An empirical result can also be obtained by maximising the correct classification rate on the validation set.

5.4.4.1 Influence of Image Bandwidth

The lower frequencies within an image typically contain information from which the presence of defects of reasonable size (i.e. spanning several pixels) can be identified by the naked eye, while the higher frequencies contain finer details such

as the material grain structure and small surface scratches. In order to assess the effect of the higher frequency components on the classification performance, a bandwidth sensitivity analysis was conducted. The goal of this study was to determine to what extent the presence of the higher frequencies affected the performance of the PCA algorithm to correctly categorise images to the classes outlined in Section 5.3.1.5 and to determine if there exists a preferential bandwidth for the pre-processing filter which maximises the probability of correct classification for the type of defects used in this study.

In all cases the DC component of the images was removed since it does not convey any information regarding the nature of the defects. The lower cut-off frequency of the image bandpass filter was set to $0.01 \frac{F_S}{2}$ which ensured that the majority of the lower frequency content above DC was retained while the upper cut-off was varied from $0.1 \frac{F_S}{2}$ to $\frac{F_S}{2}$, where F_S is the sample frequency. For convenience in discussing the results, the following text will assume the lower cut-off to be negligibly greater than DC and define the bandwidth as the same as the value of the upper cut-off frequency. Thus, the frequency range between $0.01 \frac{F_S}{2}$ and $0.1 \frac{F_S}{2}$ will be referred to as having a bandwidth of $0.1 \frac{F_S}{2}$ although it will be understood that the actual value is $0.09 \frac{F_S}{2}$.

A data set containing $K = 15$ images of each of the four classes was arbitrarily selected to train the system and the same images were used for each frequency interval. This is assessed and refined in the next section. The bandwidth was varied in increments of $0.05 \frac{F_S}{2}$, resulting in a total of 19 frequency sets between the stated bandwidth limits. Note: Given the limited nature of the data set, there is an inherent limit to the dimensionality (i.e. number of principal components) that can be used reliably. This is due to the limited number of data points available due to the finite number of images available. This causes the covariance matrix to become poorly conditioned and hence, following its inversion during calculation of the Mahalanobis distance, tend towards a singular result. It was found experimentally that if N images per class were used, then the covariance matrix would be poorly defined if more than $N - 1$ principal components were used. Therefore at each frequency interval, the system was tested using from 1 up to 14 principal components producing a total of 266 data sets. The procedure for the bandwidth sensitivity study can be summarised by the following pseudo-code:

- For bandwidth < max(bandwidth)
 - i. Perform pre-processing on all images at given bandwidth
 - For number of PCs < max(number of PCs)
 - ii. Run training algorithm

iii. Calculate classification rate for each binary classifier

iv. Increment number of PCs

Loop

v. increment bandwidth

Loop

The mean correct classification rate across the four binary classifiers using the validation set is shown in Figure 5.12 for one to fifteen principal components. The result indicates that the classification rate is relatively insensitive to changes in bandwidth (albeit with a marginal decline as more frequencies are retained). The maximum variation is approximately $\pm 5\%$ and it is in fact far more influenced by the level of dimensionality. On this basis it can be concluded that, for the given image set, the bandwidth choice is relatively unimportant and so $0.1 \frac{F_s}{2}$ was used for this classification system. The dimensionality of the data set, however, is likely to play a more important role in terms of performance.




Figure 5.12: Bandwidth sensitivity: Mean correct classification rate obtained using the validation set for varying dimensionality (1 to 14 PCs).

5.4.4.2 Sensitivity to Volume of Training Data

This section investigates the sensitivity of the classification rate to the number of training images provided. As outlined in the previous section, it was found that if N images per class were used, then in order to ensure the covariance matrix is not poorly defined, a maximum of $N - 1$ principal components can be used. Thus the investigation was limited to between 1 and $N - 1$ principal components with N varying from 10 to 30. A lower limit of 10 was selected so as to allow sufficient scope for intra-class variations to be accounted for and an upper limit of 30 used so as to allow a 10 image validation and 5 image test set.

The classification rate has been broken down, this time, so as to allow the analysis of the individual classification rates for each particular class of images. In subsequent graphs, the legend *ND, HC...etc* indicates the classification rate obtained by that particular binary classifier. If there are $D = 4N$ images in the training image set, then there will exist $D - 4$ useful principal components. Thus, as the volume of training images used increases, so too will the number of useful principal components. In order to allow comparison between the different results, the classification performance is compared against the percentage of the total variance of the data set retained by the principal components rather than the actual number of principal components themselves. Therefore it is possible to say that using a particular number of principal components which will result in a total of $P\%$ of the overall variance being retained, a classification rate of C will be achieved.

The procedure for this study is summarised by the following pseudo-code:

- i. Perform pre-processing on all images at given bandwidth
For number of training images < max(number of training images)
For number of PCs < max(number of PCs)
 - ii. Run training algorithm
 - iii. Calculate classification rate for each binary classifier
 - iv. Increment number of PCsLoop
 - v. Increment number of training images
- Loop

Figure 5.13 (a) to (d) shows the results of this study. It can be seen that, as the percentage of the variance covered increases above $\sim 50\%$, the classification performance actually decreases. This suggests that the data present in the higher principal components (lower contribution to the total variance) is less important in terms of identifying the type of defect and may be more concerned with the finer


details within an image. It was observed that a variance of approximately 40-50% resulted in the highest classification rate and this generally corresponded to using two principal components. Figure 5.14 (a) shows the influence of the training volume on the classification rate for two principal components. It can be seen that there is no distinct optimum for each class, so a trade-off is required which provides acceptable performance across the four classes and hence maximises the performance of the ensemble of classifiers. As the *defect/no-defect* classification is the most important, it is best to pick a training volume that first maximises this, and then subsequently all others as a secondary goal. It can be seen that using more than 14 images results in 100% classification of the validation set for the ND class. This is a highly significant result and indicates that the primary goal set out in the specification may be achievable providing the result can be shown to be sufficiently generalised. Using 24-27 training images gives the best performance across the three remaining classes and 25 was selected as a suitable overall operating point.

The bandwidth sensitivity study was repeated using 25 training images (rather than 15 as previously) and the results are shown Figure 5.14 (b). The results are broadly in agreement with the earlier study and the conclusions drawn on the relative importance of the bandwidth are therefore still valid.


The effectiveness of this result alone in terms of tuning the performance is dubious since overtraining the system is a danger and could result in a system which is poorly generalised. It is of critical importance to identify how the performance of the system varies with the images that are presented for training purposes while keeping all other system parameters constant. This *generalisation* issue is the subject of Section 5.4.4.3.

5.4.4.3 Generalisation

It is important to have a representative collection of images which encompasses all the variations that may occur for a particular type of defect. If images containing particular features found in a subset of a given defect class are not included as training data then such features would be expected to be eliminated when test data is expressed in terms of the resultant feature vector. Put simply, if a feature isn't in the training images then it cannot be retained in any future test images constructed from the resulting feature vector. Such approximations may then lead to classification errors and it is therefore important to determine to what extent the classification performance is influenced by the actual data presented for training.




(a)




(b)


Figure 5.13: Correct classification rate against the percentage of variance covered by the principal components for the following binary classifiers: (a) *no defect*, (b) *horizontal crack* (c) *corrosion pit* and (d) *diagonal crack*.




(c)



(d)



(a)



(b)


Figure 5.14: (a) The mean classification rate, using the validation set, as a function of the number of training images. (b) shows the results of the earlier bandwidth sensitivity using the new training volume of 25 images. It can be seen that there are no significant differences to the earlier results.

This is the process of determining a classifier's generalisation and the 'leave-one-out' technique was employed. As the name suggests, this process involves using N-1 of the N class images for training while generally reserving one image for testing purposes and is regarded as a good method for testing the sensitivity when only a limited set of images exists [102]. The downside however is that the process has to be repeated N times. Following the results of the training volume study, a data set of 26 images was produced allowing 25 images to be used for training in each of the 25 training combinations. However, the classification rate of the validation set was calculated rather than of the image being left out. In order to isolate the sensitivity between classes, the process was performed for each of the four classes in turn while keeping the other three image sets constant.


- i. Perform pre-processing on all images at given bandwidth
- ii. Define a set of N training sets that each have only 1 image different
- For Index < N
- iii. Set current training set = training set (Index)
 - For number of PCs < max(number of PCs)
 - iv. Run training algorithm
 - v. Calculate classification rate for each binary classifier
 - vi. Increment number of PCs
 - Loop
- vii. Increment Index
- Loop

Figure 5.15 (a) to (d) shows the mean classification performance for a given number of PCs across all combinations of the training images with an error bar of \pm one standard deviation. It can be seen that there is an element of cross-sensitivity between classes in all cases, with variations in the training images of any one class affecting (to varying degrees) the classification performance of the remaining classes. This is reasonable since any change to the training set will have an impact on the feature vector that is generated. In most cases, however, the error bar is considerably less than the maximum of approximately \pm 2.5%. Each classifier profile is very similar in the four plots and this is reasonable given that only 1 image across the four classes is changed each time.

It has therefore been shown that the inter-class cross-sensitivity is relatively small in proportion to the mean classification rate and therefore does not present a serious concern at the given dimensionality.

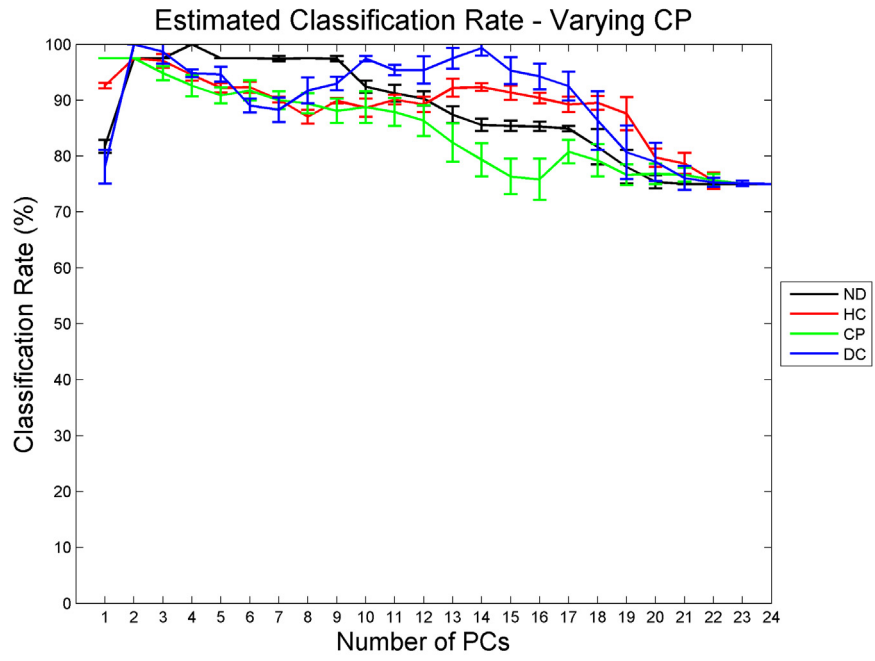


(a)




(b)

Figure 5.15: Generalisation. Mean classification rate with error bars showing one standard deviation for (a) varying ND, (b) varying HC, (c) varying CP and (d) varying DC training images.



(c)



(d)

5.4.5 Performance of Individual Binary Classifiers

The performance of all of the binary classifiers is high (>95%) for 2 PCs. The mean performance of the ND classifier, across all of the generalisation tests, was 97.5% which, although is marginally short of the target of 100% outlined in the specification, is still a very good result. Given that the dimensionality of the data has been reduced to 2, the resulting weight vectors obtained using Equation (5.13) representing all of the training images can be viewed as a 2D scatter plot. This is shown in Figure 5.16 and is a key validation of the algorithm. The figure is colour coded for each class as follows: Black for the *no defect* class, red for the *horizontal crack* class, green for the *corrosion pit* class and blue for the *diagonal crack* class.

It can be seen that there is very good separation of the four classes within the feature space and consequently, it should be expected that it would be possible to design a correspondingly good classifier. There is a strongly elliptical pattern to all classes other than the ND class, validating the decision not to use the Euclidean distance as a metric (which would attempt to define a circular decision boundary). The spread of HC images exhibits significant variation along the x-axis, with much less along the y-axis, indicating that it is most significantly affected by the first principal component. Conversely, the DC images have greater variation along the y-axis and hence, the second principal component. The CP class images occupy the




Figure 5.16: 2D scatter plot showing the location of the images from each of the four classes within the 2D feature space.

‘middle-ground’, with broadly equal variation in the direction of both principal components. Above all, however, it can be seen that the ND class is distinct from the other classes and has a relatively low variation in both dimensions. This strongly supports the classification performance observed by the ND binary classifier. The importance of this result, in terms of showing that there is a clear distinction between the defined classes, cannot be overstated and, in contrast with the other approaches discussed in Section 5.2.3, provides a solid basis for confidence in the classifier performance.

5.4.6 Combining Multiple Classifiers

5.4.6.1 Introduction

Up until now, the result of each binary classifier has been taken in isolation. The basic system could operate using a single ND binary classifier to determine the defect/no-defect case as required to meet the primary goal. However, by combining the outputs from all available classifiers it may be possible to determine a more specific classification of the image. This Section will outline the important background theory on classifier combination and then some combination algorithms will be implemented on the binary classifiers that have been developed. The concepts outlined here are not, however, restricted to binary classification, as will become clear in Section 5.5.

Consider two groups of classifiers. The first group contains individual classifiers, each of which can correctly classify a proportion of a particular set of test images while also generating some classification errors. The images that are misclassified, however, are relatively classifier specific. The second group of classifiers is similar except, in this instance, they are all relatively consistent in terms of which images they can correctly and incorrectly classify. However, this group of classifiers is affected by random noise. Intuitively it can be seen that there is value in combining the results of the classifiers within each group to harness the complementary information and thereby minimise the classification errors. In the first instance this takes advantage of the uniqueness of the misclassifications associated with each classifier and in the second instance the random error can be reduced through averaging across all classifier outputs. In recent years, classifier combination has enjoyed a period of significant interest within the field of pattern recognition. The overall aim of the combination, however, is ultimately the same: to achieve a greater classification performance from the ensemble than can be achieved by any of the base classifiers separately.

The plethora of combination schemes have been categorised in different ways by a variety of different authors. Kuncheva [103] defined two types of combination: *selection* and *fusion*. Classifier *selection* involves defining specific classifiers to cover a restricted portion of the feature space and then subsequently selecting the most appropriate classifier based on the features present in the particular input image. This draws parallels with having a panel of experts, each of which has detailed knowledge of their own field of study, but whom remain relatively ignorant of the fields of study of their colleagues. Classifier *fusion*, on the other hand, assumes that each classifier has knowledge of the entire feature space and can be thought of as taking the consensus decision of a group of equally qualified experts. Kittler et al [104], in outlining a theoretical framework for classifier combination, defined two classes of combination problem for the purposes of analysis. The first involves the situation where each classifier in the ensemble uses the same representation of the input data while the way in which this is processed by each classifier varies. For instance, two classifiers each supplied with the same input data but one calculating the Euclidean distance to a particular point in the feature space while the other calculates the Mahalanobis distance to the same point. The second class of problem involves the situation where each classifier extracts data from the input pattern that is unique to that classifier. While all images in the data set, for the system described in Section 5.4, are expressed in terms of the same feature set, the actual data input to each of the four classifiers, namely the Mahalanobis distance to a particular class centroid, is different and hence it is to Kittler's second class of system that the binary combination problem belongs. Techniques for combining classifiers can be divided into either trainable or non-trainable, depending on whether they require further training beyond that required by each individual classifier or not respectively. Trainable combiners can lead to better classification performance, however non-trainable algorithms are simpler to implement and, with lower requirements on the total volume of training data required, are popular for systems where the volume of training data is limited, such as in this case. The precise algorithm used for combination depends to a large extent on the nature of the base classifiers used and the output that they generate.

5.4.6.2 Combination Algorithms

The combination of a number of one-vs.-all binary classifiers is a problem which has no universally applicable optimum solution. The simplest, and some of the most popular, non-trainable methods include the *product*, *sum*, *mean*, *median*, *max*, *min* and *majority vote* rules.

The product and sum rules involve multiplying and adding the output probabilities for each classifier respectively and assigning the pattern to the class which has the greatest value across the ensemble. The mean and median rules assign the pattern to the class which has the greatest mean or median respectively. The max and min rules select the classifier which has either the greatest (most optimistic) or least (most pessimistic) confidence, amongst the ensemble, in its classification of the input pattern respectively. The max rule is susceptible to misclassifications arising from base classifiers which have been overtrained and consequently have an artificially high confidence for all input patterns. Both the min and max rules rely on the fact that the confidence outputs are correctly scaled relative to each other so as not to allow one classifier to appear more or less confident than the rest of the ensemble. Given the training strategy adopted in this system, relative scaling is not an issue since each classifier is subject to the same process. The *majority voting* scheme is slightly different in that it uses hard outputs and can be used to combine the outputs of each classifier to form a single prediction. For the binary OVA approach, this equates to the ‘winning’ class having a value of one and the rest having zero. Clearly, however, this leads to the possibility of obtaining a split decision where more than one classifier claims a particular image as a member or, alternatively, where no classifier claims the image as a member. In this instance, either an unknown classification can be returned or the image can be assigned to a class either according to the *a priori* class probabilities if they are different or at random if they are the same.

5.4.6.3 Combination Theory

Using Kittler’s analysis for a Bayesian classifier, given an input pattern (in this case a particular distance measurement), an image should be assigned to the class which has the greatest *a posteriori* probability. For a system consisting of R classifiers each outputting a measurement x_i for $i = 1, \dots, R$ the input image will be classified to class c_j if:

$$P(c_j|x_1, \dots, x_R) = \max_{k=1}^m [P(c_k|x_1, \dots, x_R)] \quad (5.16)$$

where k is the class index. In order to calculate the *a posteriori* probabilities required by the decision rule shown in Equation (5.16) the joint probability density functions $P(x_1, \dots, x_R|c_k)$ must be known. The calculation of this for a multivariate situation such as this is relatively difficult and so it is necessary to introduce a number of assumptions to make the analysis simpler. It was shown in [104] that by assuming that either

- i. the R classifiers are statistically independent or

ii. in addition to this independence, the *a posteriori* probability calculated by the R classifiers does not vary greatly from the *a priori* probability then, from Equation (5.12), the product and sum rules can be derived as Equation (5.17) and Equation (5.18) respectively, for an m class system.

$$P^{-(R-1)}(c_j) \prod_{i=1}^R p(c_j|x_i) = \max_{k=1}^m \left[P^{-(R-1)}(c_k) \prod_{i=1}^R p(c_k|x_i) \right] \quad (5.17)$$

$$(1-R)P(c_j) + \sum_{i=1}^R p(c_j|x_i) = \max_{k=1}^m \left((1-R)P(c_k) + \sum_{i=1}^R p(c_k|x_i) \right) \quad (5.18)$$

If it is assumed that the priors for each class are equal (based on the fact that there is an equal volume of training data for each class) the prior terms in Equation (5.17) cancel out yielding the following [102]:

$$\prod_{i=1}^R p(c_j|x_i) = \max_{k=1}^m \left[\prod_{i=1}^R p(c_k|x_i) \right] \quad (5.19)$$

While the assumptions made are unlikely to be strictly true, particularly for large numbers of classifiers in the case of statistical independence, they are ones from which, in many applications, useful results can still be obtained. Furthermore, it can be shown that many of the typical combination mechanisms can be derived by further simplifications of the product and sum models [104]. It is interesting to analyse the two basic combination mechanisms, since it allows the behaviour of the derived combination algorithms to be understood more fully. In the first instance, consider the product rule. By inspection of Equation (5.17) it can be seen that if one or more classifiers incorrectly predicts a very low probability of class membership, then this will result in the combined probability of membership also being very low. This necessitates that very good estimates of the conditional probabilities should be available. In contrast, the nature of the sum rule makes it much less sensitive to single spurious outputs from a particular base classifier. The error factor is defined as the factor by which the true probability varies from the estimated probability. It was shown in [104] that the error factor for the product and sum rules could be given by Equation (5.20) and Equation (5.21) respectively.

$$ErrorFactor_{product} = 1 + \sum_{i=1}^R \frac{e_{ki}}{P(c_k|x_i)} \quad (5.20)$$

$$ErrorFactor_{sum} = 1 + \frac{\sum_{i=1}^R e_{ki}}{R \sum_{i=1}^R P(c_k|x_i)} \quad (5.21)$$

where e_{ki} is the extent to which the estimated *a posteriori* probability deviates from the true probability, i.e. $\hat{P}(c_k|x_i) = P(c_k|x_i) + e_{ki}$. Since the *a posteriori* probabilities must be less than one, the product rule error factor actually shows that the individual classifier errors will be amplified. However using the sum rule, each time for the class which is most probable, the summation term will potentially be greater than unity (because it is the outputs from different classifiers that are being summed) therefore the error is damped. It can be concluded that in cases where the estimates of the class-conditional probabilities are likely to contain significant estimation error, combination algorithms based on the sum rule are likely to perform better than those based on the product rule, due to the error dampening effect. Since this is likely to be true of the derivation of the densities as outlined in Section 5.4.2, it was concluded that the sum rule and its derivatives were most likely to be the best choice for the combination algorithm in this case.

Starting from the sum rule, the max rule can be derived by first assuming that the summation term of the sum rule can be approximated by the maximum of the *a posteriori* probabilities:

$$(1-R)P(c_j) + \sum_{i=1}^R p(c_j|x_i) = \max_{k=1}^m \left((1-R)P(c_k) + \sum_{i=1}^R p(c_k|x_i) \right)$$

$$(1-R)P(c_j) + R \max_{i=1}^R (p(c_j|x_i)) = \max_{k=1}^m \left((1-R)P(c_k) + R \max_{i=1}^R (p(c_k|x_i)) \right)$$

Next, by assuming that each class has equal prior probabilities, the expression can be simplified to yield the *max* rule as shown in Equation (5.22):

$$\max_{i=1}^R (P(c_j|x_i)) = \max_{k=1}^m \max_{i=1}^R [P(c_k|x_i)] \quad (5.22)$$


This rule states that an image should be assigned to class, c_j , if the maximum *a posteriori* probability of the image belonging to this class, across all the R classifiers, is equal to the maximum *a posteriori* probability found in any of the classes across all the R classifiers. For the less mathematically minded this can be restated as ‘if the probability of the image being class c_j (as determined by any of the R classifiers) is the greatest, then assign the image to c_j ’. For further information on the derivations presented here, the reader is directed to [104].

5.4.6.4 Binary Fusion Results


Binary combination was performed using the 'max' rule by taking each image and identifying which classifier had the highest probability of the image being a class member (i.e. "Who is the most confident in their prediction"). If this maximum probability is greater than 50% then the image is classified as a member of that particular class; if it is not then this means that all classifiers have returned a higher probability of the image being a non-member than a member and hence the image, in this instance, is defined as *unclassifiable*.

The confusion matrices for the binary classifier are shown for both the validation and test sets in Table 5.1 and Table 5.2 respectively, where U denotes unclassifiable. The performance equates to a classification rate of 95% and 85% respectively. It can be seen that on the validation set, two of the images are unclassifiable indicating that none of the four binary classifiers indicated a probability of class membership of greater than 50% for that image. For the test set, there are two no-defect images which are unclassifiable and one horizontal crack image. The fact that there are no incorrect classifications is encouraging since this shows that no other binary classifier is incorrectly 'confident' about its classification. Also, the classifier output for the correct class may be the highest out of the 4 classifiers but if it is less than 50% it will still be classified as *unclassifiable* because it is more likely to be a non-member. For all five unclassifiable images, shown in Figure 5.17, this was found to be the case with the Mahalanobis distance of each being shortest to a member of its own class. However, the distribution of known class images was such they had, in most instances, a low probability of membership to any of the classes. Hence, all binary classifiers deemed these images to be most probably non-members resulting in an *unclassifiable* assignment. Note: The 2D FFT data in this figure is presented post-compression and hence has had the PCA algorithm applied. This explains the bipolar magnitude values arising from the zero mean requirement of Equation (5.3).

Direct comparison between the fused classifier and the individual classifiers is difficult because it is not equating like for like (i.e. the individual system separates only one class while the fused system separates four). Thus, the two systems must be taken as ones that each fulfil a different need: the individual ND classifier to determine simply if a defect exists in an image and the fused binary ensemble to determine the nature of the defect. A more appropriate comparison for the binary ensemble classifier would be with a multi-class ensemble classifier. Such a system will be considered in the following section.




	Mem (%)	non (%)
ND	44.8	56.2
HC	49.1	51.9
CP	17.4	82.6
DC	0	100




(a)

(b)




	Mem (%)	non (%)
ND	12.2	87.8
HC	46.9	53.1
CP	30	70
DC	0	100




(c)

(d)




	Mem (%)	non (%)
ND	0	100
HC	4.1	96.9
CP	0	100
DC	0	100




(e)

(f)

Figure 5.17: Binary system: unclassifiable images. Original images are shown above their filtered reconstruction using 2 principal components. The 2D FFT corresponding to the reconstruction is shown in the right-hand column: (a-f) No Defect data, (g-h) Horizontal Crack data and (i-j) Corrosion Pit data. The probabilities returned by each classifier are shown in the centre of the figure.




	Mem (%)	non (%)
ND	0	100
HC	4.1	96.9
CP	0	100
DC	0	100




(g)

(h)



	Mem (%)	non (%)
ND	0	100
HC	0	100
CP	0.52	99.4
DC	0	100



(i)

(j)

Table 5.1: Binary Fusion: Confusion matrix using validation set

Actual\Predicted	ND	HC	CP	DC	U
ND	9	0	0	0	1
HC	0	10	0	0	0
CP	0	0	9	0	1
DC	0	0	0	10	0

Table 5.2: Binary Fusion: Confusion matrix using test set.

Actual\Predicted	ND	HC	CP	DC	U
ND	3	0	0	0	2
HC	0	4	0	0	1
CP	0	0	5	0	0
DC	0	0	0	5	0

5.5 Multi-class Sensors to Enhance Probability of Detection

5.5.1 Overview

In Section 5.3.3, It was discussed that for a multi-class system, the multiple non-member classes could be described by a single non-member class with a probability density function comprising of a mixture of Gaussians, thereby essentially reducing the problem to one involving two classes. While lending itself well to the simple binary classifier implementation there is additional information content available within the non-member distributions which could be harnessed. Using Figure 5.8 as an example, and noting from Section 5.4.2 how the density functions are defined, it can be seen that for each of the four classes, four probability density functions are defined, thus yielding a total of 16 distributions. It would be expected that the Mahalanobis distances of images, belonging to the four-class data set, to a particular class of interest will fall into four distributions with a distinct mean and variance corresponding to each class. Assuming that appropriate training data is used, the distributions obtained in this way from the training procedure could be, loosely, described as the characteristic of a *sensor* for this particular class type (e.g. corrosion pit). New images could then be assigned a probability of membership to the four known classes derived from this sensor profile. As there are four types of defect that are of interest, and hence four classes, this provides four different ‘sensor’ measurements of the probability of an image belonging to any given class. Observing Figure 5.8 and taking c_1 , c_2 , c_3 and c_4 to be the *no defect*, *horizontal crack*, *corrosion pit* and *diagonal crack* image classes respectively, it can be seen that this ‘sensor’ has reasonably good separation of the *no defect* and *diagonal crack* classes from their remaining three classes however the *horizontal crack* and *corrosion pit* classes show a high degree of overlap. While there may be varying degrees of overlap between the distributions, there is still information contained there-in relating to the class membership probability for a given input Mahalanobis distance which is discarded in the binary classification system outlined in Section 5.4. Furthermore, by defining a set of sensors which provide a probability measure for each class, it is possible to apply a data fusion algorithm to the resulting collection of soft outputs to potentially enhance the probability of detection of the ensemble. No further statistical assumptions are made here that differ from that of the binary case, the key difference is the manner in which the non-member data is utilised.

Using the same training data as in Section 5.4, the distributions for the four class sensors are shown in Figure 5.18 (a) to (d) for a bandwidth of $0.1 \frac{FS}{2}$ and with a

dimensionality of two. A logarithmic scale has been used for clarity. It can be seen that the distributions of the data is far from ideal with several instances of overlapping PDFs, however this exposes the raw limit on the resolvability of an image to a class using the 2D FFT as an input. The performance of the individual sensors at classification as well as the performance of the fused ensemble is dealt with in the following section.

The theoretical analysis of the combination algorithms outlined in Section 5.4.6 is equally valid in the multi-class instance. As with the binary system, the max rule is applied this time to the collection of four multi-class sensors however, since each sensor provides an estimate of the probability of membership of all known classes, the product rule can also be used.

5.5.2 Performance of Sensors Approach

5.5.2.1 Individual and Fused Classifiers

The confusion matrices for the four individual sensor are provided in Table 5.4 through to Table 5.7. It can be seen that each exhibits good classification performance for their primary class (e.g. the ND class is the primary class of the ND sensor etc) with this in general deteriorating for the remaining three classes, as would be expected from the distributions shown in Figure 5.18. The ND, HC and CP sensors all perform more poorly than the binary ensemble classifier. The DC sensor returns 1 misclassification (CP to HC) on the validation set which in terms of correct classification rate is higher than the binary ensemble. However, in terms of the misclassifications, the latter can be considered to perform better. This may be said because only an unclassifiable result can be returned by the binary ensemble where-as the DC sensor actually incorrectly assigns a class.

The confusion matrices for the validation and test sets using the max rule are given in Table 5.8 and Table 5.9 respectively and for the product rule in Table 5.10 and Table 5.11 respectively. The classification rates are also summarised in Table 5.3 for convenience. It can be seen that both rules work well, with the product rule surprisingly providing a 100% classification of both the validation and test image sets. This is surprising, primarily due to the relatively 'unforgiving' nature of this combination method in that any classifier which incorrectly outputs a close to zero probability for a particular case will strongly influence classification in that instance. Furthermore, it would be expected that such performance would only be possible where a strong assumption of conditional independence (i.e. the result of one


classifier would not change the belief in any of the other classifiers) could be made. Given the nature of the sensor definitions, this is counter-intuitive.

The correct classification performance of the max rule on the validation set is comparable with the binary system in that all but two images are labelled correctly. Analysis of the results of the test set however indicates that the max rule outperforms the binary system and yields only one incorrect classification compared with three images identified as unclassifiable by the binary system.


Comparing the performance of the binary and the multi-class approaches on the independent test set, the product rule applied to the multi-class data provides the best classification performance, followed by the max rule on the same data and lastly the combined binary classification system. There is, however, not a great deal of difference between the three results.

Classifier	Validation (%)	Test (%)
ND (Individual)	72.5	65
HC (Individual)	80	65
CP (Individual)	90	85
DC (Individual)	97.5	100
Max	95	95
Product	100	100

Table 5.3: Summary of classification rates for multi-class sensors




(a)

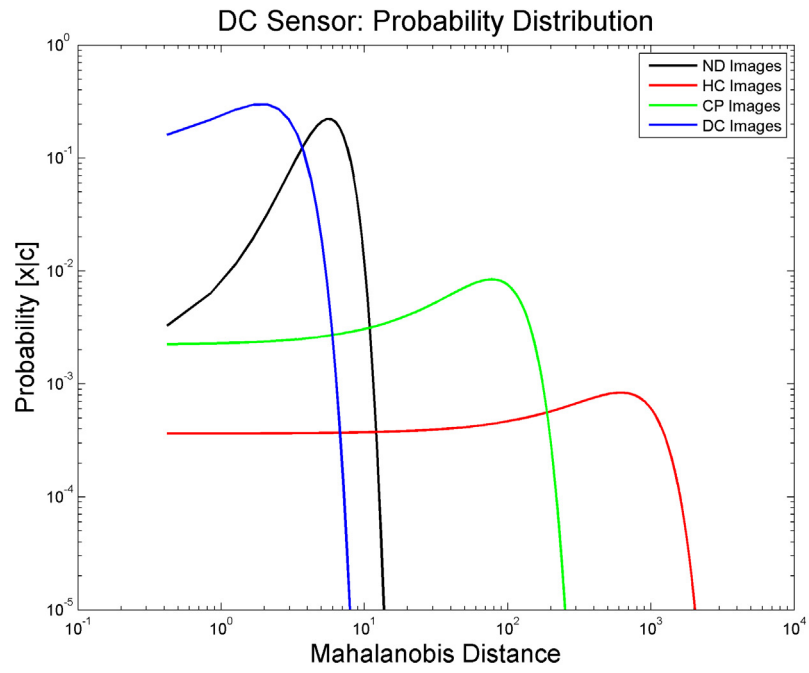


(b)

Figure 5.18: Probability density functions for each class of images using two principal components and an image bandwidth of $0.1 \frac{S}{2}$. (a) the *no defect* sensor, (b) the *horizontal crack* sensor, (c) the *corrosion pit* sensor and (d) the *diagonal crack* sensor.



(c)



(d)

Table 5.4: Confusion matrix for ND sensor (Validation Set)

Actual\Predicted	ND	HC	CP	DC
ND	9	0	0	1
HC	0	4	6	0
CP	0	0	6	4
DC	0	0	0	10

Table 5.5: Confusion matrix for HC sensor (Validation Set)

Actual\Predicted	ND	HC	CP	DC
ND	10	0	0	0
HC	0	10	0	0
CP	2	0	3	5
DC	1	0	0	9

Table 5.6: Confusion matrix for CP sensor (Validation Set)

Actual\Predicted	ND	HC	CP	DC
ND	9	0	0	1
HC	0	10	0	0
CP	1	0	9	0
DC	2	0	0	8

Table 5.7: Confusion matrix for DC sensor (Validation Set)

Actual\Predicted	ND	HC	CP	DC
ND	10	0	0	0
HC	0	10	0	0
CP	0	1	9	0
DC	0	0	0	10

Table 5.8: Multi-Class Fusion: Confusion matrix for max rule combination on validation set

Actual\Predicted	ND	HC	CP	DC
ND	10	0	0	0
HC	0	10	0	0
CP	1	0	9	0
DC	1	0	0	9

Table 5.9: Multi-Class Fusion: Confusion matrix for max rule combination on test set

Actual\Predicted	ND	HC	CP	DC
ND	5	0	0	0
HC	0	5	0	0
CP	0	0	5	0
DC	1	0	0	4

Table 5.10: Multi-Class Fusion: Confusion matrix for product rule combination on validation set

Actual\Predicted	ND	HC	CP	DC
ND	10	0	0	0
HC	0	10	0	0
CP	0	0	10	0
DC	0	0	0	10

Table 5.11: Multi-Class Fusion: Confusion matrix for product rule combination on test set

Actual\Predicted	ND	HC	CP	DC
ND	5	0	0	0
HC	0	5	0	0
CP	0	0	5	0
DC	0	0	0	5

5.5.2.2 Explanation for Misclassifications

Figure 5.19 shows the scatter plot of the training images provided earlier, with the validation images (square marker) and test images (cross marker) superimposed. The colour coding is as previously outlined, however with the following additions. Shown in light blue are the images that were misclassified by the binary ensemble classifier and shown in purple are the images that were misclassified by the max rule approach to the multi-class sensors. One image is indicated with both a purple square outline and a light blue fill indicating that it was incorrectly classified by both methods. As the product rule produced no errors, there is no index corresponding to it. What is clear from the placement of these images in relation to their co-members is that, in the cases where an error has occurred, the Mahalanobis distance of the images is shortest to its parent class. It appears that the misclassifications arise when the distances are converted to probability estimates based on the distributions of Mahalanobis distances obtained during the training process. This can be clearly seen by considering the CP image misclassified by the max rule as being a member of the ND class. In this instance, the CP sensor returns a strong probability that the image belongs to the ND class, overriding the less confident, but correct, predictions made by the other three sensors. Observation of the probability density functions for the CP sensor, shown in Figure 5.18(c), shows




Figure 5.19: Scatter plot showing the weights of the training, validation and test images. It can be seen that they all follow a consistent pattern of distribution.

that the means of both the ND and CP classes are relatively close together and, consequently, because the Mahalanobis distance of the image is larger than expected for a CP image it appears to strongly favour the classification of it as a ND member. If further training data were used, this would most likely just increase the standard deviation of the CP class, and consequently the probability of the image, given the Mahalanobis distance, would always be higher for the ND class. Since the result of this sensor is strongly in favour of an ND classification, it overrules the remaining three sensors when the max fusion rule is applied and results in the misclassification. It is for the same reason that the binary combination system returns the image as *unclassifiable* since no classifier returns a >50% probability of the image being a member of the parent class.

5.6 K-Nearest Neighbour Classification

5.6.1 Overview

The K-nearest neighbour (KNN) method is generally regarded as a good benchmark for classifier performance [105]. The primary reason for this is that, in the case where $K = 1$, the error is bounded at the lower extent by the Bayes probability of error and, as the number of samples tends to infinity, at the upper extent by *twice* the Bayes' error. This is summarised in Equation (5.23), where R is error rate, R^* is the Bayesian error rate and c is the number of classes [106].

$$R^* \leq R \leq R^* \left(2 - \frac{c}{c-1} R^* \right) \quad (5.23)$$

This system uses some distance metric (Euclidean, Mahalanobis etc) relating an image to each of the training images and then selects the K closest images (i.e. shortest distances). The image is then assigned to the class which the majority of the K images belong to. A suitable value of K is usually determined following a training process.

The KNN algorithm replaces the procedure outlined in the statistical framework of Section 5.3.3 however the pre-processing and dimensionality reduction steps remain the same. In order to allow comparison with previous results, the same bandwidth and dimensionality was used.

The Mahalanobis distance is the distance to the centre of a particular *distribution* rather than a simple point-to-point measurement and so the Euclidean distance is more commonly used when the KNN algorithm is employed. Following the pre-

processing stage, the Euclidean distance of each validation and test set image to each image in the training set is calculated. The K instances which have the smallest distance are then extracted and the classes of the corresponding training images recorded. The image under investigation is then assigned to the most frequent class within this K subset. This is summarised by the following pseudo-code, letting *test_val* denote a vector containing all validation and test data and *training* denote a vector containing all training data:

- i. Perform pre-processing on all images at given bandwidth and dimensionality
 - For $i < \text{length}(\text{test_val})$
 - For $j < \text{length}(\text{training})$
 - ii. Calculate Euclidean distance between *test_val*(i) and *training*(j).
 - iii. Increment j
 - Loop
 - iv. Sort distance in descending order and select K smallest values (nearest images)
 - v. Identify the classes of these images and assign *test_val*(i) to the class that is most frequently within the K subset.
 - vi. Increment i
 - Loop

5.6.2 Results

The results in Figure 5.20 show the performance of the KNN algorithm on both the validation and test image sets for increasing K value. It can be seen that for $1 < k < 4$, a classification rate of 100% is achieved for both the training and test sets however for $k > 5$ the classification rate on the validation set fluctuates and the performance on the test set reduces to 95%.

Comparing the merits of the Bayesian approach of Section 5.3.3 with the KNN method presented here it can be seen that both provide the opportunity of achieving 100% classification rates. The KNN neighbour method is simple and takes no specific account of the statistical spread of the data set. The Bayesian approach, on the other hand, is complex with significant calculation overheads however it benefits from the ability to assign a level of confidence to its class prediction. To obtain such a confidence measure within a KNN framework would require a sufficiently high value for K so as to ensure a good resolution of the probability measure. Given the results obtained here, this would reduce the classification rate and therefore a trade-off emerges between simplicity, performance and functional outputs. It was deemed that the confidence




Figure 5.20: K-nearest neighbour results on validation and test sets for increasing value of K. Classification rate of 100% is achieved on the validation set for $1 < K < 4$ which results in an equivalent rate for the test set. Increasing K beyond 5 reduces the performance on the validation set and causes the classification rate of the test set to reduce to 95%.

assignment was sufficiently valuable as an output as to warrant the added computational expense.

5.7 Performance Comparison Using Real Defects

5.7.1 Creating Samples and Image Set

Thus far, the proof of concept has been achieved using idealised defects. The ‘cracks’ are slots with well defined edges, the ‘corrosion pits’ are drill holes with well defined circular shape and the ‘no defect’ images, in most instances, are free from surface aberrations. Real defects are unlikely to be so amenable to distinguishing between classes. Access to samples containing genuine defects is limited, particularly in the sort of numbers that are required for training a classifier. In order to achieve a more realistic and stern test for the pattern recognition system, a number of defects were artificially corroded into a carbon steel test plate. Each defect was created using a mould placed on the metal surface and containing saline water. A positive electrode was placed in the solution and the steel plate connected to earth, thereby creating an electrochemical cell. A current of 2 A was applied for

a period of approximately 10-12 hours causing corrosion of the parent metal within the region of the mould. Two varieties of mold were used: a cylinder for a circular *corrosion pit* (CP) and a slot for a *slot* defect (SL). For simplicity only one orientation of slot was used and so, coupled with the no-defect (ND) images, this yields a three class system. As the nature of the corrosion cannot be controlled further than the rate at which it occurs, the resulting defects have less well defined edges than the machined defects, with particular regions within the mould being more corroded than others. A sample from each of the three classes is given in Figure 5.21. A total of 50 images per class were captured using 7 different defects (per class) yielding a total image set of 150 images. Each image is uniquely different from the rest due to either the type of defect it included, the lighting conditions or the location of the defect within the image frame. The image set was divided into a 30 image training set (of which the optimum volume would be determined as per Section 5.4.4.2), 10 image validation set and a further 10 image test set.

5.7.2 Training and Optimisation

The procedure outlined in Section 5.4.4.1 to Section 5.4.4.3 was once again applied, this time to the new image set. As was concluded previously, the image bandwidth does not play a significant part in the resulting classification rate however, in this case, a bandwidth of $0.15\frac{F}{2}$ was marginally superior. A good compromise between the classification rate of each of the three classes was obtained using a combination of $N = 20$ training images while reducing the dimensionality to two as shown in Figure 5.22. Although this training set size does not yield the maximum ND classification rate ($N = 19$ images), it was selected because the SL classification rate increased significantly for a marginal decrease in ND performance. The generalisation procedure was conducted, again using the leave-one-out method, for each of the three classes. The results are shown in Figure 5.23(a), (b) and (c) with the error bars indicating \pm one standard deviation. The suitability of using two principal components is evident once again. The size of the generalisation error is larger than that obtained using the synthetic defect images at $\pm 5\%$. This is most probably due to the greater intra-class variability within the training set.

Having identified the most appropriate system parameters, the system can be tested using the method outlined previously. The results are presented and discussed in Section 5.7.3.



(a)



(b)



(c)

Figure 5.21: Corrosion defect images: (a) *No defect*, (b) *slot* and (c) *corrosion pit* images on a steel plate. Both (b) and (c) were generated using an electrochemical cell.






Figure 5.22: Classification rate, using the validation set, as a function of the number of training images for each of the three classes.




(a)

Figure 5.23: Generalisation to (a) ND Class, (b) SL class and (c) CP class.



(b)



(c)

5.7.3 Classification Results

5.7.3.1 Individual Binary Classifiers

The performance of the ND and CP binary classifiers is significantly poorer than for the idealised image set although both still produce a 70% correct classification rate. The ND rate falls some way short of that required to achieve the primary goal. The SL classifier performs better yielding approximately 95% correct classification rate. The reason for this can be traced to the spatial distribution of the data set shown in Figure 5.24. It is clear that there is significant overlap in the ND and CP classes meaning a perfect decision boundary is impossible. The separation of the SL images from the rest is better although there are instances of cross-over. This can be explained using the image in Figure 5.21(a). The ND class images typically contain a level of surface content that was not present in the idealised images. This gives rise to an increased frequency content in all directions of the 2D FFT which, previously, was the defining characteristic of the CP class. The more regular edges of the SL class result in a better distinction between this and the other classes.

5.7.3.2 Binary Fusion




Figure 5.24: 2D scatter plot of the three classes of corrosion images within the feature space.

A binary combination method using the ‘max’ rule was applied to the new classifiers and the confusion matrices for the validation and test sets shown in Table 5.13 and Table 5.14 respectively. Given that the validation and test image sets contain 10 images per class, the performance translates to correct classification rates of 63% and 60% respectively. What is interesting to note, and in contrast with the previous image set, is that on 9 occasions a misclassification occurs meaning that in these instances, a classifier is incorrectly confident about its result. Previously, there were no incorrect classifications, instead images were either correctly classified or deemed unclassifiable.

5.7.3.3 Multi-Class Sensors

The confusion matrices are shown in Table 5.15 through to Table 5.21 and the results summarised in Table 5.12. The performance of the individual, multi-class sensors is generally a little poorer than the binary ensemble with the ND sensor, SL sensor and CP sensor having a 56%, 53% and 53% correct classification rate respectively on the validation set.

Neither the max or product rules have a significant effect on the classification rate on the validation set however result in better performance on the test set. On balance, the product rule appears to be the best classification option however given the results of the generalisation study, there would most probably be very little difference in overall performance between the binary ensemble, max and product rule systems.

Classifier	Validation (%)	Test (%)
ND (Individual)	43	56
SL (Individual)	53	53
CP (Individual)	53	53
Max	50	66
Product	53	73

Table 5.12: Summary of classification rates for multi-class sensors on real images

Table 5.13: Binary Fusion: Confusion matrix using validation set

Actual\Predicted	ND	SL	CP	U
ND	8	0	1	1
SL	0	9	0	1
CP	4	4	2	0

Table 5.14: Binary Fusion: Confusion matrix using test set.

Actual\Predicted	ND	SL	CP	U
ND	8	0	2	0
SL	0	10	0	0
CP	4	6	0	0

Table 5.15: Confusion matrix for ND sensor (Validation Set)

Actual\Predicted	ND	SL	CP
ND	4	0	6
SL	0	7	3
CP	8	0	2

Table 5.16: Confusion matrix for SL sensor (Validation Set)

Actual\Predicted	ND	SL	CP
ND	2	0	8
SL	0	10	0
CP	1	5	4

Table 5.17: Confusion matrix for CP sensor (Validation Set)

Actual\Predicted	ND	SL	CP
ND	2	3	5
SL	4	4	2
CP	0	0	10

Table 5.18: Multi-class Fusion: Confusion matrix for max rule combination on validation set

Actual\Predicted	ND	SL	CP
ND	3	2	5
SL	0	8	2
CP	6	0	4

Table 5.19: Multi-class Fusion: Confusion matrix for max rule combination on test set

Actual\Predicted	ND	SL	CP
ND	8	0	2
SL	0	9	1
CP	3	2	5

Table 5.20: Multi-class Fusion: Confusion matrix for product rule combination on validation set

Actual\Predicted	ND	SL	CP
ND	6	2	2
SL	0	5	5
CP	7	0	3

Table 5.21: Multi-class Fusion: Confusion matrix for product rule combination on test set

Actual\Predicted	ND	SL	CP
ND	8	0	2
SL	0	7	3
CP	4	0	6

5.7.3.4 K-Nearest Neighbour

The validity of the conclusion of the previous Section is supported by the results of the KNN method shown in Figure 5.25 where the classification rate lies in the range of 55-65% and 30-55% for the validation and test sets respectively, depending on the value of K. This is a significant reduction compared with the 95-100% achieved on the idealised image set.




Figure 5.25: KNN applied to corrosion defects for varying value of K.

5.8 Discussion

In the course of this Chapter, several implementations of a statistical classifier have been developed using 2D FFT data, extracted from images captured by the on-board camera, as its input. The system implemented is based on the Eigenface approach and uses principal component analysis to reduce the dimensionality of the data set using a set of features. The results are summarised in Table 5.22 and Table 5.23 for synthetic and real defects respectively. The goal of the approach was to develop a system that primarily could distinguish between images that contain defects and those that do not. A classification rate of 100% was specified and was achieved using a multi-class sensor approach using the product rule on both the validation and test sets. The individual binary ND classifier yielded 97.5% which, although marginally short of the required performance, is still a very good result. The advantage of this system is in its greater simplicity (with no fusion algorithm required) however at the expense of a less specific assignment of the image class. Both results are backed up by very good separation of the underlying data set allowing a high degree of confidence to be placed in the result. A more realistic data set, generated by artificially corroding steel samples, yielded significantly poorer results. Again, the product rule performed best with a correct classification rate of 73% with the individual binary ND classifier yielding 70%. Analysis of the class distributions in Figure 5.24 lead to the conclusion that the poor performance is due to the fundamental lack of distinction between the classes rather than a failure in the approach itself. This was particularly evident between the ND and CP classes. It has therefore been shown that a system based on the Eigenfaces method can be applied to NDE however, like any system, the performance will vary depending on the quality of the input data.

It was shown that for the idealised data set, where the defects have regular edges and the parent metal has little or no surface aberrations visible in the captured images, excellent classification performance can be achieved. The multi-class sensor system using the product rule has been shown to achieve a performance comparable to a KNN algorithm. Furthermore, because of the fundamental statistical basis of this system, a level of confidence could be attributed to the classification which, if emulated under a KNN approach, would require a value of K that would actually reduce the classification rate. However, this capability comes at the expense of significantly greater computational complexity. The performance of the system using real defects was, again, comparable to the results obtained by a KNN approach. Therefore, where good separation can be achieved and where a statistical confidence is more valuable than computational complexity, a multi-class

sensor approach is preferable over a simple KNN algorithm. For NDE applications, it is felt that the hard decision output (class label) of a classifier such as KNN is less appropriate and that more value is gained by having some concept of how confident the system is in that classification.

For both data sets, the image bandwidth was found to have less of an effect on the classification rate than the overall level of dimensionality that the data set was reduced to. Realistic defect geometry, such as a fatigue crack, is likely to require a different approach to that outlined in this Chapter. Principally, the pre-processing stage will need to be modified to account for the higher frequencies of interest signifying such defects. The classification methodology however could be used without modification since provided the input data has a statistical difference, the proposed system will be able to identify this and thus classify accordingly.

The lack of availability of truly independent defect images is an issue which is common to image classifier development. In truth, the only way to ensure very good generalisation is to have a large number of unique images in the training set with further unique images in the validation and test sets. Practically this has not been, and rarely is, possible and therefore it must be concluded that the performance demonstrated here is likely to be skewed in favour of a correct classification. However, the comparison of the system performance against the KNN algorithm yields a valuable result in terms of relative performance against a standard benchmark method.

Classifier	Idealised	
	Validation (%)	Test (%)
Binary Fused	95	85
ND (Individual)	72.5	65
HC (Individual)	80	65
CP (Individual)	90	85
DC (Individual)	97.5	100
Max	95	95
Product	100	100
KNN	100	100

Table 5.22: Synthetic Data: Summary of classification rates for multi-class sensors

Classifier	Real	
	Validation (%)	Test (%)
Binary Fused	63	60
ND (Individual)	43	56
SL (Individual)	53	53
CP (Individual)	53	53
Max	50	66
Product	53	73
KNN	53	66

Table 5.23: Real Data: Summary of classification rates for multi-class sensors

Chapter 6

Eddy Current Array

6.1 Introduction

Eddy current testing (ECT) is a commonly used, non-contact means of inspecting conducting materials [37, 38]. In contrast with air-coupled ultrasound, where the majority of the incident pressure wave is reflected by any metal-air boundary, the flux linkage mechanism upon which eddy current testing is founded is much less sensitive to the existence of such boundaries. A plethora of probe types and configurations abound, varying from simple, single- or dual-coil systems, to arrays [107, 108, 109, 110] and ultimately to largely lab-based systems such as the high precision, cryogenically cooled SQUID probes [111, 112]. The use of array technology, in particular, is appealing for automated inspection as it removes the need for a mechanical scanning system to position the probe, while providing enhanced detection capabilities through the flexibility of tailoring the array configurations to the inspection task.

This Chapter describes the development of an eddy current array (ECA) payload for deployment on a robotic pipe inspection system. The novelty of the work is primarily in the integration of the two systems, however the techniques used to improve the sensitivity are also considered to be novel. The Chapter starts with an overview of ECT theory and then establishes the optimum frequency and defect metric for greatest probe sensitivity. FEM is then used to relate the defect metric to actual defect geometries. An overview of the system and imaging algorithm is provided and finally, the performance is compared against a commercial system.

6.2 Background

6.2.1 Eddy Current Theory

Passing an alternating current (AC) through a coil generates a changing magnetic field which induces eddy currents, of the same frequency, in any conducting specimen that intersects it. As the induced eddy currents are also AC they, in turn, establish their own magnetic field which, according to Lenz's law, is of equal magnitude but opposite phase to that of the field that established it. For non-ferromagnetic materials, such as aluminium and stainless steel, this results in a partial reduction of the primary field. As the coil inductance, L , is directly proportional to the number of turns, N , and the magnetic flux density, Φ , this reduction in field causes a reduction in the impedance of the coil relative to its in-air value. The quantity $N\Phi$ is termed the 'flux-linkage' and the relationship is given in Equation (6.1).

$$L \propto N\Phi \quad (6.1)$$

For ferromagnetic materials, such as carbon steel, although the underlying phenomenon is the same, the effect is somewhat different. Due to the ferrite nature of the material, there is a high concentration of magnetic flux at the coil-metal boundary which *increases* the flux linkage and hence impedance of the coil. In both cases, the presence of a defect interrupts the magnetic field, as shown in Figure 6.1, thereby modifying the eddy currents induced in the specimen and consequently changing the impedance of the coil. As the impedance has both real and reactive elements, the presence of a defect can result in amplitude and phase changes in the measured coil voltage.

In the absence of any other variables, designing an inspection system to detect these changes would be relatively trivial. However, inevitably, there are several other factors that also influence the coil impedance. Primarily, these include [113]:

- i. Excitation frequency, f
- ii. Electrical conductivity of the specimen, σ
- iii. Magnetic permeability, μ (for ferromagnetic materials)
- iv. Lift-off distance between the coil and the specimen
- v. Coil size (number of turns, diameter)
- vi. Temperature




Figure 6.1: (a) Qualitative example showing the field of a coil on a conducting medium subject to AC excitation and (b) showing that the presence of a defect alters the magnetic field. The difference in the magnetic field in the two examples results in a difference in measured coil impedance.




Figure 6.2: Theoretical variation of skin depth with frequency, for aluminium and 304 grade austenitic stainless steel.

The excitation frequency is perhaps the most obvious factor which affects the impedance of the coil with higher values directly increasing the reactive component. Furthermore, the current in a conducting medium is concentrated to a layer around the outer surface whose thickness, or depth, is inversely proportional to the frequency of the current passing through it. Thus, as the frequency increases, the electrons are forced to occupy a smaller area and consequently the resistance of the coil increases. This phenomena is known as the ‘skin effect’ and the skin depth (δ) of a material, given by Equation (6.2), is defined as the depth at which the current density reduces to $\frac{1}{e}$ (37%) of its surface value.

$$\delta = \frac{1}{\sqrt{\pi f \mu \sigma}} \quad (6.2)$$

In terms of NDE, a consequence of the skin effect is that, by increasing the frequency of the excitation current, the eddy current field can be concentrated to the near-surface of the test specimen thereby producing a high sensitivity to small, surface-breaking defects. For sub-surface defects, a suitably low excitation frequency is required to ensure that the field interacts with the defect sufficiently to cause a change in coil impedance. The skin depth of ferritic materials such as steel is much less than that of austenitic stainless steel due to the inverse-root relationship with magnetic permeability, as shown by Equation (6.2). Also since the conductivity also obeys an inverse-root relationship, materials of high conductivity such as aluminium and copper will have a lower skin depth than other non-ferromagnetic materials with low conductivity, such as stainless steel. The variation of skin depth with frequency for aluminium and 304 grade austenitic stainless steel is given in Figure 6.2.

The influence of magnetic permeability and coil dimensions can be related back to the direct proportionality of the inductance and the flux-linkage outlined in Equation (6.1). Thermal effects, if not compensated for, can lead to significant repeatability errors and care must be taken in the design of eddy current systems to minimise the electrical self-heating of coils due to normal operation.

6.2.2 Probe Design

6.2.2.1 Coil Configuration

Typically, coils are configured as either absolute, differential or reflection probes. An absolute probe consists of a single coil with the drive signal also acting as the sensing signal. Such probes are particularly sensitive to lift-off and probe wobble which can lead to spurious indications of defects. Differential probes consist of a

pair of coils which usually each form one half of a balanced Wheatstone bridge. The detection signal, in this case, is the out-of-balance voltage from the bridge circuit. In this configuration, the effect of lift-off is cancelled out, provided that each coil experiences it to the same degree, however the sensitivity of both coils is still inversely proportional to the distance from the conducting surface. The differential probe is fundamentally a relative system and simply detects a *difference* in the conditions under either of its coils. A limitation arising from this, for the purposes of crack detection, is that if the defect length exceeds the coil spacing then only the leading and trailing edges of the crack will be detected. Reflection probes consist of electrically separate coils: one to induce the eddy currents and one or more sensing coils to detect variations in the induced field. Such probes allow flexibility in terms of optimisation of the driver/pick-up coils thereby offering the potential for high sensitivity to very small defects and their performance can be equivalent to that of the preceding configurations.

The use of an array is an attractive solution in many automated applications for three primary reasons:

- i. Increases the speed of inspection
- ii. Eliminates the need for mechanical scanning
- iii. Provides greater data accuracy.

The latter is true because the fixed reference system of the array provides more accurate information on the position of the coils. The following sections detail the development of three array configurations consisting of varying numbers of differentially connected coils.

6.2.2.2 Design of Coil

Coils are typically wound on a rigid former made from a dielectric (electrically insulating) material with the centre of the former referred to as the core. A probe may contain a high permeability, ferromagnetic element in its core, termed 'ferrite-cored', or may be 'air-cored', with the latter often being used to refer to anything other than the former. Having much the same effect as a ferromagnetic plate, the inclusion of a ferrite core at the centre of a coil winding concentrates the magnetic flux and, therefore, increases the impedance. Because the flux density is greater, defects cause a bigger reduction and therefore ferrite-cored probes are typically more sensitive than their air-cored counterparts. The particulate structure of ferrite means it is electrically non-conducting which also ensures that there are no unwanted eddy current losses within the core itself.

The coil design, for this system, was influenced by the work of Thompson [114] and Friedrich [81]. In [114], it is concluded from a series of experimental tests that, for flaw sensitivity, the coil winding should be positioned between 1.5 mm and 2.0 mm from the end of the ferrite rod nearest the inspecting surface. Furthermore, it is recommended that, where space is not at a premium, the length of the rod should be as large as possible (minimising the winding-core ratio) in order to maximise the inductance. In [81], a single coil pair, phase detection system achieved good results using 400 turn coils of 34SWG wire with former inner and outer diameters of 5 mm and 10 mm respectively. This successful model was used for the differential eddy current array. The coil was wound around a PVC former which provided the recommended coil position and the 1.5mm diameter by 15 mm length ferrite rod that was used allowed a winding-to-core length ratio of approximately 0.46. The impedance and phase of a sample of seven such coils were measured, in air, over a 10 to 500 kHz frequency range (encompassing the core optimum operating frequency of 125kHz) at 500 Hz intervals using a Hewlett Packard 4194A impedance analyser. The resistance and inductance were calculated and are shown in Figure 6.3 (a) and (b). In Figure 6.3 (a), the skin-effect induced, frequency dependence of the coil resistance is quite clearly evident and from (b) the coil reactance can be seen to vary slightly with frequency due to capacitive effects within the coil.


6.3 Analysis of Resistive, Inductive and Capacitive Circuit

Independently, a coil has a relatively limited response to a defect, however this can be enhanced through the use of a resonant configuration such as the RLC network shown in Figure 6.4. It is important to characterise the behaviour of this circuit prior to the design of the system and to develop expressions for the measurable quantities in terms of the basic circuit elements. The following section highlights the key relationships of interest drawn from the detailed derivation provided in Appendix B.


6.3.1 Fundamentals of an RLC Inspection Circuit

The impedance, Z , of a circuit consisting of the series combination of a resistor, inductor and capacitor is given by Equation (6.3).

$$Z = R + j\omega L - \frac{j}{\omega C} \quad (6.3)$$



(a)



(b)

Figure 6.3: (a) Resistance and (b) reactance of a sample of seven coils measured in-air over a 500 kHz bandwidth.




Figure 6.4: Circuit diagram of resistor-inductor-capacitor circuit

where R is the resistance, ω is the angular frequency (rad/s), L is the inductance of the coil (H) and C is the capacitance (F). The magnitude and phase of the impedance is thus:

$$|Z| = \sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C} \right)^2 \right]} \quad (6.4)$$

$$\theta = \text{atan}\left(\frac{\omega^2 LC - 1}{\omega CR} \right)$$

Resonance occurs when the impedance of the circuit is minimum and hence when,

$$\frac{\omega^2 LC - 1}{\omega C} = 0 \Rightarrow f_{res} = \frac{1}{2\pi\sqrt{LC}} \quad (6.5)$$

At resonance, f_{res} , the phase, θ , is 0° and the magnitude of the impedance is purely real. The magnitude of the coil voltage, V_L , for a unity voltage input may be calculated using the voltage divider rule and simplified to Equation (6.6):

$$|V| = \sqrt{\frac{L^2 \omega^4 C^2}{\omega^2 C^2 R^4 + \omega^4 L^2 C^2 - 2\omega^2 LC + 1}} \quad (6.6)$$

The phase, ϕ , of the coil voltage is given by Equation (6.7).

$$\phi = \frac{\pi}{2} - \text{atan}\left(\frac{\omega^2 LC - 1}{\omega CR} \right) \quad (6.7)$$

From Equation (6.7) it can be seen that, at resonance, the phase of the voltage across the inductor will be 90° relative to the supply voltage.


The phase difference between two coils, having inductances L_1 and L_2 , is given by:

$$\Delta\phi = \text{atan}\left(\frac{\omega^2 L_2 C - 1}{\omega CR}\right) - \text{atan}\left(\frac{\omega^2 L_1 C - 1}{\omega CR}\right) \quad (6.8)$$

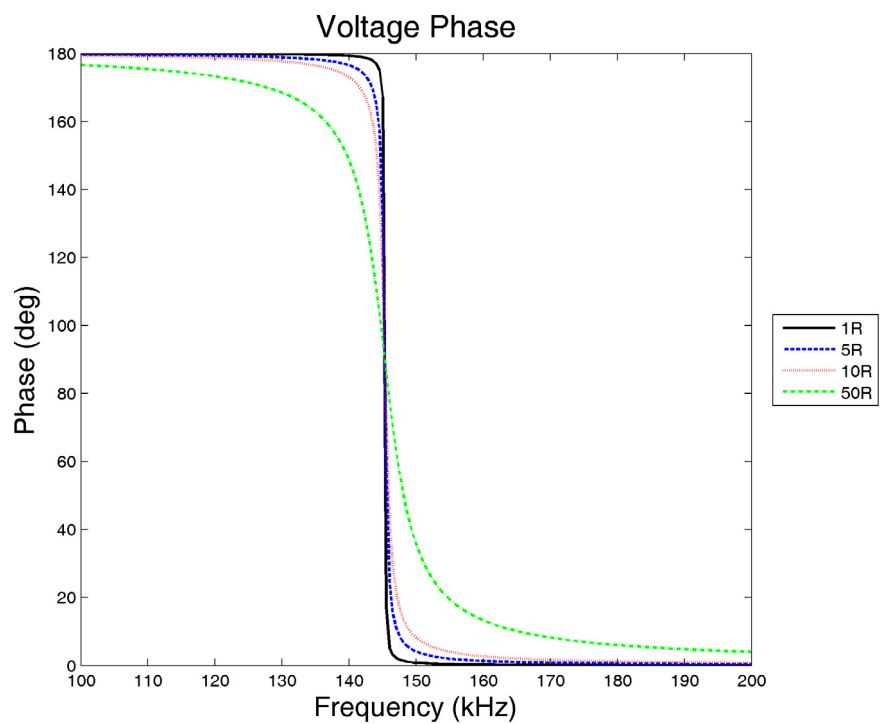
It is this equation that is fundamental to the phase measurement system implemented.

6.3.2 Effect of Series Resistance

The resistor in an RLC circuit essentially acts as damping which controls the rate at which the capacitor charges and discharges and controls the bandwidth over which the phase transitions from 360 to 0 degrees, as shown in Figure 6.5 (a) and (b). For a hypothetical circuit having zero resistance (and ideal reactive components), this bandwidth will tend to zero and the transition will occur entirely at the resonant frequency as illustrated in Figure 6.7. Considering the phase *difference* of any two coils, it can be seen that this will either be 0 degrees (when both coils are identical) or 360 degrees otherwise. Since all eddy current testing is based on a defect generating a change in impedance in the coil then this provides a potentially very sensitive metric for detecting defects. Practically it is not possible to have a circuit with zero resistance and hence the sensitivity to defects will be reduced due to the slower rate at which the phase transitions with frequency. Furthermore, instead of there being a step change in phase difference, a frequency response that rises and falls about a maximum value is obtained, as shown in Figure 6.6 (a) for $R = 1 \Omega$. The maximum phase difference will be related to the relative spacing of the resonant frequencies and hence related to the change in inductance ΔL . As the change in inductance increases then an increasing frequency range will appear in which the coils are in anti-phase ($\Delta\phi = 180^\circ$) as shown by comparing Figure 6.6(a) and (b) for $\Delta L = 100\mu\text{H}$. The minimum inductance change at which such a condition occurs determines the dynamic range of the system, beyond which it saturates and further changes cannot be detected. Since the resistance of a coil is dependent on frequency then, in theory, we can improve the sensitivity of the eddy current system by operating at a lower frequency and taking advantage of the lower RF resistance of the coil in spite of the increased skin depth in the specimen. This has the consequence, however, of increasing the uncertainty of the defect location within the test material since the range of depth being covered has increased. If the resistance is increased then the dynamic range can be increased at the expense of sensitivity as shown in Figure 6.6 (b) for $R = 50 \Omega$.




(a)




(b)

Figure 6.5: Effect of series resistance on (a) magnitude and (b) phase of the coil voltage for $L = 1.2 \text{ mH}$ and $C = 1 \text{ nF}$.



(a)



(b)

Figure 6.6: Dependence of eddy current sensitivity and dynamic range on series resistance (a) $R = 1 \Omega$ and (b) $R = 50 \Omega$, for $L = 1.2 \text{mH}$ and $C = 1 \text{nF}$.




Figure 6.7: Effect of change in inductance for RLC circuit with R tending to 0Ω for $L = 1.2 \text{ mH}$ and $C = 1 \text{ nF}$.

6.3.3 Optimum Frequency Selection

In Section 6.3.2, it was established that, for a given value of series resistance and frequency, there is range of ΔL values over which the phase *difference* will vary in magnitude from 0 to 180° . Beyond a particular value of ΔL , the phase difference will saturate. It follows then that, within this range of inductance, there must exist at least one optimum frequency at which the phase difference is maximised.

Given Equation (6.8), the optimum frequency will occur when the derivative of $\Delta\phi$ with respect to frequency is zero, corresponding to a stationary point in the gradient:

$$\frac{\partial}{\partial \omega} \Delta\phi = \frac{\partial}{\partial \omega} \left[\text{atan} \left(\frac{\omega^2 (L + \Delta L) C - 1}{\omega C R} \right) - \text{atan} \left(\frac{\omega^2 L C - 1}{\omega C R} \right) \right] = 0$$

$$\frac{\partial}{\partial \omega} \Delta \phi = \frac{-\frac{2L}{R} + \frac{\omega^2 LC - 1}{\omega^2 CR}}{1 + \frac{(\omega^2 LC - 1)^2}{\omega^2 C^2 R^2}} + \frac{\frac{2(L + \Delta L)}{R} - \frac{\omega^2 (L + \Delta L)C - 1}{\omega^2 CR}}{1 + \frac{(\omega^2 (L + \Delta L)C - 1)^2}{\omega^2 C^2 R^2}} = 0$$

Solving for ω and taking the positive, real solution the optimum frequency is found to be:

$$\omega_{opt} = \sqrt{\frac{-(2L + \Delta L) + CR^2 + \sqrt{\Delta L^2 + 16L\Delta L + 16L^2 - 2CR^2(2L + \Delta L) + (CR^2)^2}}{2CL(L + \Delta L)}}$$

Neglecting the higher order terms (CR^2) this reduces to Equation (6.9)

$$f_{opt} \approx \frac{1}{2\pi} \sqrt{\frac{-(2L + \Delta L) + \sqrt{\Delta L^2 + 16L\Delta L + 16L^2}}{2CL(L + \Delta L)}} \quad (6.9)$$

which is real for $L + \Delta L \geq 0$ and is independent of R . The error in this approximation increases with the square of resistance and linearly with capacitance. Taking the limit $\Delta L \rightarrow 0$ in Equation (6.9) yields Equation (6.10) which will be recognised as the expression for the coil resonant frequency.

$$\lim_{\Delta L \rightarrow 0} f_{opt} \approx \frac{1}{2\pi} \sqrt{\frac{-(2L) + \sqrt{16L^2}}{2CL(L)}} = \frac{1}{2\pi\sqrt{LC}} \quad (6.10)$$

The approximate solution is shown in Figure 6.8 against ΔL for $L = 1.2$ mH and $C = 1$ nF along with the error in the approximation for $R = 1, 50, 100$ and 500Ω . The optimum frequencies predicted by Equation (6.9) compare well with those in Figure 6.6. The non-linearity in the error due to R can be readily observed. Thus for small values of CR^2 , Equation (6.9) can be taken to approximate the true optimum frequency.

6.3.4 Suitable Defect Metrics

An ideal defect measurement metric is one that changes rapidly with changes in the conditions of interest. Typically, single frequency, differential eddy current inspection either involves measurement of the out-of-balance voltage or the phase difference between a pair of coils. An alternative implementation, using multiple frequencies, involves the measurement of the resonant frequency of the coil, deriving an indication of a defect based on a shift from the normal on-metal resonance. The problem with single frequency systems is that, in addition to the saturation issue outlined in Section 6.3.2, it is unlikely that the inspection frequency




Figure 6.8: Optimum frequency as a function of ΔL for $C = 1$ nF and $L = 1.2$ mH. Secondary axis shows the error due to the approximation for $R = 1, 50, 100$ and 500Ω .

will ever be the most sensitive for a given inspection. The advantage of the resonant frequency method is the fundamental relationship between the resonant frequency and coil inductance. However, for all but the most heavily damped circuits, the phase difference at the optimum frequency will be relatively large, and hence more sensitive, for quite a small shift in resonance. An alternative metric has been identified that provides the benefits of phase detection without the problem of saturation.

Consider the two phase responses in Figure 6.9, corresponding to a hypothetical two-coil situation where one coil is over a defect ($L = 1.2$ mH) and the other is on defect-free metal ($L = 1.5$ mH). As the change in inductance, and hence resonant frequency, is proportional to the size of a defect, it can be seen that for the detection of very small defects the difference in resonance will be very small and hence the resolution of the measurement system must be very high; the phase difference method suffers from a similar stipulation, albeit to a lesser degree, for realistic levels of damping. The most sensitive metric is one that combines the phase information across a range of frequencies and corresponds to calculating the area between the two responses in Figure 6.9. More formally this phase integral, PI, may be written as shown in Equation (6.11).

$$PI = \int_{f_1}^{f_2} \Delta\phi dF \quad (6.11)$$

By making use of all phase information available, small changes in coil inductance produces a relatively large change in PI. A further advantage of this method over a conventional, single-frequency phase difference approach, is in its relative insensitivity to the series resistance of the resonant circuit. As was shown in Section 6.3.2, the resistance has a damping effect and therefore increasing values reduce the size of phase-difference for a given LC configuration. As illustrated in Figure 6.10 (a), for a low resistance (10 Ω) the phase transition occurs over a short bandwidth and the peak phase difference reaches a relatively high value. Conversely, for a higher resistance (100 Ω) the phase transitions much less rapidly and the peak phase-difference reduces by 70%. In both instances, however, this reciprocal nature leads to only a 2% variation in PI using Equation (6.11) and, consequently, a much more stable defect metric.

If both coils are positioned over identical material, then the relative phase difference will be zero for all frequencies. Assuming that the resistance is constant, as the defect depth approaches zero the change in inductance will tend to zero and the optimum frequency will tend to the resonant frequency of the coil for zero lift-off, f_{zero} . The response of the coil to a 100% through thickness defect can be modelled by the coil response in air, thus the second frequency limit occurs at the resonant frequency of the coil in air, f_{air} . As the change in inductance due to the presence of ferromagnetic and non-ferromagnetic materials is of opposite polarity, the required frequency sweep range in order to guarantee at least one optimum operating point is given by Table 6.1. In reality, the resistance will change with the defect size and furthermore, since the PI metric uses data from all frequencies that yield a non-zero phase difference, a wider frequency sweep range will be required. For the purposes of this system, it was found that a sweep range of 100 to 200 kHz was suitable.

Table 6.1: Optimum frequency sweep range

Metal Type	Frequency Range
Non-ferromagnetic	$f_{air} < f < f_{zero}$
Ferromagnetic	$f_{zero} < f < f_{air}$




Figure 6.9: Phase response of two coils for varying frequency. The large change in area between the curves for small inductance changes highlights the superiority of this metric compared with phase difference, $\Delta\Phi$, or resonant frequency, Δf , measurement methods shown.




Figure 6.10: Sensitivity of the RLC series resistance on the PI and single frequency phase difference defect metrics.

6.4 Finite Element Modelling

6.4.1 Introduction

Having established a relationship between the measurable defect metric and the coil parameters, the logical progression is to extend this to relate to the actual defect geometry. The use of finite element modelling (FEM) for the characterisation of eddy current coils for NDE is particularly attractive since it allows the response to a variety of defect geometries to be explored that would otherwise be mathematically intractable to solve. The system can be modelled using a variety of FEM software packages including COMSOL Multiphysics [115], Vector Fields Software [116], Quickfield [117] and Maxwell 2D [118] or semi-analytical packages such as EXTENDE CIVA [119]. The advantage of the latter is that it is relatively fast compared with FEM techniques and can model complex geometry more efficiently. As the novelty of the ECA is in its integration with the robotic platform, it was deemed sufficient to consider only simple geometry. Furthermore, since speed was not critical and due to the well-established knowledge-base within the research group, a FEM approach was adopted. As with all numerical approximations, there is an element of error introduced and a trade-off exists between the size of this error and the size of the model; with larger models requiring more time to execute. Furthermore, as it does not produce a closed-form solution, the response to varying geometrical parameters can only be obtained by collating the results from a multitude of individual models. This can be a time consuming but, ultimately, necessary process.

6.4.2 Basic Model

6.4.2.1 Overview

The COMSOL ‘Azimuthal Induction Currents’ application mode of the AC/DC module was used and, for reasons of computational efficiency, a 2D axi-symmetric model was defined. Constraining the model to two dimensions also limits the modelled defects to axi-symmetric geometry however for the simple analysis that will be investigated in this instance, the approximation is justified. The basic model consists of five elements, or ‘sub-domains’ as shown in Figure 6.11: the ferrite rod, wire, metal block, defect and the free-space region which encompasses the entire model. The z-axis acts as the axis of symmetry around which the model is revolved. While the use of the inherent symmetry provides a convenient means to simplify the problem, it does also limit the nature of the defects that can be modelled to

those that are similarly axi-symmetric. The free space region has a magnetic insulation boundary condition which sets the magnetic potential to zero along this boundary. The defect sub-domain is modelled as a free space region having the same material properties as the free-space sub-domain minus the magnetic boundary condition. The number of turns, N , the geometry of the defect and all material properties can be configured, allowing a plethora of inspection scenarios to be simulated. Table 6.2 shows the properties assigned to each of the sub-domains. For an aluminium block, the conductivity is $3.7 \times 10^7 \text{ S m}^{-1}$ and the relative permeability is one. In order to realistically model the system, and thereby achieve an accurate prediction of the coil inductance, each coil loop must be modelled as a separate sub-domain and the skin-effect accounted for. The input current is defined as 170 mA, equivalent to the maximum output current from a buffer amplifier [120]. This affects the magnitude of the eddy currents in the plate however does not affect the flux density/linkage and consequently does not affect the inductance of the coil (neglecting thermal effects). The total current in each coil loop consists of two components: the input current (I_0) and the induced current, the latter of which is calculated within the model. As the helical nature of the coil winding cannot be represented geometrically in a 2D plane, a series of global constraints are applied to ensure the current in each conductor is constant. An ordinary differential equation (ODE) is defined for each loop voltage, V_N , which adapts so that the loop current, I_N , equals the input current. The loop current is defined as:

$$I_N = \int_S J_\phi \quad (6.12)$$

where J_ϕ is the total current density in the sub-domain of area S . The total voltage across the coil is given by the sum of the loop potentials as defined in Equation (6.13).

$$V_{total} = \sum_{n=1}^N V_n \quad (6.13)$$

Assuming the capacitive effects to be negligible, the coil inductance, L , and resistance, R , are given by the imaginary and real parts of the impedance respectively:

$$L = \text{imag} \left(\frac{V_{total}}{2\pi f I_0} \right) \quad (6.14)$$

$$R = \text{real} \left(\frac{V_{total}}{2\pi f I_0} \right) \quad (6.15)$$

where f is the excitation frequency (Hz). For a N -turn coil wound around a cylindrical ferrite core, the inductance can be calculated using Equation (6.16).

$$L = \frac{K\mu_o\mu_{rod}N^2A_e10^4}{l} \quad (6.16)$$

where, K is the inductance factor, μ_o is the permeability of free-space, μ_{rod} is the rod permeability, A_e is the cross-sectional area of the core and l is the length of the core taken up by the coil winding. The rod permeability differs from the material relative permeability in that it describes the increase in inductance between an air-cored and ferrite rod cored coil rather than the actual relative permeability of the material; this is typically much higher. The inductance factor takes account of the extent to which length of the rod is occupied by the current carrying conductor and is provided by the core manufacturer. For the selected core, having $K = 1.6$ (corresponding to a winding-core ratio of 0.46), $\mu_{rod} = 48$, $A_e = 0.01770 \text{ cm}^2$ and $l = 1.5 \text{ cm}$, the calculated inductance is 1.82 mH compared with a simulated prediction of 1.4 mH. This discrepancy is possibly due to the fact that Equation (6.16) doesn't take account of the wire gauge which is 36 SWG (0.2 mm diameter).




Figure 6.11: Basic model consisting of a coil wound around a ferrite core mounted on top of a metal block containing a defect. The number of turns, nature of the defect and all material properties can be configured.

Property/ Sub-domain	Relative Permeability μ_r	External Current Density J_ϕ^e (A m ⁻²)	Conductivity σ (S m ⁻¹)	Relative Permittivity ϵ_r
Ferrite Core	2300	0	0	1
Metal Block	μ_{Block}	0	σ_{Block}	1
Defect	1	0	0	1
Coil	1	0	5.7×10^7	1

Table 6.2: Properties of model sub-domains

6.4.2.2 Mesh Sensitivity Analysis

In finite element modelling, a mesh is a discretisation of the model geometry into a series of small elements of some simple shape for the purposes of computation. As with any such process, the accuracy of the resulting model is dependent on the level of approximation and hence the mesh element size. Within COMSOL it is possible to specify several parameters to the *free mesh* algorithm that specify the nature of the resulting mesh. The *maximum element size* defines, unsurprisingly, the maximum size of any element that may be used and therefore is the minimum resolution of the model. It is possible to assign different values to particular sub-domains and edges in order to increase/decrease the local resolution. In the absence of explicit instructions, the meshing engine determines the element size throughout the various regions of the model, ensuring that this is always below the defined maximum. The *element growth rate* defines the rate at which the size of the elements can increase from one element to the next. The greater this value, the larger the potential variation in size between adjacent elements. This parameter is particularly important where there is a significant difference between the largest and smallest elements used in the model.

A mesh sensitivity analysis is important in order to obtain a quantitative estimation of the scale of the meshing error. The basic model shown in Figure 6.11 is used for the study, since it contains a defect in the centre of the metal block and therefore is likely to have the greatest meshing requirements. Initially, the default growth rate of 50% is kept constant and the element size varied from 7.5×10^{-5} to 5×10^{-3} . Taking

the finest mesh to exemplify the best approximation to the inductance, L_{best} , the ‘error’ in the coil inductance, L_{error} , is calculated as the difference in inductance for a given mesh size relative to this.

$$L_{error} = \frac{L - L_{best}}{L_{best}} \times 100 \quad (6.17)$$

where L is the calculated inductance for a model with a maximum mesh size of 0.1, 0.25, 0.5, 0.75, 1.0, 2.5 and 5.0 mm. The mean error and standard deviation is calculated for 15 different models in which the position of the defect relative to the axis of symmetry is varied. Table 6.3 shows that the percentage error due to the maximum mesh size is relatively small at <1.22%. This is understandable given the fact that the finer regions of geometry will automatically have a denser mesh applied by the meshing algorithm.

Max. Element Size (mm)	Mean Error (%)	Standard Deviation (%)
0.1	0.16	0.0047
0.25	0.72	0.0128
0.5	1.05	0.0155
0.75	1.16	0.0139
1.0	1.22	0.0226
2.5	1.22	0.0214
5.0	1.22	0.0214

Table 6.3: Mean and standard deviation of error due to max. element size

Growth Rate (%)	Mean Error (%)	Standard Deviation (%)
1	0.02	0.0052
5	0.13	0.0100
10	0.51	0.0109
25	2.40	0.0391
50	6.86	0.1515

Table 6.4: Mean and standard deviation of error due to growth rate

Thus the maximum element size will only be significant in the less densely populated regions of the model geometry, such as the upper-right region of the

model boundary sub-domain. A value of 7.5×10^{-4} was selected since this ensured an error of only $\sim 1.1\%$ with an acceptable standard deviation.

The mean error was subsequently calculated using this new maximum element size of 7.5×10^{-4} for element growth rates of 1%, 5%, 10%, 25% and 50% relative to an assumed best approximation value of 0.5%. It can be seen from Table 6.4 that for growth rates of 10% and below the mean error is less than $\sim 0.5\%$ where-as for 25% and 50% it is significantly larger, with a significant increase in the standard deviation. A growth rate of 5% was deemed to be a suitable compromise between inductance error and model size (and hence execution time).

6.4.3 Lift-off

The response of the coil to lift-off from an aluminium surface is shown in Figure 6.12 for three frequencies: 100 kHz, 150 kHz and 200 kHz. In this instance, the original model is modified such that the defect width is that of the metal block width, and consequently as the depth of this is increased, the air-gap (and hence lift-off) is increased. The inductance follows a first-order response from a contact inductance of 1.07 mH tending towards an in-air value of 1.4 mH. As would be expected, the gradient is steepest for lower levels of lift-off indicating that even small differences in the height of two coils above a surface will result in a relatively large change in inductance and hence phase difference. The effect of frequency is relatively small on the predicted inductance, however there is a significant variation in the coil resistance between the three frequency intervals, resulting from the skin effect within the coil. The variation due to the defect is evident at each frequency interval and also follows a first order response. Figure 6.13 (a) shows the qualitative distributions of induced current density for a coil placed directly on top of an aluminium surface while (b) shows the same but for a lift-off of 3 mm. It can be seen that the current density within the metal block reduces when there is lift-off resulting in a lower opposing magnetic field and hence a larger coil inductance.

6.4.4 Sensitivity to Lateral Position and Radius

As discussed in Section 6.2.1, the presence of a ferrite core concentrates the magnetic flux and hence results in the existence of a region of peak sensitivity along the probing face of the coil. Any defect within this region which is sufficiently large that it interrupts the eddy currents in the conducting material would be expected to have a more significant impact on the measured coil inductance than in any other region. The sensitivity profile of a coil to the lateral position of a 6 mm deep defect

is quantified in Figure 6.14 at 200 kHz for three defect radii: 1 mm, 3 mm and 6mm. Noting the axi-symmetric constraint, once again, for each defect size the *inside radius* is varied from 0.1 mm to 12 mm in 0.1 mm increments. It is evident, from the inductive response, that a peak exists for a defect positioned at 0.8mm from the centre of the coil in all cases and that for a radius greater than 10mm the response of the coil is negligible. As would be expected, the larger the width of the defect the less conducting area for the eddy currents to occupy and thus the greater the relative change in coil inductance. As the radius of the ferrite core is 0.75 mm, the region of maximum sensitivity lies beneath the boundary between the core and the winding. The response of the coil resistance is more complex. For a 1mm defect, the maximum occurs at 0.8 mm, in agreement with the inductive results, however as the width is increased the distinction of this point reduces becoming a local maximum for a 3 mm width and ultimately just a transition point in the rate of change of the coil resistance for a 6 mm width. Figure 6.15 (a) and (b) shows the qualitative eddy current distribution for a 1mm wide slot, of the same depth, positioned at a radius of 0.8 mm and 6 mm respectively.

Figure 6.16 shows the effect of varying the radius of a 1 mm, 3 mm and 6 mm deep defect from 0.1 mm to 15 mm, again in increments of 0.1 mm. As the defect size increases beyond that of the coil itself, the model tends to that of a lift-off model and consequently the coil parameters approaches that predicted by Figure 6.12 for a lift-off equal to the defect depth. Figure 6.17 (a) and (b) shows the qualitative eddy current distribution for a 1.8x3 mm deep and a 10x3 mm deep defect respectively.

6.4.5 Sub-Surface Defect

In some instances, a defect, such as erosion in the inner surface of a pipe or delamination in a multi-layered structure, does not emanate from the inspection surface. To assess the performance on sub-surface defects, a 1 mm thick cavity was modelled for varying depths below the inspection surface. As there is a conducting layer above the defect and in direct contact with the coil, eddy currents will be generated in this region meaning the reduction in coil inductance due to the void will be much less significant than for the previous cases. This is exemplified in Figure 6.18 where the change in inductance can be seen to be very small and sensitivity confined to a region less than 0.5 mm below the surface. The change in resistance is more significant with a change of approximately 20 Ω however with a similar level of depth sensitivity. As for the previous examples, the qualitative eddy current distribution is shown, this time for a frequency of 100 kHz in Figure 6.19




Figure 6.12: Lift-off characteristics as predicted by COMSOL simulation for 100kHz, 150kHz and 200kHz.

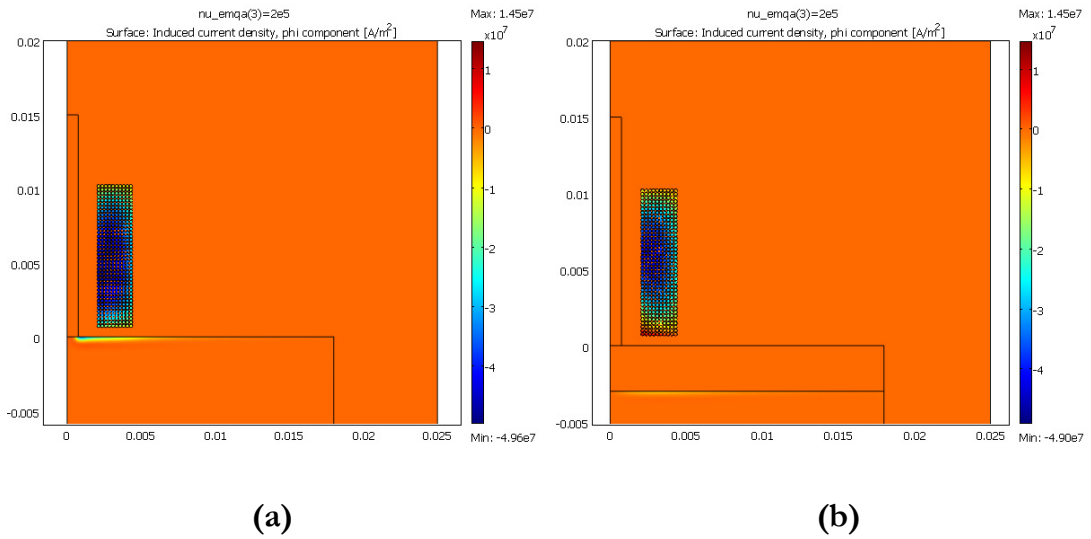


Figure 6.13: (a) Qualitative distribution of eddy currents at 200 kHz for a probe in contact with an aluminium plate and (b) for a lift-off of 3mm.




Figure 6.14: Effect of position of a slot from the centre of the probe on the coil inductance and resistance for three widths of defects at 200 kHz: 1mm, 3mm and 6mm at 200kHz.

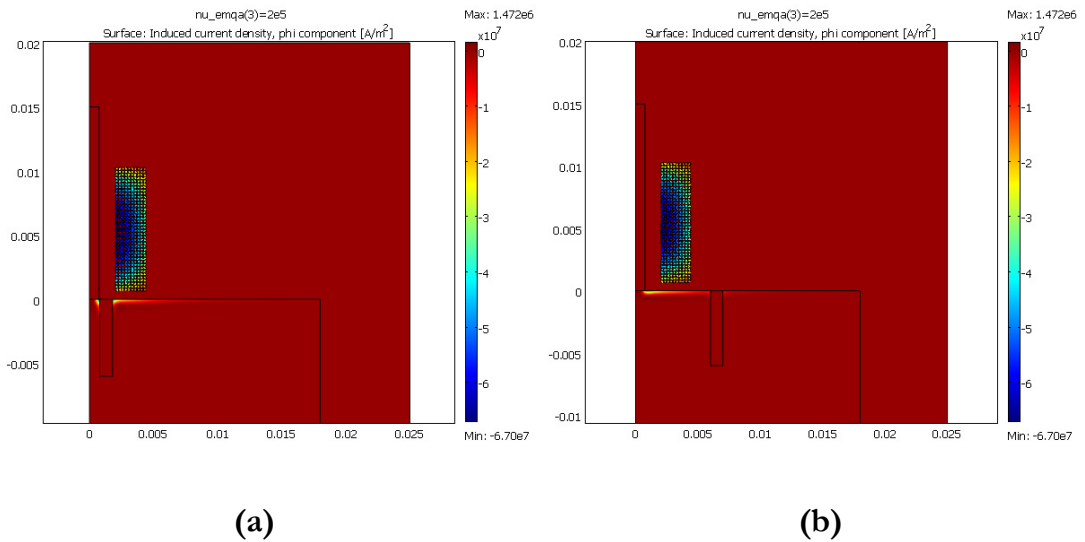


Figure 6.15: (a) Qualitative distribution of eddy currents at 200 kHz for 1 x 6mm defect positioned at 0.8 mm from centre of the coil. (b) shows the induced currents for the same defect at 6mm where the inductance has almost reached its ‘on aluminium’ value.




Figure 6.16: Effect of the radius of a 10mm deep hole from the centre of the probe on the coil inductance and resistance at 200 kHz.

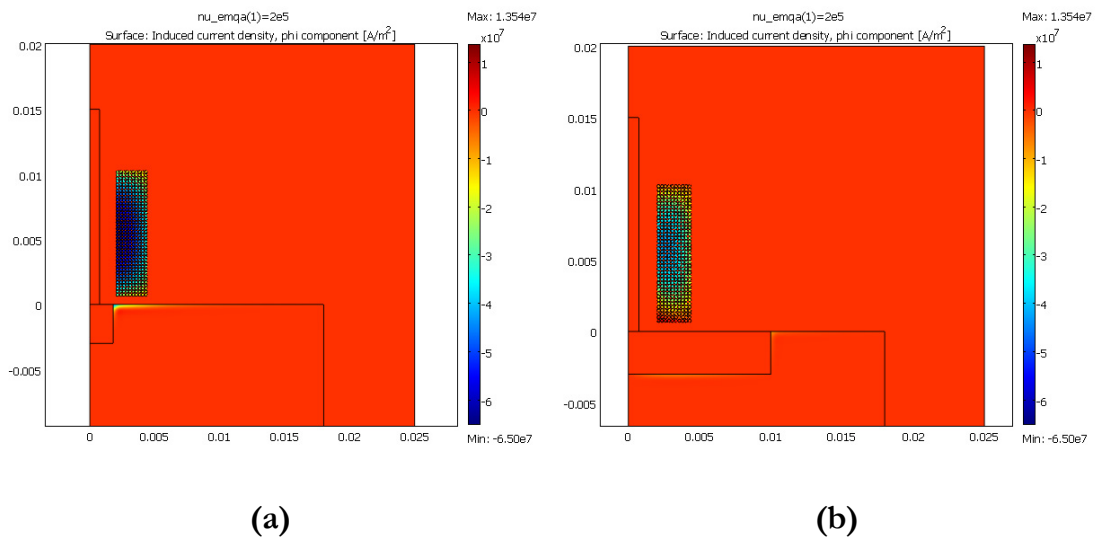


Figure 6.17: Qualitative distribution of eddy currents at 200 kHz for coil centred on a (a) 1.8 mm radius x 3 mm deep defect a (b)10 mm radius x 3 mm.





Figure 6.18: Effect of the depth of a sub-surface cavity from the probing surface on coil inductance and resistance at 100 kHz.



(a)

(b)

Figure 6.19: Qualitative distribution of eddy currents at a probing frequency of 100kHz for an intermediate defect positioned at (a) 0.1mm and (b) 0.5mm below the probing surface.


(a) and (b) for a cavity located 0.1 mm and 0.5 mm beneath the probing surface respectively. Given the ideal nature of the model and the limited change in inductance (and hence resonant frequency), it can be concluded that by using the frequency range 100-200 kHz, it is unlikely that such a defect could be detected in practice.

6.4.6 Depth Detection Using Frequency Sweep

In this section, the theoretical analysis outlined in Section 6.3.1 is combined with the FEM results by inputting the coil resistance and inductance predicted by the model into Equation (6.8) for a fixed circuit capacitance of 1 nF. This is performed for different defect geometries. Figure 6.20 (a) shows that by increasing the defect depth while keeping the radius fixed at 5 mm, the peak phase difference increases and the frequency at which this occurs varies. This is consistent with the analysis in Section 6.3. As predicted, the greater the series resistance, the lower the peak phase difference and wider the range over which the phase difference is non-zero. Figure 6.20 (b) shows the computed PI against defect depth for various defect radii in aluminium. The effect of the fixed RLC resistance is shown for a value of 10 ohms (solid line) and for 100 Ω (dashed line). There is a clear relationship between PI and defect depth, given a fixed radius, and the effect of series resistance can be seen to be minimal. However, it can be seen that, for a given value of PI, there exists multiple solutions for the defect geometry meaning that precise defect sizing cannot be obtained from a single probe position. Similar results are evident in Figure 6.21 for 304 stainless steel ($\sigma = 1.39 \times 10^6$ and $\mu_r = 1.008$) albeit with lower magnitudes, consistent with that expected from the significantly lower material conductivity. A two-stage imaging process where-by a first pass estimates the defect dimensions and the second stage re-processes the data to calculate the depth based on this may be possible although this would be complicated in the, quite likely, situation where a defect straddles more than one coil. *A priori* information of the defect surface dimensions, possibly derived from a visual payload, could be used in conjunction with the PI to identify the depth of the defect. The combination of visual and ECA data in this way results in a system which has greater NDE value than each sensor provides individually. The concept of sensor fusion will be explored in greater detail in Chapter 6.


Having derived a sound theoretical basis for a frequency agile ECA system for defect sizing, Section 6.5 will now outline the development of the system.

Theoretical Phase Difference for Varying Defect Depths



(a)

Theoretical PI for Varying Defect Geometry



(b)

Figure 6.20: (a) Phase versus frequency plots for different hole depths in aluminum and (b) the phase integral (PI) plotted against hole depth for varying defect diameters. In (b), the solid and dashed lines represent a RLC fixed resistance of 10Ω and 100Ω respectively.




Figure 6.21: Theoretical phase integral (PI) plotted against defect depth for varying defect diameters in 304 austenitic stainless steel. The solid and dashed lines represent a RLC fixed resistance of 10 Ω and 100 Ω respectively.

6.5 System Overview

6.5.1 Introduction

The eddy current array payload consists of five functional units, as shown in Figure 6.22: signal generation, drive signal conditioning, coil selection, received signal conditioning and digital signal processing (DSP). The hardware can be separated into four modules: the main motherboard, a combined DSP/function generator (DSP/FG) unit, a received signal conditioning board and the array probe unit. The DSP/FG board is a multi-purpose, flexible processing unit developed within CUE for mobile robotics applications and docks into the main motherboard. It contains a DSPIC33F family microprocessor [121] interfaced with an AD5930 programmable sweep function generator [122] and performs both the signal generation and DSP functions. The DSP/FG board slots into the main eddy current array motherboard containing the electronics that provide the drive signal conditioning. The eddy current coils are mounted on the separate array probe unit which contains a bank of multiplexers (MUXs) that determine which coils are energised. The probe is connected to the main motherboard via a 10-wire custom interface which allows it to be situated remotely from the main electronics housing. As the MUXs are mounted on the probe itself, the interface is standard, regardless of the number of coils, and therefore different probe arrangements can be connected to the motherboard by simply defining the number of coils in the main GUI and eliminating the need for any hardware modifications. The received signal conditioning board provides a hardware filter and amplification stage for the coil signals and, as with the DSP/FG board, docks directly into the main motherboard. The loop is completed with a direct connection between the output of the received signal conditioning board and the analogue-to-digital conversion (ADC) input on the DSP/FG unit. The final, digital signal processing, stage then takes place on the microprocessor. The eddy current array motherboard forms a layer of the main robot electronics stack and interfaces through the primary signal bus. The communications with the robot central processor occurs over an RS232 interface. Separate ground and power sources are used for digital and analogue components throughout, to minimise the effect of noise in the measurement system.

6.5.2 Signal Generation

The AD5930 programmable sweep function generator is programmed to produce a frequency sweep from 100-200 kHz in 1 kHz increments. As the signal generation and sampling is performed on the same unit (i.e. the DSPIC processor), each

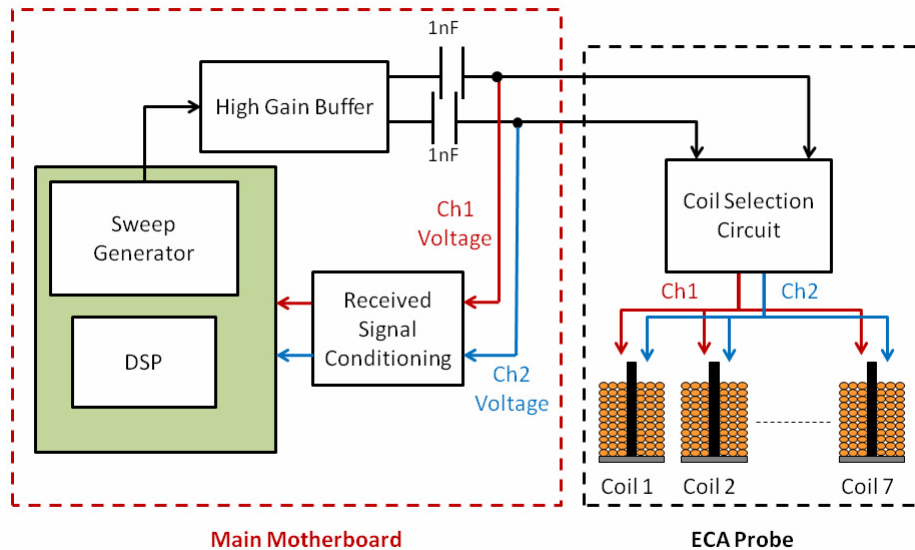


Figure 6.22: System overview separated into the functional units implemented on the RPC electronics stack and on the probe.

frequency interval can be restricted so that it is output only for the length of time taken by the processor to sample the signal, thereby keeping power consumption to a minimum. The frequency sweep parameters are set upon initialisation of the main payload software which occurs at start up. All communications between the DSPIC processor and the AD5930 takes place over a serial peripheral interface (SPI) in a master-slave arrangement, with the latter being the slave. The burst period and frequency increment occur in unison through transitioning the GPIO connected to the CNTRL pin of the FG board either low-high (start output/increment frequency) or high-low (stop output). The output signal is a $3.3 V_{pk-pk}$ DC biased to 1.65 V with a maximum current of 4 mA.

6.5.3 Drive Signal Conditioning

Each coil is connected in an LC series resonant circuit and forms one arm of a balanced Wheatstone bridge. The relatively low maximum current output of the function generator stage is insufficient to drive such a circuit with a reasonable degree of sensitivity and therefore the signal is first fed through a LMH6718 programmable gain buffer [120], capable of producing 170mA with a bandwidth of 110 MHz from a single +5 V supply. When considering how best to excite the coils, a trade-off is required between using a current that is high enough to generate a strong field in the test specimen but which is adequately small that the thermal effect of coil expansion on the coil inductance can be neglected. Such thermal

effects are important to avoid, in order to prevent time varying measurement errors. This issue is more pertinent in ferrite-cored probes which can also suffer from magnetic hysteresis effects due to over-excitation. A variable, ‘tuning’, resistor is connected in series with the output of the buffer, as shown in Figure 6.23, to allow the current to be adjusted. The buffer is configured in inverting mode with a gain of one and the input biased to 2.5V using resistors R3 and R4 to conform with the single supply operation. This bias is removed by the capacitor preceding the coil to ensure there is no additional heating of the inductive element due to the DC excitation.

6.5.4 Coil Selection

The two coil drive signals, ‘drive-signal 1’ and ‘drive-signal 2’, form part of the 10-wire probe interface, along with the communications, power and ground. Typical MUX current ratings are in the region of tens of mA and consequently cannot be used to directly drive the eddy current coils of a resonant circuit. Instead, a dual-bank, 16 channel multiplexer [123] is used to drive the activation LED of an opto-isolated MOSFET [124] as shown in Figure 6.24. Bank A controls the switching of drive-signal 1 to any of the coils in a particular probe (maximum of 16) and bank B performs the same task for drive-signal 2. The 0.7 ohm on-resistance of the MOSFET device ensures there is minimal unwanted damping effect on the LC circuit and has the advantage that the gate signal is completely separate from the switched signal thus minimising potential cross-talk. Since each coil can be connected to either of the drive signals, it is necessary to have a dedicated MOSFET for each channel. A dual chip package is used to minimise the form factor of the probe electronics.

The switching process is controlled, for reasons of simplicity of integration with the main robot software, by the main robot processor through the I2C interface however this operation could ultimately be shifted to the DSPIC processor in order to produce a self-contained unit.

6.5.5 Received Signal Conditioning

The received signal conditioning hardware is a two-stage buffer-filter combination, shown in Figure 6.25, that docks into the EC array motherboard. While it is not strictly correct to refer to the coil signal as a ‘received’ signal as in the usual context of a reflection probe (since it is actually the same as the coil drive signal), it is helpful here to use such terminology to distinguish the stage at which the signal

conditioning occurs. The unity gain buffer has a 1.65V bias applied to the non-inverting input in order to produce a DC offset output signal suitable for subsequent input into the ADC stage. As the inductive qualities of the coil harbour the high possibility for electrical interference from the environment, and following good design practice for the front end of a sampling system, a lowpass filter stage bandlimits the signal to the range of 0 to 300kHz, set by components R_{12} , C_2 . The upper limit is selected so as to ensure a relatively flat response over the sweep range.

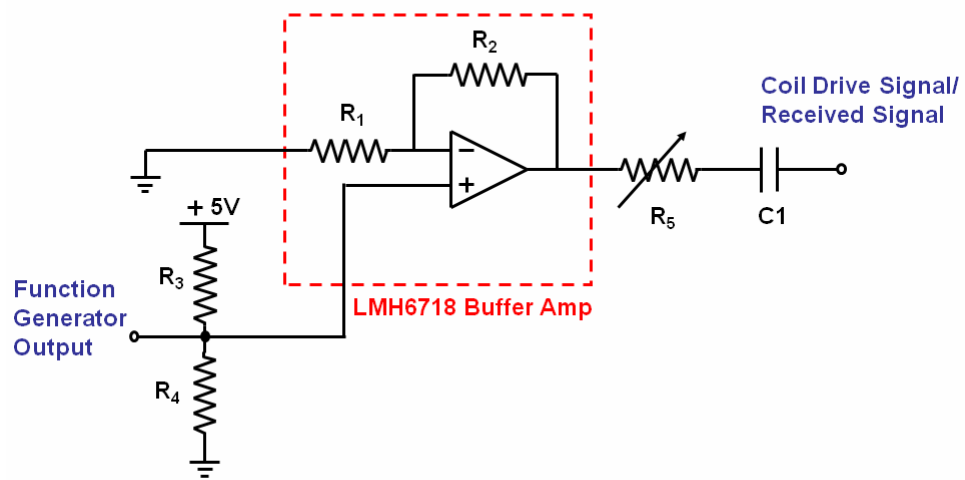


Figure 6.23: Coil drive circuit diagram consisting of a unity-gain buffer amplifier feeding into the resistive and capacitive elements of the RLC circuit.

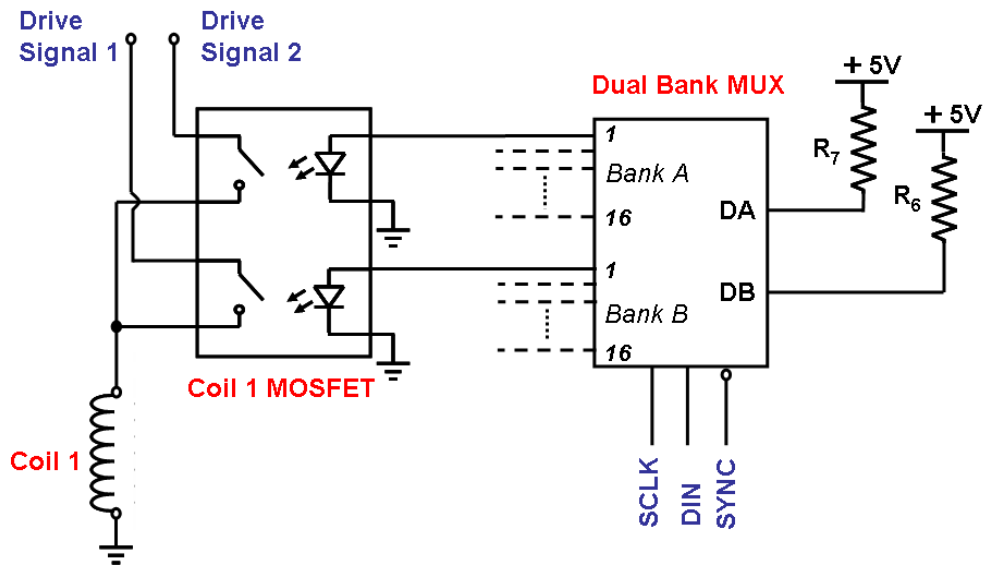


Figure 6.24: Coil selection circuit with the connections for channel 1 illustrated.




Figure 6.25: Received signal conditioning circuit.

6.5.6 Digital Signal Processing (DSP)

The DSP unit comprises of three stages: data capture, upsampling and phase measurement.

Data capture is performed using the integrated ADC chip of the DSPIC33F processor, configured for 10-bit operation. Each eddy current signal is fed into a separate ADC channel, both of which are sampled simultaneously when instructed by the eddy current program. The inputs to the ADC have built-in protection diodes which prevent damage occurring in the event of a temporary over-voltage. Each sample is stored in a 16-bit integer format within the 256 byte ADC sample register and, consequently, only 128 data-points are contained within. The particular DSPIC used has a 2.2 Msps ADC yielding a Nyquist frequency of 1.1 MHz. Given the 100 to 200 kHz sweep range, this equates to between 11 and 5 sample points per period respectively. Clearly such a low number of sample points is insufficient for an accurate phase comparison system and would yield considerable rounding error. It is therefore necessary to perform upsampling.

Upsampling is the process of increasing the sample rate of a signal in post-processing. After a signal has been captured using an ADC at a sample rate F_s , then, for a desired sample rate of NF_s , $N-1$ zeros are inserted between the original samples as shown in Figure 6.26 for $N = 2$. This process is known as ‘zero-padding’ and intuitively it can be seen that, since the time window of the signal has not changed but there are now N -times more sample points, it has resulted in an increase in the sample rate. When a signal is sampled at a frequency F_s , spectral repetitions appear at integer intervals of the Nyquist rate of $\frac{F_s}{2}$. While the process of zero-padding does not change the frequency content of the signal, since the sample rate has now been increased by a factor, N , it does result in the appearance of spectral repetitions below the new Nyquist rate of $\frac{NF_s}{2}$. Furthermore, the addition of zeros in the signal, on its own, is not particularly helpful for the purpose of gaining more points of reference to do phase measurements from. In order to remove these aliased frequency components and return the signal to something similar to the original, an ‘interpolating’ filter is used to filter out frequency components greater than $\frac{F_s}{2}$. In the time domain, this low-pass filter interpolates the zero samples inserted between the original points, thereby recovering the original signal but sampled at a higher rate. A multi-rate system ensures the most efficient upsampling system and therefore a different filter is needed for each upsampling stage. This process is shown in Figure 6.27 for a typical captured signal upsampled by a factor of two




Figure 6.26: Depiction of upsampling process. Signal is upsampled by a factor of two yielding an effective increase in the sample rate.




Figure 6.27: Signal captured by ADC before and after upsampling by a factor of two.

The original signal captured by the ADC is first fed into a peak detector function which determines how many periods of the signal were captured. In order to allow averaging of the subsequent phase measurements to be performed, it is desirable to have several periods of each waveform available. However, there is a trade-off between storing this additional data and utilising that memory for upsampling a more limited number of periods and hence increasing the resolution of each of the individual phase measurements. Due to the finite memory resources available on the chip, the maximum size of each upsampled signal buffer is 1030 samples. Although the data size doubles with each upsampling stage, the removal of the group delay and limiting the maximum number of periods that are stored mean that a maximum of four upsampling stages can be performed if necessary. A heuristic approach was employed to define a maximum of 9 periods and thus, if the original data contains more than this, only those are retained. In all cases the data is upsampled at 2.2 MHz, 4.4 MHz and 8.8 MHz with an optional 17.6 MHz stage. In real terms this equates to three stages of upsampling for frequencies between 100 and 140 kHz and four for frequencies from 141 to 200 kHz. It will be obvious that sampling a frequency sweep will generate a number of sample points per period which decreases as the frequency increases. This in turn results in a decreasing phase measurement resolution. While the upsampling regime does not solve this problem, it does ensure that the average number of point per period over the entire sweep is higher and therefore the average resolution of the phase measurement is greater.

Following the initial capture and peak detection stage, the following steps are repeated each time the signal is upsampled with the only difference being the filter that is used. The signal is first scaled in software to occupy the full 16-bit integer resolution and zero padded, before being shifted through the digital low-pass filter. Each filter is an equi-ripple low pass implementation ensuring a sharp cut-off frequency and passband and stop band attenuations of 0 and 30 dB respectively. Following this, the filter group delay is removed and the process repeated for either three or four times as is appropriate. The filter length and group delay of each filter is shown in Table 6.5. The magnitude response of the 4.4 MHz filter is shown in Figure 6.28.

Once the signal has been processed to ensure there are a sufficient number of data points per period, the phase measurement algorithm is invoked. As the location of the peaks will have been affected by the upsampling process, the same peak detection routine is applied for a second time to ensure that spurious peaks are not registered. The average period across both signals is calculated and used to calculate

the phase difference between the peaks of signals one and two. The phase is then wrapped to -179 to +180 degrees by adding/subtracting one period length as necessary. The average of all the phase values for this particular frequency interval is then stored as the resulting phase measurement. Prevention measures are employed to ensure overflow does not occur and to prevent excessive rounding errors which could result in incorrect measurements.

Table 6.5: Group delay for each multi-rate filter.

Filter Frequency (MHz)	Filter taps	Group Delay
2.2	16	8
4.4	32	16
8.8	64	32
17.6	151	75.5

6.5.7 Phase Measurement Error

The accuracy of the software phase detection system was quantified at 100 kHz, 150 kHz, 200 kHz, 250 kHz and 300 kHz using a 3V sinusoidal input to the ADC from a dual channel signal generator. The phase was set on the signal generator and the EC array software routine, outlined in Section 6.5.6, triggered. Five samples were taken at each 5 kHz frequency increment over the full range of phase difference and the mean error calculated. A normal distribution was assumed and is shown in Figure 6.29 for the tested frequencies. The observed bias is due to the rounding down preference adopted in the algorithm. Aside from a phase overlap issue at phases close to $\pm 180^\circ$, which were omitted in this instance, the magnitude of the mean error in the 100 kHz to 200 kHz range is less than 1 degree, indicating good accuracy from the software algorithm. As the frequency is increased there is a general trend towards greater mean error coupled with a greater standard deviation due to the reduction in data points available per period within the operational frequency range of the ECA system however, the error is acceptable.




Figure 6.28: Active bandpass filter frequency response for 4.4MHz upsampling stage. The Direct-Form FIR implementation passband is between 0 and 300kHz and the stop frequency is 400kHz.




Figure 6.29: Phase measurement error for five frequency intervals. The mean error of the hardware/software system is shown to be <1% for the frequency range 100 to 200 kHz.

6.6 Transducers

6.6.1 Linear Array

Initial research centred on the development of a linear array and was focussed on the inspection of a flat aluminium plate. The geometry and arrangement of the array both play a critical role in the ultimate performance of the device. Considering defects whose geometry in one dimension is significantly greater than in another, such as fatigue cracks, consistent successful detection at the most basic level necessitates that the inter-coil spacing either be smaller than the lesser dimension of the defect or a multi-row array be employed where no line across the probe face fails to cross a coil. The corollary of the former is that there can be no mechanical support structure to hold the coils and consequently the only realisable solution is to have multiple rows with each offset such that the centre of each coil in row $n + 1$ is equidistant between each pair of coils in row n . The minimum requirement is two rows and, as each row is added in the scanning direction, additional rows allow greater inspection coverage per probe position (and hence speed) but no greater defect detection capability. To minimise the multiplexing requirements, a dual row design was adopted as shown in Figure 6.30.

6.6.2 Circular Array

An alternative design, termed the ‘circular array’, was explored which differs from the linear design by having a central coil with six satellite coils at 30° intervals. In this instance, the important design parameter is the radius of the satellites relative to the central coil, again with the stipulation that no line must exist across the probe face which does not intersect with at least a single coil. An exploded diagram of the model and the final realised design is shown in Figure 6.31 (a) and (b) respectively. The advantage of this probe over its linear counterpart is in situations where there are a series of points of interest and where a larger field of inspection at each of these points is valuable.

6.6.3 Pipe Inspection Array (Fusion Probe)

For pipe inspection scenarios it is necessary to have a probe that conforms to the diameter of the pipe under test. A single probe, containing spring loaded coils would provide the greatest flexibility however, for simplicity, a fixed 50 mm diameter probe was developed. The design itself is similar to a linear array wrapped in a semi-circle around the diameter and consists of 11 coils as shown in

Figure 6.32 (a) and (b). Situated in front of the array is a camera mounting set at a 33 mm offset from the pipe surface to ensure that the full diameter of the pipe is captured in an image frame. This dual mode sensor is fundamental to the fusion process that will be discussed in Chapter 7. The sensor unit has two bearings which allow to it to be scanned in the axial direction of the pipe powered by a linear actuator. In this way it is possible to capture an image of the pipe and then scan the eddy current array over the same area, before ultimately fusing the two data sets.






Figure 6.30: Linear array comprising of a 16 coil probe in a staggered two row arrangement.




(a)




(b)

Figure 6.31: Circular array consisting of a central coil with 6 satellite coils. (a) Exploded solid model showing key components and (b) fabricated probe.



(a)



(b)


Figure 6.32: Pipe inspection array from two perspectives (a) and (b).

6.7 Results


6.7.1 Coil Performance

The magnitude and phase responses of the coils were obtained using a Hewlett Packard HP4194A impedance analyser and the calculated resistance and inductance compared against that predicted by FEM. Figure 6.33 (a) and (b) show the results for a coil on a piece of aluminium containing no defects over a frequency range of 100 to 200 kHz. It can be seen that the trend is reasonably similar, albeit with the FEM predicting higher values of both resistance and inductance than was found experimentally. The difference in the latter is 0.153 mH and remains relatively constant throughout the frequency range. However, the rate at which the resistance of the coil increases with increasing frequency is lower in reality than is expected from the simulation leading to a difference in resistance of 3.9 Ω and 18.3 Ω at the lower and upper limit respectively. The response of the coils to lift-off follows much the same trend with, in general, the FEM over-estimating the true value as shown in Figure 6.34. The rate at which the resistance decreases with increasing lift-off is lower in reality than predicted from FEM and the relative difference at each frequency is also smaller leading to an intersection between the measured and theoretical responses for 100 and 150 kHz. Once again, the predicted inductance variation shows good consistency with the experimental results while showing a relatively fixed offset across the 0 to 5 mm lift-off range.

It is likely that the discrepancy in resistance as a function of frequency is the result of inadequate modeling of the skin effect in each conductor. A solution to this would be to implement a finer mesh however this is at the expense of considerably higher simulation time. This is supported by Figure 6.35 which shows spatial quantisation of the induced eddy currents in the outer conductors. It is concluded that for fine coil diameters such as are used, a semi-analytical approach may be more favourable and result in more accurate results. The coil resistance as a function of frequency is likely to be further affected by the assumption that the number of coil windings along the model z-axis is consistent which, in practise, is very difficult to achieve without professional orthocyclic winding facilities. Professional manufacture of coils is therefore recommended for future sensors.




(a)




(b)

Figure 6.33: (a) Resistance and (b) inductance of a coil against frequency compared with that predicted from FE modelling.



(a)



(b)

Figure 6.34: (a) Resistance and (b) inductance of a coil for varying degrees of lift-off compared with that predicted by FE modelling.




Figure 6.35: Induced current distribution showing quantisation in the outer conductors.

6.7.2 Eddy Current Array Imaging

6.7.2.1 Imaging Algorithm

A multitude of imaging algorithms have been proposed for eddy current systems however these are typically designed for absolute or reflection probe configurations and typically involve the calculation of the absolute coil parameters [125, 126, 127, 128]. The following algorithm differs from this in that it is designed specifically to address the issue of imaging using an array of differentially connected coils.

The location of the centre of each coil is fixed relative to the local, or ‘probe’, reference system. As both the linear and circular arrays are designed for the inspection of flat surfaces, the z co-ordinate of each coil is constant and consequently can be ignored; for the pipe inspection array, this simplification does not hold. Figure 6.36(a) depicts the circular array situated such that coil 5 is placed over a defect. The footprint of each coil is approximated by N pixels. As each coil is wound by hand, there is an inevitable variation in their construction and hence performance. Small errors in the number of turns and factors such as non-uniformity of wire placement around the former will result in an unwanted phase variation with frequency between any coil pair. It is therefore necessary to perform a calibration measurement on a sound specimen of metal prior to the start of the scan. The result of this process is a set of phase vectors, $\Delta\phi_{cal_{m,n}}$, for each coil pair (m, n) , which define the ‘ambient’ response of the coil-pairs for each frequency

interval. This vector is then used to correct the corresponding result in all subsequent probe locations, by direct frequency subtraction using Equation (6.18).

$$PI_{m,n} = \sum_{k=1}^K [\Delta\phi_{measured_{m,n}}(k) - \Delta\phi_{cal_{m,n}}(k)] \quad (6.18)$$

where, $\Delta\phi_{measured_{m,n}}(k)$ is the phase difference vector for coil pair (m, n) at frequency interval, k , and $m \neq n$.

The imaging algorithm is represented in Figure 6.36(b) for two positions, x_j and x_{j+1} along the x-axis. In both cases, coil 5 is situated over a defect, represented in the diagram as a black rectangle. As the system is differential and not absolute, it is not possible to know, directly, under which coil the defect is present. However, it would be expected that every result that involves this coil would result in a differential phase response. Thus, by combining the information from all coils of the array, the most likely location of the defect can be determined.

The size of the defect can also be determined to a level of accuracy proportional to the size of the coil and relative coil spacing.

The phase integral is calculated for each pair of M coils using Equation (6.18) and the average of all values involving a particular coil, $P_{coil_m}(j, k)$ for $m = 1...M$, is then assigned to the N pixels representing that coil footprint equally, as shown in Figure 6.36(b). For a given probe location, (j, k) , the phase data mapped to the pixel footprints of each coil is transformed from the probe reference system to the image reference system using Equation (6.19).

$$\begin{bmatrix} x_{image} \\ y_{image} \end{bmatrix} = \begin{bmatrix} x_{probe}^{coil} \\ y_{probe}^{coil} \end{bmatrix} + \begin{bmatrix} x_{image}^{probe} \\ y_{image}^{probe} \end{bmatrix} \quad (6.19)$$

where, $\begin{bmatrix} x_{probe}^{coil} \\ y_{probe}^{coil} \end{bmatrix}$ and $\begin{bmatrix} x_{image}^{probe} \\ y_{image}^{probe} \end{bmatrix}$ are the 2D position of the coil relative to the probe

axis system and the position of the centre of the probe relative to the image axis system respectively. The appropriate pixels in the image plane are then updated by adding the mapped values to the existing pixel value. Finally, each pixel is averaged by the number of samples that it received, to account for multiple measurements of the same region.

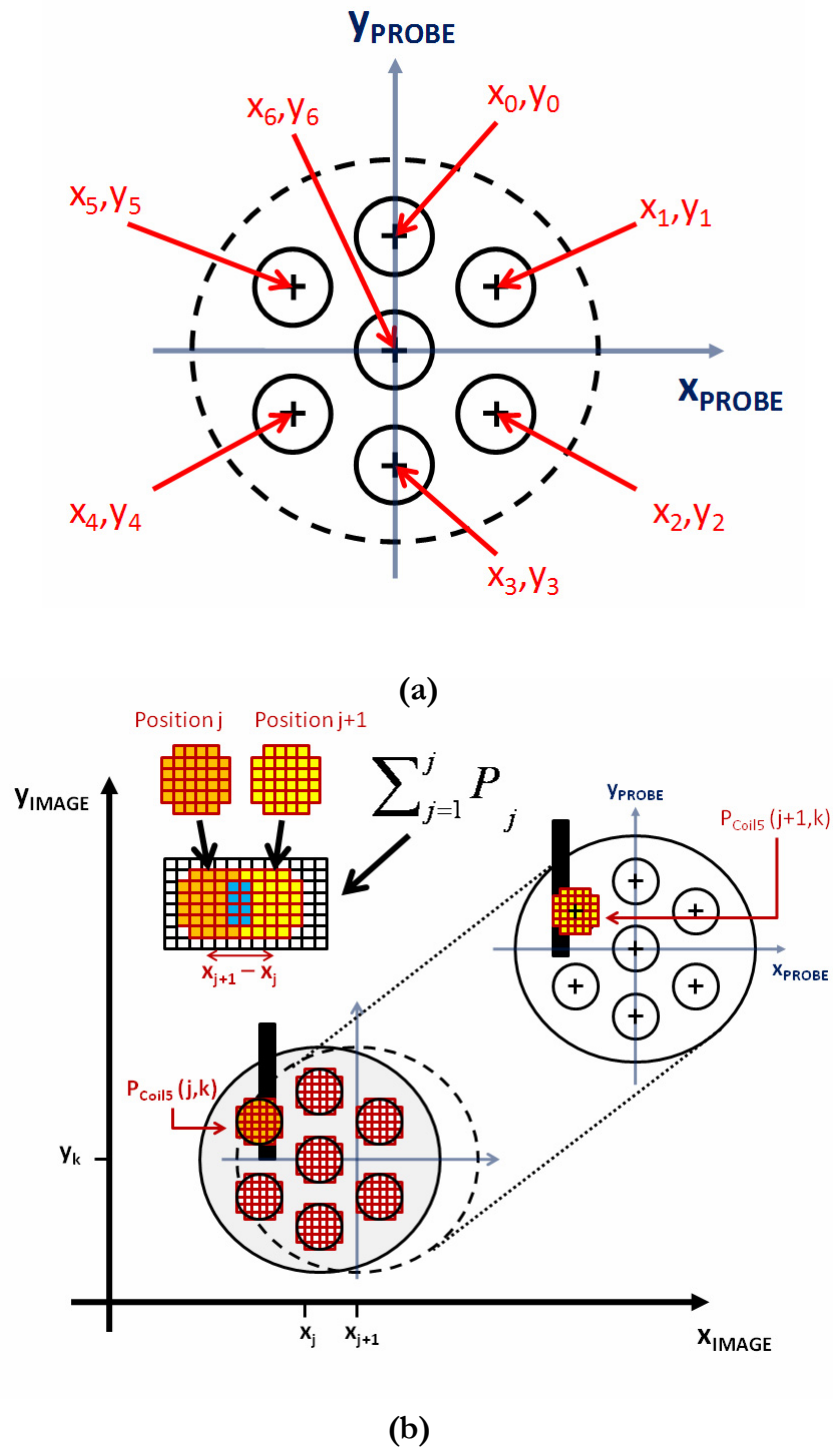


Figure 6.36: (a) Schematic of circular probe showing the relative location of the centres of each coil relative to the probe reference system. (b) diagrammatic representation of the imaging algorithm combining the phase data, $P_{coil_5(j)}$ and $P_{coil_5(j+1)}$, from two probe locations, j and $j+1$.

This process produces three distinct levels within the resulting image:

- Noise floor - representing the coil response to areas of no defect.
- Defects - representing the coil response to defects.
- Unscanned areas - where no data has been captured.

The latter can be minimised, and in fact eliminated, through judicious choice of scanning resolution.

6.7.2.2 Imaging Simulation

A simulation algorithm was developed, in MATLAB, in order to assess the effect of different defect configurations on the resulting images. As the imaging system is independent of the probe used, for simplicity and validation of the later fusion probe results in Chapter 6, the linear/curved array coil configuration was modelled for a seven coil probe. This is shown in Figure 6.37 and comprises of a two dimensional surface, corresponding to the ‘unwrapped’ circumference of the pipe (x-axis), containing two 10 mm diameter defects modelled using a 1 mm mesh. The array is scanned along the axial length of the pipe in the y-direction. The model itself consists of three elements: the pipe (including defects), the array and the image matrix. The pipe is defined by an $N \times N$ matrix of pixels with a value of one assigned to those pixels that represent a defect and a value of zero assigned elsewhere. The array is similarly modelled, this time as an $M \times N$ array, with the pixels corresponding to the coil footprint defined as having a value of one. For simplicity, it is assumed that the sensitivity profile of the coil is uniform. As with a real inspection, the array template is moved along the pipe in the scanning direction, in this case in 2 mm increments for speed. At each interval, the pixels of each coil are multiplied with the coincident pixels of the pipe and the sum across the coil footprint calculated. This is equivalent to calculating the overlapping region indicated in Figure 6.37. If the coil is entirely over a defect-free region, this total will equal zero. If any part of a defect overlaps the coil footprint, this will result in a total value which is proportional to the extent to which the defect covers the coil. For each pair of coils, the difference in the total is calculated and added to the pixel values, corresponding to the coils, of the image matrix, in the same way that the phase difference is assigned in the real imaging algorithm. This process can be likened to the increase in the phase difference (in this case pixel difference) as a defect encompasses more of a coil compared with another coil that is unaffected by the defect.

The result of the simulation is shown in Figure 6.38 where two defects regions are indicated by peaks in the image matrix. As defect 1 and 2 are identical and cover




Figure 6.37: Imaging model. The difference in the area of overlap between each coil pair is used to determine the value added to each pixel under the coil footprint in the image matrix.

each coil to the same extent, the peak magnitude associated with each defect is the same. The shape of the defects is not well represented due to the relative size of the coils and defects, as well as the approximation of the circular shape by the 1 mm mesh. The ‘noise’ regions are an artefact of the differential system and are an inherent part of the algorithm as it is not possible to know under which coil the defect exists, as discussed in Section 6.7.2.1. Figure 6.39 shows the effect of moving defect 1 by negative 2 mm along the x-axis such that the area of overlap with coil 0 is reduced while the overlap with coil 6 is increased. The result of this is that the peak pixel value of the defect 1 indication is now less than for defect 2 and there is an increased pixel value associated with coil 6. This can be thought of as ‘leakage’ due to the fact that the detection of defect 1 is now less dominated by the response of coil 0 and increasingly by coil 6. This effect leads to the conclusion that, for an unknown inspection situation where it cannot be assured that any defects will be positioned favourably, the pixel intensity cannot be guaranteed to indicate the defect depth. Rather, the intensity of the pixels in the area surrounding the defect must also be considered. Scanning at a finer resolution and in the x and y dimensions would solve this however the inspection time would be much longer.




Figure 6.38: Simulated output of the imaging algorithm, where coils 0 and 1 are positioned symmetrically over the defects. The resulting pixel values of the image matrix indicating the defect are equal.

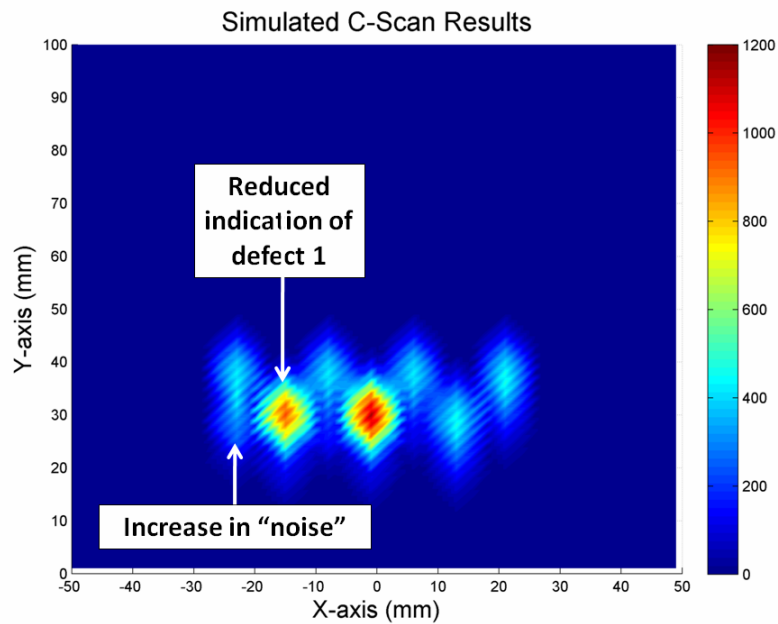


Figure 6.39: Simulated output of the imaging algorithm, where the centre of defect 1 is shifted by -2 mm along the x-axis. The resulting pixel values of the image matrix, indicating this defect, are reduced and there is an increase in the pixel values corresponding to coil 6.

6.7.3 C-Scan Results

6.7.3.1 Comparison with Commercial System

The imaging algorithm outlined in Section 6.7.2 was applied experimentally to a scan of a section of aluminium containing 16 mm diameter manufactured drill-holes of depths: 1.1 mm, 0.9 mm, 0.7 mm, 0.5 mm, 0.2 mm, 0.1 mm and 0.05 mm as shown in Figure 6.40. The circular probe was positioned at 132 positions in a regular grid covering a region of 150 x 400 mm within which the defects were located. Following the results of the imaging simulation, the grid was designed such that the centre of the probe would coincide with the centre of the defects, at the corresponding location. Figure 6.41 (a) shows the results from the inspection with the average PI (i.e. corrected for the number of samples per pixel) plotted on the z-axis and the x and y axes referring to the corresponding axes of the plate. It can be seen that the four largest defects can be clearly identified with the 0.2 mm defect also detectable but quite close to the noise floor. As predicted, the value of PI decreases with defect depth meaning that the relative size of the defects can be identified from the results.

In order to benchmark the performance of the system against an industry standard device, the same defects were inspected using the commercially available Olympus NDT OMNIScan MX ECA [49]. This 32 coil device uses similar diameter coils and has an operational frequency range of 20 Hz to 6 MHz depending on the type of probe used.



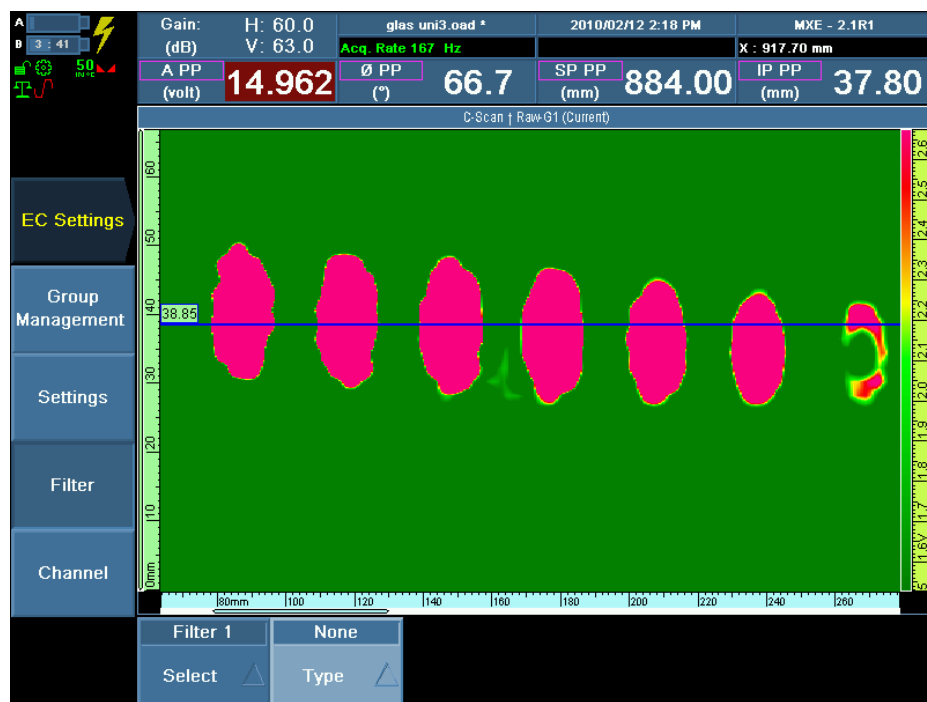


Figure 6.40: Aluminium test plate containing seven defect with depths from right-left of 1.1 mm, 0.9 mm, 0.7 mm, 0.5 mm, 0.2 mm, 0.1 mm and 0.05 mm.



(a)



(b)

Figure 6.41: (a) C-scan image generated from data obtained using the ECA system at 132 regular locations on an aluminium test plate. (b) shows the same area of plate inspected using the commercially available OlympusNDT OMNIScan MX ECA system (largest defect shown on right).

The use of a frequency sweep is possible only in a piecewise fashion, using this system, with post-processing then needed to combine the C-scans from each frequency interval into a single resultant image. Typically, a lower frequency than the ECA developed in this Chapter is used (5-30 kHz for absolute probe and 5-100 kHz for reflection probe) which minimises the unwanted lift-off effects. The results of a scan using a manual probe fitted with an encoder are shown in Figure 6.41 (b) and in this instance, due to the direction in which the probe was scanned, the largest defect appears on left. It can be seen that all seven defects are clearly detectable

6.7.3.2 Using Multi-Frequency Data to Reduce Noise


In Section 6.3.1, a relationship for the phase difference between 2 coils was derived. Setting $L_2 = L + \Delta L$, where ΔL is the change in inductance, Equation (6.8) becomes Equation (6.20):

$$\Delta\phi(\omega) = \text{atan}\left(\frac{\omega^2(L + \Delta L)C - 1}{\omega CR}\right) - \text{atan}\left(\frac{\omega^2 L_1 C - 1}{\omega CR}\right) \quad (6.20)$$


The use of frequency sweep excitation allows this known functional relationship with frequency to be exploited to fit a curve to the experimental data and in so doing significantly reduce the noise. Such a technique could not have been used if single excitation had been used since this relationship could not be determined. Assuming that L , R and C are known for all K frequencies within the sweep range, the mean squared error (MSE) can be calculated for each value of ΔL using Equation (6.21):

$$MSE = \frac{1}{K} \sum_{k=0}^K (\Delta\phi_{theoretical}(k) - \Delta\phi_{theoreticalexperimental}(k))^2 \quad (6.21)$$

By selecting the value of ΔL that minimises the MSE, a best estimate of ΔL can be produced for each coil pair, at each probe position. The data from Figure 6.3 (a) was used to correct the resistance as a function of frequency and ΔL was varied from 0 to 1 mH in 0.01 mH increments. Figure 6.42 (a) shows a phase measurement captured containing significant noise and the associated best-fit data yielded by the above process. Re-processing the data from the plate scan shown in Figure 6.41 (a) using this technique yields the result in Figure 6.42 (b). It is immediately apparent that the noise levels have been reduced with many areas having zero value due to the best fit occurring for $\Delta L = 0$. Consequently, the 0.2mm defect is now much more clearly visible.



(a)



(b)

Figure 6.42: (a) Best-fit to experimental data. (b) Re-processed plate scan using best-fit correction.

6.8 Discussion

The novelty of this Chapter is primarily in the integration of the ECA with the robotic platform. Also, the theory presented on reducing the sensitivity of the system to the choice of inspection frequency, noise at the inspection frequency and changes in the circuit resistance using the phase integral approach along with the estimation of the optimum frequency for phase detection could not be found in any literature reviewed and is therefore also considered as a contribution. The use of a frequency sweep excitation has the additional advantage over a single frequency excitation of being able to exploit the known functional relationship between the phase difference and excitation frequency of two coils to reduce the noise.

The ECA system developed has been tested on a variety of defect depths and found to produce results consistent with those expected through simulation. This has been benchmarked against an industry standard device and found to detect five out of the seven defects found using this system. The diameter of the defects used are of the order of that of the coils however the detection of 0.2 mm depth is considered reasonably good considering the fabrication facilities available although falls short of the 0.1 mm depth defined in the specification set out in Chapter 3. The coil diameter can be scaled down to allow smaller diameter defects to be detected however this will require more sensitive electronics to detect the signals resulting from the lower generated field they produce. This is unlikely to present a significant design challenge, particularly in light of the fact that commercial systems allow the use smaller coils, however this would need to be evaluated more thoroughly. The poorer performance in comparison with the commercial system is believed to be a combination of the variability of the hand-wound coils and the quality of the electronics used. It is recommended that future systems incorporate precision, orthocyclic coils and that professional standard fabrication facilities used to create the PCBs. The significant advantage of the in-house system, however, is that it can be integrated into a lightweight, low power robotic platform. Although the OMNIScan MX ECA was capable of detecting the two smallest defects tested (0.1 mm and 0.05 mm), the associated electronics comes at a significantly greater size (244 mm x 182 mm x 57 mm), weight (1.2 Kg) and cost (£20-30k). Additional noise is introduced in the custom ECA through the limitation of a 10bit ADC; the Moore's Law nature of microprocessor development means that currently better ICs are available which could address this problem. As all of the electronics have been developed using relatively limited in-house facilities, there is an inevitable compromise in terms of the achievable signal-to-noise ratio. Although the current speed for complete data capture, processing and communications is approximately

50 seconds for a 7 coil (21 pair) system, the dominant overhead is the phase detection at each frequency interval. With the luxury of more development time, it is conceivable that this could be reduced by an order of magnitude. Given that each coil pair has 100 frequency intervals, the total number of equivalent single frequency measurements for a 7 coil probe is 2100. This equates to an interval acquisition rate of 42 Hz which, when considering the developmental nature of the system, is reasonably high. Operational speed could be significantly increased through the use of additional channels and parallel measurement however such work is beyond the scope of this research.

The FEM approach was found to have limitations for the given wire gauge manifested by spatial quantisation in the outer conductors of the winding. This led to an inadequate model of the skin effect resulting in a noticeable discrepancy in the resistance vs. frequency characteristic. Furthermore, the poor quality coil windings noted earlier also contributed to a deviation in the model of the resistance as a function of lift-off. This necessitated the implementation of a calibration procedure however an element of error was also introduced here. As the frequency response of the coil varies with defect size (and frequency of peak phase difference changes), the subtraction of a calibration measurement taken on sound metal on a frequency-by-frequency basis becomes less accurate with increasing defect size. Again, through precision construction of the coils, to ensure consistent windings per layer, the variation between coil performance will be much lower and therefore the error introduced through calibration will be much less significant.

An upshot of the differential nature of the inspection system is that, for each probe position, there are two possible solutions that generate the same output data. Consider the scenario shown in Figure 6.43, where there exists a crack under coils 4 and 5 of the circular probe, termed the 'defect coils'. As data is collected for each coil-pair, any results which involve either of the defect coils will produce a higher value of PI than are produced by any other pair of coils. Crucially, this includes the result from the two defect coils themselves. As the name suggests, the differential probe is sensitive to *differences* in conditions between coils but if two coils experience the same conditions, whether this be over a defect or over sound metal, the result will be similarly small. Consequently, defect scenario 1, shown in Figure 6.43 (a) is indistinguishable from defect scenario 2, shown in Figure 6.43 (b) for a single probe position. This limitation may be overcome in one of two ways. Firstly it can be assumed that, in general, as the expected defects are most likely to be small, it can be assumed that the largest area is non-defective. This naturally leads to a different limitation whereby very large defects would be assumed to be smaller than

they actually are. The second solution involves using the data from subsequent probe positions to update the prediction of the defect size and shape. Considering the situation depicted in Figure 6.43 (a) where a defect is under coils 4 and 5, if the probe is shifted to the right and the phase information calculated, two new possibilities can be defined as shown in red in Figure 6.44 (a) and (b). At this point, the data from the previous step can be compared with the current results and consistencies identified. Had defect scenario 2 been the correct form, then it would be expected that, for the given change in probe position, the coils in the green region of the probe would be affected instead of the coils in the red regions. However, the predicted response of the new probe position to defect scenario 1 is consistent and therefore it can be concluded that this is the correct solution. This process involves greater computational complexity however provides greater SNR over the solution outlined earlier as no phase contributions are assigned to coils which do not result in the expected consistent result for defect location.

A list of parameters that can affect the performance of any eddy current system was outlined in Section 6.2.1 and of particular applicability to industrial scenarios is the variation in magnetic permeability. It was identified in Chapter 2 that all grades of austenitic steel have some remnant ferromagnetic properties. Cold work can lead to a partial transformation from austenite to the ferromagnetic martensite phase resulting in variations in the magnetic permeability. Due to the local nature of such variations, the effect of cold work could not be calibrated for and hence this represents a fundamental limitation of the eddy current modality and not of any particular inspection system itself.




Figure 6.43: (a) Defect scenario 1 and (b) defect scenario 2 are indistinguishable in a differential phase system.

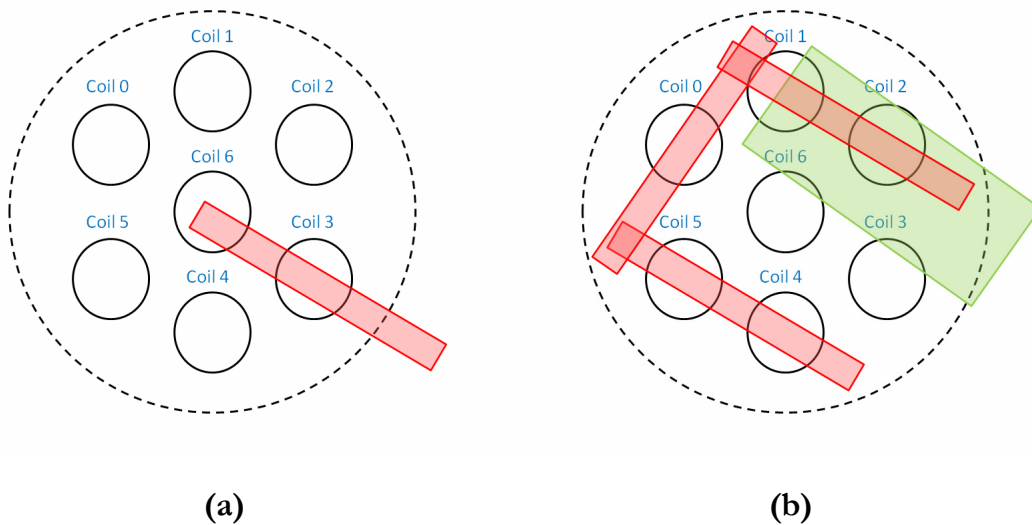


Figure 6.44: (a) Defect scenario 1 and (b) defect scenario 2 where, in both instances, the probe has been shifted to the right from those shown in Figure 6.43. In (b), if the defect was of the type of defect scenario 2, the coils affected would be in the area shown in green rather than in the area shown in red. Thus, there is no consistent progression for defect scenario 2 leading to the conclusion that the defect is of the form of defect scenario 1.

Chapter 7

Multi-Sensor Data Fusion

7.1 Introduction

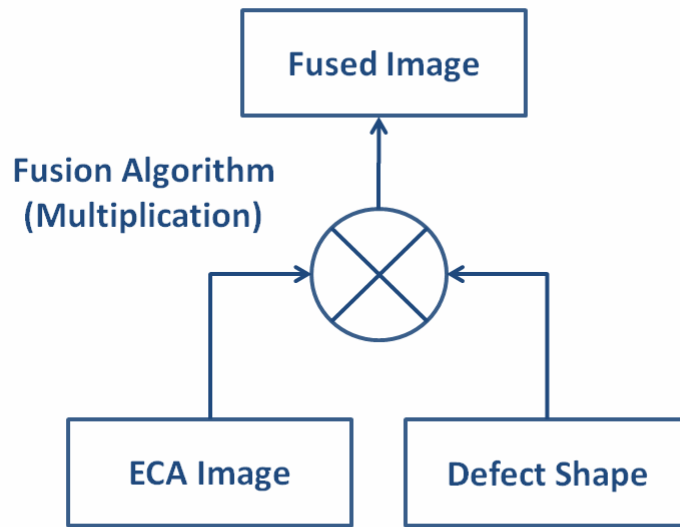
Due to the short wavelength associated with optical inspection, the spatial resolution of a typical camera far exceeds the resolution realisable from a discrete array of non-optical sensors. However, the increased level of detail exacerbates the problem of distinguishing true defects from defect-free surfaces. In Chapter 5, the eddy current array was shown to be capable of detecting defects of varying size and ultimately providing an estimation of the defect depth. However, to achieve the required fidelity to accurately determine the shape of a defect requires an unrealisably small coil diameter and impractically high number of elements. The corollary of this is to combine data from both sensors into a single inspection result and thereby gain the benefits of both. Following an indication of likelihood from the pattern classifier, the area would then be scanned using the eddy current array and finally an appropriate shape extraction algorithm would be employed and fusion performed. Such data fusion [129, 130] ties together the three main threads of this research and is the subject of this Chapter.

The Chapter begins with an outline of a framework for on-board fusion for any type of defect while limiting the specific implementation to corrosion pit defects. Future work will involve extending this to develop the algorithms to process other types of defect. The Chapter concludes with two examples of on-board data fusion on a sample containing two artificial corrosion pits (drill holes) and a slot, both in 304 austenitic stainless steel.

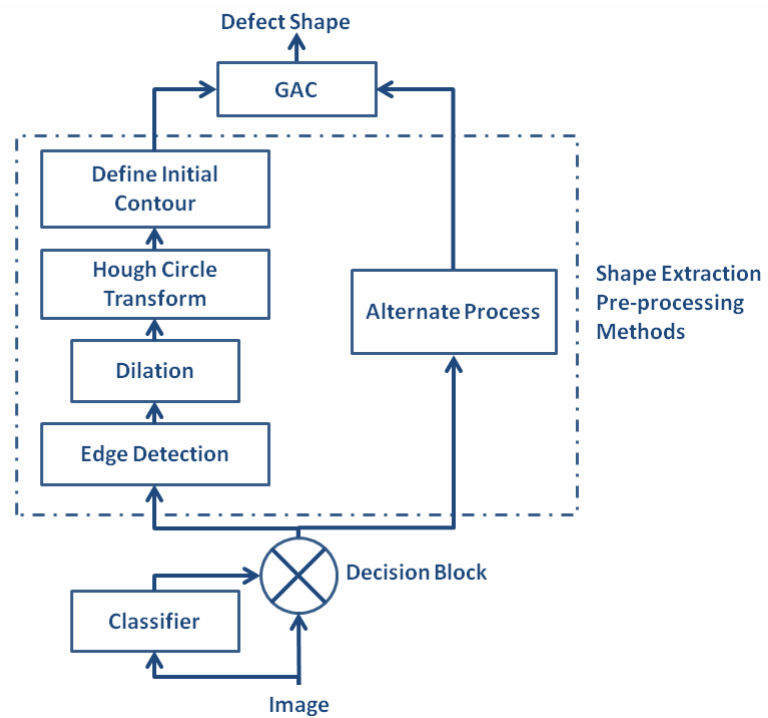
7.2 Development of a Data Fusion Framework

7.2.1 Scope

Designing an all-encompassing data fusion system for any defect geometry would consist of a significant research undertaking in itself and the scope within this research project must therefore be more limited. The aim of the data fusion is simple: to obtain a better understanding of the defect by the combination of the visual and ECA data. The available inputs to the fusion algorithm are the ECA image, the predicted defect class based on the output from the classification system and finally, the raw image itself from which the defect shape can be determined. Two data sets are often combined using either a Bayesian inference or pixel-level approach [129]. A Bayesian system would involve determining the uncertainty in the defect shape resulting from the shape extraction process along with the uncertainty in the size of the defect from the eddy current array and then combining the two data sets using Bayes' theorem in the same manner as was used in Chapter 5. The main difficulty with this approach lies in quantifying the uncertainty in the defect shape for complex geometries. A pixel-level approach may be implemented in a number of ways including addition, subtraction, multiplication, division and weighted versions of each of the two data sets. Such systems require perfect registration and are algorithmically very simple to implement. The use of Bayesian techniques in the application presented in this thesis is rather involved and a pixel level approach is favoured with the former discussed in more detail as part of the programme of future work. The fusion will combine the defect shape with the corresponding depth indication from the ECA as shown in Figure 7.1(a). This lends itself well to a pixel-level multiplication combination since one wishes an eddy current indication of depth within an area suspected as being a defect to register a strong output while eddy current indications outwith (that cannot be visually observed) to register a small output. Adopting a modular approach, a framework for such a system can be established that involves use of the full complement of data available to the robotic inspection system and which can be tailored in future as research in this area progresses. For simplicity, the shape extraction process will be designed primarily for broadly circular defects (simulating corrosion pits) however the fusion will also be tested on slots for completeness. A block diagram of the shape extraction process for corrosion pits is given in Figure 7.1(b). The output from the classifier is fed into the decision block which implements an appropriate pre-processing method based on the estimation of the defect class to yield a first guess at the defect shape.



(a)



(b)

Figure 7.1: Functional block diagram of (a) the data fusion and (b) the shape extraction processes.

Only circular defects will be implemented in this case and therefore only one shape extraction method is shown. All methods feed the initial shape (or contour) estimate into the geometric active contour (GAC) block, discussed in Section 7.2.2.1, which then adapts to arbitrary shapes and ensures a closed contour solution.

7.2.2 Overview of Procedure for Artificial Corrosion Pits

The shape extraction process used for corrosion pits involves four separate stages of image processing: edge detection, dilation, Hough transform and GAC processing. The use of GAC has a large bearing on the processing stages that precede it and hence the discussion in the following sections will be taken in the reverse order in which they are applied with GAC being outlined first followed by the Hough transform as used for circles.

MATLAB was used to perform the processing due to the availability of image processing toolboxes for the key operations (GAC, Hough transform etc). The task of porting this functionality so that the data can be entirely processed ‘on-board’ is one that would require a significant investment of time, however the hardware resources have been put in place, during the design process, to cater for this. This final stage of automation is recommended as an area for future work.


7.2.2.1 Geometric Active Contours

Active contours is an image processing technique that seeks to identify internal boundaries within an image [131-139]. The original active contour algorithm, or ‘snakes’, modelled a contour as a spline [137]. A set of points defined around the target object then evolve to encompass it in an energy minimisation process. The energy function consists of three terms: an internal energy representing the weighted effects of spline bending and stretching, an image energy constituting the factor driving the spline to the image features (such as edges and lines) and finally, an external constraint energy that ensures the spline ends near a local minimum. The shape of interest should represent a minimum of such an energy function with each iteration yielding a set of spline points having collectively a lower value of the energy function and thus a better match to the target feature. The main limitations of this parametric approach are the dependency of the result on the initial contour and the difficulty associated with accounting for merging contours. Geometric active contours (GAC) was introduced to address these limitations and employs the level set method [138, 139]. Instead of explicitly defining a contour using splines, in GAC it is defined implicitly as the zero-level set of a level set function. The level



set function is a 3D surface and so the contour is a 2D slice through it at the zero-level. Rather than evolving the contour directly, as in the parametric formulation, the level-set function is evolved according to a particular partial differential equation (PDE) and consequently so is the contour. This process is shown in Figure 7.2 with the level-set function being shown in the left-hand column and the associated contour shown in the right-hand column. An initial guess of the shape is provided in (a), superimposed on an artificial image, and the initial level set function in (b) is then calculated based on this starting contour. Following this, the operation then essentially operates in reverse, with the contours in (b) and (c) being extracted as slices at the ‘zero-level set’ (x-y plane at $Z = 0$) of their associated level-set function. In this way, it can be seen that it is possible for the contour to separate and/or merge in a way that would be very difficult to achieve in a parametric system, yielding an important advantage over the ‘snakes’ method. Many variants of the technique exist, with differing methods for implementing the curve evolution, however the method presented in [131] provides a suitably automated implementation. Many GAC systems require the level-set function to be periodically re-initialised to a signed distance function (a function that is differentiable at all necessary points) in order to ensure numerical stability during evolution. However, this implementation defines an energy function which penalises deviations from this function (similar to the internal energy term of the snakes algorithm) and hence is inherently more stable. The second, ‘external’, energy term is one that drives the zero-level to the image features.

The main advantage of GAC is that it allows arbitrary shapes to be detected and is thus much more suitable for NDE applications. For each specific type of defect, having its own characteristic shape, it may be that in some cases a far simpler algorithm would provide similar results at a lower computational expense. However, the advantage of the GAC approach is that it can be used on *any* defect and therefore is far more flexible. Thus rather than implementing many different shape extraction algorithms, each with their own performance characteristics, GAC provides this functionality within a single system.



As stated, in order to be able to employ the algorithm an initial estimate of the contour must be provided which is then modified during the search for boundaries. This can be defined as the outer extremes of the image frame however, it was found that this did not always guarantee segmentation of two defects due to the effect of lighting and surface curvature. Significantly better results were achieved by first identifying the centre points of the defects and then expanding out to find the boundaries.





(a)



(b)



(c)



(d)

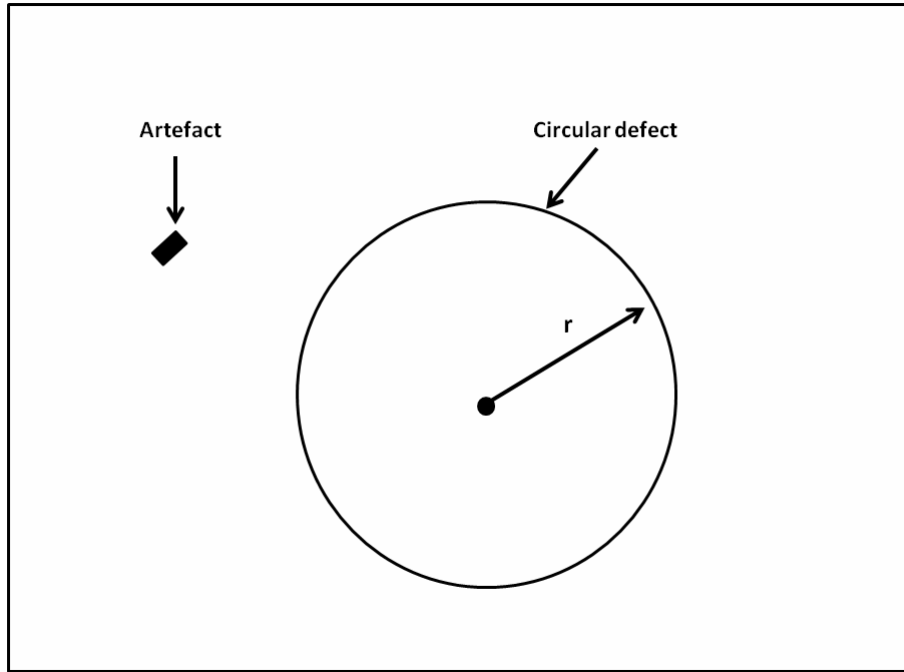
Figure 7.2: Evolution of the level-set function. (a) Shows the initial contour on the image from which the shape is to be extracted, (b) shows the associated level-set function and contour and (b) and (c) show the evolved functions and contours for 100 and 300 iterations respectively.

In order to achieve this, and given the *a priori* knowledge of the defect class (in this case a circular defect), an edge detection process is used followed by a Hough transform to identify the centres of any circles having a radius, r . The initial contour is then defined as a circle of radius $\frac{r}{2}$ centred at these points, from which point the active contour algorithm determines the precise shape of the defect. In this way, the process allows defects that are not perfectly circular to be extracted. The scale of the captured image is reduced for this process in order to increase the processing speed of the active contour algorithm and to ensure pixel-pixel registration for the subsequent data fusion.


7.2.3 Hough Transform

The Hough transform is best known for linear edge detection applications where each point on an edge can have a range of straight lines of gradient, m , and intercept, c , defined that intersect it. Each of these lines defines a single point in Hough space. For points lying on the same straight line there will be one particular point in the Hough space (and hence one value of m and c) that is dominant and this corresponds to the true edge parameters. The same theory can be adapted for detecting circles. Figure 7.3 (a) shows a binary image containing a circular defect and an additional unrelated artefact generated from a hypothetical edge detection filter. The resolution of the image is such that the circle is represented by 13 pixels (0 to 12) and the unwanted artefact by 1 pixel.

Using *a priori* knowledge that the artefact of interest has a radius, r , for each edge pixel, a circle of radius, r , is defined with the pixel as the centre point. This is shown in Figure 7.3 (b) for three edge pixels, 1 to 3, defining the circle and for edge pixel 13 defining the unwanted artefact. It can be seen that the circles defined on the circular defect all intersect at its centre-point. An accumulator is used to count the number of times each pixel is intersected by one of these circles. Extrapolating this to use all of the detected edge points, it can be seen that a ‘hot-spot’ will appear at the true defect centre and a threshold can be set to pick it out from the accumulator. For practical edge detection however, the results are imperfect and a degree of uncertainty must be tolerated. This is particularly the case for circular shapes on curved surfaces, viewed in only two dimensions, where the perspective results in an elliptical representation. Ideally a perspective transformation step would precede any shape extraction and is a recommended area of future work. The detection of edges belonging to a circle can, to some extent, be improved by enhancing (or thickening) the edges prior to the application of the Hough algorithm. This can be done through a morphological process called ‘dilation’.



(a)



(b)

Figure 7.3: (a) Image containing a circular defect and an unrelated artefact. (b) Illustration of the Hough transform.

The dilation operator effectively thickens the binary edges by the size of a given structuring element (e.g. a disk). The net effect of doing this prior to the Hough transform is that there is greater tolerance to imperfect circularity however, this is at the expense of greater uncertainty in the precise location of the centre. This can be overcome by an additional validation step which discards any centres that are within a small number of pixels of each other. In practice, the correct radius is not known in advance and consequently a search process must be employed to find the correct value. This process in itself can be problematic yielding multiple solutions however, for the purposes of demonstrating the principle, this will be overlooked in order to focus on the more important issue of data fusion.

Image registration is inherent through the use of the fusion probe array, outlined in Chapter 5, and is illustrated in Figure 7.4. Due to the axial offset between the EC coils and the array, following the capture of an image using the camera, the array is driven forward through a distance $\text{LensOffset} - \text{CamScope}_y$ to take it to its start position. Data capture can then commence with the array being moved through a distance of $2 \times \text{CamScope}_y$.

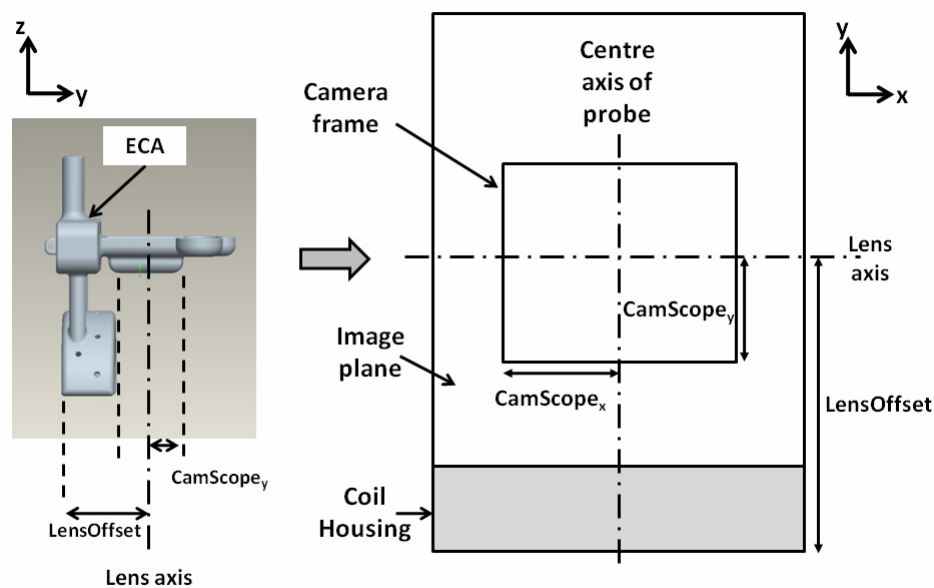


Figure 7.4: Image registration is solved by the fixed link between the camera and the ECA of the fusion probe. After an image is captured, the array is simply driven forward by a distance $\text{LensOffset} - \text{CamScope}_y$ and the EC scan started from there.

7.3 Data Fusion on Real Data

The following section demonstrates the principle of the data fusion process through a step-by-step example of an inspection on a 50mm diameter section of 304L austenitic stainless steel pipe containing two 10mm through-thickness, circular defects. The fusion probe fitted with seven coils was used for the inspection, allowing the same control electronics to be used as for the circular array probe used in Chapter 6. Figure 7.5 shows an image captured using the built-in camera module. Both defects are identical however defect two is situated in the centre of the x-axis of the frame and defect 1 is situated approximately 12mm to its left and slightly lower with respect to the y-axis. The effect of the curvature can be clearly seen in the elliptical appearance of defect 1. The effect of a Sobel edge detection filter and dilation using a two pixel radius, disk-shaped structuring element is shown in Figure 7.6. Having reduced the resolution of the original data by a factor of 9, for the computational efficiency of the GAC algorithm, the necessity for this step is evident by the imperfect circular shape of the defects. Without this dilation process, the peak magnitude in the Hough accumulator would be smaller and more distributed, making centre detection difficult.

The Hough space representation of the dilated image is given in Figure 7.7 with two regions of local maxima indicated corresponding to the most likely centres of the two defects. There is less certainty surrounding the position of defect 1 (manifested as a lower peak value) which is to be expected, given its elliptical representation. The regional maxima (pixels of the same intensity whose value is greater than those of its boundary) of the accumulator is calculated and those maxima that exceed a specified threshold are returned as possible centres. The threshold, t , for this process is specified as $t = \pi r$, which equates to a pixel having at least 180° of the circle at a radius, r , from it. Figure 7.8 shows the locations of possible centres superimposed on the scaled grayscale image. Two clusters can be identified and the mean of each is defined as the true centre. The initial contour is generated for each defect, using this data, and fed into the active contour algorithm in two separate steps. The results of the shape extraction process are shown in Figures 7.9 and 7.10 at 20, 100, 200 and finally 300 iterations for both defect 1 and 2 respectively. While yielding slightly smaller results than in reality, the algorithm performs well at adapting to the basic shape of both defects and is stable by 300 iterations. A region of high intensity pixels around defect 2, caused by reflection of the on-board light source from the rear-inner surface of the pipe, causes the boundary to be pulled towards the centre at this point, however the adaption to the remaining boundary is relatively good.




Figure 7.5: 26mm x 36mm image of 304L austenitic stainless steel pipe containing two 10mm through thickness defects.

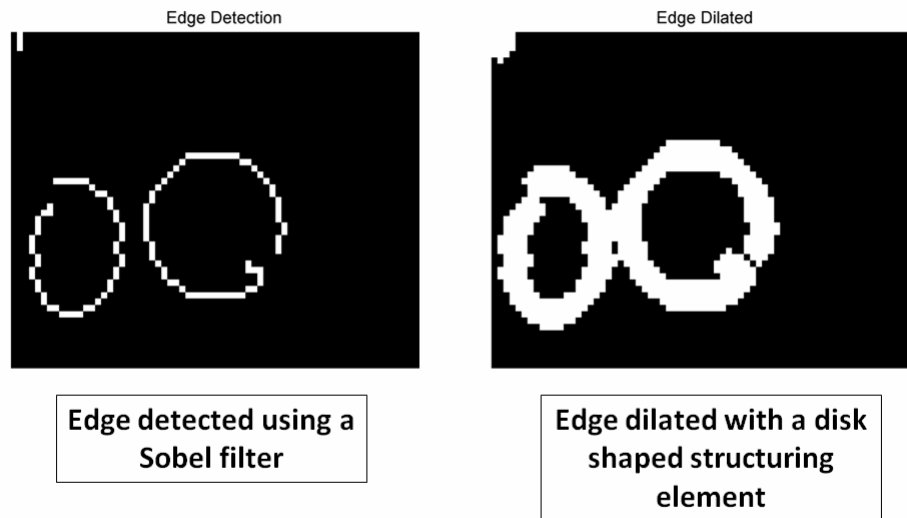


Figure 7.6: Edge processing stage. This involves applying a Sobel filter to the grayscale image followed by the application of a morphological dilation operation using a disk shaped structuring element having a 2 pixel radius.




Figure 7.7: The Hough space representation of the edge detected image shows the locations where the hough circles intersect, indicating the likely centres of a circle defined by the edges.




Figure 7.8: Location of centres that have accumulator values that exceed the threshold and have different values.

Having extracted the shape of the defect from the original image, this can now be fused with the eddy current image. Combining both of the extracted contours while setting all pixels within each to logic one and all pixels outwith to logic zero yields the defect mask shown in Figure 7.11. The results of the eddy current scan of the pipe are given as a 2D representation in Figure 7.12 and wrapped to the diameter of the pipe in Figure 7.13. Although there is a clear indication of the existence of some form of defect through the presence of two peaks above the noise level, the shape information has been poorly represented. This is consistent with the results of the imaging simulation for the same defect scenario in Chapter 5, and is to be expected given the relative dimensions of the coils and the defects.

Fusion is performed, through the multiplication of the ECA data by the defect mask. This equates to a probabilistic system where the uncertainty associated with the camera data set is zero. The result of this is to combine the two data sets yielding zeroes everywhere except within the area defined by the defect mask. Fusion is performed on a pixel-by-pixel basis as given by Equation (7.1), where i and j are the pixel co-ordinates and $I_{ECA}(i,j)$, $I_{mask}(i,j)$ and $I_{fused}(i,j)$ are the ECA, defect mask and fused image values, respectively, at pixel (i,j) .

$$I_{fused}(i,j) = I_{ECA}(i,j) \times I_{mask}(i,j) \quad (7.1)$$

This fusion strategy was adopted due to its computational simplicity, however an alternative Bayesian method is outlined in Chapter 7 as part of the programme for future work.

It can be seen from Figure 7.14 that the result of this yields a far more representative indication of the type of defect with the 2mm region between the two defects now visible. The correct registration of the two data sets is evident in the accuracy of the mapping of the two peaks to the defect mask. That being said, the inaccuracy of the size of the defect mask means that some data is lost. A limitation of data fusion in this way lies in the binary nature of the defect mask. If part of a defect is obscured in the image, and therefore undetectable by the shape extraction algorithm, then the fused data will be zero at this point. Developments of the methodology will be discussed as part of a future work programme in Chapter 7. A criticism of this form of fusion is that any shape of defect that lies over a noise region of the eddy current image may appear to be defective as all other regions are masked to zero. This issue can be illustrated by introducing a registration error into the previous data set such that the camera data coincides with noise pixels in the ECA data.




Figure 7.9: Active contour algorithm on defect 1.




Figure 7.10: Active contour algorithm on defect 2.




Figure 7.11: Defect mask, extracted from camera image, defined as logic one in the region within the identified contours and logic zero elsewhere.




Figure 7.12: 2D representation of the eddy current array data from the fusion probe on the test specimen.




Figure 7.13: Eddy current data wrapped to pipe circumference with no fusion performed.




Figure 7.14: Result of the fusion of camera image and eddy current array data wrapped to circumference of the pipe.




Figure 7.15: Effect of registration mismatch on the fusion result.

Figure 7.15 demonstrates that, in this scenario, while the outline of the defect is still obvious the magnitude is significantly lower than for the correctly registered image and therefore the interpretation should be straightforward.

A more challenging inspection task is the 2.5mm wide by 1.5mm deep, two-angled slot, again in 304L austenitic stainless steel, shown in Figure 7.16. The shape extraction algorithm method was applied again, however in this case the initial contour was manually created using MATLAB so that, when overlaid on the original image, its boundary was within the region of the slot. As the edges are more poorly defined in the case where a defect is not through-thickness, a contrast enhancing ink was applied to the surface of the slot. Without this, the active contour algorithm performed very poorly even when a good initial contour was defined. The defect mask generated by this process is shown in Figure 7.17 and it can be seen that the shape of the circumferential region of the slot has been detected reasonably well, however the remainder is less well represented. In particular, the length of the angled section is underestimated by several millimeters and the width at the junction of the two slots is only a few pixels wide. The ECA results, in Figure 7.18, are much more poorly defined than for the larger, through-thickness defects. However two peaks can be identified, corresponding to the circumferential section of the slot.




Figure 7.16: 26mm x 36mm image of 304L austenitic stainless steel pipe containing slot coated in a contrast enhancing agent.




Figure 7.17: Defect mask generated by manually specifying an initial contour within the region of the slot and subsequently allowing it to adapt.




Figure 7.18: Eddy current data wrapped to the circumference with no fusion performed




Figure 7.19: Fused image and eddy current array data wrapped to pipe circumference.

Although the magnitude of these are only 10% of that of the previous example, they are still discernible from the background noise level. As in the previous example, a phase response greater than the noise level is attributed to the image where any coil is in contact with the defect, which in this instance, occurs over a larger axial extent. The fusion process provides a more representative illustration of the defect geometry, as shown in Figure 7.19, however the inaccuracies of the shape extraction process are evident in the reduction in peak intensity of the fused ECA data compared with that of Figure 7.18. This is due to the mask being smaller than it should be and therefore being zero in the pixels corresponding to the maximum total PI.

7.4 Summary

In this Chapter, a framework for processing and fusing eddy current array and visual images has been presented, built around the GAC algorithm. The techniques outlined represent a first-attempt at data fusion using the full complement of data available to the robot and are not intended to be regarded as an optimum solution. Good results have been achieved on relatively ideal synthetic defects (drill holes) where contrast within the image was good and the ECA gave strong indications of defect presence. A more comprehensive system must be developed in order to fully handle different defect geometries and the suitability of the framework must then be assessed based on its overall performance. The potential for further work within this area is extensive and the framework presented here will allow for a variety of methodologies to be implemented. The current system implements a pixel-level combination technique which is made possible by the inherent registration provided by the fusion probe. This, rather rudimentary, fusion technique is ideal as a means of proving the concept of the integrated inspection system, however, does not take account of any statistical variation in the two data sets. As part of the programme for future work, a Bayesian fusion system will be suggested which has very close similarities with the approach adopted in Chapter 5 and could be considered to be the statistically optimum method of fusion.

Chapter 8

Conclusions and Future Work

8.1 Conclusions

8.1.1 Robotic Platform

Previous research within CUE had focussed on the concept of miniature magnetic wheeled robots however the limitation of this technology to ferromagnetic surfaces, coupled with their unsuitability for anything other than large diameter pipes precluded their use in this instance. The result of the present research has been to create an entirely new robotic platform specific to pipe inspection applications. An iterative design methodology was adopted yielding, in the first instance, a linear pipe crawler capable of locomotion along straight sections of pipe. This model was refined and enhanced in a second design step yielding the lighter-weight Mark II system, featuring a range of kinematic abilities including horizontal/vertical motion and movement around bends having a radius of 4 times the pipe diameter or greater. The fundamental design of both Mark I and II platforms centres on two interconnected, semi-circular collars whose separation is controlled by a stepper motor and which constituted the method of generating axial motion. The semi-circular structure was used for its strength, facility for easy insertion onto a pipe, ability to overcome pipe hangers and the ease with which it could be manufactured. The optimised collar design of the Mark II RPC represents the minimum size solution for the given pipe diameter using such a mechanical morphology. With the advent of rapid prototyping facilities offering the ability to create highly complex geometrical shapes, more compact solutions may now be

possible that would allow the form-factor to be reduced in ways not previously possible. This was explored with the Mark IIr platform, yielding promising results.

A modular approach was taken to the design of the on-board electronics meaning that much of the hardware could be used for both the Mark I and II mechanical platforms. Four modules were used, three of which formed the main processing stack involved in robot control with the fourth being specific to the NDE payload. The central processor consists of a 400MHz embedded computer running Linux interfaced with an 8-bit microcontroller for low-level control. Significant additional capacity exists over and above that required for the robot control and payload management creating the ability for substantial processing to take place on-board. Communications occurs over an 802.11g wireless connection thereby allowing a fully wireless inspection platform and circumventing practical issues concerned with the use of an umbilical in the inspection environment. The power consumption of the system in standby mode is 3.6 W which, using the 11.1 V, 910 m Ah, Li-Po battery leads to a 3 hour standby battery life. Due to an element of mechanical backlash in the actuators used for gripping, during vertical locomotion, the effect of gravity requires that one set of grippers be energised at all times in order to maintain traction. This results in the highest power demand and thus the biggest drain on battery capacity. For continual vertical locomotion the battery lifetime reduces from 3 hours to approximately 50 minutes however, this may be increased through the use of better quality motors with equivalent backward drive force but lower backlash.

In Chapter 1, the key motivations for remote inspection and deployment of an intelligent inspection platform were identified as facilitating inspection of areas where access is restricted, increasing the efficiency of inspection and increasing the accuracy of the quantitative results. The platform that has been developed has achieved the first and, by providing a much quicker means of conducting the inspection than is required to send a human into a processing cell, also the second goal. The specification outlined in Chapter 3 required the ability to travel along pipes with a nominal diameter of 50 mm and allow for inline variations of ± 10 mm. This has been achieved through the inclusion of intermeshing grippers. The speed of the platform was regarded as a secondary design consideration behind kinematic ability and mechanical stability. Thus, for vertical ascent on a straight section of pipe, the speed equates to 4.7 mm s^{-1} however as the cornering mechanism requires a smaller step-size to be adopted, this reduces to 1.1 mm s^{-1} around bends. The significant challenge of maintaining traction on austenitic stainless steel (and other non-ferromagnetic materials) has been overcome through the application of

mechanical adhesion whilst still ensuring a low form-factor design. Furthermore, a fail-safe mechanism has been implemented through high backward drive force motors to prevent traction loss in the event of a power failure. One limitation that was not addressed, however, was the inclusion of a system for circumferential motion. Such kinematic ability is essential for inspecting typical facilities so that the robot can orient itself correctly in order to pass pipe hangers. While ultimately necessary for practical inspection, this additional degree of freedom presented an engineering obstacle that detracted from the key research objectives concerned with NDE sensing and could not therefore be justified within the given time-scale.

Given the limited resources and knowledge-base within CUE regarding the radiation hardening of materials, it was deemed to be beyond the scope of this research. While ultimately necessary for a multi-use inspection platform, this would represent no additional contribution to the field over and above that of the development of the platform and consequently was omitted from the specification. This approach allowed the focus of the research to be on the creation of the novel platform coupled with the sensing and fusion capabilities while ultimately acknowledging this basic limitation.

Two NDE payloads were integrated with the robotic platform: a camera and an eddy current array. The camera served a dual purpose role, providing the input for both defect classification and geometric shape extraction, and is discussed next.

8.1.2 Visual NDE Pattern Classification

A statistical pattern recognition system was created using 2D FFT representations of the NDE images as inputs and yielding a probability of class membership as its output. The system was first developed using a set of idealised images featuring synthetic defect geometry that was congruent with the generation of unique frequency spectra and latterly using real defects. Four classes were defined for the development system, corresponding to images having no defects, corrosion pits and horizontal or diagonal slots (with respect to the image frame). This was condensed to three for the case of real defects, with only a single orientation of slot being considered. These (more realistic) defects were created by corroding carbon steel within circular and rectangular moulds of varying dimensions. In both cases, the classification system procedure was the same and is based on the eigenfaces technique used in facial recognition systems. Principal component analysis was used to extract the eigenvectors ordered according to descending value of the variance. By constructing a feature vector consisting of a subset of these and re-

expressing the original 2D FFT data solely in terms of this subset, the size of the data set was reduced to only two dimensions for both the ideal and real defect cases. The Mahalanobis distance of each image was finally calculated to each of the known classes, assuming Gaussian distributions and concludes the classifier training procedure. As new images are gathered, the frequency data is first calculated and then expressed in terms of the feature vector before finally the probability of class membership is determined using the calculated Mahalanobis distance and the known distributions. For each class, a set of distributions, and therefore a set probability estimates, are obtained. For both data sets a validation set consisting of 10 images per class was used to fine-tune parameters such as image bandwidth, number of training images and size of the feature vector and a test set consisting of 5 and 10 images per class used, respectively, to characterise the classifier performance. Several classifier implementations were investigated: individual binary, binary fusion, individual multi-class and multi-class fusion. The performance on both data sets was benchmarked against the K-nearest neighbours algorithm in order to establish a benchmark to assess their performance.

A statistical classifier, such as this, has an important fundamental difference over other classification systems, particularly artificial neural networks, in that it is based on a measurable difference in the actual statistics of the data set. Thus, in contrast with neural networks where the result is input in advance and a set of weights adjusted to achieve the desired result, there is no design bias involved in this system, and therefore it can be confidently concluded that the classifier is responding to quantifiable underlying statistical differences in the data.

The goal of the approach was to develop a classification system, in MATLAB, that could primarily distinguish between images containing defects and those that do not. For the idealised data set it was found that excellent classification performance could be achieved. A classification rate of 100% was specified and was achieved using a multi-class sensor approach using the product rule on both the validation and test sets. The individual binary ND classifier yielded 97.5% which, although marginally short of the required performance, is still considered to be a very good result. The greater computational simplicity (no fusion) of the individual binary approach comes at the expense of a less specific assignment of the image class. Both results are backed up by very good separation of the underlying data set allowing a high degree of confidence to be placed in the result. Significantly poorer performance was achieved using the real data set generated by artificially corroding steel samples. The product rule, once again, performed best with a correct classification rate of 73% with the individual binary ND classifier yielding 70%.

Analysis of the class distributions for this data set indicated that the poor performance was due to the fundamental lack of distinction between the classes rather than a failure in the system itself. This was particularly evident between the ND and CP classes.

The multi-class sensor system using the product rule was shown to achieve a performance comparable to a KNN algorithm on both data sets. In addition to this, because of the fundamental statistical basis of this system, a level of confidence could be attributed to the classification which, if emulated under a KNN approach, would require a value of K that would actually reduce the classification rate. The advantage, however, of the KNN algorithm is that it has significantly less computational overhead associated with it and is therefore potentially more suitable for on-board, real time processing. Overall, it can be concluded that, where good separation can be achieved in the underlying data and where a statistical confidence is more valuable than computational complexity, a multi-class sensor approach is preferable over a simple KNN algorithm. For NDE applications, it is felt that the hard decision output (class label) of a classifier such as KNN is less appropriate and that more value is gained by having some concept of how confident the system is in that classification. Given the high level of processing capabilities designed into the platform the computational expense is not considered to be too significant, particularly since the training process can be done offline.

For both data sets, the image bandwidth was found to have less of an effect on the classification rate than the overall dimensionality of the data set. The results of the real image set highlights a key limitation of the system. The pre-processing stage is based on the 2D FFT and, as real defects are unlikely to have straight edges, frequency components are likely to appear at a range of orientations. Furthermore, regions of high contrast, such as dirt deposits, on the surface will generate frequency content in the same way that a defect of the same shape would, resulting in false positives. Lastly, due to the translation invariant property of the 2D FFT, distributed regions of high contrast will generate frequency content in a wide range of orientations, mimicking the defining characteristic of a corrosion pit class. The latter two points can be easily addressed through the incorporation of a treatment phase, such as dye penetrant with an associated developer, prior to the pre-processing stage in order to generate a monochrome image. Adoption of such a system would largely eliminate the effect of contrast regions within images having no defect and thus would offer a significantly improved correct classification potential for this high priority class. However, this is at the expense of the

incorporation of a system for dispensing these fluids in an appropriately controlled manner. For defects such as a fatigue crack, it is likely that a different approach to pre-processing would be required. Principally, the subject of differentiating between scratches and cracks would need to be investigated as well as a suitable filtering strategy to account for the higher frequencies of interest signifying such defects. The classification methodology however could be used without modification since provided the input data has a statistical difference, the proposed system will be able to identify this and thus classify accordingly.

In practice, it is difficult to obtain the large volume of unique images required for a truly unbiased performance assessment. Consequently it is acknowledged that the results reported are likely to be skewed in favour of a correct classification. While this may be true, the relative comparison of the system against a recognised benchmark such as the KNN yields a valuable result in itself.

Ultimately, this research has demonstrated that the Eigenfaces technique for pattern recognition can be successfully applied to NDE images. Further research is, however, required in order to address the limitations that have been identified.

8.1.3 Eddy Current Array

An integrated eddy current array has been developed that uses a frequency sweep (100-200kHz) as the excitation signal. The motivation for the use of such a drive signal lay in the observation that the excitation frequency that maximised the phase difference between two coils (where one was positioned over a defect) varied with the actual geometry of the defect itself. An equation specifying the optimum frequency for a phase detection systems, in terms of the RLC components of the drive circuit, has been established. Due to the variation in conditions of the underlying material, a single frequency system is unlikely to be operating at the optimal frequency and so a swept frequency approach is desirable. A theoretical framework established a relationship between defect geometry and the integrated phase difference of the coil voltages across the frequency range. For small defects, while the phase difference at each frequency interval may also be small, the accumulated difference across the frequency range yields a metric with a relatively large change and thus, in theory, a more sensitive system. For both aluminium and 304 austenitic stainless steel, it was found that, given the diameter of a circular defect, the depth could be determined and vice-versa. In contrast to other multiple frequency methods, a significant advantage of using this defect metric was in the relative insensitivity to the drive circuit series resistance. It was shown that for a

given change in inductance, the peak phase difference reduced by 70% while the PI reduced by only 2%. The use of a frequency sweep excitation has the additional advantage over a single frequency excitation of being able to exploit the known functional relationship between the phase difference and excitation frequency of two coils and was shown to reduce noise in an image of an aluminium plate.

The aim of the ECA research was to develop a differential sensor array that could be integrated with the robotic pipe inspection system and was capable of detecting defects down to 0.1 mm in depth. The ECA was found to be capable of detecting a 0.2 mm deep defect which falls short of the required performance. A standard commercial system as tested on the same specimen and was able to detect all test defects (down to 0.05 mm). However, the hardware associated with this device constituted a size, weight and cost that was not feasible for adaption within a lightweight, low power, mobile inspection platform. Furthermore, access to the raw data for post-processing is not available. Given the theoretical basis underpinning the integrated system, it is concluded that the reduced sensitivity is largely associated with the quality of the electronic design. The limitation of the 10 bit sampling accuracy coupled with the relatively limited in-house manufacturing facilities mean that achieving comparable SNR with that of a system produced on a commercial budget is unrealistic. However, the significant advantage of an integrated solution, as opposed to an adapted off-the-shelf solution, lies firstly in the flexibility to tailor to an inspection task and secondly in the ability to access data for further processing. The latter is particularly important where the evaluation of different imaging algorithms may be desirable. The speed of the system, given its developmental nature, is relatively high and equates to 42 single frequency acquisitions and phase calculations per second. As the dominant overhead is the phase detection algorithm, given further development and refinement of the program code, the speed could be increased considerably. The use of additional channels for parallel processing would increase this further and is suggested as part of the future work programme.

As the coils were wound manually, there was an inevitable variation in performance and consequently a calibration procedure was necessary. As the frequency response of the coil is related to the defect size, a frequency-by-frequency correction becomes less accurate as the defect size increases (and the frequency of the peak phase difference changes). It was concluded, therefore, that the only way to substantially improve the result was to limit the variation in coil response through higher precision construction ensuring consistent numbers of total windings and consistent number of windings per layer.

The full coil was modelled using a FEM package. Limitations in this model arose due to the small wire gauge and were manifested by spatial quantisation in the outer conductors of the winding. This led to an inadequate model of the skin effect resulting in a noticeable discrepancy in the resistance vs. frequency characteristic.

The differential nature of the ECA was established for its suitability to mobile robotic applications. However, this restricted the applicability of more common imaging techniques that utilise impedance calculations and therefore a custom algorithm was developed. One of the difficulties arising from imaging the phase difference of a differential pair of coils lies in the fact that it is not possible, from a single probe position, to know under which coil a defect lies. Thus, the measured phase difference must be assumed to belong to either coil and therefore added to the corresponding pixels of the resulting image. However, it would be expected that if a coil lies over a defect then any phase comparison between this and a coil on defect-free metal would yield a difference. An unavoidable uncertainty, however, arises from the fact that the same result would be obtained in the former situation as would be in the inverse case, where all coils but one are affected by a defect. This leads to two possible solutions for a given probe location. The first solution to this is to assume that, given the size of the array, a defect is unlikely to occupy the majority of the surface and therefore it should always be assumed that the defect is the smaller of the two solutions. This, however, leads to errors where in fact such large defects *are* encountered. The second, more computationally demanding, solution involves using the data history from a series of probe positions to update the prediction of the defect size and shape. Such a technique would result in greater SNR since no phase contributions would be added to pixels corresponding to coil locations that do not have a defect.

The fundamental limitations of the ECA system are primarily the limitations of the eddy current modality. Local variations in magnetic permeability and conductivity in an otherwise consistent material represent an inspection issue that cannot be calibrated for and will therefore, in all probability, result in a false positive for structural defects. The fusion of ECA data with another modality (e.g. visual) should provide a means of overcoming this.

8.1.4 Data Fusion and the Integrated Inspection System

The data fusion implementation was shown to yield a valuable improvement in the clarity of the defect over either data set individually. The shape extraction system

was built around the geometric active contours algorithm in order to provide a very flexible means of detecting arbitrary defect shapes. In order to achieve best performance it was found that an initial starting contour should be supplied to the GAC algorithm and this was implemented for circular defects. The initial contour, in this instance, was generated using the Hough transform to identify the location of the centre of any circles having a radius, r . Following this, an initial contour is estimated as a circle having a radius $\frac{r}{2}$ and centred at these points. Once a stable solution has been achieved for each defect, a combined binary mask is created representing the shape of any identified defect regions within the image frame.

A pixel-level multiplication approach was used to combine the defect mask and the eddy current array, with image registration provided intrinsically through the use of the fusion probe. The multiplication algorithm is basically a probabilistic approach with the uncertainty of the camera set to zero and thus, any areas not suspected by the camera as having a defect, are taken as defect free. Good results were achieved for the two circular defects with a significantly better correlation between the appearance of the fused data and the actual defect than was evident from the eddy current data alone. The presence of the defect is confirmed through strong indications from the eddy current array at the precise location and has been shown to be a distinguishable characteristic, that can minimise the risk of false positives due to irrelevant surface aberrations. However, the effect of surface curvature was such that the contrast in certain areas of the defect was poor and subsequently the adaption to the defect shape suffered at these locations. The fusion process was repeated for a 2.5mm wide by 1.5mm deep, two-angled slot in the same pipe material. It was found that if an initial contour was manually defined to lie within the region of the slot, a reasonably good representation of the defect shape could be extracted. The size of the defect relative to that of the coils meant that the peak magnitude of the eddy current image in the location of the defects was only 10% of the previous result, although still discernible from the background noise level. The result of the data fusion was consequently poorer than for the well defined through-thickness holes, however still yielded a more accurate representation of the shape than the ECA data taken in isolation. Having said that, the inaccuracies of the shape extraction process coupled with the severity of the multiplication operator resulted in attenuation of the eddy current data in certain regions.

An issue with the data fusion mechanism is the effect of surface curvature on the shape extraction process which tends to distort the natural defect shape. It was shown that this effect caused circular objects to appear elliptical and, consequently, so to was the extracted defect mask. This can be addressed through the

implementation of a matrix transformation using the known lens position relative to the pipe surface and the lens' optical parameters. The need to detect the defect centres in order to define the initial contour complicates the process for unknown radii of defects and therefore requires a repeated implementation of the Hough process for radial increments. The primary issue with the method of combination is the zero uncertainty assumption of the camera. An extension is therefore proposed in the future work section to address this.

This research has taken a significant step towards developing an extensive inspection system that can be used for real industrial pipe inspection applications. The complete inspection system is shown in Figure 8.1 performing an inspection of a 50mm diameter, 304L austenitic stainless steel pipe and depicted as a functional block diagram in Figure 8.2. All sensing capabilities rely on the robotic platform to deliver them to the desired location and provide a stable structure on which to perform the inspection. Each payload provides a result that can either be assessed individually or fed into the data fusion process. A fusion probe has been developed to provide a means of directly registering the eddy current and visual data sets. Thus, having arrived at the inspection point, the camera is used to capture an image of the pipe and, following this, the eddy current array is scanned over the area of the captured image frame. The process of combining the sensor results links all of the units of this research and yields a single unified inspection result, that enhances the clarity of the defect. Following the defect classification process, an additional processing stage is required, to extract the shape of the defect. This has been demonstrated convincingly for a 304L austenitic stainless steel pipe containing two 10mm diameter through-thickness holes.

8.2 Future Work

The suggestions for further work can be broadly characterised into one of the following two divisions: robot design and NDE payloads. The robot design is further sub-divided into mechanical and electronic design and these along with each of the sensing payloads will be addressed in turn.

8.2.1 Robot Design

The primary goal of the robotic research within CUE is to produce systems that are capable of operating in real industrial environments and that would be adopted by the end-users. Consequently, it is of critical importance that the RPC be developed to incorporate a means of circumferential motion.




Figure 8.1: Mark IIr RPC inspecting a 50mm diameter, 304L austenitic stainless steel pipe using the fusion probe.

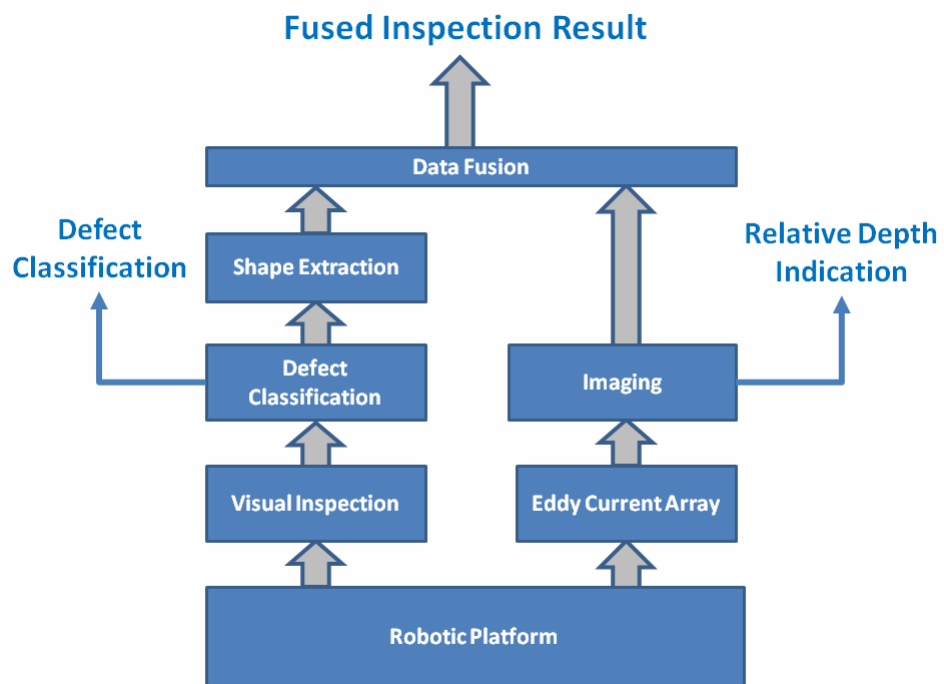


Figure 8.2: Integration of the inspection system with each ascending level building on the functionality of the previous level.

As outlined earlier, such kinematic ability is essential in order to allow the robot to correctly orient itself for passing pipe hangers and inspection of the full extent of the pipe surface. A possible solution that minimises the effect on the overall chassis design, would be to create a bespoke gripper fitted with omni-directional wheels [140]. The advantage of this would be two-fold, with the ability to travel around the circumference being complemented by a means of travelling along straight sections of pipe much more quickly. Although this increases the kinematic redundancy, as discussed in Section 3.3.2, this may be advantageous where long sections of pipe must be traversed. By modifying the gripper rather than the structure of the chassis, the means of obstacle avoidance is preserved.

In addition to this, the gripping force should be increased to allow heavier payloads to be carried as well as increasing the reliability of the traction mechanism and facilitating low power battery operation through judicious motor control. This is most effectively achieved through the use of more powerful motors, rather than the passive spring system outlined in section 4.4.1.1, which is more complex. A key feature of the motors would be a low mechanical backlash together with a backward drive force sufficient to support the weight of the robot and payload. The results of the FEA have been supported by the observed performance, with the structural strength more than capable of handling the forces exerted on it. Given these results, there is scope to increase the motor force, as indicated, while still meeting the minimum displacement criterion.

With the most significant proportion of the development complete, the modularity of the stack electronics becomes less critical. The integration of these into a single commercially manufactured board would minimise the form-factor and allow them to be encased in a shielded housing. This would provide a more robust system with less risk of the associated interconnect becoming entangled on the surrounding infrastructure as well as providing a degree of ‘weather-proofing’. At present, the bandwidth of the communications is not really utilised to any significant degree. Addition of a real-time video link, ideally mounted on a 2 DOF platform, would be invaluable for robot control and as a navigational tool. Furthermore, this would allow visual inspection of adjacent pipework thereby increasing the power of the system as an inspection tool. The performance effect of the (potentially) severe multi-path from densely packed pipework should also be assessed. One area in particular that has not been investigated within this project has been that of accurate positioning however, this is currently a very active area of research within CUE. The use of trilateration is limited by the lack of available access ports and the desire to minimise the number that are opened at any one time. Laser tracking

technology [141] is a possible solution with the RPC being fitted with a reflective beacon and the laser tracker automatically following the reflected signature. The accuracy of such systems is reputed to be in the sub-millimeter range over tens of meters and algorithms are generally built-in to allow a degree of failure of LOS.

8.2.2 NDE Payloads

It was discussed in Chapter 5 that where a unique frequency spectra could be obtained, good classification performance could be achieved. The deterioration of this for more representative images suggests that a pre-treatment phase should be introduced prior to capturing images. In so doing, the effect of dirt deposits would be less detrimental to the classification performance since these would be largely masked by the application of the developer fluid. Furthermore, the effectively monochrome image rather than the highly reflective, grayscale image of the pipe surface would yield an input that was significantly less sensitive to the effect of the surface curvature on the level of light reflection. While the use of such a technique cannot change the intrinsically less uniform edges found in real defects, the performance of the defect/no-defect classification would be increased. As the cost of an error is greater for this classification compared with that of sub-categorising which class of defect may be present, this would be a significant improvement. However, the inclusion of such a process would increase the overall time required to perform an inspection. Finally, the vast processing capacity of the main processor leads to the obvious idea of implementing the classifier on-board the RPC. The advantage of this would be to obtain a greater level of autonomy, with the inspection system then able to operate unsupervised until a defect is encountered.

The main improvement in terms of the eddy current array payload involves the creation of a set of precision coils with a more consistent frequency response across the working range. Coupled with the production of higher quality drive and sensing electronics, this would achieve a much higher SNR and potentially allow smaller defects to be detected. The inclusion of a simultaneous drive and receive signal capture would allow the phase of each coil to be detected relative to the input signal. In so doing, the polarity of the phase shift could be determined and consequently the location of the defect could be identified to a particular coil. This would greatly improve the resulting image generated by eliminating this positional uncertainty in a computationally efficient way. Given the luxury of more time, more robust and time efficient software could have been developed. The use of the digital filtering capacity of the DSP unit led to a self-contained payload system

however, by capturing the data and then processing it on the higher memory, main processor, the accuracy of the phase calculation could be increased leading to a lower rounding error. Finally, the speed of the system should be increased through the implementation of parallel coil excitation and processing.

Given the statistical nature of the defect classification system and combination, the obvious extension is to apply this to the data fusion process. By quantifying the uncertainty in the shape extraction process along with the uncertainty in the size of the defect from the eddy current array, the two data sets can be combined using Bayes' theorem in exactly the same manner as was used to enhance the classification results of the multi-class sensor ensemble. The main difficulty with this lies in quantifying the uncertainty in the defect shape. For simple geometry such as circles, this can be achieved relatively simply through repeated measurement of the defect diameter. However, for more complicated geometry which cannot be characterised by a single metric, the complexity of the problem increases with the number of varying dimensions. The significant advantage of such a fusion approach, however, is that it would represent the statistically optimum solution. An additional suggestion is to implement the functionality of the data fusion process on-board in order to achieve true, semi-autonomous operation. Given the extensive processing capabilities designed into the platform, this will be a trivial, albeit time-consuming, process.

The concluding suggestion involves the integration of a new sensor payload in the form of an integrated ultrasonic phased array system. With advances in electronic technology increasing the processing capabilities within a given form-factor, the development of a phased array controller suitable for mobile robot applications will soon be realised. This will combine the benefits of robotic inspection with a highly flexible, precision sensing system.

8.3 Technology Transfer

This section is intended to be a footnote to the thesis and will briefly outline some of the key points that must be addressed before this system can be usefully used by the industrial end-user.

The system in its current shows a high level of potential for inspection use however it is not yet suitable for deployment in a nuclear processing cell. The design and kinematic functionality of the system is good however the addition of circumferential motion must be pursued. The issue of radiation hardening must

also be given a more thorough analysis and a review of the material choices for all mechanical and electrical components undertaken (particularly to remove any Chlorides). The results of the pattern recognition system are promising and some dedicated research on this could yield a system that could then be implemented on the robotic platform. In its present state, while the theoretical basis is sound, none of the classifiers are regarded as being robust enough for field use and require a more rigorous verification process, ideally with images from the cells in which the inspections would typically take place. The integrated eddy current array as an inspection payload also shows significant promise and ultimately only requires ruggedising and professional quality fabrication before it can be usefully used for NDE. As the demand for remote inspection in industry increases it is expected that external pipe inspection systems, such as that outlined in the thesis, will form a meaningful part of a regular inspection programme.

Appendix A

Eddy Current Array

A.1 Analysis of Resistive, Inductive and Capacitive (RLC) Circuit

The following section provides a derivation of the phase difference for a RLC circuit as used in Chapter 5.

A.1.1 RLC Impedance

The impedance (Z) of the circuit consisting of a series combination of a resistor, inductor and capacitor (RLC), shown in Figure A.1 is given by Equation (A.1).

$$Z = R + j\omega L - \frac{j}{\omega C} \quad (\text{A.1})$$

where R is the resistance, ω is the angular frequency (rad/s), L is the inductance of the coil (H) and C is the capacitance (F). Expressed in polar form:

$$Z = \sqrt{\left[R^2 + \left(\omega L - \frac{1}{\omega C} \right)^2 \right]} \angle \text{atan} \left(\frac{\omega L - \frac{1}{\omega C}}{R} \right)$$
$$Z = \sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C} \right)^2 \right]} \angle \text{atan} \left(\frac{\omega^2 LC - 1}{\omega CR} \right)$$

Thus, the magnitude and phase are given by:

$$|Z| = \sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C} \right)^2 \right]} \quad (\text{A.2})$$
$$\theta = \text{atan} \left(\frac{\omega^2 LC - 1}{\omega CR} \right)$$




Figure A.1: Circuit diagram of resistor-inductor-capacitor circuit (RLC)

Resonance occurs when the impedance of the circuit is minimum and hence when,

$$\frac{\omega^2 LC - 1}{\omega C} = 1 \Rightarrow f_{res} = \frac{1}{2\pi\sqrt{LC}}$$

And at resonance, $|Z| = R$ and $\theta = 0^\circ$.

The voltage across the eddy current coil may be calculated using the voltage divider rule as follows:

$$V_L \angle \theta = V_s \angle 0 \cdot \left(\frac{j\omega L}{Z} \right) = V_s \angle 0 \cdot \left(\frac{j\omega L}{R + j\omega L - \frac{j}{\omega C}} \right)$$

$$H = \frac{V_L \angle \theta}{V_s \angle 0} = \left(\frac{j\omega L}{R + j\omega L - \frac{j}{\omega C}} \right) = \frac{\omega L \angle \frac{\pi}{2}}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C} \right)^2 \right]} \angle \text{atan} \left(\frac{\omega^2 LC - 1}{\omega CR} \right)}$$

Separating the real and imaginary parts of H yields:

$$\begin{aligned} \operatorname{Re}\{H\} &= \frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \cos\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right] \\ \operatorname{Im}\{H\} &= \frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \sin\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right] \end{aligned}$$

with the magnitude of the voltage given by Equation (A.3).

$$|V| = \sqrt{\operatorname{Re}\{H\}^2 + \operatorname{Im}\{H\}^2} \quad (\text{A.3})$$

The magnitude of the coil voltage, V_L , for a unity voltage input, V_s , is given by:

$$|V| = \sqrt{\left[\frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \cos\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right] \right]^2 + \left[\dots + \frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \sin\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right] \right]^2}$$

Expanding and simplifying, the following expression is obtained:

$$|V| = \sqrt{\frac{L^2 \omega^4 C^2}{\omega^2 C^2 R^4 + \omega^4 L^2 C^2 - 2\omega^2 LC + 1}} \quad (\text{A.4})$$

The phase of the coil voltage is given by ϕ :

$$\phi = \operatorname{atan}\left(\frac{\operatorname{Im}\{H\}}{\operatorname{Re}\{H\}}\right) = \operatorname{atan}\left(\frac{\frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \sin\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right]}{\frac{\omega L}{\sqrt{\left[R^2 + \left(\frac{\omega^2 LC - 1}{\omega C}\right)^2\right]}} \cos\left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right]}\right)$$

$$\phi = \operatorname{atan}\left[\tan\left(\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega CR}\right)\right)\right]$$

$$\phi = \frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 LC - 1}{\omega C R}\right) \quad (\text{A.5})$$

From Equation (A.5), it can be seen that, at resonance, the phase of the voltage across the inductor will be 90° relative to the supply voltage. The phase difference between two coils, L_1 and L_2 , is given by Equation (A.6).

$$\begin{aligned} \Delta\phi &= \left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 L_2 C - 1}{\omega C R}\right) \right] - \left[\frac{\pi}{2} - \operatorname{atan}\left(\frac{\omega^2 L_1 C - 1}{\omega C R}\right) \right] \\ \Delta\phi &= \operatorname{atan}\left(\frac{\omega^2 L_2 C - 1}{\omega C R}\right) - \operatorname{atan}\left(\frac{\omega^2 L_1 C - 1}{\omega C R}\right) \end{aligned} \quad (\text{A.6})$$

References

- [1] Frost & Sullivan, “Trends in the global NDT equipment market”, 18th August 2009
- [2] WNA Report, “The New Economics of Nuclear Power”, World Nuclear Association
- [3] Health and Safety Executive, “Report of the Investigation into the Leak of Dissolver Product Liquor at the Thermal Oxide Reprocessing Plant (THORP), Sellafield, Notified to HSE on 20 April 2005”
- [4] “Sellafield radioactive leak to cost £300m”, Paul Brown, The Guardian (UK), Monday 13th June 2005.
- [5] Frost & Sullivan, “Non-Destructive Test Equipment Segments - An Overview”, 6th September 2006
- [6] Frost & Sullivan, “Non-destructive Test Equipment Market - Consolidating Towards a Digital Future (Part 1)”, 10th March 2004
- [7] M Friedrich, W Galbraith and G Hayward, “Autonomous Mobile Robots for Ultrasonic NDE”, Proceedings of IEEE International Ultrasonics Symposium, Vancouver, pp 902-905, October 2006.
- [8] Hargreaves, M. L., “Digital Processing of Ultrasound Signals Back-scattered from Coarse Grained Austenitic Stainless Steel”, PhD Thesis, University of Keele, 1988
- [9] Whitaker, J. & Jessop, T., “Ultrasonic Detection and Measurement of Defects in Stainless Steel-A Literature Survey”, Br. J. Non-Destr. Test., 1981, 23, 293-303
- [10] Tomlinson, J., Wagg, A. & Whittle, M. “Ultrasonic inspection of austenitic welds” British Journal of Non-Destructive Testing, 1980, 22, 119-127
- [11] Hudgell, R., Seed, H. “Ultrasonic longitudinal wave examination of austenitic welds” British Journal of Non-Destructive Testing, 1980, 22, 78-85
- [12] Silk, M., Lidington, B., Hammond, G., “A time domain approach to crack location and sizing in austenitic welds”, British Journal of Non-Destructive Testing, 1980, 22, 55-61
- [13] M. Sitti and R.S. Fearing, “Synthetic Gecko Foot-Hair Micro/Nano-Structures for Future Wall-Climbing Robots”, IEEE Proceedings-International Conference on Robotics and Automation, Vol 1, pp 1164-1170, Sep 2003.

- [14] C. Menon, M. Murphy and M. Sitti, "Gecko Inspired Surface Climbing Robots", IEEE International Conference on Robotics and Biomimetics, pp 431-436, August 2004.
- [15] S. Kim, A. T. Asbeck, M. R. Cutkosky and W. R. Provancher, "SpinybotII: Climbing Hard Walls with Compliant Microspines", IEEE Proceedings-12th International Conference on Advanced Robotics, pp 601-606, July 2005.
- [16] Z.Y. Qian, Y.Z. Zhao, Z. Fu and Y. Wang, "Fluid Model of Sliding Suction Cup of Wall-climbing Robots", International Journal of Advanced Robotic Systems, Vol 3, No 3, pp275-284, 2006.
- [17] D. Longo and G. Muscato, "A modular approach for the design of the Alicia3 climbing robot for industrial inspection", Industrial Robot: An International Journal, Vol 31, No 2, pp 148-158, 2004.
- [18] Pelrine, R.; Prahlad, H.; Eckerle, J.; Kornbluh, R. & Stanford, S. Electroadhesion, Patent number 7773363, 2007
- [19] www.sri.com/rd/electroadhesion.html
- [20] Hagen Schempf, Edward Mutschler, Alan Gavaert, George Skoptsov, William Crawley, "Visual and Nondestructive Evaluation of Live Gas Mains Using the Explorer Family of Pipe Robots", Journal of Field Robotics 27(3), 217-249 (2010), © 2010 Wiley Periodicals, Inc.
- [21] Yunwei Zhang, Guozheng Yan, "In-pipe inspection robot with active pipe-diameter adaptability and automatic tractive force adjusting", Mechanism and Machine Theory 42 (2007), 1618-1631.
- [22] Zheng Hu and Ernest Appleton, "Dynamic characteristics of a novel self-drive pipeline pig", IEEE transactions on robotics, vol 21, no 5, pp781-789 october 2005
- [23] Gálvez, J.; De Santos, P. & Pfeiffer, F. "Intrinsic tactile sensing for the optimization of force distribution in a pipe crawling robot", IEEE/ASME Transactions on mechatronics, 2001, 6
- [24] Suzumori, K.; Miyagawa, T.; Kimura, M. & Hasegawa, Y., "Micro inspection robot for 1-in pipes", IEEE/ASME transactions on mechatronics, 1999, 4, 286-292
- [25] R K Stanley, "Recent advances in robotic coiled tubular assessment", Insight, Vol 49, No9, September 2007.
- [26] Koichi Suzumori, Shuichi Wakimoto, Masanori Takata, "A miniature inspection robot negotiating pipes of widely varying diameter", Proceedings of the 2003 IEEE International Conference on Robotics and Automation
- [27] P. Chatzakos, Y.P. Markopoulos and K. Hrissagis, "On the development of a modular external-pipe crawling omni-directional mobile robot", Industrial Robot: An International Journal, Vol 33, No 4, pp291-297, 2006.

- [28] Changhwan Choi, Seungho Jung, Seungho Kim, “Feeder Pipe Inspection Robot with an Inch-worm Mechanism Using Pneumatic Actuators”, *International Journal of Control, Automation and Systems*, vol 4, no 1, pp 87-95, February 2006
- [29] Saltaren, R.; Aracil, R.; Reinoso, O. & Scarano, M. Climbing parallel robot: a computational and experimental study of its performance around structural nodes *Robotics, IEEE Transactions on*, IEEE, 2005, 21, 1056-1066
- [30] R. Aracil, R. Saltaren, O.Reinoso, 'Parallel robots for autonomous climbing along tubular structures', *Robotics and Autonomous Systems*, Vol 42, pp 125-134, 2003.
- [31] M. Tavakoli, M.R. Zakerzadeh, G.R. Vossoughi and S. Bagheri, “A hybrid pole climbing and manipulating robot with minimum DOFs for construction and service applications”, *Industrial Robot: An International Journal*, Vol 32, No 2, pp 171-178, 2005.
- [32] M. Almonacid, R. J. Saltaren, R. Aracil, O.Reinoso, “Motion Planning of a Climbing Parallel Robot”, *IEEE Transactions on Robotics and Automation*, Vol 19, No 3, pp485-489, June 2003.
- [33] R. Aracil, R. Saltaren, O.Reinoso, “A climbing parallel Robot”, *IEEE Robotics & Automation Magazine*, pp 16-22, March 2006.
- [34] Xia-Shan Gao, Deli Lei, Qizheng Liao, Gui-Fang Zhang, “Generalized Stewart-Gough Platforms and Their Direct Kinematics”, *IEEE Transactions on Robotics*, Vol 21, No 2, pp 141-151, April 2005.
- [35] Halmshaw, R., “Non-Destructive Testing.”, Edward Arnold(UK), 1991,, 1991, 323
- [36] Kupperman, D. & Reimann, K., “Ultrasonic Wave Propagation and Anisotropy in Austenitic Stainless Steel Weld Metal”, *IEEE Transactions on Sonics and Ultrasonics*, 1980, 27, 7-14
- [37] McMaster, R.; McIntire, P. & Mester, M. *Electromagnetic testing: eddy current, flux leakage, and microwave nondestructive testing* American Society for Nondestructive Testing, 1986
- [38] *Non-Destructive Testing Handbook, Volume 4 - Electromagnetic Testing*. ASNT.
- [39] Haldane, R., Yin, W., Strangwood, M., Peyton, A. & Davis, C., “Multi-frequency electromagnetic sensor measurement of ferrite/austenite phase fraction--Experiment and theory”, *Elsevier*, 2006, 54, 1761-1765
- [40] Hansen, J. “The eddy current inspection method. Part 2. The impedance plane and probes”, *Insight Non-Destructive Testing and Condition Monitoring*, BINT, 2004, 46, 364-365

- [41] Hansen, J. "The eddy current inspection method. Part 4. Applications, practical testing and advanced concepts", *Insight*, 2004, 46, 480-483
- [42] McGonnagle, W. "Nondestructive testing", Gordon and Breach, 1966
- [43] Yin, W. & Peyton, A., "Thickness measurement of non-magnetic plates using multi-frequency eddy current sensors", *NDT & E International*, Elsevier, 2007, 40, 43-48
- [44] Hansen, J. "The eddy current inspection method: Part 3. Instrumentation and applications", *Insight, Non-Destructive Testing and Condition Monitoring*, BINDT, 2004, 46, 414-416
- [45] "Vector 22 - Flaw Detection and Material Property Evaluation in Metal Products", 2005, General Electrical Measurement & Control.
- [46] "Phasec 3 Series - General Applications", 2008, General Electrical Measurement & Control.
- [47] US-454 series, Uniwest.
- [48] ZETEC MIZ-21B, Zetec Inc.
- [49] OMNIScan MX ECA, Olympus NDT Inc.
- [50] "Multiscan MS 5800 High Speed, Multichannel, multifrequency NDE System", P/N ZDS0171 Rev 02/06, 2005, Zetec Inc.
- [51] "TC7700 Remote Data Acquisition Unit", P/N ZDS0169 Rev 2/06, 2005, Zetec Inc.
- [52] "ECVision - Advanced Analysis Software for Eddy Current Arrays", Rev 12/07, Zetec Inc.
- [53] "Flexible Eddy Current Array Probe for Weld Inspections", Rev 8/07, Zetec Inc.
- [54] Xu, L.; Krzyzak, A. & Suen, C., "Methods of combining multiple classifiers and their applications to handwriting recognition", *IEEE Transactions on Systems Man and Cybernetics*, 1992, 22, 418-435
- [55] Lim, S.; Lee, K.; Byeon, O. & Kim, T., "Efficient iris recognition through improvement of feature vector and classifier", *ETRI journal*, Citeseer, 2001, 23, 61-70
- [56] Turk, M. & Pentland, A., "Eigenfaces for recognition", *Journal of cognitive neuroscience*, MIT Press, 1991, 3, 71-86
- [57] Lyons, M.; Budynek, J. & Akamatsu, S., "Automatic classification of single facial images", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1999, 21, 1357-1362
- [58] Ahanathapillai, V.; Soraghan, J.; Hamilton, D. & Morison, G., "Echocardiographical Sequence Analysis for the Diagnosis of Heart Wall Damage", *Digital Signal Processing*, 2007 15th International Conference on, 2007, 155-158

- [59] Schwender, H.; Zucknick, M.; Ickstadt, K. & Bolt, H., "A pilot study on the application of statistical classification procedures to molecular epidemiological data", *Toxicology Letters*, Elsevier, 2004, 151, 291-299
- [60] He, Q. & Wang, J., "Principal component based k-nearest-neighbor rule for semiconductor process fault detection", *American Control Conference*, 2008, 2008, 1606-1611
- [61] Vailaya, A.; Figueiredo, M.; Jain, A.; Zhang, H.; Technol, A. & Alto, P., "Image classification for content-based indexing", *IEEE Transactions on Image Processing*, 2001, 10, 117-130
- [62] Wang, G. & Liao, T. W., "Automatic identification of different types of welding defects in radiographic images", *NDT & E International*, 2002, 35, 519 - 528
- [63] Gayer, A.; Saya, A. & Shiloh, A., "Automatic Recognition of Welding Defects in Real-Time Radiography", *Ndt International*, Elsevier, 1990, 23, 131-136
- [64] da Silva, R.; Caloba, L.; Siqueira, M. & Rebello, J., "Pattern Recognition of Weld Defects Detected by Radiographic Test", *Ndt & E International*, Elsevier, 2004, 37, 461-470
- [65] Wang, Y., Sun, Y., Lv, P. & Wang, H., "Detection of Line Weld Defects Based on Multiple Thresholds and Support Vector Machine", *NDT & E International*, Elsevier, 2008, 41, 517-524
- [66] Zhao, W.; Chellappa, R.; Phillips, P. & Rosenfeld, A., "Face recognition: A Literature Survey", *Acm Computing Surveys (CSUR)*, ACM, 2003, 35, 399-458
- [67] Osuna, E.; Freund, R. & Girosi, F., "Training Support Vector Machines: An Application to Face Detection", *cvpr*, 1997, 130
- [68] Ganster, H.; Pinz, A.; Rohrer, R.; Wildling, E.; Binder, M. & Kittler, H., "Automated melanoma recognition", *IEEE Transactions on Medical Imaging*, Citeseer, 2001, 20, 233-239
- [69] Lee, T.; Ng, V.; Gallagher, R.; Coldman, A. & McLean, D., "Dullrazor: A software approach to hair removal from images", *Computers in Biology and Medicine*, Elsevier, 1997, 27, 533-543
- [70] Zharkova, V. & Jain, L., "Introduction to pattern recognition and classification in medical and astrophysical images", *Artificial Intelligence in Recognition and Classification of Astrophysical and Medical Images*, Springer, 2007, 1-18
- [71] www.hologic.com/en/breast-screening/imagechecker/, Accessed 27th December 2010

- [72] Astley, S. & Gilbert, F., "Computer-aided detection in mammography", *Clinical radiology*, Elsevier, 2004, 59, 390-399
- [73] Brem, R. & Schoonjans, J., "Radiologist detection of microcalcifications with and without computer-aided detection: a comparative study", *Clinical radiology*, Elsevier, 2001, 56, 150-154
- [74] Yu, S. & Guan, L., "A CAD system for the automatic detection of clustered microcalcifications in digitized mammogram films", *Medical Imaging, IEEE Transactions on*, IEEE, 2002, 19, 115-126
- [75] Meinel, L.; Stolpen, A.; Berbaum, K.; Fajardo, L. & Reinhardt, J., "Breast MRI lesion classification: Improved performance of human readers with a backpropagation neural network computer-aided diagnosis (CAD) system", *Journal of Magnetic Resonance Imaging*, Wiley Online Library, 2007, 25, 89-95
- [76] Karssemeijer, N., "Stochastic model for automated detection of calcifications in digital mammograms", *Image and Vision Computing*, Elsevier, 1992, 10, 369-375
- [77] Ahanathapillai, V.; Soraghan, J.; Hamilton, D. & Morison, G., "Echocardiographical Sequence Analysis for the Diagnosis of Heart Wall Damage", *Digital Signal Processing*, 2007 15th International Conference on, 2007, 155-158
- [78] Ahanathapillai, V.; Soraghan, J. & Sonecki, P., "Delaunay Triangulation Based Image Enhancement for Echocardiography Images", *17th European Signal Processing Conference (EUSIPCO 2009)*, 2009
- [79] Turmon, M.; Pap, J. & Mukhtar, S., "Statistical Pattern Recognition for Labeling Solar Active Regions: Application to SOHO/MDI Imagery", *The Astrophysical Journal*, IOP Publishing, 2002, 568, 396
- [80] Priyantha, N., "The cricket indoor location system", PhD Thesis, 2005
- [81] Friedrich, M., "Design of miniature mobile robots for non-destructive evaluation", PhD Thesis, University of Strathclyde, 2007
- [82] PVC-CAW/PVC-HSV/PVC-MZ, rev 01/94, Simona AG
- [83] Miniature Linear Motion Series: PQ12, PQ12f, Firgelli Technologies Ltd, www.firgelli.com
- [84] Non-captive Stepper Motor, 25NCLA-B01, Motion Control Products Ltd, www.motioncontrolproducts.co.uk
- [85] www.gasketing.net/rubber.htm
- [86] Xiang, G.; Pownuk, A.; Kosheleva, O. & Starks, S., "Von Mises Failure Criterion in Mechanics of Materials: How to Efficiently Use it Under Interval and Fuzzy Uncertainty", *Fuzzy Information Processing Society, 2007. NAFIPS'07. Annual Meeting of the North American*, 2007, 570-575

- [87] Vicat Softening Temperature, DIN 53460
- [88] Texas Instruments. L293, L293D: Quadruple Half-H Drivers, rev B, November 2004.
- [89] Gumstix Connex 400xm, GS400J-XM, Gumstix Inc.
- [90] Robostix, BRD00019, Gumstix Inc.
- [91] Easy-Radio ER400TRS, 433-434MHz FM Transceiver, LPRS Ltd, March 2003
- [92] "Specification of the Bluetooth System: Experience More", BLUETOOTH SPECIFICATION Version 4.0
- [93] IEEE Std 802.11g - 2003.
- [94] www.intel.com/technology/comms/wusb/index.htm
- [95] Theodoridis, S.; Koutroumbas, K. & Smith, R., "Pattern Recognition", 3rd edition, Academic Press
- [96] Jain, A.; Duin, R. & Mao, J., "Statistical pattern recognition: A review", IEEE Transactions on pattern analysis and machine intelligence, 2000, 22, 4-37
- [97] Kittler, J., "Combining classifiers: A theoretical framework", Pattern Analysis & Applications, Springer, 1998, 1, 18-27
- [98] Schwarzer, G.; Vach, W. & Schumacher, M. "On the misuses of artificial neural networks for prognostic and diagnostic classification in oncology Statistics in medicine, John Wiley & Sons, 2000, 19, 541-561
- [99] Pierce, S.; Worden, K. & Manson, G. "Evaluation of neural network performance and generalisation using thresholding functions", Neural Computing & Applications, Springer, 2007, 16, 109-124
- [100] Berkhin, P., "A survey of clustering data mining techniques", Grouping Multidimensional Data, Springer, 2006, 25-71
- [101] MacKay, D. "Information theory, inference, and learning algorithms", Cambridge Univ Pr, pg 347, 2003
- [102] Bishop, C. "Neural networks for pattern recognition", Oxford University Press, USA, 1995 , pg23, 375, 387
- [103] Kuncheva, L., "Combining pattern classifiers: methods and algorithms", Wiley-Interscience, 2004
- [104] Kittler, J., H. M. D. R. & Matas, J., "On combining classifiers", IEEE Transactions on Pattern Analysis and Machine Intelligence, 1998, 20, 226-239
- [105] Nabney, I., "NETLAB: algorithms for pattern recognition", Springer Verlag, 2002, pg 105
- [106] Cover, T. & Hart, P. "Nearest neighbor pattern classification", IEEE Transactions on Information Theory, 1967, 13, 21-27

- [107] Uesaka, M.; Nakanishi, T.; Miya, K.; Komatsu, H.; Aoki, K. & Kasai, K., “Micro eddy current testing by micro magnetic sensor array”, IEEE Transactions on Magnetics, 1995, 31, 870-876
- [108] Chen, X. & Ding, T., “Flexible eddy current sensor array for proximity sensing”, Sensors and Actuators A: Physical, Elsevier, 2007, 135, 126-130
- [109] Uesaka, M.; Hakuta, K.; Miya, K.; Aoki, K. & Takahashi, A., “Eddy-current testing by flexible microloop magnetic sensor array”, IEEE Transactions on Magnetics, 1998, 34, 2287-2297
- [110] Grimberg, R.; Udpa, L.; Savin, A.; Steigmann, R.; Palihovici, V. & Udpa, S., “2D eddy current sensor array”, NDT & E International, Elsevier, 2006, 39, 264-271
- [111] Muck, M.; Korn, M.; Welzel, C.; Grawunder, S. & Scholz, F., “Nondestructive evaluation of various materials using a SQUID-based eddy-current system”, IEEE Transactions on Applied Superconductivity, 2005, 15, 733-736
- [112] Kreuzbruck M.v., Allweins K., Ruhl T., Muck M., Heiden C., Krause H.-J., Hohmann R., Defect detection and classification using a SQUID based multiple frequency eddy current NDE system, IEEE Transactions on Applied Superconductivity, 11(1):1032--1037, 2001.
- [113] Halmshaw R., "Non-Destructive Testing", 2nd edition, 1991
- [114] Thomson, J.M., An Eddy Current Array Instrument for Offshore Application, PhD Thesis, University of Strathclyde, 1991,
- [115] www.comsol.com/products/acdc
- [116] Vectorfields (Opera 2D and Opera 3D), Cobham Plc, www.cobham.com/about-cobham/avionics-and-surveillance/about-us/technical-services/kidlington/products.aspx.
- [117] Quickfield, Tera Analysis Ltd, www.quickfield.com
- [118] Maxwell 2D, ANSYS, Inc., www.ansoft.com/products/em/maxwell
- [119] EXTENDE CIVA, www-civa.cea.fr
- [120] LMH6718, Dual, High Output, Programmable Gain Buffer, National Semiconductor Corp., January 2003
- [121] Microchip. dsPIC DSC High-Performance 16-Bit Digital Signal Controllers, rev 1 edition, 2005.
- [122] AD5930, Programmable Frequency Sweep and Output Burst Waveform Generator, Analogue Devices Inc., Rev 0.
- [123] ADG725, 16-/32 Channel, Serially Controlled, 4 Ohm, 1.8V to 5.5V, $\pm 2.5V$, Analog Multiplexers, Analogue Devices Inc., Rev. A
- [124] ASSR-1420, General Purpose, Form A, Solid State Relay (Photo MOSFET), Avago Technologies, July 2007

- [125] Velez, J.; Linhares, F.; Borges, A.; Xiong, H.; Saunders, N. & Beck, M., “An overview of electromagnetic inductance tomography: description of three different systems”, *Meas. Sci. Technol*, 1996, 7, 261-271
- [126] Tamburrino, A. & Rubinacci, G., “Fast methods for quantitative eddy-current tomography of conductive materials”, *IEEE Transactions on Magnetics*, 2006, 42, 2017-2028
- [127] Zorgati, R.; Duchene, B.; Lesselier, D.; Pons, F.; et Recherches, D. & EDF, S., “Eddy current testing of anomalies in conductive materials. I. Qualitative imaging via diffraction tomography techniques”, *IEEE Transactions on Magnetics*, 1991, 27, 4416-4437
- [128] Fukutomi, H.; Huang, H.; Takagi, T. & Tani, J., “Identification of crack depths from eddy current testing signal”, *IEEE Transactions on Magnetics*, 1998, 34, 2893-2896
- [129] Gros, X., “NDT data fusion”, Arnold Publishers, 1997
- [130] Blum, R. & Liu, Z. *Multi-sensor image fusion and its applications* CRC, 2006
- [131] Gui, C. & Fox, M. “Level set evolution without re-initialization: A new variational formulation”, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005. CVPR 2005, 2005, 1
- [132] Shih, F. & Zhang, K. *Locating object contours in complex background using improved snakes* *Computer Vision and Image Understanding*, Elsevier, 2007, 105, 93-98
- [133] Wang, L.; Li, C.; Sun, Q.; Xia, D. & Kao, C. *Brain MR image segmentation using local and global intensity fitting active contours/surfaces* *Medical Image Computing and Computer-Assisted Intervention--MICCAI 2008*, Springer, , 384-392
- [134] Wang, S.; Lim, K.; Khoo, B. & Wang, M. *An extended level set method for shape and topology optimization* *Journal of Computational Physics*, Elsevier, 2007, 221, 395-421
- [135] Chan, T. & Vese, L. *Active contours without edges* *IEEE Transactions on image processing*, Citeseer, 2001, 10, 266-277
- [136] Li, C.; Huang, R.; Ding, Z.; Gatenby, C.; Metaxas, D. & Gore, J. *A variational level set approach to segmentation and bias correction of images with intensity inhomogeneity* *Medical Image Computing and Computer-Assisted Intervention--MICCAI 2008*, Springer, 2008, 1083-1091
- [137] Kass, M.; Witkin, A. & Terzopoulos, D. *Snakes: Active contour models* *International journal of computer vision*, Springer, 1988, 1, 321-331
- [138] Osher, S. & Sethian, J. *Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations* *Journal of computational physics*, Elsevier, 1988, 79, 12-49

- [139] Chan, T. & Vese, L. Active contours without edges IEEE Transactions on image processing, Citeseer, 2001, 10, 266-277
- [140] Omniwheel, www.kornylak.com/wheels/omniwheel.html
- [141] Leica Laser tracker, www.leica-geosystems.co.uk/en/Leica-Absolute-Tracker-AT901_69047.htm