

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

School of Geography

**Sub-pixel mapping of rural land cover features from fine
spatial resolution remotely sensed imagery**

By

Matthew W. Thornton

Doctor of Philosophy

September 2006

UNIVERSITY OF SOUTHAMPTON

Abstract

FACULTY OF ENGINEERING, SCIENCES AND MATHEMATICS
SCHOOL OF GEOGRAPHY

Doctor of Philosophy

SUB-PIXEL MAPPING OF RURAL LAND COVER FEATURES FROM FINE
SPATIAL RESOLUTION REMOTELY SENSED IMAGERY

By Matthew W. Thornton

Mapping rural land cover features, such as trees and hedgerows, for ecological applications is a desirable component of the creation of cartographic maps by the Ordnance Survey and inclusion in geographic database systems such as OS Mastermap®. Based on the phenomenon of spatial dependence, sub-pixel mapping can provide increased mapping accuracy of such features. A simple pixel swapping algorithm for super-resolution sub-pixel mapping was applied to predicted class proportions derived from a soft classification of simulated and real fine spatial resolution remotely sensed imagery. Input proportions were super-resolved into sub-pixels using a specified zoom factor. Sub-pixels were then iteratively swapped until the spatial correlation between sub-pixels for the entire image was maximised. The standard pixel swapping algorithm was developed to increase the accuracy with which rural land cover features were predicted. Firstly, the algorithm was modified to increase the likelihood of predicting linear features, such as hedgerows, on the basis of measured anisotropy within class proportions. Secondly, an image fusion component was integrated to enable the use of multiple datasets such as panchromatic imagery, to refine the prediction of the geometric characteristics of the predicted features. The new pixel swapping technique, increased the accuracy with which rural land cover features compared with the standard technique and substantially increased the utility of such land cover maps compared with conventional mapping techniques, such as classification.

1. Introduction	1
2. Literature Review	2
3. Methodology	3
4. Results	4
5. Discussion	5
6. Conclusion	6
7. References	7
8. Appendix	8

Contents

Abstract	ii
Contents	iii
List of Figures	vii
List of Tables	x
List of Equations	xii
Declaration of Authorship	xiii
Acknowledgements	xiv
Dedication	xv
1. Introduction	2
1.1 Land Cover	2
1.1.1 The importance of land cover mapping.....	3
1.1.2 The availability of land cover data.....	4
1.2 Principles of Remote Sensing.....	7
1.2.1 Physical basis.....	8
1.2.2 Images and Pixels	9
1.2.3 Radiometric Resolution.....	10
1.2.4 Spatial resolution	10
1.2.5 Spatial variation	11
1.2.6 Spatial dependence.....	12
1.3 Research objectives.....	13
1.3.1 Rationale.....	14
1.3.2 Aims and Objectives	15
1.3.3 Thesis Outline.....	16
1.3.4 Publications	16
1.4 Conclusion.....	17
2. Background and Methods	20
2.1 Introduction	20
2.2 Classification	21
2.2.1 Hard classification.....	21
2.2.2 Soft classification.....	23
2.2.2.1 Fuzzy <i>c</i> -means	24
2.2.2.2 Linear mixture model.....	25
2.3 Super-resolution techniques for land cover mapping	25
2.3.1 Geostatistical theory.....	26
2.3.2 Sub-pixel mapping.....	26
2.3.3 Sub-pixel mapping: pixel swapping.....	29
2.3.3.1 Simulated Annealing.....	31
2.3.3.2 Swap criteria	31
2.4 Identifying land cover features.....	32
2.4.1 Edge Detection.....	33
2.4.2 The Hough transform	33
2.4.3 Active contour models: snakes	34
2.4.4 ‘Linearisation’	36
2.5 Image fusion	37
2.5.1 Super-resolution image fusion.....	41
2.6 Accuracy assessment.....	42

2.6.1	Confusion matrix	42
2.6.2	Root mean square error (RMSE)	45
2.6.3	Pearson's product-moment correlation coefficient.....	46
2.6.4	Difference imaging	46
2.7	Conclusion.....	47
3.	Study Areas and Data.....	50
3.1	Introduction	50
3.2	Field site selection 'decision support system'	51
3.3	DSS Step 1: Image components	53
3.3.1	Field site characteristics	53
3.3.2	Feature characteristics.....	54
3.3.2.1	Width (and size).....	54
3.3.2.2	Curvature	54
3.3.2.3	Embedded objects	55
3.4	DSS Step 2: Remotely sensed imagery.....	55
3.4.1	Study area	56
3.4.2	Real imagery.....	57
3.4.2.1	Aerial photography.....	57
3.4.2.2	Satellite sensor imagery.....	59
3.4.3	Simulated imagery	60
3.5	Step 3: Field site selection.....	61
3.5.1	Simulated imagery (1).....	61
3.5.2	Simulated imagery (2).....	63
3.5.3	Real imagery.....	64
3.5.3.1	Field Data	67
3.6	Summary	68
4.	Evaluating the super-resolution pixel swapping technique	70
4.1	Introduction	70
4.2	Simulated Imagery.....	71
4.2.1	Soft proportions	71
4.2.2	Pixel swapping.....	71
4.2.3	Results	72
4.2.3.1	Simulated image 1: Diamond	72
4.2.3.2	Simulated image 2: Line.....	78
4.2.3.3	Simulated image 3: Complex scene	85
4.2.4	Discussion of results	94
4.3	Real imagery.....	96
4.3.1	Soft classification.....	96
4.3.2	Pixel swapping.....	97
4.3.3	Results	98
4.3.3.1	Field site 1 (Site ref: A).....	98
4.3.3.2	Field site 2 (Site ref: B)	102
4.3.3.3	Field site 3 (Site ref: C)	106
4.3.4	Discussion of results	110
4.4	Mathematical morphology	112
4.4.1	Field Site 1.....	113
4.4.2	Field site 2	115
4.4.3	Field site 3	117
4.4.4	Discussion of results after mathematical morphology	119
4.5	Discussion of chapter results.....	120

4.6	Conclusion.....	122
5.	Linearisation.....	124
5.1	Introduction	124
5.2	Linearisation.....	125
5.2.1	Introduction	125
5.2.2	Linearising the pixel swapping technique.....	126
5.2.3	Multiple directions.....	129
5.3	Results	131
5.3.1	Field site (spatial resolution 2.5 m).....	131
5.3.2	Field site (spatial resolution 5 m).....	139
5.3.3	Discussion of results	147
5.4	Advanced linearisation – 12 directions.....	148
5.5	Discussion of chapter.....	152
5.6	Conclusion.....	154
6.	Image fusion.....	156
6.1	Introduction	156
6.2	Image Fusion	156
6.2.1	Methodology.....	158
6.2.2	Defining the energy function.....	160
6.3	Initialisation.....	160
6.3.1	Simulating panchromatic imagery	161
6.4	Results	161
6.4.1	Field site (spatial resolution 2.5 m).....	161
6.4.2	Field site (spatial resolution 5 m).....	164
6.5	Discussion of results	166
6.6	Discussion of chapter.....	168
6.6.1	Parameterisation.....	169
6.6.1.1	Weightings.....	169
6.6.1.2	Spatial resolution.....	170
6.6.1.3	Image coregistration.....	171
6.7	Conclusion.....	173
7.	Discussion and further research	175
7.1	Introduction	175
7.1.1	Overall performance.....	175
7.1.2	Assumptions and considerations.....	176
7.1.2.1	Accuracy assessments	177
7.1.3	Limitations of the algorithm.....	178
7.2	Discussion of the model.....	179
7.2.1	Zoom factor	180
7.2.2	Exponential function window parameters.....	183
7.2.3	Anisotropy ratio	185
7.2.4	Annealing schedule	185
7.3	Future research	186
7.3.1	Prediction of direction.....	186
7.3.2	Soft classification, noise, geometric error and real imagery	186
7.4	Conclusion.....	188
8.	Conclusions.....	190
8.1	Introduction	190
8.2	Summary	190
8.2.1	Background.....	190

8.2.2	Research aims and objectives	191
8.2.3	Development and analysis	191
8.2.4	Results	192
8.3	Implications	195
8.4	Conclusion	195
	References.....	197
	Appendix 1.....	208
	Appendix 2.....	211

The first stage of the research was to identify the research objectives and aims of the study. This was done by reviewing the literature on the topic and identifying the gaps in the current knowledge. The research objectives were then formulated based on these gaps.

The second stage of the research was to develop the research design and methodology. This involved identifying the research questions and hypotheses, and determining the appropriate methods for data collection and analysis. The research design was then developed based on these considerations.

The third stage of the research was to collect and analyze the data. This involved conducting the research and collecting the data, and then analyzing the data to identify the results. The results were then presented in a clear and concise manner.

The fourth stage of the research was to discuss the implications of the findings and to draw conclusions. This involved interpreting the results in the context of the research objectives and aims, and discussing the implications of the findings for practice and theory.

The fifth stage of the research was to write the final report. This involved summarizing the findings and conclusions, and presenting them in a clear and concise manner. The final report was then submitted for publication.

List of Figures

Figure 2.1: The image fusion model (Pohl, 1999).....	38
Figure 2.2: The three levels of image fusion (Pohl and Van Genderen, 1998).....	39
Figure 2.3: The discrete wavelet frame transform.....	40
Figure 3.1: Flow diagram illustrating the decision support system. Solid lines represent direction of decision making. Dashed lines illustrate questions dependent on other areas of the system.....	52
Figure 3.2: Sample colour aerial photograph of the study area, spatial resolution 25 cm, Christchurch study area.....	58
Figure 3.3: Simulated image 1, ‘diamond’, (dimensions: 250 pixels by 250 pixels, number of classes: 2).....	62
Figure 3.4: Simulated image 2, ‘line’, (dimensions: 250 pixels by 250 pixels, number of classes: 2).....	62
Figure 3.5: Simulated image 3, ‘complex scene’, (dimensions: 250 pixels by 250 pixels, number of classes: 5).....	63
Figure 3.6: Real imagery, Site A, (dimensions: 35 pixels by 35 pixels, number of classes: 3).....	65
Figure 3.7: Real imagery, Site B, (dimensions: 46 pixels by 169 pixels, number of classes: 4).....	66
Figure 3.8: Real imagery, Site C, (dimensions: 73 pixels by 73 pixels, number of classes: 4).....	66
Figure 4.1: Soft proportions for feature 1, ‘Diamond’, (a) feature, (b) background ...	73
Figure 4.2: Hard classification, diamond feature.....	74
Figure 4.3: Super-resolved output, feature 1, (a) initial allocation, (b) super-resolved output.....	75
Figure 4.4: Subsection of simulated image 1 used in confusion matrix.....	77
Figure 4.5: Soft proportions for feature 2, ‘Line’, (a) feature, (b) background.....	79
Figure 4.6: Hard classification, line feature.....	80
Figure 4.7: Super-resolved output, feature 2, (a) initial allocation, (b) super-resolved output.....	81
Figure 4.8: Subsection of feature 2 used in confusion matrix.....	83
Figure 4.9: Class proportion images for feature 3: (a) “field 1”, (b) “field 2”, (c) “field 3”, (d) “field 4”, (e) “hedgerow”. Units: pixels.....	86
Figure 4.10: Hard classification, complex scene.....	87
Figure 4.11: Super-resolved output, feature 3, (a) initial allocation, (b) super-resolved output.....	89
Figure 4.12: Difference images, (a) Diamond shape, (b) Soft classification, (c) Pixel swapping.....	93
Figure 4.13: Proportion images for field site 1: (a) “hedgerow”, (b) “cereal”, (c) “woodland”. Units: percentage of a pixel’s area.....	99
Figure 4.14: Super-resolved output, field site 1.....	101
Figure 4.15: Proportion images for field site 2: (a) “woodland”, (b) “hedgerow”, (c) “non-ripe cereal”, (d) “ripe cereal”. Units: percentage of a pixel’s area.....	103
Figure 4.16: Super-resolved output, field site 2.....	105
Figure 4.17: Proportion images for field site 3: (a) “hedgerow”, (b) “woodland” (i.e. individual trees), (c) “ripe cereal”, (d) “non-ripe cereal”. Units: percentage of a pixel’s area.....	107

Figure 4.18: Super-resolved output, field site 3.....	109
Figure 4.19: Super-resolved output with mathematical morphology applied, field site 1.....	114
Figure 4.20: Super-resolved output with mathematical morphology applied, field site 2.....	116
Figure 4.21: Super-resolved output with mathematical morphology applied, field site 3.....	118
Figure 5.1: Predicting the direction of a linear feature from soft proportions. A constraint is imposed that a linear feature must pass through the central pixel of a moving window. In this example, a linear feature is predicted at 45° from north.	126
Figure 5.2: Exponential distance decay models, (a) Isotropic (b) Anisotropic. Legend: White = 1, Black = 0.....	127
Figure 5.3: Directions, (a) The 4 directions in a 3 by 3 window, (b) The 8 directions in a 5 by 5 window.....	129
Figure 5.4: Hard classification, SR 2.5 m.....	131
Figure 5.5: Field site, spatial resolution 2.5 m, (a) Input imagery, (b) Random allocation.....	133
Figure 5.6: Direction maps (a) Anisotropic, 3 by 3 window, 2.5 m SR, (b) Anisotropic, 5 by 5 window, 2.5 m SR. Legend (A _{ij}): Cyan – Low weight, Red – High weight.....	134
Figure 5.7: Pixel swapping output, 2.5 m SR (a) Standard, 2.5 m (b) Linearised, 3 by 3 window, (c) ‘Linearised’, 5 by 5 window.....	135
Figure 5.8: Difference maps: spatial resolution 2.5 m, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window. Legend: black – correct allocation, white – error of omission, grey - error of commission.....	137
Figure 5.9: Hard classification, 5 m SR.....	139
Figure 5.10: Field site, spatial resolution 5 m, (a) Input imagery (b) Random allocation.....	141
Figure 5.11: Direction maps (a) Anisotropic, 3 by 3 window, 5 m SR, (b) Anisotropic, 5 by 5 window, 5 m SR. Legend (A _{ij}): Cyan – Low weight, Red – High weight.....	142
Figure 5.12: Pixel swapping output, spatial resolution 5 m (a) ‘Linearised’, 3 by 3 window, (b) ‘Linearised’, 5 by 5 window, (c) ‘Linearised’, 5 by 5 window.....	143
Figure 5.13: Difference maps: spatial resolution 5 m, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window. Legend: black – correct allocation, white – error of omission, grey - error of commission.....	145
Figure 5.14: Super-resolved output. Linearised pixel swapping, 3 by 3 window, 12 directions (a) spatial resolution 2.5 m (b) spatial resolution 5 m.....	150
Figure 5.15: Difference maps. Linearised pixel swapping, 3 by 3 window, 12 directions (a) spatial resolution 2.5 m (b) spatial resolution 5 m. Legend: black – correct allocation, white – error of omission, grey - error of commission.....	151
Figure 6.1: Preliminary stage of fusion model.....	158
Figure 6.2: The fusion model.....	159
Figure 6.3: Simulated panchromatic image.....	161
Figure 6.4: Fusion output, (a) Input imagery, 2.5 m SR, (b) Random allocation, (c) Super-resolved output.....	162
Figure 6.5: Difference image, super-resolved output.....	163

Figure 6.6: Fusion output, (a) Input imagery, 2.5 m SR, (b) Random allocation, (c) Super-resolved output.....165

Figure 6.7: Difference image, super-resolved output.....166

Figure 6.8: Comparison of super-resolution output with target image: (a) Target image, (b) Linearised pixel swapping output, (c) Linearised pixel swapping output with image fusion167

Figure 6.9: Effect of weights on super-resolution output, (a) $k_1 = 0.5, k_2 = 0.5$, (b) $k_1 = 0.6, k_2 = 0.4$, (c) $k_1 = 0.75, k_2 = 0.25$, (d) $k_1 = 0.9, k_2 = 0.1$, (e) $k_1 = 1, k_2 = 0, ..170$

Figure 6.10: Effect of inaccurate coregistration of input imagery, super-resolved output (a) Perfect registration, (b) 1 pixel displacement (1.25 m), (c) 2 pixel displacement (2.5 m), (d) 3 pixel displacement (3.75 m), (e) 4 pixel displacement (5 m)172

Figure 7.1: Effect of zoom factor on super-resolved output (linearised pixel swapping, 3 by 3 window, 2.5 m SR)182

Figure 7.2: Non-linear parameter of the exponential distance decay model, linearised pixel swapping with fusion: (a) low rate of decay, (b) increased rate of decay, (c) highest rate of decay184

Figure 7.3: Size of exponential function window (a) 5 by 5, (b) 9 by 9, (c) 13 by 13. Scale: sub-pixels.....184

List of Tables

Table 2.1: Example confusion matrix involving four classes	42
Table 3.1: Specifications of aerial photography	57
Table 3.2: Quickbird multispectral and panchromatic imagery characteristics	59
Table 3.3: Simulated imagery 2.....	64
Table 4.1: Accuracy assessment, diamond feature	73
Table 4.2: Confusion matrix, hard classified diamond feature.....	74
Table 4.3: Confusion matrix, simulated image 1, (a) initial allocation, (b) super-resolved output	76
Table 4.4: Confusion matrix, subsection of feature 1, (a) hard classification, (b), soft classification, (c), super-resolution	77
Table 4.5: Significance test, Hard classification & Pixel swapping	78
Table 4.6: Accuracy assessment, simulated image 2	79
Table 4.7: Confusion matrix, hard classied line feature.....	80
Table 4.8: Confusion matrices, feature 2, (a) Initial allocation, (b) Super-resolved output.....	82
Table 4.9: Confusion matrix, subsection of feature 2, (a) Hard classification, (b) Soft classification, (c) Pixel swapping.....	83
Table 4.10: Significance test, hard classification & pixel swapping	85
Table 4.11: Accuracy assessment, feature 3	87
Table 4.12: Confusion matrix, hard classification, complex scene	88
Table 4.13: Confusion matrices, feature 3, (a) Initial allocation, (b) Super-resolved output	90
Table 4.14: Confusion matrices, subsection of complex scene	91
Table 4.15: Accuracy assessment, field site 1	100
Table 4.16: Confusion matrix, field site 1	102
Table 4.17: Accuracy assessment, field site 2	104
Table 4.18: Confusion matrix, field site 2.....	106
Table 4.19: Accuracy assessment, field site 3	108
Table 4.20: Confusion matrix, field site 3.....	110
Table 4.21: Confusion matrix, super-resolved output with mathematical morphology applied.....	115
Table 4.22: Confusion matrix, super-resolved output with mathematical morphology applied.....	117
Table 4.23: Confusion matrix, super-resolved output with mathematical morphology applied, field site 3	119
Table 5.1: Confusion matrix, hard classification, 2.5 m SR.....	132
Table 5.2: Super-resolved output confusion matrices, 2.5 m spatial resolution, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window.....	136
Table 5.3: Significance tests, (a) PS & LPS, 3 by 3 window, (b) PS & LPS, 5 by 5 window	138
Table 5.4: Confusion matrix, hard classification, 5 m SR	140
Table 5.5: Super-resolved output confusion matrices, 5 m spatial resolution, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window	144

Table 5.6: Significance tests, (a) PS & LPS, 3 by 3 window, (b) PS & LPS, 5 by 5 window	146
Table 5.7: 12 possible directions in a 3 by 3 window, Scale: pixels	149
Table 5.8: Confusion matrices, linearised pixel swapping, 12 directions, spatial resolution: 2.5 m.....	150
Table 5.9: Confusion matrices, linearised pixel swapping, 12 directions, spatial resolution: 5 m.....	151
Table 6.1: Confusion matrix, super-resolved output (2.5 m input SR).....	163
Table 6.2: Significance test, linearised pixel swapping and linearised pixel swapping with fusion	164
Table 6.3: Confusion matrix, super-resolved output (5 m input SR).....	165
Table 7.1: Example processing time for a 3-class classified image, Size: 100 pixels by 100 pixels, zoom factor: 5.....	176
Table 8.1: Summary of results, spatial resolution 2.5 m.....	193

List of Equations

Equation 1: Pixel swapping, calculating A_i	10
Equation 2: Exponential distance decay function	11
Equation 3: McNemar's significance test	12
Equation 4: Root Mean Square Error (RMSE)	13
Equation 5: Pearson's product-moment correlation coefficient	14
Equation 6: Anisotropy correction	15
Equation 7: Image fusion, stage 1	16
Equation 8: Image fusion, stage 2	17
Equation 9: Image fusion, energy function	18

Acknowledgements

I would like to take this opportunity to express my thanks to everyone that has helped me during the course of this PhD. First and foremost, to my supervisors. Professor Peter Atkinson has been invaluable: as a source of immense knowledge; fantastic support and guidance; faith and criticism, and all (usually) with a smile. Without Pete, none of this would have been possible. Dr. David Holland of the Ordnance Survey has been incredibly helpful as a further source of input, both in terms of supervision of the work and as a critic, often with a new and unconsidered viewpoint.

I am extremely grateful to both the School of Geography and the Ordnance Survey for co-funding this research. The administrative staff in the School of Geography were very supportive and efficient with the day-to-day things and the computing staff, in particular Dr. Charlie Kerr, provided top quality computing resources which enabled the processing and analysis of data. The Ordnance Survey provided me with some wonderful datasets to work with, as well as offering a wealth of expertise in various parts of the work.

I am also very grateful to the people who assisted in my fieldwork: Ilse Steyl, Jana Fried, Nick Hamm, Karen Anderson, and Peter Gething. Peter also did a remarkable job of putting up with me in the Postgraduate offices. I would also like to thank Hampshire County Council for the use of a differential GPS system during fieldwork, as well as each of the landowners who kindly granted me access to their land.

I would also like to acknowledge some people who provided me with technical assistance: Robert Maunder, for his incredible C++ knowledge, Professor Mark Nixon and Dr. Matt Wilson.

Finally, I am indebted to my friends and family: to my parents for their faith and confidence in me; to my brother for always being sceptical, and to all my friends for putting up with me throughout.

List of Abbreviations

AA	Automobile Association
ANN	artificial neural network
ASCII	American standard code for information interchange
AVHRR	advanced very high resolution radiometer
AVIRIS	airborne visible-infrared imaging spectrometer
CASI	compact airborne spectrographic imager
CORINE	Co-Ordinating Information on the European Environment
DIS	data and information system
DSS	decision support system
DVD	digital versatile disc
DWFT	discrete wavelet frame transform
ETM	enhanced thematic mapper
FCM	fuzzy <i>c</i> -means
FTP	file transfer protocol
GCP	ground control point
GIS	geographic information system
GPS	global positioning system
GRE	ground resolution element
HNN	Hopfield neural network
IDW	inverse distance weighting
IFOV	instantaneous-field-of-view
IGBP	International Geosphere Biosphere Programme's
IHS	intensity-hue-saturation
LCMGB	Land Cover Map of Great Britain
LiDAR	light detection and ranging
ML	maximum likelihood
MLP	maximum likelihood probabilities
MM	mixture model
MSS	multispectral scanning system
NIR	near infrared
NLP	non-linear parameter
NOAA	National Oceanographic and Atmospheric Administration
OS	Ordnance Survey
pan	panchromatic
PCA	principal components analysis
PCS	principal component substitution
pdf	probability density function
PSF	point spread function
RAC	Royal Automobile Club
SAR	synthetic aperture radar
SPOT	Système Pour L'Observation de la Terre
SSSI	Site of Special Scientific Interest
TM	thematic mapper

Chapter 1

Introduction

with individuals who are familiar with the principles of the research process

and Cover

of the research process, including the role of the researcher and the role of the research participant. The research process is a complex and multi-stage process that involves the identification of a research problem, the formulation of research objectives, the design of a research study, the collection and analysis of data, and the interpretation of the results. The role of the researcher is to guide the research process and to ensure that the research is conducted in a rigorous and ethical manner. The role of the research participant is to provide the data that is needed to answer the research question. The research process is a collaborative effort between the researcher and the research participant, and it is essential that both parties understand their respective roles and responsibilities.

Chapter 1

1. Introduction

In this thesis, a super-resolution pixel swapping technique for mapping land cover is presented, with a specific emphasis on the mapping of small rural features, such as hedgerows and trees, from fine spatial resolution remotely sensed imagery.

In this preliminary chapter, land cover is defined and the value and importance of its mapping is described and discussed. The availability of land cover information, such as remotely sensed imagery, is then discussed, including an introduction to the key principles of remote sensing and the process of classifying remotely sensed imagery for the creation of land cover maps. A common practical limitation of classifying remotely sensed imagery (mixed pixels) is then introduced and the use of super-resolution techniques as a tool to minimise the effect of mixed pixels and enable more accurate mapping of land cover is introduced. The chapter then presents a rationale for the research and closes with a statement of the objectives of the research presented in this thesis.

1.1 Land Cover

Land cover is an intrinsic component of many environmental systems and represents a key variable in scientific research. Land cover is defined as ‘the description of the physical nature of the land surface, for example, vegetation, buildings, water or bare soil’ (Wyatt *et al.*, 1993). Land cover plays an important part in systems such as the water cycle (Congalton *et al.*, 1998), where a heavily vegetated area such as a forest will, for example, affect the speed at which water is returned to its source (in contrast to a recently deforested area) and the climate cycle, for example, the role of different land cover objects within the creation of regionalized climate (Kerr and Ostrovsky, 2003).

The term land cover should not be confused with the term 'land use'. Land use, as defined by (Campbell, 1996), is "the use of land by humans, usually with emphasis on the functional role of land in economic activities". In contrast, land cover could be considered to be the physical cover of the land surface. Land use is not directly observable and although it might be possible to infer land use from independent land cover objects (for example, a large area of trees (land cover objects) might suggest a land use of forest) there will be significant uncertainty in this prediction (Lambin *et al.*, 2001). For example, each of the four examples of land cover (vegetation, buildings, water, bare soil) in the definition by Wyatt *et al.* (1993), above, may represent many different contrasting land uses. 'Vegetation' land cover could represent either agricultural land use in the form of cereal fields or alternatively domestic land use in the form of a sports playing field. 'Buildings' land cover might suggest a residential housing estate or a commercial business park. 'Water' land cover might suggest a human-made reservoir for recreational purposes or an agricultural irrigation system. 'Bare soil' land cover might suggest agricultural land use or a commercial quarry. Accordingly, combining land cover types contextually, by considering the characteristics of each land cover individually and in aggregate may allow accurate prediction of land use.

1.1.1 The importance of land cover mapping

Information on land cover is used in many different applications. Inventories of land cover data, such as the International Geosphere Biosphere Programme's (IGBP) Data and Information System (DIS) (Belward *et al.*, 1999), are used to determine and develop land management procedures. Repeated land cover mapping through the creation of temporal sequences facilitates change detection and monitoring (Stefanov *et al.*, 2001). Indeed, the use of remote sensing data is of great utility to global leaders involved in the Kyoto protocol, a worldwide discussion on global climate change (Rosenqvist *et al.*, 2003). Advances in land cover planning can be effected by the study of the interaction of land cover types for the development and implementation of land use policies and strategies around human existence.

“Three decades of research have given us a far more nuanced view of Earth’s “resources”. Our concerns now center on understanding how the land surface helps govern the primary cycles of energy, water, and carbon, and how the variability of those cycles may be affected by our own activities.”

(Editorial, 2001)

To assist our understanding of the Earth’s surface, mapping land cover and in particular, land cover objects, is important (Foody, 2002b). For example, the ability to locate individual buildings within a scene would be useful for town planning, or in the context of this research, the ability to locate and extract small rural features would allow for their accurate cartographic mapping. The creation of land cover maps dates as far back as 1086 with the creation of the Domesday Book; a collation of parish records of land cover to national scale. Since the 1890s, the Ordnance Survey has mapped the whole of Great Britain creating maps at a variety of scales. Greater and more accurate information on land cover objects would be of great utility to the Ordnance Survey, for inclusion within modern-day maps of the UK.

1.1.2 The availability of land cover data

Due largely to advances in technology, the availability of land cover data in recent years has increased. The availability of land cover data can be described in two ways: 1) what data have been acquired and 2) how are these data dispersed.

Remote sensing is a well-used source of land cover information, most commonly acquired in the format of a remotely sensed image, from ground-based, aerial-based or satellite-based survey (Cihlar, 2000). Ground-based methods of data acquisition, such as field survey, provide potentially accurate land cover information over a small area. Such methods are often inappropriate for use in particular types of application due to the large acquisition times, the small area covered and the expense of such data relative to that provided by other sources of remotely sensed data (Veitch *et al.*, 1995). Field survey is frequently used as a source of remotely sensed data. For example, airborne remote sensing provides data, such as aerial photography,

considerably faster than field survey and these data can be provided at a fine spatial resolution in a variety of formats, relatively inexpensively (Congalton *et al.*, 1998). Additionally, many commercial firms are able to provide aerial photography “on demand”, i.e., they are able to purposively collect data in a specified area as required by the needs of the application (Mullerova *et al.*, 2005).

There are many examples of cartographic mapping that involve the integration of multiple sources of land cover data. For example, the Land Cover Map of Great Britain (LCMGB: <http://www.ceh.ac.uk>) is a digital dataset that classifies the whole of Great Britain into 25 land cover types (Fuller *et al.*, 1994). It was generated from satellite sensor imagery, mainly Landsat Thematic Mapper (TM) imagery, and is created approximately every 10 years. These data are also being integrated into CORINE (Co-Ordinating Information on the European Environment), a mapping initiative set up by the European Environment Agency to provide a Europe-wide map of land cover in 44 classes (<http://science.ceh.ac.uk/subsites/CORINE>). Such maps and mapping initiatives rely heavily on the plentiful availability of accurate land cover information. In recent years, there has been a great deal of investment in the use of satellite-based remote sensing (Loveland and DeFries, 2004). Satellite-based remotely sensed imagery is often preferable to aerial imagery since the imagery acquired from the satellite sensor can be provided at local to global scales, with greater frequency of repeat coverage enabling, for example, the assembly of temporal datasets covering large areas of the Earth’s surface (Gregoire *et al.*, 2003).

Satellite sensors launched in the last 10 years are now providing imagery at fine spatial resolutions, such as Système Pour L’Observation de la Terre (SPOT), (5 m), Ikonos multispectral (4 m), Quickbird multispectral (2.6 m), Ikonos panchromatic (1 m) and Quickbird panchromatic (0.6 m). Quickbird multispectral satellite sensor imagery, for example, is available at a spatial resolution of 0.6 m, which uses Quickbird panchromatic imagery to increase the spatial resolution of the standard multispectral imagery via a pansharpening fusion technique (Vijayaraj *et al.*, 2006). Remotely sensed data are available in an ever-burgeoning array of formats, from panchromatic, multispectral and hyperspectral image formats to imagery acquired using laser scanning, such as Light Detection and Ranging (LiDAR), and radar scanning, such as Synthetic Aperture Radar (SAR).

There have been many advances in technology, both in satellite remote sensing and image processing, which, when combined with the relaxation of legislation concerning existing and new data collection (Aplin *et al.*, 1997) and the commercialisation of remote sensing, have resulted in the proliferation of the availability and use of remotely sensed imagery for land cover mapping. Consequently, the technologies and techniques used in the creation of land cover maps are continually developing, as further advances in the types of maps and the technology that underpins the creation of these maps is achieved. In addition, the ways in which the data that these maps provide is disseminated is constantly changing. For example, the Ordnance Survey maintains a database of feature-level geometric information on many different types of land cover objects across the majority of the United Kingdom, called OS MasterMap©. This database can be used to generate highly specialised maps where, for example, only the information requested by the user is displayed on the map. One such implementation is Digimap (<http://www.digimap.com>), an internet-based service, where highly specialised maps can be generated instantly. For example, an agency in charge of water services could use the Digimap service to generate a map for a predefined area that displays only rivers and field boundaries, which could be used to determine access routes to their water supplies, for example, across private land. The growth in use of the Internet and electronic communications and developments in data storage technology have greatly influenced the availability and distribution of remotely sensed data (Su *et al.*, June, 2003). Broadband internet connections have enabled the transmission of small numbers of data quickly and easily. Moderate numbers of data can be recorded to a DVD inexpensively. Improved electronic storage capabilities now mean that a single data centre can store terabytes of data, which can be distributed efficiently via the Internet or private FTP sites to authorised users.

In the last two years, remotely sensed imagery has been placed into the public spotlight as several commercial applications became available. Foremost is the Google Earth service (<http://earth.google.com>), a software application that provides near global coverage of the Earth using remotely sensed imagery. Users are able to pan and zoom within the application and add pinpoints to favourite places or get directions to popular attractions (Butler, 2000). Similar services are those that

combine information from a Geographic Information System (GIS) to provide a combination of satellite sensor imagery, aerial photographs and cartographic maps, such as Google Maps (<http://maps.google.com>), Multimap (<http://multimap.com>) and Streetmap (<http://streetmap.co.uk>). Motoring organisations in the UK, such as The AA (<http://theaa.com>) and the RAC (<http://www.rac.co.uk>) utilise this technology and further extend the functionality of these maps, by offering comprehensive route planner services, providing users of these services with detailed directions and maps of the route between any two or more locations, with near European coverage. Such services represent only a small sample of the proliferation of using remotely sensed data in general purpose applications but are representative of the rapid expansion of the use of these data from the academic domain through to the commercial and public domains. As the availability of land cover information continues to increase and the use of land cover object databases, such as OS MasterMap or Google Earth develops, so the requirement for techniques to process these data to provide information on all kinds of land cover objects, including hedgerows and trees, for use in such databases also increases.

1.2 Principles of Remote Sensing

Earlier sections of this chapter evaluated the importance of land cover maps and considered the potential benefits of using remote sensing to create these maps. The following section describes how remote sensing can be used to assist in the creation of land cover maps. Initially, the theory and physical basis of remote sensing are described. Methods for overcoming some of the practical limitations are presented and then some key principles relevant to use of remotely sensed imagery, such as the effect of spatial resolution and mixed pixels, are introduced. The aims and objectives of this research project precede a conclusion.

1.2.1 Physical basis

(Campbell, 1996) defines remote sensing as:

“the practice of deriving information about the Earth’s atmosphere and surface using images acquired from an overhead perspective, using electromagnetic radiation reflected or emitted from the Earth’s surface”.

All physical objects on the surface of the Earth reflect electromagnetic radiation at certain frequencies within the electromagnetic spectrum. Remote sensing equipment is able to monitor and measure reflected radiation and reflectance data, recorded in the sensor’s instantaneous-field-of-view (IFOV), that are considered to be representative of the land cover type at the land-level. A remotely sensed image, however, is not a map of land cover; it is a collection of measured radiation values. It is *variation* in measured radiation, that is, variation in brightness per waveband, across an image that facilitates prediction of land cover and allows identification of features within an image, for example, using knowledge of the way in which measured radiation interacts with the environment and the reflectance characteristics of different types of land cover. Such information is commonly acquired during field study on the ground, for example, using techniques such as field spectroscopy to record the reflectance characteristics of different land cover types (Baath *et al.*, 2002). For example, the red and near infrared (NIR) portion of the electromagnetic spectrum is commonly used to map vegetation, since chlorophyll in plants absorbs radiation in the red wavelengths and the internal scattering properties of leaves increases reflectance in the NIR wavelengths of the electromagnetic spectrum (Kerr and Ostrovsky, 2003). Such information enables classification and mapping of many land covers within a remotely sensed image.

Remote sensing systems measure reflected radiation in sets of wavebands. A waveband can be defined by its starting position and width in relation to the electromagnetic spectrum. For example, the Landsat TM sensor measures in seven distinct wavebands across the spectrum (blue, green, red, near infrared, middle infrared, thermal infrared, middle infrared) in the wavelengths (excluding the thermal

band) from 0.45 μm to 2.35 μm (with the Enhanced Thematic Mapper (ETM) providing an additional 15 m spatial resolution panchromatic image), whereas the Airborne Visible-Infrared Imaging Spectrometer (AVIRIS) has 224 narrow wavebands in approximately the same wavelengths as the Landsat sensor but due to the narrower wavebands, produces a finer spectral resolution. The decision on an appropriate sensor to use should be made based on the requirements of the project, dependent primarily on the features of interest and the amount of detail required (Ju *et al.*, 2005). For example, a project monitoring the growth of tropical forests in Africa might be suited to Landsat TM imagery which offers moderate spatial and spectral resolution but wide spatial coverage, whereas a project monitoring the growth and decline of individual species of tree in a tropical forest in Africa would require greater spectral detail, for which AVIRIS may be more suitable, due to its finer spectral resolution.

1.2.2 Images and Pixels

A remotely sensed image is a collection of discrete pixels containing data on measured reflectance of the Earth's surface. Pixels are the smallest element of the image and represent an output device to which a remotely sensed observation is assigned. The Ground Resolution Element (GRE) of the sensor is the actual area that a sensor "sees" in making an observation. The number of pixels in an image will depend on the IFOV and the swath (the width of the image on the ground) of the sensor. In theory, each pixel should be representative of the reflectance from its own area and it should, therefore, be possible to use the image to identify the land cover on the ground. In reality, however, several factors affect this ability. These include the spatial resolution, available wavebands (spectral resolution), completeness of the imagery, as well as the make and model and the point spread function (PSF) of the sensing equipment (Campbell, 1996), in addition to several practical limitations, which affect the quality of the acquired imagery. Such limitations include atmospheric interference (e.g., cloud cover), geometric inconsistencies (e.g., movements of the sensor platform, relief variability and the curvature of the Earth) and technical problems with the remote sensing system. It is common practice to post-

process remotely sensed imagery using a variety of correction procedures to minimise the effect of such limitations (Ju *et al.*, 2005).

1.2.3 Radiometric Resolution

Radiometric resolution is a term used to describe the sensitivity of a sensor to the magnitude of electromagnetic energy it receives, that is, it is a measure of its ability to discriminate slight differences in energy (Campbell, 1996), commonly referred to as (levels of) “brightness”. The finer the radiometric resolution of a sensor the more sensitive it is to detecting small differences in reflected or emitted energy and the more levels of brightness are in the resultant image (Richards and Jia, 2006).

The maximum number of brightness levels available depends on the number of bits used in representing the energy recorded. Thus, if a sensor used 8 bits to record the data, there would be 2^8 digital values available, ranging from 0 to 255, whereas an 11 bit sensor would have 2^{11} values ranging from 0 to 2047 and would have a finer radiometric resolution (Richards and Jia, 2006).

For fine spatial resolution, the sensor tends to have a small IFOV. A consequence of a small IFOV, however, is that the amount of energy detected decreases as spatial resolution increases which leads therefore to a decrease in radiometric resolution. To maintain both a fine spatial resolution and radiometric resolution would require a broader wavelength range detected for a particular channel or band, which would decrease the spectral resolution of the sensor (Forshaw *et al.*, 1983). Conversely, coarser spatial resolution would allow increased radiometric and/or spectral resolution (Tso and Mather, 2001). Thus, the three types of resolution must be balanced against the desired capabilities and objectives of the sensor.

1.2.4 Spatial resolution

The spatial resolution of an image is a term used to describe the size of the area on the ground from which the sensor receives electromagnetic radiation in each pixel (and is therefore approximately equal to the GRE, most commonly referred to as the *support*,

that is, the size of the area from which sample data are recorded). A quantitative definition of spatial resolution is not generally possible, since it can be interpreted in a variety of ways (Atkinson and Aplin, 2004). In most cases, the geometric properties of the sensor are used instead. For example, the Landsat TM has a support of 30 m by 30 m. Pixels in the images have dimensions of 30 m by 30 m and accordingly the spatial resolution is said to be 30 m. The Landsat Multispectral Scanning System (MSS), for example, has a support of 79 m by 79 m, but due to the design of the imaging device, has an overlap of 11 m between pixels, so that each pixel represents an area of 57 m by 79 m. Spatial resolution is a fundamental consideration in any research involving remotely sensed imagery, since it affects several facets of the resultant imagery, such as the scale of spatial variation within the imagery and the frequency of mixed pixels. Generally, it is true that the finer the spatial resolution of the imagery the smaller the ground objects that can be distinguished within it (but equally more imagery is required to cover larger areas on the ground) (Atkinson and Aplin, 2004).

1.2.5 Spatial variation

Selection of an appropriate remote sensing scanner in order that the resultant imagery is suitable for the intended research is important, since land cover types identifiable at one spatial resolution may not be identifiable at a different spatial resolution (Woodcock and Strahler, 1987) due to the effect of spatial variation. For example, mapping a large fairly homogeneous area such as desert would require a sensor such as the National Oceanographic and Atmospheric Administration Advanced Very High Resolution Radiometer (NOAA AVHRR) which in local coverage mode has a spatial resolution of 1.1 km, whereas mapping a small area with high frequency spatial variation in land cover, such as many urban areas, might require a sensor with a finer spatial resolution. Depending on the intended use of the imagery, a spatial resolution of 5 m, as available from the SPOT Panchromatic sensor, could be used to extract individual features, such as buildings, from urban areas.

Spatial variation is an important consideration when creating maps of land cover. Spatial variation describes the variation between land cover types within a remotely

sensed image and is intrinsically linked with the spatial resolution of the image. The concept of the L-resolution and H-resolution models of spatial resolution (Strahler *et al.*, 1986; Woodcock and Strahler, 1987) describes the effect of spatial resolution on the scale of spatial variation observable within remotely sensed imagery. In the L-resolution case, features of interest are smaller than the spatial resolution of the imagery and cannot be resolved spatially. In the H-resolution case, features of interest are larger than the spatial resolution of the imagery and can be resolved. The L-resolution case describes a frequent problem with remotely sensed imagery, that of *mixed pixels*. Mixed pixels occur where a single pixel contains more than one land cover type. Mixed pixels are one of the most common sources of error in image classification. In general, the propensity for mixed pixels to occur increases as the spatial resolution of the imagery becomes coarser. In some cases, therefore, increasing the spatial resolution of the imagery may reduce the number of mixed pixels. However, increasing the spatial resolution beyond the H-resolution case may lead to oversampling, where variation within land cover types occurs. For example, an area of land cover at a coarse spatial resolution may be labelled “woodland” but at a finer spatial resolution, individual species of tree and areas of bare soil or grass within the forest may become discernible, and pixels that should ideally be labelled “woodland” may be incorrectly labelled “bare soil” or “grass”. The effect of mixed pixels on the interpretation of a remotely sensed image therefore depends on the spatial resolution of the remotely sensed imagery, the sizes of features to be resolved and the objectives of the research being carried out.

In most cases, finer spatial resolutions will reveal greater spatial variation than coarser spatial resolutions. However, this is wholly dependent on the types of land cover, the spatial arrangement of the land cover and the spatial frequency of the land cover types within the image.

1.2.6 Spatial dependence

Geostatistics is a set of techniques for the analysis of spatial data, first presented in (Matheron, 1965) and popularised in (Journel and Huijbregts, 1978), primarily as a set of techniques to assist in the mining industry. The cornerstone of geostatistics is the

phenomenon of spatial dependence. Spatial dependence is the tendency for proximate observations to be more alike than those further apart (Curran and Atkinson, 1998; Curran *et al.*, 1998) and can generally be described by the variance and semi-variance methods of the geostatistical paradigm. The phenomenon of spatial dependence is critical, since it enables the prediction of the relationship between two or more points in both discrete and continuous probability distributions. For example, by assuming spatial dependence, it is possible to make predictions on the spatial arrangement of land cover within a remotely sensed image. Super-resolution techniques are one example of this and are commonly implemented in a mapping objective to overcome the difficulties associated with mixed pixels in the classification of remotely sensed imagery.

Super-resolution techniques generally seek to adjust the spatial resolution of remotely sensed imagery so as to optimise the scale of spatial variation for the mapping objective. Such adjustments are made by increasing the spatial resolution of the imagery via a process known as downscaling. Estimated proportions of land cover classes within pixels (e.g., from a soft classification) can be used to perform a classification on the imagery at a finer spatial resolution (at the sub-pixel scale) and by assuming spatial dependence within and between pixels, the accuracy with which land cover classes are mapped can be increased by maximising spatial correlation between pixels such as to increase the spatial clustering within and between pixels (Atkinson, 2005).

1.3 Research objectives

Developments within the field of remote sensing, particularly as a result of technological advances, have led to a significant increase in the use of remotely sensed imagery for the creation of land cover maps. In particular, remotely sensed imagery at finer spatial resolutions than were previously available (from sources such as those described in Section 1.1.2), has increased the accuracy with which such maps can be created.

1.3.1 Rationale

To increase our understanding of, and ability to manage, the natural and built environments of the Earth, it is important to have information on land cover. Rural linear land cover features, such as hedgerows and pathways, are an important constituent of rural biodiversity. Hedgerows, in particular, support a rich diversity of plants, insects, birds and mammals and form an integral part of the UK Biodiversity Action Plan (UKBAP, <http://www.ukbap.org.uk>), cited as a “priority habitat”. Over the last 40 years, up to half of the UK’s hedgerows have been removed (McCollin, 2002), which is known to have had a profound effect on many different components of the UK’s rural biodiversity (Robinson and Sutherland, 2002). Accordingly, accurate maps of these rural linear features would be of great utility to the organisations involved in the protection of UK hedgerows.

Land cover information is available from many different sources, but in recent years, a common source of such information has been remote sensing. Technological advances in remote sensing have enabled fast and relatively inexpensive collection of large amounts of data whilst developments in image processing have enabled accurate and semi-automated analysis of these data. One of the most important technological developments has been the ability to acquire remotely sensed imagery at finer spatial resolutions than was previously available. Fine spatial resolution remotely sensed imagery is often used to extract more information from a scene than would have been possible at coarser spatial resolutions. Information on small land cover objects is important for use in many applications, such as land use planning and management, GIS databases and change detection, particularly in the case of monitoring and maintaining the UK’s natural hedgerows and rural areas. With the increased availability of fine spatial resolution remotely sensed imagery, the identification of land cover features smaller than was previously possible becomes feasible. However, the existence of mixed pixels in such remotely sensed imagery remains the most common source of error in land cover maps, and despite advances in classification techniques, such as the development of soft (or fuzzy) classification, a solution to the problem is still required (Foody, 2002b).

1.3.2 Aims and Objectives

In most cases, the ability to create accurate land cover maps relies on the availability of accurate feature-level geometric information on the location and position of individual land cover types. There are many techniques in common use for the creation of land cover maps, starting with classification techniques, which are covered in greater detail in Chapter 2. However, as finer spatial resolution remotely sensed imagery became available, so did the effect of mixed pixels increase (Foody, 2002b), since small land cover artefacts that were not visible at coarser spatial resolutions, occupy appreciable proportions of pixels in fine spatial resolution remotely sensed imagery.

To make best use of fine spatial resolution remotely sensed imagery for land cover mapping, techniques that actively allow for the effect of multiple land cover types within single pixels (mixed pixels) are required. Super-resolution mapping is one such set of techniques. Based on the phenomenon of spatial dependence and using information on the proportions of individual classes within pixels, super-resolution mapping has already been seen to map land cover with increased accuracy over conventional classification methods (Atkinson, 1997; Verhoeve and De Wulf, 2002).

The objective of the research presented in this thesis was to develop an accurate technique for mapping small features in fine spatial resolution imagery using super-resolution techniques. The aims of this research were:

- Evaluate the potential of super-resolution techniques for accurate mapping of land cover and land cover features, in particular, linear features
- Define the important characteristics of land cover features and investigate the extent to which super-resolution mapping can be used to map these characteristics
- Develop an algorithm that can map land cover more accurately than standard mapping techniques, such as hard classification, with appropriate consideration to those characteristics described previously

The rationale for the objective is presented in 1.3.1. In order to test the development of the algorithm presented in this thesis, land cover features such as hedgerows and trees were used as target features and standard accuracy assessment techniques were used in evaluating the super-resolution technique for the objective.

1.3.3 Thesis Outline

The thesis is organised accordingly: Chapter 2 sets the contextual background to this research. Firstly, the motives and key concepts that form the rationale of this research are described. Secondly, a review of existing methods and techniques from the literature is presented, including a description of each of the methods used in this research. Chapter 3 introduces each of the datasets used in this research for evaluating and testing the techniques. Chapter 4 introduces the super-resolution pixel swapping technique in a practical sense, by evaluating its performance when applied to simulated and real remotely sensed imagery. Chapters 5 and 6 represent the primary contribution of this research. Chapter 5 presents a pixel swapping technique for identifying and mapping rural features, including linear features, with greater accuracy than was previously possible. Chapter 6 presents an image fusion framework, to integrate multiple sources of data. In Chapter 7, a discussion of the important issues that arose during this research is provided. Chapter 8 closes the thesis with conclusions.

1.3.4 Publications

The research presented in this thesis has resulted (in part or in full) in the following publications (or submissions) and conference presentations:

Peer-reviewed journal papers:

Thornton, M.W., Atkinson, P.M., and Holland, D.A. Super-resolution image fusion for sub-pixel mapping of rural land cover features from fine spatial resolution remotely sensed imagery. *To appear as part of a journal special edition arising from the RSPSoc Annual Conference, Cambridge, 2006.*

Thornton, M.W., Atkinson, P.M., and Holland, D.A., 2007. Sub-pixel mapping for extracting rural linear land cover features from fine spatial resolution remotely sensed imagery. *Computers & Geosciences*, in press.

Thornton, M.W., Atkinson, P.M., and Holland, D.A., 2006. Sub-pixel mapping of rural land cover objects from fine spatial resolution satellite sensor imagery using super-resolution pixel swapping. *International Journal of Remote Sensing*, 27, 473-491.

Conference proceedings/presentations:

Super-resolution mapping of linear land cover features from remotely sensed imagery. *International Symposium on Remote Sensing of Environment (ISRSE)*, Hawaii, 2003.

Sub-pixel identification of fine spatial resolution land cover objects from Quickbird™ satellite sensor imagery using super-resolution mapping techniques. *Remote sensing and Photogrammetry Society (RSPSOC) Annual Conference*, Aberdeen, 2004.

Simulated annealing for super-resolution land cover mapping from fine spatial resolution remotely sensed imagery. *Remote sensing and Photogrammetry Society (RSPSOC) Annual Conference*, Portsmouth, 2005.

1.4 Conclusion

Remotely sensed imagery is a common source of information for creating land cover maps. A remotely sensed image, as a set of discrete pixels, is a record of reflected electromagnetic radiation of an area of the Earth's surface. Using knowledge of the reflectance characteristics of the Earth's surface, many objects can be identified from the radiation they reflect. Traditional techniques achieve this by classifying remotely sensed images into a set of discrete land cover classes and representing these classes visually in a map of land cover. In recent years, the availability of remotely sensed imagery in general has increased, but notably technological advances have led to the availability of remotely sensed imagery at fine spatial resolutions.

Fine spatial resolution remotely sensed imagery generally enables the production of more accurate classification maps. However, the derisory effect of mixed pixels on

classification accuracy is still apparent. Using super-resolution techniques, such as sub-pixel mapping based on the phenomenon of spatial dependence, pixels can be classified with a greater accuracy than standard classification techniques.

This introductory chapter has discussed the basic principles behind remote sensing and image analysis for predicting land cover on the ground and stated the objectives of the research in this thesis. Chapter 2 continues in more detail by presenting and analysing the key set of techniques, which form the basis of this research project.

Chapter 2

Chapter 2

Background and Methods

The first part of the chapter discusses the background of the study, including the importance of understanding the role of the environment in the development of children's language and the need for a comprehensive approach to the study of language development. The second part of the chapter describes the methods used in the study, including the design of the study, the participants, and the data collection procedures.

The study was designed to investigate the role of the environment in the development of children's language. The participants were 100 children aged 2-5 years, who were recruited from a community center. The data collection procedures included observations of the children's language use in natural settings, as well as structured interviews with the parents. The results of the study are discussed in the following sections.

Chapter 2

2. Background and Methods

2.1 Introduction

The methodology required for super-resolution mapping, the primary goal of this research, is a process whereby (i) input remotely sensed imagery is classified (using one or more of a range of available soft classification techniques), to provide input data for (ii) a super-resolution mapping technique. During the course of this research, the functionality of an existing super-resolution pixel swapping technique was extended to produce accurate land cover maps of small rural features, such as hedgerows and trees.

In this chapter, the methods and techniques used in this research project are reviewed. Initially, a review of the concept of classification is presented, followed by an analysis of two soft classification methods used in this research project. Following the review of classification techniques, super-resolution mapping is introduced and a background to existing techniques is presented. The pixel swapping technique used in this research is then presented and analysed.

During this research, the super-resolution pixel swapping technique was modified in two key ways: (i) 'linearisation' was performed to increase the accuracy with which linear features were predicted and, (ii) an image fusion component was integrated to further increase the accuracy of prediction and refine feature-level delineation of feature characteristics. Accordingly, the background to the role of feature identification and image fusion and existing techniques are presented.

The final component of this chapter reviews techniques that were used for assessing the accuracy of the methods used in this research. Various aspects of accuracy assessment are discussed and the methods employed are reviewed. The chapter closes with conclusions.

2.2 Classification

Classification is one of the most common objectives for which remotely sensed imagery are applied (Foody, 2002b). Land cover classification is the process of converting remotely sensed imagery (spectral domain) into a representation of land cover on the ground (spatial domain), often in the form of a land cover map (Tso and Mather, 2001). This is achieved by assigning pixels (containing digital numbers) to land cover types (land cover classes) (Campbell, 1996). Classification can be divided broadly into two types – hard classification and soft classification.

2.2.1 Hard classification

Hard classification comprises a large proportion of all available classification techniques (Foody, 2002a; Foody, 2002b). Traditional hard classification techniques are formulated under a key principle, that is, a pixel is assigned to a single class (Atkinson and Aplin, 2004). Common techniques for performing hard classification are the maximum likelihood, minimum-distance-to-means and parallelepiped methods (Campbell, 1996).

The maximum likelihood (ML) classifier is perhaps the most commonly implemented approach to hard classification. Classes within an image are defined on the basis of measured training data based on Bayes' Theorem. ML classification proceeds by estimating the histogram of each class in multivariate feature space from training data. The mean and variance-covariance are then estimated which characterizes the “likelihood”. The maximum posterior probability of a pixel belonging to a class is then calculated by normalizing the product of the likelihood and the prior. A pixel is

then allocated to the class with which it has the highest posterior probability- the maximum likelihood (Atkinson and Lewis, 2000). In other words, the ML classifier measures the mean and variance-covariance between bands per class, such that the probability of a pixel belonging to a class can be estimated from its spectral response.

The ML classifier is often described as the most accurate of the three main conventional methods of hard classification (Foody, 1990; Congalton, 1991; Campbell, 1996; Foody, 2002b), though the accuracy can depend on an appropriate sample size, accurate training data to ensure the class characteristics, as determined by mean vector covariance matrices, are representative and that the data are accurately represented by a normal distribution. The ML classifier is also fairly statistically and computationally intensive, though given the availability of technology for image processing, as well as field spectroscopy for collection of training observations, the ML classifier is the most widely used hard classification technique.

Classification techniques often make assumptions about the data that they are being applied to. For example, the ML classifier is based on the assumption that the prior probability in the information class has a Gaussian (normal) distribution (Mather, 1999) which, in reality, is rarely the case; generalising the distribution of a complex land cover class, which often incorporates unpredictable deviations from normality, may be inappropriate.

Hard classification techniques commonly make the assumption that individual pixels are discrete units (on the ground), that is, they are pure pixels containing a single land cover class. In reality, where land cover classes are not discrete units and intergrade across many pixels, this assumption is not justifiable (Foody, 2002a). Pixels often contain a mix of more than a single land cover type. These pixels are called *mixed pixels* and are a fundamental consideration in remote sensing research. (Foody, 2002a; Foody and Mathur, 2006) suggest that alternative methods for hard classification, such as evidential reasoning or neural networks, are no more or less suited to classification than conventional hard classification, similarly due to the questionable reliability of the assumptions inherent in these techniques. However, a well-

established set of classification techniques, so-called soft classifiers, provides a possible solution to the problem of mixed pixels.

2.2.2 Soft classification

In a remotely sensed image, the spatial frequency of land cover and the frequency of sampling (factors such as the spatial resolution of the sensor, effects of the point-spread function and to a lesser extent, pre-processing corrections) will likely result in mixed pixels (pixels containing more than one class) (Gebbinck, 1998; Choi and Lee, 2001; Tso and Mather, 2001; Hochberg and Atkinson, 2003). For example, a pixel might exhibit 60% forest, 30% urban and 10% water yet using the maximum likelihood classifier it would be assigned to a single class (in this example, it would be assigned to forest). In most cases, this is inappropriate.

Soft classification (sometimes referred to as fuzzy classification) techniques evolved as a response to the problem of mixed pixels, where, on the assumption that a pixel is a linear combination of the spectral responses of the land cover within it, the proportions of each class within a pixel can be estimated and a proportional representation of each class is given, such that, instead of assigning a pixel to a single class (a one-to-one relationship), a pixel is assigned to multiple classes (a one-to-many relationship) (Tso and Mather, 2001).

According to (Tso and Mather, 2001), soft classification techniques fall into three categories: fuzzy set theory, artificial neural networks and spectral mixture analysis. Fuzzy set theory algorithms work on the existence of fuzzy partitions between classes, which can be used to construct fuzzy membership values for each class. A straightforward implementation of fuzzy set theory is a 'softened' version of the traditional hard maximum likelihood classifier known as the fuzzy maximum likelihood classifier (Wang, 1990). It uses the fuzzy mean and fuzzy covariance matrix, and outputs fuzzy membership grades for each pixel (Tso and Mather, 2001). Earlier work on fuzzy methods led to the development of the fuzzy *c*-means clustering algorithm (Bezdek *et al.*, 1984) which may be used for either unsupervised (Key *et al.*, 1989) or supervised (Cannon *et al.*, 1986; Foody and Cox, 1994) classification. Fuzzy rule bases are an extension to the fuzzy set membership functions and can be

used to generate a series of rules on how to unmix pixels in feature space (Wu *et al.*, 2001). (Ho *et al.*, 1994) suggest a series of techniques for combining multiple classifiers to maximise classification accuracy by using different classifiers based on decision rules.

Artificial neural networks (ANNs) may be used to soft classify remotely sensed imagery (Bernard *et al.*, 1997). The network is trained initially with mixed pixels using pixel proportions. Input vectors generate a membership value for each class. This information is then repeatedly fed forward and back propagated through the network such that it learns the characteristics of the classes and accordingly assigns pixels to multiple classes based on their spectra. In most cases, ANN classification is more accurate than traditional classifiers (Foody, 1996; Atkinson and Tatnall, 1997).

Two of the most commonly used soft classification techniques are the fuzzy *c*-means and the linear mixture model, both of which are used in this thesis and, therefore, are described in more detail below.

2.2.2.1 Fuzzy *c*-means

The fuzzy *c*-means clustering algorithm was first presented by (Bezdek *et al.*, 1984), as a more robust clustering algorithm than traditional algorithms, such as the ISODATA method (Tso and Mather, 2001) and as a means of overcoming the false clustering problem (Acqua and Gamba, 2001). Initially, the number of clusters (classes) in an image is chosen, and a $c \times n$ sized fuzzy covariance matrix is created accordingly, where c is the number of clusters and n is the number of pixels. Then, the mean value of each cluster is calculated and using a membership weighting exponent and a distance metric, such as the Mahalanobis distance, the membership of each pixel in the covariance matrix to each cluster is updated iteratively (Campbell, 1996). The membership weighting exponent, m , can be varied such that the fuzziness (hardness) of the classification can be controlled, that is, pixels can be “members” of many classes or only a few. The fuzzy *c*-means classifier has been used extensively in a variety of different applications (Cannon *et al.*, 1986; Nguyen and Cohen, 1993; Foody and Cox, 1994; Bastin, 1997; Foody, 1998).

2.2.2.2 Linear mixture model

Spectral mixture analysis uses pure pixels (or “endmember” pixels), which are homogeneous pixels, that is, pixels that are known to contain a single class. Each endmember may be found in feature space, around which a polyhedron can be drawn – all possible mixtures are found within this polyhedron in feature space (Settle and Drake, 1993). An example of spectral mixture analysis is the linear mixture model (Adams *et al.*, 1985; Settle and Drake, 1993).

Mixture modelling assumes that within a pixel, the spectral responses of each class are mixed linearly in proportion to the area that each class represents. Therefore, the spectral response of each pixel is a linear combination of the endmember spectra of each class and contains information on the proportion of each class found within that pixel. Spectral endmembers for each class can then be used to unmix the pixels and predict the proportions of each class within each pixel (Adams *et al.*, 1985; Settle and Drake, 1993; Foody and Cox, 1994; Bastin, 1997; Manolakis *et al.*, 2001; Shimazaki and Tateishi, 2001; Settle, 2002).

Soft classification comprises a set of alternative techniques to hard classification by providing information on the number of classes within a pixel and the proportion of the pixel that each class occupies. However, the information provides only the proportions of each class: the location of these proportions is unknown. Super-resolution techniques are commonly used to map the location of class proportions within pixels.

2.3 Super-resolution techniques for land cover mapping

Remotely sensed imagery can be used, through classification and a variety of post-classification techniques, to create land cover maps. Some of these techniques draw on the basic concepts and theory of geostatistics, in particular, the phenomenon of spatial dependence, and whilst the common geostatistical tools, such as the variogram are not used in this research, a brief description of geostatistical theory is presented.

After a brief review of geostatistical theory, advanced techniques for land cover mapping are presented, including the concept of sub-pixel mapping, which ultimately leads to the super-resolution techniques used in this research: pixel swapping and, subsequently, an updated version based on simulated annealing.

2.3.1 Geostatistical theory

Geostatistics can be defined as the study of phenomena that fluctuate in space (Deutsch and Journel, 1998). Geostatistical methods first appeared to assist in spatial prediction in the mining industry (Matheron, 1971). However, these methods are increasingly being applied to a diverse range of fields, including geological, atmospheric and environmental sciences and more recently, geostatistics has been used in remote sensing applications (Curran and Atkinson, 1998; Curran *et al.*, 1998). Geostatistics offers a collection of tools that have been developed to aid the understanding and modelling of spatial variability. One use of such information would be for the prediction of unknown values within a sample (Isaaks and Srivastava, 1989). In the prediction of unknown values, the phenomenon of spatial dependence (that is, proximate values are more alike than those further apart) is commonly used. For example, a geostatistical tool such as the variogram can be used to describe the relationship between pairs of known values in a given space. This information can then be used to interpolate values at locations for which no sampled data exist (Lloyd and Atkinson, 2004). Collectively, geostatistical theory, and its tools and techniques can be used to model across a set of data, as well as providing estimates as to the uncertainty of the model and its predictions.

2.3.2 Sub-pixel mapping

(Atkinson, 1997) used an assumption of spatial dependence within and between sub-pixels to map the location of land cover classes within pixels. This was achieved by using the output from a soft classification, which provided information on the proportions of each class in each pixel. The assumption was valid for re-creating the

layout and areal coverage of the land cover, and the algorithm produced an acceptable degree of accuracy for semi-natural land cover. However, the technique compared sub-pixels to pixels and so the complex mixing in the data caused the simple technique to suffer from problems.

(Verhoeve and De Wulf, 2002) used similar assumptions as (Atkinson, 1997), but formulated them as a linear optimization problem. The algorithm was applied to synthetic imagery and a Système Pour L'Observation de la Terre (SPOT) multispectral image of Sahelian wetlands. Compared with traditional hard classification techniques, acceptable results were produced, but the method was similarly impaired to that of (Atkinson, 1997) in that sub-pixels were compared to pixels, rather than to other sub-pixels, causing the appearance of linear artefacts.

(Zhan *et al.*, 2002) developed the assumptions of (Atkinson, 1997) and (Verhoeve and De Wulf, 2002) and applied the technique to two of the four mixed pixels scenarios suggested by (Fisher, 1997), the boundary pixel and intergrade pixel. They used a two-stage methodology. Initially, the maximum likelihood (ML) classifier was used to derive maximum likelihood probabilities (MLP) of class membership and then, secondly, applied an inverse distance-weighting predictor to interpolate a probability surface at the sub-pixel scale with the MLP values for the central pixel and neighbourhood pixels at the pixel scale. The image was classified according to the interpolated MLP values at the sub-pixel scale. However, comparison between pixels was only made on sub-pixels in either the current sub-pixel or the corners of neighbouring pixels which reproduced the scale and linearity limitations that arose in the (Atkinson, 1997) and (Verhoeve and De Wulf, 2002) methods.

Further work on the Verhoeve and De Wulf (2002) method was carried out using genetic algorithms (Mertens *et al.*, 2003). Working on the basis of spatial dependence on both simulated and real imagery, the genetic algorithm is a search and replace optimisation algorithm based on the principles of natural selection and genetics. Based on a fitness value, sub-pixels were iteratively compared with other sub-pixels, such that the “fitter” sub-pixels were used to influence the selection process and vary the arrangement of sub-pixels within an image until spatial dependence between sub-pixels was optimised. When compared with traditional hard classification techniques,

the genetics method was quantitatively more accurate on both the simulated and real imagery.

(Aplin and Atkinson, 2001) developed a technique for sub-pixel mapping that operated on a “per-field” basis. Most sub-pixel mapping techniques classify within sub-pixels independently from ancillary data (e.g., vector data) that may be available, whereas the per-field technique grouped individual sub-pixels of similar land cover based on available Ordnance Survey polygons. Their method required the use of two datasets, the Ordnance Survey Land-Line (a vector dataset of land cover polygons) and Compact Airborne Spectrographic Imager (CASI) imagery. Four land cover classes were selected from the imagery and a supervised ML classification was then performed using the Mahalanobis distance measure. A soft classification was also applied. The two datasets were then integrated by rasterising the vector Land-Line data and then relating the two as ASCII files. Per-field classifications were carried out on these data, performed on both the hard classified and soft classified imagery. For the hard classified imagery, the land cover class from the CASI image was assigned to the (new) Land-Line polygons. For each polygon, the area covered by each class was summed, and the modal land cover class was assigned to the polygon as a whole. For the soft classified image, the Mahalanobis distance of each pixel and the Land-Line polygon within each CASI pixel were ranked, with land cover labels being assigned based on the shortest distance and largest proportional cover.

Artificial neural networks (ANNs) can be used for the creation of land cover maps. (Paola and Schowengerdt, 1995) have previously established the benefits of using an artificial neural network (ANN) over conventional classification methods such as the ML classifier. (Tatem *et al.*, 2001) used a Hopfield neural network (HNN) to map the location of land cover classes within pixels based on the output of a soft classification. The HNN has been shown to be able to predict the location of class proportions to produce a sub-pixel scale land cover map for use in land cover target identification and for mapping of sub-pixel scale land cover features on both simulated and real imagery (Tatem *et al.*, 2002b; Tatem *et al.*, 2002a; Tatem *et al.*, 2003). (Mertens *et al.*, 2004b) have applied the use of neural networks and wavelet coefficients to sub-pixel mapping and (Nguyen *et al.*, 2005) extended the use of the HNN to include LiDAR data for the creation of super-resolution maps using digital surface models.

Other sub-pixel techniques have enabled multiple-scale land cover change detection (Braswell *et al.*, 2003), land cover mapping, target detection from unknown sub-pixels using spatial filtering (Ashton, 1998) and sub-pixel land cover mapping using mixture analysis (Ju *et al.*, 2003), the linear mixture model (Manolakis *et al.*, 2001), indicator geostatistics (Boucher and Kyriakidis, 2006), spatial attraction models (Mertens *et al.*, 2006) and shoreline mapping (Muslim *et al.*, 2003; Muslim *et al.*, 2006).

Making more realistic assumptions about the contents of pixels and providing a mechanism for minimising the effect of mixed pixels lead to increases in classification accuracy. In most cases, the sub-pixel techniques described above demonstrated significantly increased classification accuracy over conventional hard classification methods.

2.3.3 Sub-pixel mapping: pixel swapping

(Atkinson, 2004b) proposed a super-resolution pixel swapping algorithm. Land cover proportions from pixels in a remotely sensed image were transformed into sub-pixels and allocated randomly. The goal was to maximise spatial correlation by varying the spatial arrangement of the sub-pixels within each pixel – the proportional values of the pixels (and quantity of sub-pixels allocated to each class) remained fixed. A distance weighted function of each sub-pixel and its neighbours was predicted that determined the position and attractiveness of the sub-pixel within the pixel. Once this information was stored, sub-pixels within pixels were initially allocated randomly, and the algorithm then compared sub-pixels. The most attractive and least attractive sub-pixels in each pixel were compared. If the attractiveness of the least positive sub-pixel was less than that of the most attractive neutral sub-pixel in the same pixel, then the land cover classes were swapped (otherwise no change was made).

The pixel swapping algorithm works as follows. Within each iteration, the attractiveness (A_i) of each sub-pixel for a particular class k is predicted as a distance-weighted function of its neighbours:

$$A_i = A_k(\mathbf{x}_i) = \sum_{j=1}^m \lambda_{ij} z_k(\mathbf{x}_j) \quad (1)$$

where m is the number of neighbours, $z_k(\mathbf{x}_j)$ is the value of the class, k (now constrained to be either 0 or 1), at the j^{th} pixel location, \mathbf{x}_j , and λ_{ij} is a distance-dependent weight predicted as:

$$\lambda_{ij} = \exp\left(\frac{-h_{ij}}{a}\right) \quad (2)$$

where h_{ij} is the distance between pixel \mathbf{x}_i (for which the attractiveness is desired) and the neighbour \mathbf{x}_j and a is the non-linear parameter of the exponential distance-decay model (here the exponential model is used, but this is not mandatory).

The algorithm is run for a fixed number of iterations or until no swaps are made. In (Atkinson, 2004b), the algorithm operated quickly and efficiently on simulated imagery and the basic, simulated shapes, on a visual comparison with the input shapes, were mapped and reformed accurately.

The pixel swapping algorithm is a simple and efficient super-resolution technique. It uses the exponential distance-decay model in ranking the attractiveness of sub-pixels within pixels. This model allows control over the way in which sub-pixels are ranked, through the use of the non-linear parameter, which controls the gradient and descent of the curve of the model. A further parameter to be set by the investigator is the distance at which sub-pixels surrounding the target sub-pixel are included (the bandwidth). For example, a bandwidth setting of two utilises the sub-pixels immediately surrounding the sub-pixel in question, and the sub-pixels immediately surrounding those. That is, a moving window with the target sub-pixel at the centre.

2.3.3.1 Simulated Annealing

Simulated annealing is an optimisation algorithm that is used to perturb a set of input data until an objective, defined by some predetermined criterion, is met (Deutsch and Journal, 1998). Simulated annealing is similar in concept to the pixel swapping technique described above, yet simulated annealing is based on a more rigorous statistical background (Tso and Mather, 2001).

Within a pixel swapping framework, simulated annealing can be used to distribute sub-pixels randomly within pixels and subsequently make swaps between sub-pixels based on conditional probability (changes in energy). Over a series of iterations, a temperature parameter, starting high, is gradually cooled until 'frozen'. At each iteration, the prediction of the swap is compared with the objective function. A swap is retained when the predicted swap results in a move towards the objective function. At higher temperatures, the probability of a swap being made is greater, where large differences in energy are allowed. As the system cools, only small increases in energy are allowed and accordingly, fewer swaps are made, with the potential for incorrect swaps decreasing over the course of the annealing schedule. The schedule lasts until repeated iterations do not lower the objective function (i.e. no swaps are made) or when a pre-specified minimum value is met. The effect of the annealing schedule is to introduce error into the system to avoid getting stuck in a local minimum.

2.3.3.2 Swap criteria

There are a wide variety of options for defining the conditions under which a swap is accepted; that is, assessing whether a predicted swap is suitable or not.

In this research, the criteria are initially defined, thus:

One target and one background sub-pixel (i.e., a 1 (target) and a 0 (background)) within a pixel were selected randomly. A_i values for each sub-pixel were calculated as per the pixel swapping method described above. The sub-pixels were then swapped and the A_i values were recalculated. If the swap resulted in an increase in A_i for the

target (1) sub-pixel, then the swap was retained. If the A_i attractiveness decreased for the target (1) sub-pixel then an annealing schedule was used to determine whether the swap was retained. In Chapter 6 of this thesis, where image fusion was used to combine multiple sources of data, the criteria were updated. Complete details are provided in Chapter 6.

2.4 Identifying land cover features

The objective of the research presented in this thesis was to develop a super-resolution mapping technique for the mapping of rural land cover features. An emphasis was placed on the mapping of *linear* rural land cover features, such as hedgerows (although other land cover features, such as trees, were also of great interest). Commonly used concepts in existing linear feature extraction techniques aided in the development of a “linearised” version of the pixel swapping technique and so a brief review of these common feature extraction techniques is now presented. However, the techniques described here were not used directly in this research. The linearised pixel swapping technique is described in detail in Chapter 5.

Feature extraction can be viewed as the objective of finding a set of vectors that represent an observation (Lee and Landgrebe, 1993; Choi and Lee, 2001). Defining a feature is a complex task since features can have many dimensions and varied characteristics, which often causes problems (Kovalev and Petrou, 1996). In simple terms, a feature can be a physical object (e.g., a tree, hedgerow, or building), an area (e.g., a field or city or geographical region) or alternatively a land cover class (e.g., within feature space) (Choi and Lee, 2001; Kuo and Landgrebe, 2002). Features extracted from images can be linear or areal: linear features, for example, are roads, pathways, hedgerows or any linear artefact found within the scene, such as the edge or boundary of a field; areal features are, for example, buildings, trees or fields (Mayer, 1999; Shi and Zhu, 2002).

2.4.1 Edge Detection

Edge detection (or enhancement) is a generic term that covers a range of feature extraction techniques for finding edges based on spatial filtering or templating. A commonly used edge detector is the Canny operator (Canny, 1986; Heath *et al.*, 1997). An image is first smoothed by Gaussian convolution and a two-dimensional first derivative operator is then applied to create a gradient magnitude image, with regions of high spatial first-derivatives forming along edges found within it. Subsequently, non-maximal suppression is tracked along the ridges, whereby any pixels not contained within the ridge are set to zero. This is controlled by two threshold values, T_1 and T_2 , where $T_1 > T_2$. Tracking commences on a point on a ridge with a value higher than T_1 , and continues until the ridge falls below T_2 . This thresholding is used to ensure that noisy edges are not broken up into a series of edge fragments. Careful selection of suitable thresholds is required in order to minimise an inevitable trade-off between extraction of all edges in the image and an increase in noise. Two other edge detectors are the Roberts Cross operator and the Sobel operator (Heath *et al.*, 1997; Richards and Jia, 1999), which consist of a pair of convolution masks (Roberts Cross is 2x2 pixels, Sobel is 3x3 pixels) where the second mask is identical to the first mask, except rotated through 90 degrees. Each mask is applied separately to the image and then combined to find the magnitude and orientation of gradients at each point. This produces a smoothed image with edges that have been both enhanced, and in some cases, widened (by several pixels).

2.4.2 The Hough transform

A well-known technique for detecting lines in images is the Hough transform (Hough, 1962). The Hough transform invariably receives as its input, the output from an edge detection algorithm. Hough presented the transformation procedure as a method to detect the presence of groups of collinear or almost collinear figure points in pictures for the extraction of complex patterns. Initially developed for extracting patterns from charged particles in bubble tubes and then in astrophysics for charting stars, the Hough transform has been developed by researchers (Duda and Hart, 1972; Leavers,

1993; Stewart, 1999) as an efficient line detector for remotely sensed imagery. It is particularly suited to noisy imagery, such as synthetic aperture radar (SAR) (Skingley and Rye, 1987). Any single point (e.g. X_i, Y_i) has an infinite number of lines passing through it. This point can be represented uniquely based on the distance from the origin (p) and the angle (θ) in slope-intercept (parameter) space. All points in image space can be transformed into lines in parameter space which are then divided up into a number of discrete *accumulator* cells which count each time a line is projected into it. Peaks in these cells mark the equations of the collinear points in $x y$ space. This is the classical Hough transform. However, a problem exists when lines become vertical since the parameter space becomes unbounded. Therefore, a normalised Hough transform is more commonly used.

The Hough transform is efficient for simple shapes such as lines, curves and circles. If shapes become mathematically complex, as Duda and Hart (1972) point out, line matching can be computationally intensive when there are a large number of points. It is possible, however, to adjust the parameters, such that the cells used to search for lines are bounded to a specified area using the angle-radius parameters. The Hough transform has also been applied to SAR imagery for the extraction of faint, thin lines, such as horse tracks in coniferous woodland (Skingley and Rye, 1987). The Hough transform, however, pays no regard to contiguity within the image, such that it would be possible for completely unrelated or meaningless groups of points to influence the line and relies heavily on thresholding (i.e., the choice of what is a “high” value in accumulator cells) which introduces bias and operator error into the results. (Davies, 1987) also presents a difficulty with the Hough transform, whereby objects that possess more than one straight edge are not detected with optimal sensitivity. The solution to the problem is a further parameterisation of the transform, which consequently requires further computation.

2.4.3 Active contour models: snakes

Active contour models (also called ‘snakes’) are “an energy minimizing spline guided by external constraint forces and internal forces that pull it toward features such as

lines and edges” (Kass *et al.*, 1988). (Nixon and Aguado, 2002) use the analogy of using a balloon to ‘find’ a shape, i.e. the balloon is placed outside the shape, enclosing it. The air is then taken out of the balloon, making it shrink and the edges of the shape are detected when no more air can be released from the balloon. Snakes, therefore, define a set of points that enclose the target feature. Snakes are advantageous over other feature extraction techniques, since image data are integrated into the process, knowledge-based constraints can be used, and desired contour properties can be achieved, all in a single process (Gunn and Nixon, 1997). Furthermore, the predicted edges are not restricted to straight edges as contouring techniques allow the prediction of the edges of curvilinear features.

(Abd-Alamageed *et al.*, 2003) describe common problems with statistical pressure snakes. Specifically, they note that the shape must exhibit a strong edge in order for it to be detected and accordingly causes difficulties for detecting shapes in “weak gradient fields” – i.e., where the transition between target and background is not pronounced. Furthermore, these techniques assume that low-order statistics, such as the mean and the standard deviation, represent the target area accurately. Snakes also require complex parameterisation. The authors present a generalised non-parametric pressure snake based on the work of (Parzen, 1962). Initially, they estimated a probability density function (pdf) of the target and background and implemented a Gaussian kernel function, followed by a k -point moving average filter, to reduce noise from the pdf. When executed, the snake ‘collapsed’ around the target within five frames, using a 30 control-point snake. Increased accuracy was achieved by increasing the number of control-points, although this resulted in a trade-off between speed and accuracy.

(Gunn and Nixon, 1994; Gunn and Nixon, 1997) have researched extensively the use of snakes, and, in work designed to overcome the problem of complex parameterisation and initial placement of the snake, developed a dual active contour system. A dual contour system is applied by placing two snakes on the image, one outside the feature, as in the conventional snake technique, followed by an additional snake *within* the feature. The outer snake is then contracted to fit the outer edge, whilst the inner snake is expanded outwards to fit the inner edge. Neither snake was programmed with any preference about which direction to move in. That is, they were

simply allowed to acquire their natural shape. This reduced the need for parameterisation and initialisation, and allowed a comparison between the two contours' energies, enabling the rejection of weak local minima. In each test case, the dual contour technique performed more accurately than conventional single contour methods, especially when there was noise or complex minima in the image, since the dual contour technique was less sensitive to such factors.

2.4.4 'Linearisation'

In this research, a linearisation technique was developed to increase the accuracy of predicted linear features in the super-resolved output of the pixel swapping technique. A summary of the linearisation technique is now provided. The technique is described in full in Chapter 5.

The linear features of primary interest in this research are hedgerows found in rural areas of the UK countryside. In addition to information in the literature, fieldwork was undertaken to investigate and define the properties and characteristics of such hedgerows. This fieldwork is described in greater detail in Chapter 3. Three key characteristics of hedgerows, which were used to assist in the selection of suitable study areas and data, were established: variable size (width and length), curvature (the tendency for the geometric orientation of hedgerows to change along its width and length) and the existence of embedded objects (trees and other land cover artefacts were commonly found *within* hedgerows). Each of these properties represented important considerations when developing the linearisation technique.

Existing feature extraction techniques (above) were reviewed to provide information on common practices in the detection of linear features. It occurred, however, that individual techniques, such as those described above, were not configured optimally for the detection of small linear features, such as hedgerows. The linearisation technique was developed such that an estimation of the existence of linear features (and the associated geometric orientation of such features, often referred to in this research as 'direction') in remotely sensed imagery, based on the input soft

proportions, was used to increase the likelihood of predicting contiguous linear features in the super-resolved output.

It is worth noting, however, that linear *features* are not explicitly *extracted*. That is, the term ‘features’ in this research is generally used to infer the prediction of features as contiguous pixels (in a visible sense) as defined by their classification. For example, a land cover class labelled “hedgerow” defines a feature, or set of features, which are “hedgerows”. However, given the super-resolved output of the pixel swapping technique, creating *actual* objects is relatively straightforward, using a simple raster-vector conversion, as performed in standard GIS applications.

2.5 Image fusion

The use of image fusion techniques for increasing the accuracy of mapping land cover classes was investigated. A discussion and review of existing image fusion technique follows.

Image fusion (or image merging, image integration or multi-sensor data fusion) is a concept comprising a set of techniques to combine multiple sources of data (Wald, 1999). The goal of image fusion is to integrate complementary information from multisensor data such that the new images are more suitable for the purpose of human visual perception and computer-processing tasks (Li *et al.*, 1995; Piella, 2003). Through the combination of multiple sources of data, it is possible to extract more information than from each source independently (Chavez, 1991) – ‘ $1 + 1 = 3$ ’ (Pohl and Van Genderen, 1998). For example, a common implementation of image fusion is to merge multispectral remotely sensed imagery with panchromatic remotely sensed imagery (Li *et al.*, 2002). Panchromatic imagery often has a finer spatial resolution than multispectral imagery, which is useful in object recognition, but lacks the spectral information for use in applications such as mapping land cover (Chibani and Houacine, 2002b; Chibani and Houacine, 2002a). By fusing complementary imagery, a multispectral image with the spatial resolution of the panchromatic image can be produced (Liu, 2000). Data fusion is not limited to simply merging two images

together; ancillary data, such as topographic maps, GPS information or other geophysical information can be used in fusion (Harris and Murray, 1989).

Before any image fusion technique is performed, input images must be geometrically co-registered to ensure that the input images are both representative of the same area on the ground (Pohl, 1999). The commonly used technique for co-registration is to resample one image to match the other, through the use of ground control points (GCPs) which are sites that are easily distinguishable on each input image (Campbell, 1996). Registration accuracy has a direct result on the accuracy of the fused output and should be accounted for in any analysis. Figure 2.1 illustrates the fusion process.

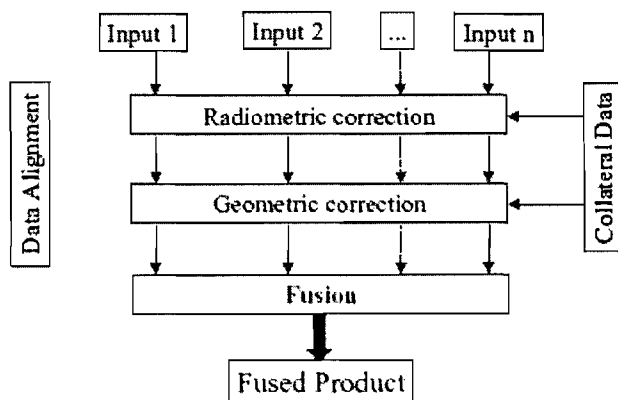


Figure 2.1: The image fusion model (Pohl, 1999)

Image fusion techniques are often divided into three categories, dependent on the stage at which fusion takes place (Li *et al.*, 2002):

1. Pixel level – individual pixels are merged
2. Feature level – groups of pixels (“features”) are merged
3. Decision level – complete images are merged

which can be visualised as shown in Figure 2.2:

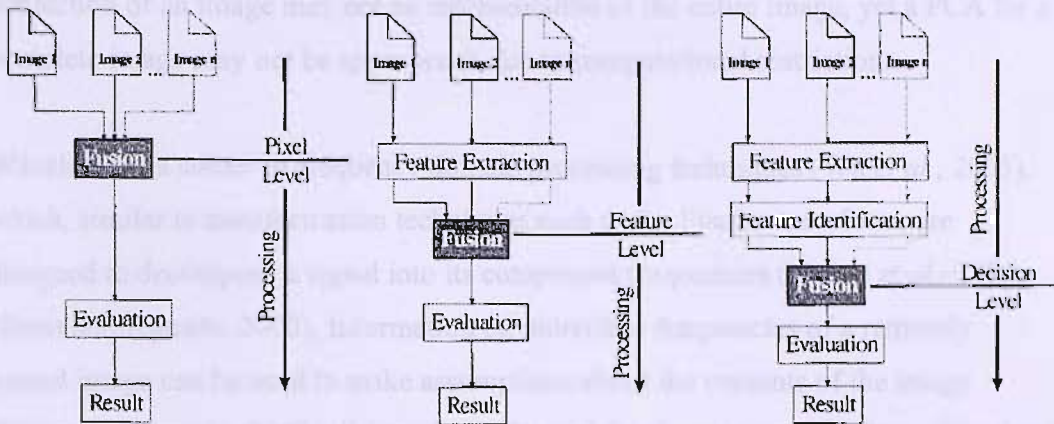


Figure 2.2: The three levels of image fusion (Pohl and Van Genderen, 1998)

One of the common image fusion techniques is Intensity-Hue-Saturation (IHS) (Pohl, 1999; Tu *et al.*, 2001; Chibani and Houacine, 2002a; Simone *et al.*, 2002). In simplest terms, one of the image components (I, H or S) is replaced with a similar component from the image to be fused. Three channels of the input data set are transformed from RGB to IHS colour space (Gillespie *et al.*, 1986; Pohl and Van Genderen, 1998). The resultant image is a linear combination of the input multispectral channels and fused panchromatic channel (Campbell, 1996). IHS is suitable for many applications, although since the initial transformation resamples the image, there can occasionally be a small loss of spectral information.

Principal components analysis (PCA) can be used for image fusion. A standard PCA is often used to identify the optimum linear combination of the original channels that account for maximum variation of pixel values in an image (Campbell, 1996). The linear combination is calculated for each pixel by summing each available channel multiplied by a specified coefficient. The coefficients, estimated by a complex process, account for maximum variation within the data set. In a typical PCA, a set of coefficients is created for each input band (PC1 – PC n). Each set of coefficients yields progressively less information. A simple method for using PCA for fusion is to replace PC1 with the image to be fused (Chavez *et al.*, 1991), which is often referred to as principal component substitution (PCS), since PC1 contains information that is common to all bands of the multispectral image, whereas the other PCs tend to

contain spectral information specific to individual bands. A drawback to the PCA approach is that the output is affected by the data within it, and a PCA for a subsection of an image may not be representative of the entire image, yet a PCA for a complete image may not be appropriate due to computational restrictions.

Wavelets are a subset of frequency domain processing techniques (Wu *et al.*, 2005), which, similar to transformation techniques such as the Fourier transform, are designed to decompose a signal into its component frequencies (Nunez *et al.*, 1999; Nixon and Aguado, 2002). Information on individual frequencies of a remotely sensed image can be used to make assumptions about the contents of the image (Nixon and Aguado, 2002). (Li *et al.*, 2002) used the discrete wavelet frame transform (DWFT) to merge Landsat and SPOT remotely sensed imagery. Input images were decomposed into a DWFT representation at the same spatial resolution. A selected sub-band of the multispectral image was then fused with the remaining sub-bands in the panchromatic image, and an inverse DWFT transform performed to produce a one band “fused” output. Figure 2.3 visualises the methodology:

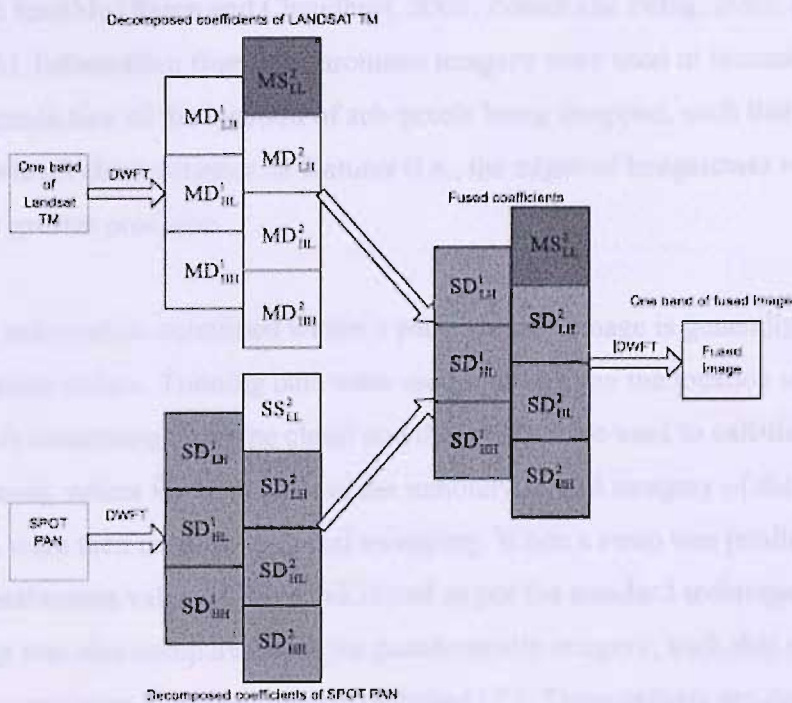


Figure 2.3: The discrete wavelet frame transform

Results from the DWFT fusion technique were demonstrated to often be more accurate when compared with other fusion techniques such as IHS and PCA (Li *et al.*, 2002).

2.5.1 Super-resolution image fusion

The techniques described above represent only a few fusion techniques in common use. Pixel level techniques represent the majority of fusion techniques, although the accuracy of any fusion technique is reliant on the accuracy of geometric co-registration. In this research, a fusion technique was developed and incorporated to increase the accuracy of the super-resolved output of the linearised pixel swapping technique. A summary of the fusion technique is now provided. The technique is described in full in Chapter 6.

The fusion technique used in this research is a pixel level technique. The successful use of such techniques in a super-resolution context has previously been demonstrated to be feasible (Rajan and Chaudhuri, 2002; Zomet and Peleg, 2002; Nguyen *et al.*, 2006). Information from panchromatic imagery were used to increase the accuracy of the prediction of the location of sub-pixels being swapped, such that the prediction of high-level characteristics of features (i.e., the edges of hedgerows) were predicted with greater precision.

The information contained within a panchromatic image is generally given as a set of intensity values. Training data were used to inform on the location of pure pixels (i.e., pixels containing only one class) and these data were used to calculate average class intensity values for each class in the remotely sensed imagery of the study area. These data were then used during pixel swapping. When a swap was predicted, the attractiveness value (A_i) was calculated as per the standard technique. However, the swap was also compared with the panchromatic imagery, such that new criteria for the acceptance of a swap were established (P_i). These criteria are described in full in Chapter 6.

2.6 Accuracy assessment

Accuracy assessment is an important component of any research. Accuracy assessment provides information on the accuracy and validity of the techniques used and feedback on areas of the methodology that may require improvement. In remote sensing land cover classification, accuracy assessment should ideally focus on the extent to which the predicted land cover at a certain point agrees with ground data collected to support the research.

2.6.1 Confusion matrix

There are many different methods for performing accuracy assessment. One of the most commonly used tools is the confusion (or error) matrix. This is a $K \times K$ square array, where K is the number of classes in the image, and is a tabular representation of the relationship between two sets of observations from the classified area – the “actual” ground data (reference) and the prediction. Table 2.1 illustrates an example confusion matrix.

Table 2.1: Example confusion matrix involving four classes

		Predicted				Total
		Grass	Cereal	Hedgerow	Road	
Reference	Grass	29	5	8	1	43
	Cereal	7	46	4	3	60
	Hedgerow	3	4	21	0	28
	Road	1	0	1	28	30
	Total	40	55	34	32	161

From the confusion matrix, it is possible to derive several indicators of “accuracy”. The “overall accuracy” (% correct) is calculated by dividing the sum of the main diagonal by the total number of observations (e.g., $(29+46+21+28)/161=77\%$). In this example, the confusion matrix can be used to indicate that approximately 77% of the

image area is classified correctly, or that there is a 0.77 probability that each pixel is classified correctly. Two further statistics that can be calculated from the matrix are the producer's accuracy (PA based on errors of omission) and the user's accuracy (UA based on errors of commission). The PA informs on how many pixels of a certain class in the ground data are correctly classified, and is calculated by dividing the class cell by the row sum of that class. For example, in the case of grass, the PA is calculated as $29/43=67.45\%$. The UA informs on the probability that a classified pixel actually represents that class on the ground (Tso and Mather, 2001). It is calculated by dividing the class cell value by the column total, for example, in the case of grass, it is $29/40=72.4\%$.

The confusion matrix provides a simple and efficient method of accuracy assessment for hard classification and super-resolution techniques. When comparing two confusion matrices, it would be useful to know if differences in overall accuracy were significant. Significance testing is a well-documented problem in the remote sensing literature and several techniques exist for this purpose. One of the most commonly used techniques is the Kappa coefficient (Campbell, 1996), although its use as a reliable metric of classification accuracy is contentious (Foody, 2002b).

The percentage correct value, provided by a confusion matrix, is often used to provide a simple metric with which to assess the accuracy of a land cover map. In this research, it was necessary to establish percentage correct values that would be deemed, in a qualitative sense, as "accurate" as a means of evaluating the techniques and drawing conclusions. This research does not attempt to provide a methodology for determining appropriate levels of accuracy; instead appropriate levels were derived from empirical testing. For real remotely sensed imagery, in this research, a value of 85% or greater was deemed as accurate. For synthetic imagery, however, where the effect of many real-world problems associated with remote sensing techniques, for example, atmospheric interference, the point spread function or the complexities of a real-world scene, were minimised or non-existent, it was anticipated that reported accuracies would be much greater and, as a consequence, differences in accuracy between the reported accuracies of different versions of the pixel swapping technique would be smaller. Therefore, accuracies of 98% or greater were deemed "accurate". Indeed, establishing a level of acceptable accuracy in the context of the

aims and objectives of this research is actually of little meaning, considering that the ultimate goal is in the development of an algorithm. It is, therefore, intended to provide a simple quantification of the performance of the algorithm. In fact, the most useful accuracy metrics are actually the differences in accuracy between different versions of the technique.

In this research, confusion matrices were constructed to evaluate the accuracy of the pixel swapping technique. It was important to evaluate the differences in reported accuracy between each version of the technique. In particular, the high level of accuracy expected (as described in the previous paragraph), meant that small differences in accuracy might be significant. To evaluate the significance of differences in predicted accuracy, McNemar's test was therefore used (Foody, 2004). The test requires a confusion matrix 2 by 2 in dimension, where the data contained within it are the values of correct and incorrect predictions from each technique to be compared. For example, the confusion matrix in Table 2.1 could be 'collapsed' thus: Correct: 124, Incorrect: 37.

McNemar's test may be expressed as (Agresti, 1996):

$$x^2 = \frac{(|f_{ij} - f_{ji}| - 1)^2}{f_{ij} + f_{ji}} \quad (3)$$

where f_{ij} is the frequency of sites lying in matrix elements i,j and the result is adjusted for continuity. The derived value (x^2) may be compared against values from a chi-square distribution to indicate its significance (Foody, 2004). In most cases, McNemar's test of significance is useful in predicting the likelihood that a difference in accuracy between, for example, two techniques, arose by chance. That is, a value that is significant at a specified level (in this research, a confidence level of 0.05% is used throughout) can be said to be unlikely to arise as a result of chance.

The interpretation of the accuracy metrics provided by the confusion matrices is important in the context of this research. The specified "acceptable" accuracy values are, in essence, arbitrary, and were derived to enable the evaluation of the techniques

and to draw conclusions on the merits of the work presented in this thesis. Indeed, the accuracy metrics used in this thesis describe the accuracy of land cover maps whereas in reality, this research was primarily interested in the delineation of individual features and, in particular, the boundaries of features, something which was difficult to analyse quantitatively. The use of difference images (below) was a useful tool in visualising error on these boundaries. McNemar’s test was applicable only where the confusion matrices were calculated as a sample of a larger population. That is, if the confusion matrices were calculated from the entire population, then it was sufficient to compare the overall accuracy values obtained from the matrix. In this situation, particularly in the case of synthetic imagery where there was an expectation of highly accurate results, the percentage correct value could be considered as both an assessment of the number of correctly allocated sub-pixels as well as the amount of error in the output. For example, if technique A was assessed as being 97% accurate (3% error) and technique B was assessed as being 98% accurate (2% error), an increase in accuracy of 1% might be considered insignificant. However, in this example, a 1% increase in accuracy also represents a decrease in error of approximately 33%, which could be considered to be significant. Where a subsection of a field site was used to create a confusion matrix, then McNemar’s test was applicable. However, it could be argued that the fieldsites used in this research, whilst identified here as discrete (that is, a population), were, in fact, a sample of a wider surrounding area and accordingly, McNemar’s test would be of interest.

2.6.2 Root mean square error (RMSE)

The root mean square error (RMSE) is the square root of the mean squared error, where the “error” is the difference between a predicted value and an observed value. In this research, the RMSE was used to estimate the overall accuracy of a soft classification. It is calculated using Tso and Mather (2001):

$$\text{RMSE} = \sqrt{\frac{\sum_k^K \sum_i^n (z_k(\mathbf{x}_i) - z_k^*(\mathbf{x}_i))^2}{n}} \quad (4)$$

where K is the number of classes, $z_k(\mathbf{x}_i)$ is the proportion of class k at the i^{th} pixel location \mathbf{x}_i of the reference (original) image and $z_k^*(\mathbf{x}_i)$ is the value of the same pixel in the soft classified image. The RMSE is presented as a percentage value and the larger the value, the more error apparent in the soft classification.

2.6.3 Pearson's product-moment correlation coefficient

The correlation coefficient (r) is a measure of how well a linear equation describes the relation between two variables X and Y . In this research, it was used to describe the amount of association between the target and the predicted proportions and informs on the precision of the prediction. It is defined as the result of dividing the covariance between the two variables by the product of their standard deviations. It is calculated per class by:

$$r_k = r_{uv} = \frac{C(uv)}{\sqrt{s^2(u) \times s^2(v)}} = \frac{C(uv)}{s(u) \times s(v)} \quad (5)$$

where $u = z_k(\mathbf{x})$ is the target image for class k , $v = z_k^*(\mathbf{x})$ is the predicted image for class k , $C(uv)$ is the covariance between u and v , s^2 is the variance and s is the standard deviation.

The coefficient ranges from -1 to 1 . A value of 1 shows that a linear equation describes the relationship perfectly and positively, with all data points lying on the same line and with Y increasing with X . A score of -1 shows that all data points lie on a single line but that Y increases as X decreases. A value of 0 shows that a linear model is inappropriate, that is, there is no linear relationship between the variables.

2.6.4 Difference imaging

A useful approach to the visualisation of error in super-resolved output was applied in this study. The accuracy assessment techniques detailed above provide statistical information on error, yet the location of the error is unknown. A *difference image* can be created to highlight the location of error. Where a suitable target image is available (i.e., a map of the actual positions of land cover, often referred to in the literature as a reference image), then the output map and the target map can be compared and combined. In this research, pixels in which there is agreement between both maps are removed, leaving only pixels in erroneous positions.

2.7 Conclusion

Image classification is commonly applied to remotely sensed imagery and is the process of assigning land cover labels to image pixels. These labels can be either derived automatically from the spectral information available (unsupervised classification) or by first inputting training data containing information on the land cover classes of interest (supervised classification). The existence of more than one land cover class in pixels (mixed pixels) can lead to inaccurate classification whereas soft classification techniques can be used to assign pixels to more than one land cover class.

Super-resolution techniques, such as pixel swapping, can be used to predict the spatial location of multiple land cover classes within pixels, using the geostatistical phenomenon of spatial dependence.

The objective of this research was to use super-resolution mapping to produce accurate land cover maps, with the emphasis being on the mapping of linear land cover features such as hedgerows. Therefore, existing linear feature extraction techniques, such as edge detectors, the Hough transform and Snakes, were reviewed to inform on the procedure of identifying linear features in remotely sensed imagery. In addition, image fusion was seen as a possible method of further increasing the accuracy of the resultant maps.

In each stage of the research methodology, accuracy assessments were used to inform on the accuracy of the classification and in validating the newly created land cover map.

The remainder of this thesis presents and analyses the findings of this research. Initially, the data used in this research is presented (Chapter 3). Chapter 4 provides a detailed description of the complete methodology of this research and an evaluation of the pixel swapping technique when applied to remotely sensed imagery. Chapter 5 describes the research carried out to 'linearise' the pixel swapping technique and Chapter 6 presents the final stage of the research: image fusion.

Chapter 3 Data

Chapter 3

Data

The data for this research was collected from a series of focus group discussions with 12 participants, all of whom were experienced teachers in primary schools in the north of England. The participants were selected through purposive sampling to ensure that they had a range of experience and perspectives on the topic. The data was collected over a period of six weeks, with each focus group lasting approximately 45 minutes. The data was then analysed using thematic analysis, which involves identifying themes in the data that are related to the research objectives. The themes identified in this research were: the importance of data in education, the challenges of data collection and analysis, and the role of data in improving teaching and learning. The findings of this research suggest that data is an essential tool for teachers, but that it can be difficult to collect and analyse. It also suggests that data can be used to improve teaching and learning, but that it must be used carefully and ethically.

The data for this research was collected from a series of focus group discussions with 12 participants, all of whom were experienced teachers in primary schools in the north of England. The participants were selected through purposive sampling to ensure that they had a range of experience and perspectives on the topic. The data was collected over a period of six weeks, with each focus group lasting approximately 45 minutes. The data was then analysed using thematic analysis, which involves identifying themes in the data that are related to the research objectives. The themes identified in this research were: the importance of data in education, the challenges of data collection and analysis, and the role of data in improving teaching and learning. The findings of this research suggest that data is an essential tool for teachers, but that it can be difficult to collect and analyse. It also suggests that data can be used to improve teaching and learning, but that it must be used carefully and ethically.

The data for this research was collected from a series of focus group discussions with 12 participants, all of whom were experienced teachers in primary schools in the north of England. The participants were selected through purposive sampling to ensure that they had a range of experience and perspectives on the topic. The data was collected over a period of six weeks, with each focus group lasting approximately 45 minutes. The data was then analysed using thematic analysis, which involves identifying themes in the data that are related to the research objectives. The themes identified in this research were: the importance of data in education, the challenges of data collection and analysis, and the role of data in improving teaching and learning. The findings of this research suggest that data is an essential tool for teachers, but that it can be difficult to collect and analyse. It also suggests that data can be used to improve teaching and learning, but that it must be used carefully and ethically.

The data for this research was collected from a series of focus group discussions with 12 participants, all of whom were experienced teachers in primary schools in the north of England. The participants were selected through purposive sampling to ensure that they had a range of experience and perspectives on the topic. The data was collected over a period of six weeks, with each focus group lasting approximately 45 minutes. The data was then analysed using thematic analysis, which involves identifying themes in the data that are related to the research objectives. The themes identified in this research were: the importance of data in education, the challenges of data collection and analysis, and the role of data in improving teaching and learning. The findings of this research suggest that data is an essential tool for teachers, but that it can be difficult to collect and analyse. It also suggests that data can be used to improve teaching and learning, but that it must be used carefully and ethically.

The data for this research was collected from a series of focus group discussions with 12 participants, all of whom were experienced teachers in primary schools in the north of England. The participants were selected through purposive sampling to ensure that they had a range of experience and perspectives on the topic. The data was collected over a period of six weeks, with each focus group lasting approximately 45 minutes. The data was then analysed using thematic analysis, which involves identifying themes in the data that are related to the research objectives. The themes identified in this research were: the importance of data in education, the challenges of data collection and analysis, and the role of data in improving teaching and learning. The findings of this research suggest that data is an essential tool for teachers, but that it can be difficult to collect and analyse. It also suggests that data can be used to improve teaching and learning, but that it must be used carefully and ethically.

Chapter 3

3. Study Areas and Data

3.1 Introduction

This third chapter of the thesis introduces the study area used in the research. To test and evaluate the super-resolution pixel swapping technique, remotely sensed data were required. There are three common sources of real remotely sensed data: field survey, using, for example, handheld field spectroscopy equipment; airborne sensors onboard aircraft, providing data such as aerial photography and airborne multispectral imagery; and spaceborne satellite sensors. As discussed in Chapter 1, field survey was the foremost source of data during much of the 20th century. Data of this type tend to be highly accurate though their use is often limited by the time required to acquire them, the comparatively small area they tend to cover and the comparatively high cost of their acquisition (Veitch *et al.*, 1995). Aerial survey provides data such as optical remotely sensed imagery considerably faster than field survey, at potentially sub-metre spatial resolutions and in multispectral or hyperspectral format. Satellite-based remotely sensed imagery is available with often global coverage from a range of sensors.

The pixel swapping technique used in this research is intuitively simple and can be applied to any imagery where the proportion of land cover classes within each pixel can be predicted. In this thesis, the technique was applied to both simulated and real remotely sensed imagery.


A ‘decision support system’ was devised to enable the selection of a set of suitable field sites to evaluate the algorithm. In this chapter, the decision support system is first described. Each stage of the system is described, starting with the requirements

of the field sites, followed by the types of imagery available for use in the research. Each of the field sites derived from applying the decision process is presented with a description of the site and supporting imagery. The chapter closes with a summary.

3.2 Field site selection ‘decision support system’

A ‘decision support system’ (DSS) is a useful tool designed to “aid the process of decision making” (Finlay, 1994). A DSS can be used to provide solutions to structured and unstructured problems where input data are both qualitative and quantitative.

The decision support system was used to assist in answering important questions about the method for selection of remotely sensed imagery and field sites for use in this research. Figure 3.1 presents a flow-diagram illustrating the decision support system.



Step 2 : Remotely Sensed Imagery

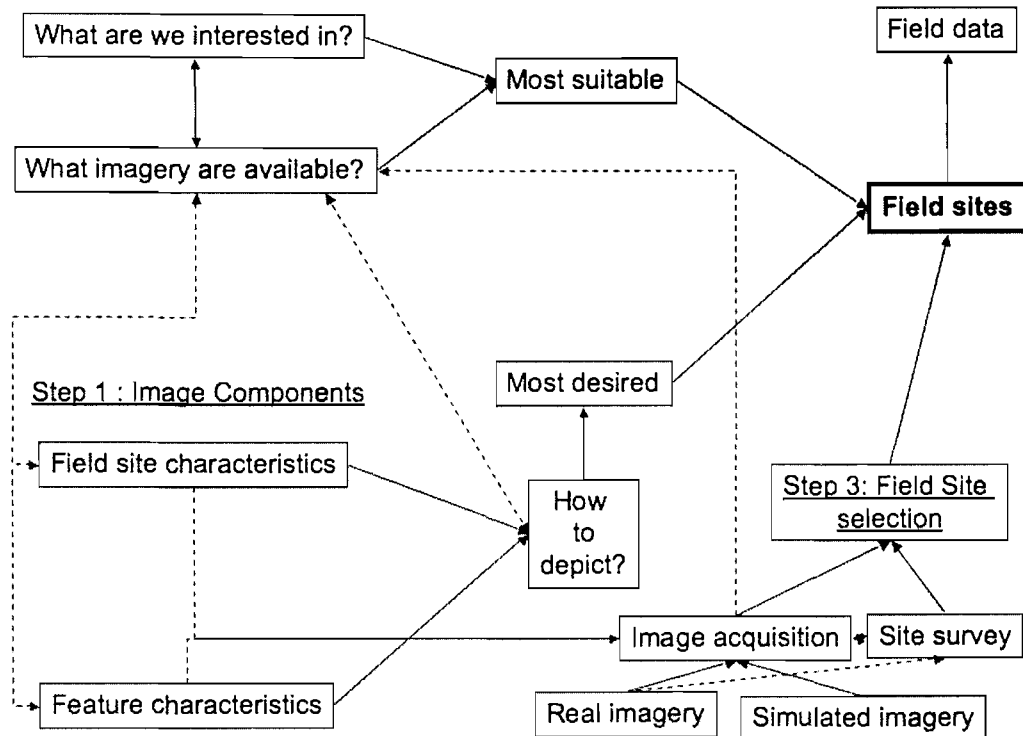


Figure 3.1: Flow diagram illustrating the decision support system. Solid lines represent direction of decision making. Dashed lines illustrate questions dependent on other areas of the system.

These questions were categorised into questions about remotely sensed imagery (constraints), and questions about image components (objectives):

Remotely sensed imagery

- What remotely sensed imagery is available?
- What characteristics of the available remotely sensed imagery are of interest?
- Which of the available remotely sensed images are most appropriate?

Field sites

- What are the important characteristics of a suitable field site?
- How can these characteristics be suitably depicted in a remotely sensed image?
- What features within the field site are of most interest?

- What are the important characteristics of these features?
- How can these characteristics be suitably depicted in a remotely sensed image?
- What affect will the characteristics of the remotely sensed image have on the field site and the features within the field site?

The answers to these questions, presented in further sections of this chapter, were used in devising a strategy for selecting field sites. In answering these questions, it was recognised that individual questions were dependent on questions elsewhere in the system (represented by dashed lines in Figure 3.1).

The following sections present the various stages of the decision process.

3.3 DSS Step 1: Image components

The first component of the decision process (lower half of Figure 3.1) was used to determine the properties of the field sites such that the sites would be suitable for use in this research. Initially, the characteristics of the overall site were considered.

3.3.1 Field site characteristics

Each field site was required to exhibit certain characteristics:

- accessible by car in order to transport bulky field equipment
- have public access (or be easy to obtain permission from the owner of the land)
- unsusceptible to damage as a result of basic field work operations
 - e.g., many of the fields in the study area were planted with cereal crops which in many cases did not provide easy access
- be of a *manageable* size
 - possible to map and measure the area in the field with field equipment
 - acceptable computational requirements

3.3.2 Feature characteristics

In addition to field site characteristics, a set of *feature* characteristics was also established as being requirements of any field site.

3.3.2.1 Width (and size)

Rural linear land cover features, such as hedgerows, tend to be between 1 m and 5 m in width (Baudry and Bunce, 2001), as opposed to urban roads, which tend to be greater than 10 m in width (Mayer, 1999; Quackenbush, 2004). The width of a linear feature is an important aspect of the feature's character. For example, in the absence of class information knowledge of the width of the feature can be used to increase the accuracy with which feature type can be predicted. For instance, ecologists might use the width of the feature (in addition to the length and height) to predict, for example, the species of vegetation within the hedgerow or the suitability of the hedgerow as a habitat for different species of fauna (Gillings and Fuller, 1998; Fuller *et al.*, 2004).

3.3.2.2 Curvature

Rural linear features, such as hedgerows, often exhibit directional variation along their length, for example, where they follow natural boundaries, such as streams and rivers, where they follow human-made linear land cover objects, such as minor roads, or where they have been planted to serve a specific purpose, such as separators between agricultural fields (Boutin *et al.*, 2002; Thenail and Baudry, 2004). The frequency of variation in curvature is diverse, from minor changes in direction (for example, where a feature has been modified to suit modern farming practices (McCollin, 2000) to major changes in direction (for example, where a hedgerow curves around the corners of a field or deviates from a straight line where it follows the natural course of a stream or river). Curvature is an intrinsic component of a feature's character, but is often not represented accurately in land cover maps created by traditional classification techniques (Kristensen, 2001; Franco, 2002; Baudry *et al.*, 2003; Gerd *et al.*, 2006).

3.3.2.3 Embedded objects

Trees and other land cover objects are often found within rural linear features such as hedgerows, which, in addition to breakpoints and localised changes in size (e.g., increased growth in small sections due to spatial variation in soil type, proximity to a water source or the use of fertilizers in adjacent fields), alters the spectral character and width of the feature at points along its length. Trees within hedgerows are a common occurrence and an important requirement of the model.

Variation in all of the above three characteristics was required when selecting a suitable set of field sites. After establishing the requirements of the field sites, it was then necessary to consider the next step of the DSS – remotely sensed imagery.

3.4 DSS Step 2: Remotely sensed imagery

The second component of the decision support system (top half of Figure 3.1) was used to select the remotely sensed imagery for use in this research. Initially, it was necessary to establish the objective of the research in order to determine a starting point. The objective of this research was to develop a super-resolution technique for the identification of linear land cover features from remotely sensed imagery, such as hedgerows and pathways. Using information from Chapter 3.3, a study area and supporting remotely sensed imagery for use in this research were selected.

The Ordnance Survey maintains an extensive library of remotely sensed imagery in a variety of formats. Remotely sensed imagery covering large areas of the United Kingdom were available for selection. The majority of these data were satellite sensor remotely sensed imagery, although in some areas, there were additional data, such as aerial photography. When selecting individual datasets from this library, these data were considered in terms of their suitability with respect to the requirements of the research and the practicality of using the imagery (for example, would the distance to

the location covered by the image inhibit site survey and field work?). Initially, a study area was selected from the available imagery.

3.4.1 Study area

The study area used in this research was in the locale of Christchurch, Dorset in South West England. Christchurch is a small suburban area with a population of approximately 45,000. Christchurch harbour is located on a natural salt marsh which is protected by Mundeford Spit, which together with Hengistbury Head are designated Sites of Special Scientific Interest (SSSI). The land use in the area surrounding Christchurch is primarily agricultural and included planted cereals, fallow fields, woodlands, small dwellings, trees, streams, hedgerows and roads. Much of the land in this area is private property. The dominant crops in the planted fields are wheat, barley and maize. Unplanted fields are used primarily for grazing animals such as sheep and cattle. The hedgerows in the study area are managed, maintained and are the direct responsibility of the owner of the land on which they are planted. Many of the pathways in the study area are Public Rights of Way as determined by the local council.

The study area was chosen primarily for the availability of remotely sensed imagery of the area, as well as the relatively short distance to the study area from Southampton University, compared with other datasets. Visual inspection of the imagery suggested many potential field sites. This was confirmed by visiting the study area. After selecting the study area, remotely sensed imagery was selected.

Remotely sensed imagery was used in two formats: *real imagery*, such as that acquired from a satellite sensor and simulated imagery, that is, synthetic imagery generated to fulfil a specific purpose.

3.4.2 Real imagery

Real remotely sensed imagery of the study area was available in two formats: aerial photography and Quickbird satellite sensor imagery.

3.4.2.1 Aerial photography

Table 3.1 presents the specifications of the colour aerial photography used in this research project, whilst Figure 3.2 presents a sample image.

Table 3.1: Specifications of aerial photography

Spatial resolution	0.25 m
Flying height:	1620 m
Lens focal length	150 mm
Other	Orthorectified with GPS and OSTN97 by Ordnance Survey 3 wavebands (red, blue, green)

Orthorectified aerial photography was used in this study. Orthorectification is a process of correcting distortions in aerial photography caused by terrain relief and camera tilt. This process involves using ground control points (GCPs) and a digital elevation model (DEM) to warp the image so that it is geometrically accurate. The resulting image is then projected onto a map projection, such as the Ordnance Survey National Grid (OSNG), to ensure that it is compatible with other spatial data. The use of orthorectified aerial photography allows for accurate measurement and analysis of land cover and features in the study area.



Figure 3.2: Sample colour aerial photograph of the study area, spatial resolution 25 cm, Christchurch study area

In Figure 3.2, the dominant land use (agricultural) of the study area is clearly visible. Areas of forest (maintained by the Forestry Commission) are also evident. The fine spatial resolution (25 cm) of the aerial photography enabled simple identification of features of interest, such as hedgerows, as well as other land cover objects such as trees and minor roads. In addition, crop growth lines and linear tractor paths (tramlines) are visible in the fields in the top right quarter. There is evidence of some shadowing in the image, particularly around the trees and areas of woodland.

3.4.2.2 Satellite sensor imagery

The satellite sensor imagery used in this research project was Quickbird panchromatic and multispectral imagery. Table 3.2 presents the specifications¹ of the imagery:

Table 3.2: Quickbird multispectral and panchromatic imagery characteristics

	Panchromatic	Multispectral
Spatial resolution	61 cm at nadir	2.44 m – 2.8 m at nadir
Swath	16.5 km (10.3 mi) width at nadir 11-bit digitization (up to 2048 levels of gray scale) Discrete non-overlapping bands	
Scene dimensions	27,552 x 27,424 pixels	6,888 x 6,856 pixels
Spectral characteristics	450 to 900 nm	Blue: 450 to 520 nm Green: 520 to 600 nm Red: 630 to 690 nm Near IR: 760 to 900 nm
Pre-processing	Radiometric corrections: <ul style="list-style-type: none"> • Relative radiometric response between detectors • Non-responsive detector fill • Conversion to absolute radiometry • Internal detector geometry Optical corrections: <ul style="list-style-type: none"> • Optical distortion • Scan distortion • Any line-rate variations • Registration of multispectral bands 	

The multispectral imagery was provided in four discrete wavebands in the same portion of the electromagnetic spectrum as the panchromatic image. The spatial resolution of the panchromatic image was approximately four times finer than the multispectral image. Each of the pre-processing steps (Table 3.2) as well as co-

¹ http://www.digitalglobe.com/product/basic_imagery.shtml

registering of the imagery (meaning that they both represented the same area on the ground) was performed by the image vendor.

In addition to real remotely sensed imagery, synthetic imagery was used.

3.4.3 Simulated imagery

Simulated (or synthetic) imagery were used in this thesis to provide a controlled environment in which to test a specific issue within a remote sensing investigation, where the effect of specific variables can be tested. In most cases, simulated imagery is used to isolate a problem or phenomenon in order to investigate it more thoroughly (Atkinson, 2004a). For example, simulated imagery can be generated to give a realistic yet simplistic representation of a typical remotely sensed image without the problems introduced by factors such as geometric error, atmospheric interference and the point spread function. By using simple simulated imagery, the understanding of how image processing techniques and algorithms operate may be increased and facilitate developments to such techniques. Additionally, the effects of these problems can be investigated by introducing simulations of the problem into the synthetic imagery in controlled amounts. For example, the effect of cloud cover or noise within an image on classification accuracy could be investigated by producing multiple versions of a simulated image, each with a different amount of simulated cloud or noise, and assessing the differences in classification accuracy between each image.

Simulated imagery is particularly suitable for super-resolution applications where the information from a soft classification, such as class proportions, are required. For example, it is possible to generate a synthetic image of a sample scene with two or more simulated land cover classes. By degrading the spatial resolution of this image mixed pixels with a proportion of each class can be created. Additionally, the starting image can be used as a target or “truth” image, if the output of the super-resolution technique matches the spatial resolution of the original simulated imagery. Such images are often free of errors such as misregistration or poor classification (Mertens *et al.*, 2004a). An accuracy assessment tool, such as the difference image, can be used

to visualise error in the super-resolved output by comparing it with the target image. The simulated imagery used in this research is described in the next section.

3.5 Step 3: Field site selection

The final stage in preparing imagery for testing the algorithm was to establish a set of simulated and real field sites with supporting remotely sensed imagery.

Two types of simulated imagery were created: (1) entirely computer-generated synthetic imagery and (2) synthetic imagery simulated from real remotely sensed imagery.

3.5.1 Simulated imagery (1)

Simulated imagery for provisional testing of the algorithm (Chapter 4) was generated using the S-PLUS software. Three images were generated. Images 1 and 2 contained simple shapes ('diamond' and 'line'). The third simulated image was a simplified representation of a typical scene in a classified remotely sensed image, containing four simulated field classes, such as cereals or grass (the actual class labels are not important), and a simple linear feature, such as a hedgerow or pathway. Figure 3.3, Figure 3.4 and Figure 3.5 present these simulated images.




Figure 3.3: Simulated image 1, 'diamond', (dimensions: 250 pixels by 250 pixels, number of classes: 2)




Figure 3.4: Simulated image 2, 'line', (dimensions: 250 pixels by 250 pixels, number of classes: 2)





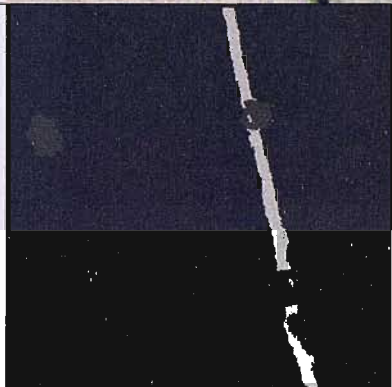
Figure 3.5: Simulated image 3, ‘complex scene’, (dimensions: 250 pixels by 250 pixels, number of classes: 5)

Each simulated image represented simplistic test data with which to evaluate the algorithm.

3.5.2 Simulated imagery (2)

Aerial photography in three bands (red, green and blue) at a spatial resolution of 0.25 m of Burton in the study area acquired on 21st December, 2000 (see section 3.4.2.1) was used to simulate a field site. A sub site (400 x 400 pixels) of the original aerial photography, which contained a simple linear feature: a hedgerow combined with four trees, one distinct from the hedgerow and three embedded within it, was selected. Using RSI ENVI, the imagery was hard classified into four unique classes (hedgerow, trees, cereal and shadow) using the Mahalanobis distance classifier, and the shadow class was removed from the image by merging it with the cereal class. This resulted in a simple 3-category classification of the input imagery (Table 3.3).

Table 3.3: Simulated imagery 2

Site reference	Simulated 2
Dimensions (pixels)	400 x 400
No. classes	3
Raw image (RGB)	
Simulated image (greyscale)	

3.5.3 Real imagery

In addition to simulated imagery, a real remotely sensed image was used. Quickbird multispectral remotely sensed imagery at a spatial resolution of 2.6 m was used as the source.

Quickbird multispectral (and panchromatic) imagery of the study area acquired on June 2nd 2001 was used to establish three field sites. The characteristics of the Quickbird sensor are described in section 3.4.2.2. After an initial visual inspection of the remotely sensed imagery, a list of approximately 10 possible field sites that were thought to meet the field site requirements was produced. A site survey of each possible field site then followed, from which five field sites were chosen. After a

second inspection of the imagery, three field sites from the remaining five were chosen. The sites chosen were deemed most suitable as per the requirements of the DSS.

After selecting three field sites, an image of each field site was produced from Quickbird satellite sensor imagery at a spatial resolution of 2.6 m using RSI ENVI. The field sites were purposefully kept small to minimise computational overheads. Each site contained at least one linear feature, such as a hedgerow. Trees were also included in each site. The field sites are presented in Figure 3.6, Figure 3.7 and Figure 3.8.




Figure 3.6: Real imagery, Site A, (dimensions: 35 pixels by 35 pixels, number of classes: 3)




Figure 3.7: Real imagery, Site B, (dimensions: 46 pixels by 169 pixels, number of classes: 4)



Figure 3.8: Real imagery, Site C, (dimensions: 73 pixels by 73 pixels, number of classes: 4)

Site A represents the most simplistic field site. The site was small with a single linear feature and three land cover classes. Site B was more complex, containing four land cover classes, including a large strip of woodland in the centre of the site. Hedgerows separated each of the fields in the site. A tree is embedded within one of these hedgerows. Site C was the most complex site where there were four land cover classes and a network of three hedgerows. Trees were embedded within each hedgerow and each hedgerow exhibited varying levels of curvature.

3.5.3.1 Field Data

Following field site selection, field data for each site were acquired between June 12th and June 23rd 2003. Since these measurements were eventually not used extensively in this thesis, concise details only are presented in this thesis. Data were acquired during June 2003 to minimise phenological and other seasonal differences between the recorded ground data and the remotely sensed imagery.

On 12th June, 2003, each field site was provisionally surveyed. Digital photographs and a simple drawn schematic with a record of any anomalies or miscellaneous features were taken. The schematic was used to record the type and position of the predominant land cover classes in the site. Positional information on the location of trees and hedgerows within it were collected using a differential GPS.

Following an initial survey of each field site, further measurements in each site were taken. On three separate days (June 14th, June 19th and June 3rd, 2003) the following sampling strategy was applied to each site. Spectral data for each class in the field site were recorded using a Milton Multiband Radiometer (MMR). For each class, spectral data in ten different locations were recorded. A differential GPS was used to record the location of these measurements. For the linear features within the field site, which were hedgerows, additional measurements were taken. The length of the hedgerow was measured with a tape measure along both edges and the location of the end points of the hedgerow were recorded with a GPS. The geometric characteristics of the hedgerows, such as width, height and length (as well as miscellaneous information such as phenological stage or anomalous data) were recorded. At 5 m intervals along the length of the hedgerow, the following measurements were taken:

- Three GPS measurements were taken: at each edge and in the centre. The edges and the centre of the hedgerow were estimated visually.
- The width of the hedgerow at the approximate top and approximate bottom was measured using a solid measuring rod.

- At the centre of the feature, a spectral measurement using an MMR was taken.

Where the hedgerow deviated from its normal course, additional measurements were made so as to characterise the hedgerow fully. Examples of the data are presented in the Appendix.

3.6 Summary

Selection of appropriate remotely sensed imagery for use in any research project is critical to the validity of the results and the conclusions reached as a consequence of these results. This chapter has described the study area and data used in this research and the process through which suitable simulated imagery was created or field sites from real satellite sensor imagery were selected. The purposes and benefits of using both real remotely sensed imagery and simulated remotely sensed imagery were discussed and each of the study areas and field sites were presented.

Chapter 4

Evaluating the super-resolution pixel swapping technique

The proposed super-resolution pixel swapping technique is evaluated using a set of test images. The results are compared with those of the conventional super-resolution techniques. The proposed technique is shown to produce better results than the conventional techniques.

The proposed technique is compared with the conventional techniques.

The proposed technique is compared with the conventional techniques. The results show that the proposed technique is superior to the conventional techniques in terms of the accuracy of the super-resolution results.

The proposed technique is compared with the conventional techniques.

Chapter 4

4. Evaluating the super-resolution pixel swapping technique

4.1 Introduction

The pixel swapping algorithm used in this research was originally developed by Professor Peter Atkinson in the School of Geography at the University of Southampton. It was presented in 2001 as a simple pixel swapping technique at the GeoComputation conference in Brisbane, Australia (Atkinson, 2001) and later appeared as an article in *Photogrammetric Engineering and Remote Sensing* (Atkinson, 2004). The objective of the research presented in this thesis was to develop the pixel swapping technique for the identification of fine rural linear features.

In order to establish the suitability of the pixel swapping technique for the objective of this research, the technique was evaluated by applying it to the simulated imagery and real satellite sensor imagery described in Chapter 3. This chapter presents the results of this preliminary evaluation of the algorithm.

The methodology for evaluating the technique consisted of three stages:

- 1) Soft classify the imagery to create soft proportions and assess the accuracy
- 2) Apply the pixel swapping technique to the soft proportions
- 3) Assess the accuracy of the super-resolved output

The methodology was first applied to simulated imagery; the results of soft classification and pixel swapping are presented and analysed (Section 4.2). The methodology was then applied to real satellite sensor imagery; the results of soft

classification and pixel swapping are presented and analysed (Section 4.3). These analyses are followed by a discussion of the overall results and key findings. The chapter closes with conclusions.

4.2 Simulated Imagery

The first stage of evaluation of the super-resolution pixel swapping technique was to apply the technique to simulated imagery. Three simple features were simulated, as described in Chapter 3.

4.2.1 Soft proportions

Soft class proportions were predicted by initially simulating each image at a fine spatial resolution and then degrading the spatial resolution of the image. Each image was simulated with dimensions of 250 pixels by 250 pixels and then degraded by a factor of five, creating input images with dimensions of 50 pixels by 50 pixels. A zoom factor of five was used in pixel swapping. The principal benefit of this action was that the spatial resolution of the super-resolved output matched the spatial resolution of the original simulated imagery enabling direct comparison between input and output.

4.2.2 Pixel swapping

The super-resolution pixel swapping technique uses the proportions information from a soft classification to predict the locations of each class within the pixels of remotely sensed imagery. A technical description of the technique is presented in Chapter 2 of this thesis. In the case of simulated imagery, soft proportions were derived from simulated mixed pixels, as described in Section 4.2.1, above. The super-resolution pixel swapping technique was applied to each of the three simulated images using a zoom factor of five.

4.2.3 Results

In this section, the results of simulating soft proportions from simulated imagery and then applying the pixel swapping technique are presented, accordingly:


- Simulated image 1: Diamond
 - Soft proportions (Figure 4.1)
 - Accuracy assessment of soft proportions (Table 4.1)
 - Super-resolved output (Figure 4.3)
 - Accuracy assessment of super-resolved output (Table 4.3)
- Simulated image 2: Line
 - Soft proportions (Figure 4.5)
 - Accuracy assessment of soft proportions (Table 4.6)
 - Super-resolved output (Figure 4.7)
 - Accuracy assessment of super-resolved output (Table 4.8)
- Simulated image 3: Complex scene
 - Soft proportions (Figure 4.9)
 - Accuracy assessment of soft proportions (Table 4.11)
 - Super-resolved output (Figure 4.11)
 - Accuracy assessment of super-resolved output (Table 4.13)

4.2.3.1 Simulated image 1: Diamond

Soft proportions in simulated image 1 were created by degrading the spatial resolution of the input imagery by a factor of five. Figure 4.1 displays the predicted soft proportions for each class in the simulated diamond feature.

(a)

(b)



Legend:

100%

50%

0%



Figure 4.1: Soft proportions for feature 1, ‘Diamond’, (a) feature, (b) background

Mixed pixels occur along the boundary between the two classes, that is, along each edge of the diamond feature and are depicted in Figure 4.1 by grey pixels. Table 4.1 shows the predicted accuracy of the soft classification.

Table 4.1: Accuracy assessment, diamond feature

RMSE	r
6.91	0.99

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a high level of accuracy.

A simple maximum likelihood classification was then applied to the diamond feature. The hard classified output is displayed in Figure 4.2.




Figure 4.2: Hard classification, diamond feature

It was evident that the hard classification of the diamond resulted in the edges of the feature taking on a “blocky” appearance.

Table 4.2: Confusion matrix, hard classified diamond feature

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	50211	391	50602	99.2
	Feature	165	11234	11399	98.5
	Totals	50376	11625	61445	
	UA (%)	99.6	96.6	% Correct: 99.1	

Overall accuracy was estimated for the diamond feature to be 99.1%. The pixel swapping technique was then applied to feature 1 using a zoom factor of five. Sub-pixels were allocated randomly and the pixel swapping algorithm ran until it

converged (i.e., no further swaps were made). Figure 4.3 shows the super-resolved output of the pixel swapping technique applied to simulated image 1.




Figure 4.3: Super-resolved output, feature 1, (a) initial allocation, (b) super-resolved output

Figure 4.3(a) was the starting point (random allocation). Sub-pixels within mixed pixels were allocated randomly. Figure 4.3(b) shows the super-resolved output. A visual inspection shows that the output closely resembled the original simulated feature. Table 4.3 presents the predicted accuracy of the super-resolved output for the initial allocation (soft classification) and super-resolved output for the whole image.

Table 4.3: Confusion matrix, simulated image 1, (a) initial allocation, (b) super-resolved output

(a)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	50237	365	50602	99.2
	Feature	365	11034	11399	96.7
	Totals	50602	11399	61271	
	UA (%)	99.2	96.7	% Correct: 98.8	

(b)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	50589	13	50602	99.9
	Feature	13	11386	11399	99.8
	Totals	50602	11399	61975	
	UA (%)	99.9	99.8	% Correct: 99.9	

Overall accuracy for the super-resolved output was estimated as 99.9%. Accuracy assessment of the soft proportions predicted an RMSE of 6.9%, but since the majority of pixels in the image were pure pixels, the effect of mixed pixels on the accuracy was reduced. To assess the effect of the mixed pixels on the accuracy, a confusion matrix (Table 4.4) was calculated for a smaller area (100 sub-pixels by 100 sub-pixels) of the original feature (illustrated by Figure 4.4).




Figure 4.4: Subsection of simulated image 1 used in confusion matrix

Table 4.4: Confusion matrix, subsection of feature 1, (a) hard classification, (b), soft classification, (c), super-resolution

(a)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	4710	150	4860	96.9
	Feature	15	5326	5341	99.7
	Totals	4725	5476	10036	
	UA (%)	99.6	97.2	% Correct: 98.3	

(b)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	4738	122	4860	97.4
	Feature	127	5214	5341	97.6
	Totals	4865	5336	9952	
	UA (%)	97.3	97.7	% Correct: 97.5	

(c)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	4853	7	4860	99.8
	Feature	7	5334	5341	99.8
	Totals	4860	5341	10187	
	UA (%)	99.8	99.8	% Correct: 99.8	

The confusion matrix in Table 4.4 for a subsection of the simulated diamond feature reveals that, as expected, the abundance of pure pixels in the image were reducing the effect of mixed pixels on the overall accuracy. Accordingly, in the subsection of the site, the overall accuracy of the hard classification was estimated as 98.3% compared with 99.8% for the whole image, a significant difference of 1.5% (Table 4.5).

Table 4.5: Significance test, Hard classification & Pixel swapping

		Pixel swapping		
		Correct	Incorrect	Total
Hard classification	Correct	10036	0	10036
	Incorrect	151	14	165
	Total	10187	14	10201

$\chi^2 = 149$. Significant at 0.05%

4.2.3.2 Simulated image 2: Line

Soft proportions in simulated image 2 were created by degrading the spatial resolution of the input imagery by a factor of five. Figure 4.5 displays the predicted soft proportions for each class in the simulated line feature.




Figure 4.5: Soft proportions for feature 2, ‘Line’, (a) feature, (b) background.

Figure 4.5 displays the predicted soft proportions for each class in the simulated line feature. Mixed pixels occur along the boundary between the two classes, that is, each edge of the line. Table 4.6 shows the predicted accuracy of the soft classification.

Table 4.6: Accuracy assessment, simulated image 2

RMSE %	r
6.59	0.99

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a high level of accuracy.

A simple maximum likelihood classification was then applied. The image is displayed in Figure 4.6




Figure 4.6: Hard classification, line feature

A visual inspection of the hard classification revealed that the edges of the feature had taken on a “blocky” appearance. The accuracy assessment of this image is displayed in Table 4.7.

Table 4.7: Confusion matrix, hard classified line feature

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	55476	300	55776	99.4
	Feature	300	5925	6225	95.1
	Totals	55776	6225	61401	
	UA (%)	99.4	95.1	% Correct: 99.0	

Overall accuracy of the hard classification was estimated at 99%. The pixel swapping technique was then applied to feature 2 using a zoom factor of five. Sub-pixels were allocated randomly and the pixel swapping algorithm ran until it converged (i.e., no

further swaps were made). Figure 4.7 shows the super-resolved output of the pixel swapping technique applied to simulated image 2.




Figure 4.7: Super-resolved output, feature 2, (a) initial allocation, (b) super-resolved output

Figure 4.7(a) was the starting point (random allocation). Sub-pixels within mixed pixels were allocated randomly. Figure 4.7(b) shows the super-resolved output. A visual inspection shows that the output closely resembled the original simulated feature. Table 4.8 presents the predicted accuracy of the initial allocation (soft classification) and the super-resolved output.

Table 4.8: Confusion matrices, feature 2, (a) Initial allocation, (b) Super-resolved output

(a)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	55334	442	55776	99.2
	Feature	442	5783	6225	92.8
	Totals	55776	6225	61117	
	UA (%)	99.2	92.8	% Correct: 98.5	

(b)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	55772	4	55776	99.9
	Feature	3	6222	6225	99.9
	Totals	55775	6226	61994	
	UA (%)	99.9	99.9	% Correct: 99.9	

Overall accuracy for the super-resolved output was predicted as 99.9%. To assess the effect of the mixed pixels on the accuracy, a confusion matrix (Table 4.9) was calculated for a smaller area (100 pixels by 100 pixels) of the feature, as illustrated by Figure 4.8.

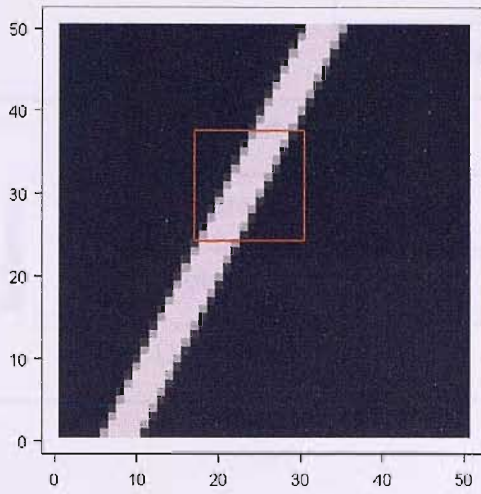


Figure 4.8: Subsection of feature 2 used in confusion matrix

Table 4.9: Confusion matrix, subsection of feature 2, (a) Hard classification, (b) Initial random allocation, (c) Super-resolved output

(a)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	7556	120	7676	98.4
	Feature	120	2405	2525	95.2
	Totals	7676	2525	9961	
	UA (%)	98.4	95.2	% Correct: 97.6	

(b)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	7500	176	7676	97.7
	Feature	177	2348	2525	92.9
	Totals	7677	2524	9848	
	UA (%)	97.6	93.0	% Correct: 96.5	

(c)

		Predicted			
		Background	Feature	Totals	PA (%)
Reference	Background	7676	0	7676	100
	Feature	0	2525	2525	100
	Totals	7676	2525	10201	
	UA (%)	100	100	% Correct: 100	

The confusion matrix in Table 4.9 for a subsection of the simulated diamond feature reveals that the abundance of pure pixels in the image were, as expected, reducing the effect of mixed pixels. Accordingly, in the subsection of the site, the overall accuracy for the hard classification was predicted at 96.5% compared with a perfect allocation of 100% for the same subsection image in the pixel swapping technique. McNemar's test was used to evaluate the significance of these results (Table 4.10).

Table 4.10: Significance test, hard classification & pixel swapping

		Pixel swapping		
		Correct	Incorrect	Total
Hard classification	Correct	9961	0	9961
	Incorrect	240	0	240
	Total	10201	0	10201

$\chi^2=242$. Significant at 0.05%.

4.2.3.3 Simulated image 3: Complex scene

Soft proportions in feature 3 were created by degrading the spatial resolution of the input imagery by a factor of five. Figure 4.9 displays the predicted soft proportions for each class in the simulated complex scene.






Figure 4.9 displays the predicted soft proportions for each class in the simulated complex scene. Mixed pixels occur along the boundary between classes that lie along the boundaries of the "grass" and on each edge of the "sidewalk". Table 4.11 shows the predicted accuracy of the soft classification.

(c) (d)



(e)



Legend:



Figure 4.9: Class proportion images for feature 3: (a) “field 1”, (b) “field 2”, (c) “field 3”, (d) “field 4”, (e) “hedgerow”. Units: pixels.

Figure 4.9 displays the predicted soft proportions for each class in the simulated complex scene. Mixed pixels occur along the boundary between classes, that is, along the boundaries of the “fields” and on each edge of the “hedgerow”. Table 4.11 shows the predicted accuracy of the soft classification.

Table 4.11: Accuracy assessment, feature 3

	RMSE %	<i>r</i>
Field 1	7.84	0.99
Field 2	6.01	0.99
Field 3	6.91	0.99
Field 4	6.12	0.99
Hedgerow	8.16	0.98

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a high level of accuracy. The greatest error was predicted in the simulated hedgerow class. This can be explained due to the hedgerow class having the fewest pixels (i.e., a “thin” hedgerow) and therefore when the spatial resolution of the original imagery was degraded, more mixed pixels in the hedgerow class were created. A maximum likelihood classification was applied to this image, as displayed in Figure 4.10.




Figure 4.10: Hard classification, complex scene

The “hedgerow” feature was predicted by the hard classification as a “blocky” feature. Accuracy assessments of the hard classification are presented in Table 4.11.

Table 4.12: Confusion matrix, hard classification, complex scene

		Predicted						
		Field 1	Field 2	Field 3	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	8950	0	0	0	136	9086	98.5
	Field 2	105	14270	0	0	94	14469	98.6
	Field 3	0	0	6440	204	15	6659	96.7
	Field 4	0	0	0	30427	115	30542	99.6
	Hedgerow	70	60	30	200	885	1245	71.0
	Totals	9125	14330	6470	30831	1245	60972	
	UA (%)	98.0	99.5	99.5	98.6	71.0	% Correct: 98.3	

Overall accuracy for the hard classification techniques was estimated at 98.3%.

However, the hedgerow class was predicted considerably less accurately than each of the other classes. The pixel swapping technique was applied to feature 3 using a zoom factor of five. Sub-pixels were allocated randomly and the pixel swapping algorithm ran until it converged (i.e., no further swaps were made). Figure 4.11 shows the super-resolved output of the pixel swapping technique applied to the third simulated image (complex scene).

(a) Initial allocation, random allocation, 255 super-resolved output (b) Super-resolved output, 255 super-resolved output




Figure 4.11: Super-resolved output, feature 3, (a) initial allocation, (b) super-resolved output

Figure 4.11(a) was the starting point (random allocation). Sub-pixels within mixed pixels were allocated randomly. Figure 4.11(b) shows the super-resolved output. A visual inspection suggests that the output closely resembled the original simulated feature. Table 4.13 presents the predicted accuracy of the initial allocation and the super-resolved output.

	Field 1	Field 2	Field 3	Field 4	Background	Total	% Correct
Field 1	2881	1	0	0	0	2882	99.9
Field 2	1	1002	0	0	0	1003	99.9
Field 3	0	0	1007	1	0	1008	99.9
Field 4	0	0	0	1033	0	1033	99.9
Background	0	1	1	2	1237	1241	99.9
Total	2882	1003	1008	1033	1241	6167	
Accuracy (%)	99.9	99.9	99.9	99.9	99.9	% Correct	99.9

Table 4.13: Confusion matrices, feature 3, (a) Initial allocation, (b) Super-resolved output

(a)

		Predicted						
		Field 1	Field 2	Field 3	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	8855	88	0	0	143	9086	97.4
	Field 2	87	14279	0	1	102	14469	98.6
	Field 3	0	0	6454	171	34	6659	96.9
	Field 4	2	1	171	30148	220	30542	98.7
	Hedgerow	142	103	33	219	748	1245	60.0
	Totals	9086	14471	6658	30539	1247	60484	
	UA (%)	97.4	98.6	96.9	98.7	59.9	% Correct: 97.5	

(b)

		Predicted						
		Field 1	Field 2	Field 3	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	9081	1	0	0	4	9086	99.9
	Field 2	1	14468	0	0	0	14469	99.9
	Field 3	0	0	6657	1	1	6659	99.9
	Field 4	0	0	0	30538	4	30542	99.9
	Hedgerow	4	1	1	2	1237	1245	99.3
	Totals	9086	14470	6658	30541	1246	61981	
	UA (%)	99.9	99.9	99.9	99.9	99.2	% Correct: 99.9	

Table 4.13 shows that the overall accuracy of the super-resolved output for the complex scene was estimated at 99.9%. This was an increase of 1.5% over the hard classification. However, attention should be drawn to the accuracy of the prediction of the hedgerow class. In the hard classification, accuracy of this class was estimated at 71.0%. In the super-resolved output, the accuracy was estimated as 99.3% a positive difference of 28.3%. Accuracy assessments were also calculated for a 100 sub-pixel by 100 sub-pixel subsection of the complex scene, which contained samples of the field 1, field 4 and hedgerow classes. These accuracy assessments are displayed in Table 4.14.

Table 4.14: Confusion matrices, subsection of complex scene

(a)

		Predicted				
		Field 1	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	3362	0	92	3454	97.3
	Field 4	0	6189	53	6242	99.1
	Hedgerow	53	92	360	505	71.2
	Totals	3415	6281	505	9911	
	UA (%)	98.4	98.5	71.2	% Correct: 97.1	

(b)

		Predicted				
		Field 1	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	3335	18	101	3454	96.5
	Field 4	2	6142	98	6242	98.3
	Hedgerow	101	100	304	505	60.1
	Totals	3438	6260	503	9780	
	UA (%)	97.0	98.3	60.4	% Correct: 99.9	

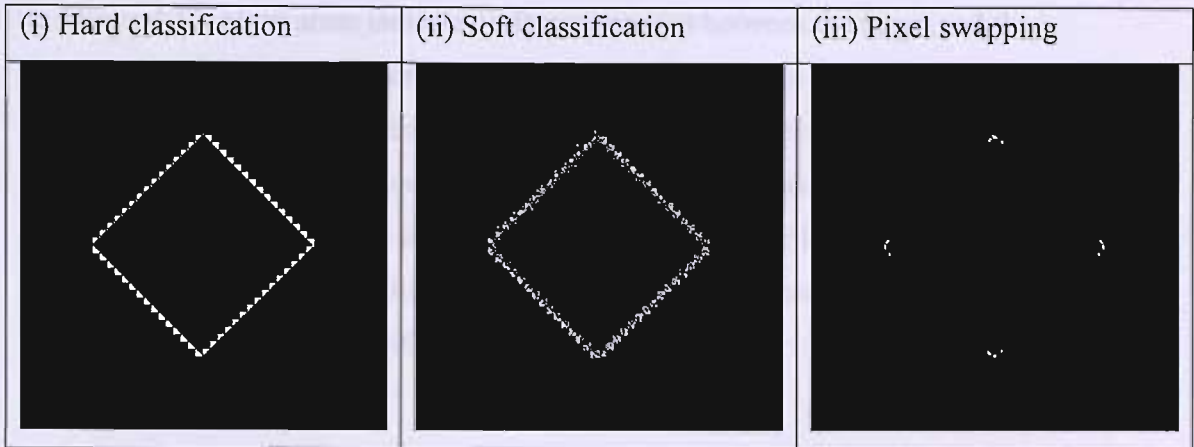
(c)

		Predicted				
		Field 1	Field 4	Hedgerow	Totals	PA (%)
Reference	Field 1	3453	0	1	3454	99.9
	Field 4	0	6241	1	6242	99.9
	Hedgerow	1	1	503	505	99.6
	Totals	3454	6242	505	10197	
	UA (%)	99.9	99.9	99.6	% Correct: 99.9	


Accuracy assessments of the three techniques to the input imagery revealed results very similar to those of the complete scene. This was not unexpected, however, as the number of mixed pixels in the subsection was still low relative to the number of pure pixels.

In addition to the confusion matrices, difference images were constructed to help visualise the error in each of the outputs. These images are displayed in Figure 4.12.

(a)



(b)



(c)




Figure 4.12: Difference images, (a) Diamond shape, (b) Soft classification, (c) Pixel swapping

In Figure 4.12, black areas indicate perfect agreement between the target and the prediction. Misclassification existed along the edges of classes at the points where classes intersected neighbouring classes, as depicted by white and grey pixels. Small areas of misclassification also existed at the topmost and bottommost ends of the hedgerow class. In each case, it was evident that considerably less error was apparent in the super-resolved output from the pixel swapping technique than either the hard classification or the soft classification.

4.2.4 Discussion of results

Degrading the spatial resolution of the simulated imagery resulted in simple features with mixed pixels. Mixed pixels (i.e., pixels with proportions of two or more classes) all occurred on or adjacent to boundaries between classes, i.e., along the edges of the diamond and line features, and between the simulated field and hedgerow classes in the third set. This was an expected outcome, since degrading the spatial resolution of a pure pixel entirely surrounded by pure pixels will always result in a pure pixel at the degraded spatial resolution. It is only where the window in the degradation process (in this case, 5 pixels by 5 pixels) covered multiple classes that mixed pixels (and accordingly soft proportions) result in the output image at the degraded spatial resolution.

Accuracy assessment of the resultant images of degrading the spatial resolution of the original simulated imagery showed very small amounts of error. In each of the three features, the RMSE was between 6% and 8%, with the exception of the hedgerow class in feature 3, where the RMSE was 8.14%. In the same class, the correlation coefficient was 0.98, and in all other classes in each feature, the correlation coefficient was 0.99. The small amounts of error in the images were not unexpected. The error occurs due to the images being reasonably small (250 pixels on each axis degraded to 50 pixels on each axis), yet the geometry of the features contained within them caused quite dense mixing within pixels (i.e., in the two class examples, proportions of 60 and 40 were often output as opposed to, for example, 90 and 10).

Since the number of pixels with proportions was low relative to the number of pure pixels, the effect of densely mixed pixels on the accuracy assessment was more acute.

Confusion matrices which were created for the soft classifications of each of the simulated images, were displayed to illustrate the benefit of using super-resolution techniques. As described in previous chapters of this thesis, soft classification does not provide any information on the locations of proportions in pixels. Therefore, the random allocation was used. In each case, the overall accuracy of the random allocation was less than that of the hard classification. This was, of course, not unexpected, since it was a random allocation. However, it was demonstrated that the pixel swapping technique could be used to predict the proper locations of the randomly allocated sub-pixels and increase the accuracy of the resultant super-resolved map.

In each of the three simulated images, the predictions made using the pixel swapping technique were very accurate. In most cases, the accuracy of the super-resolved output was 99%. In some cases, the accuracy was reported as 100%. The areas of error in the super-resolved output of each of the sites were around the edges of the classes, for example, on vertices of the diamond feature, the outermost edges of the line feature, and on the boundaries of the simulated field and hedgerow classes in feature 3. In each case, though, the super-resolved output matched the input imagery with very little error. When the accuracy of the technique was assessed in subsections of the simulated features, the accuracy of the hard classification technique was less than the accuracy for the whole image, whereas, in most cases, the accuracy of the super-resolved output was comparable to that of the whole image, indicating that the effect of mixed pixels on the hard classification was significant. That is, the pixel swapping technique demonstrated an ability to map classes within mixed pixels with greater accuracy than hard classification techniques. Difference images of each of the features, revealed the areas in which misclassification occurred. As expected, these areas were at the boundaries of classes, where mixed pixels were created in the process of degrading the spatial resolution of the original imagery.

On the basis of these results, the standard pixel swapping algorithm displayed great potential for mapping land cover classes in remotely sensed imagery. Mixed pixels

were mapped more accurately than in a traditional hard classification. The simulated imagery used in this initial evaluation identified key parts of features that would require consideration in development of the algorithm (i.e., boundaries of features, in particular, boundaries that changed direction, e.g., the four points of the diamond shape that displayed most error). The simulated imagery were simplistic and represented a useful set of test cases, without many of the complications of using real remotely sensed imagery, such as atmospheric interference or the point-spread function. Mixed pixels created by degrading the spatial resolution of the input imagery were mixed perfectly linearly, which simplified the super-resolution task. In reality, however, such perfect mixing is rarely the case. Therefore, in order to fully evaluate the pixel swapping algorithm for sub-pixel mapping, it was applied to real remotely sensed imagery.

4.3 Real imagery

The second stage of testing the pixel swapping algorithm with real imagery is now presented. The imagery used in this section is described in Section 3.4.2 of this thesis.

4.3.1 Soft classification

Quickbird imagery at a spatial resolution of 2.6 m of each of the three field sites in the Christchurch study area were soft classified to provide soft proportions for input to the pixel swapping algorithm. Two soft classification methods were used: the fuzzy *c*-means (FCM) and the linear mixture model (MM), as described in Section 2.2.2.1 and section 2.2.2.2 of this thesis. In the case of the FCM, a parameter known as the fuzzy exponent, that determines the fuzziness of the output, was set at 2. Results could, therefore, be expected to fall between minimum and maximum entropy. Values between 1.5 and 2.5 were evaluated, but did not lead to an increase in accuracy. A value of 2 is commonly found in the literature, and so was adopted in this thesis. Accuracy assessment was applied to the soft classification output, as described in

Section 2.6 of this thesis. Information on the location of pure pixels, for use within soft classification, was collected during field survey.

Soft classification results are presented as a set of greyscale images – one for each land cover class in the field site. These images are displayed from the FCM only. In these images, white represents areas completely covered by a particular class (100%); black represents zero proportion of a particular class. Greys, therefore, indicate areas of mixing between classes. Accuracy assessments of each soft classification are presented in tables showing the RMSE and correlation coefficient for each of the soft classification methods.

4.3.2 Pixel swapping

The super-resolution pixel swapping method uses the soft proportions from a soft classification to map the locations of classes within remotely sensed imagery. A technical description of the technique is presented in Chapter 2 of this thesis. At this stage, information on soft proportions were derived from two soft classification techniques, as described in Section 2.2.2 above. The super-resolution pixel swapping technique was applied to each of the three field sites using a zoom factor of five. The effect of the spatial resolution and the zoom factor were explored as part of this research (the results are presented later in this thesis). For this initial evaluation sub-pixels were allocated randomly and the pixel swapping algorithm iterated until it converged (i.e., no further swaps were made).

In this stage of the research, where the objective was to evaluate the potential of super-resolution pixel swapping for the identification of fine linear features, confusion matrices were constructed to assess the accuracy of the linear feature (i.e., hedgerow) class against all other classes.

4.3.3 Results

In this section, the results of simulating soft proportions from simulated imagery and then applying to the pixel swapping technique are presented, accordingly:

- Field site A
 - Soft proportions (Figure 4.13)
 - Accuracy assessment of soft proportions (Table 4.15)
 - Super-resolved output (Figure 4.14)
 - Accuracy assessment of super-resolved output (Table 4.16)
- Field site B
 - Soft proportions (Figure 4.15)
 - Accuracy assessment of soft proportions (Table 4.17)
 - Super-resolved output (Figure 4.16)
 - Accuracy assessment of super-resolved output (Table 4.18)
- Field site C
 - Soft proportions (Figure 4.17)
 - Accuracy assessment of soft proportions (Table 4.19)
 - Super-resolved output (Figure 4.18)
 - Accuracy assessment of super-resolved output (Table 4.20)

4.3.3.1 Field site 1 (Site ref: A)

Field site 1 was soft classified with the FCM and MM. Figure 4.13 displays the predicted soft proportions for each class in field site 1.

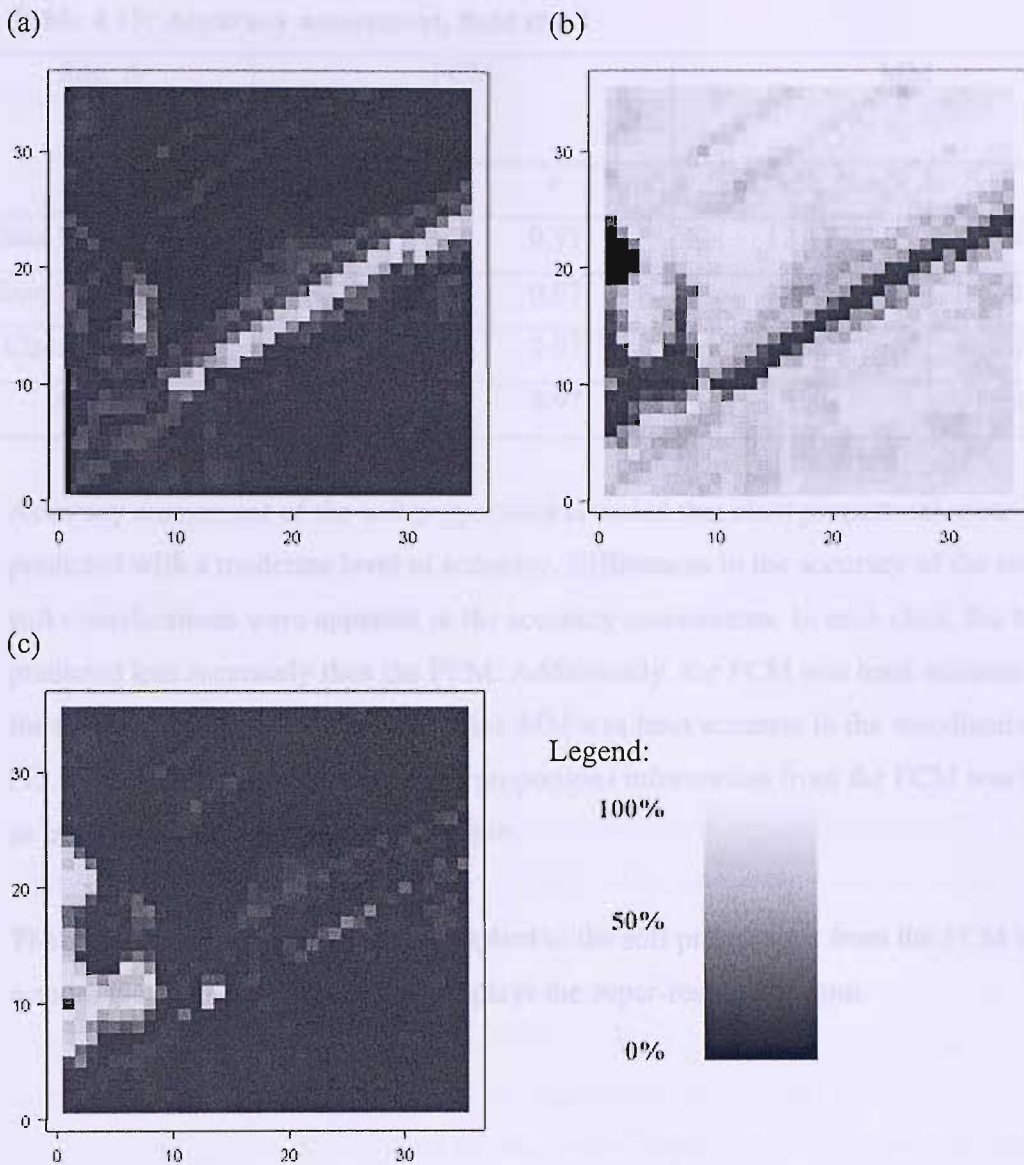


Figure 4.13: FCM Proportion images for field site 1: (a) “hedgerow”, (b) “cereal”, (c) “woodland”. Units: percentage of a pixel’s area.

Mixed pixels occurred within each land cover class, as depicted by grey pixels. Table 4.15 shows the predicted accuracy of the soft classification.

Table 4.15: Accuracy assessment, field site 1

Site: A	FCM		MM	
	RMSE %	<i>r</i>	RMSE %	<i>r</i>
Class 1 (hedgerow)	11.55	0.97	12.17	0.97
Class 2 (woodland)	11.30	0.97	15.45	0.95
Class 3 (cereal)	12.36	0.97	13.97	0.96
Average	11.73	0.97	13.86	0.96

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a moderate level of accuracy. Differences in the accuracy of the two soft classifications were apparent in the accuracy assessments. In each class, the MM predicted less accurately than the FCM. Additionally, the FCM was least accurate in the cereal class (12.36%) whereas as the MM was least accurate in the woodland class (15.45%). For the above reasons, the proportions information from the FCM was used as input to the pixel swapping technique.

The pixel swapping technique was applied to the soft proportions from the FCM using a zoom factor of five. Figure 4.14 displays the super-resolved output.

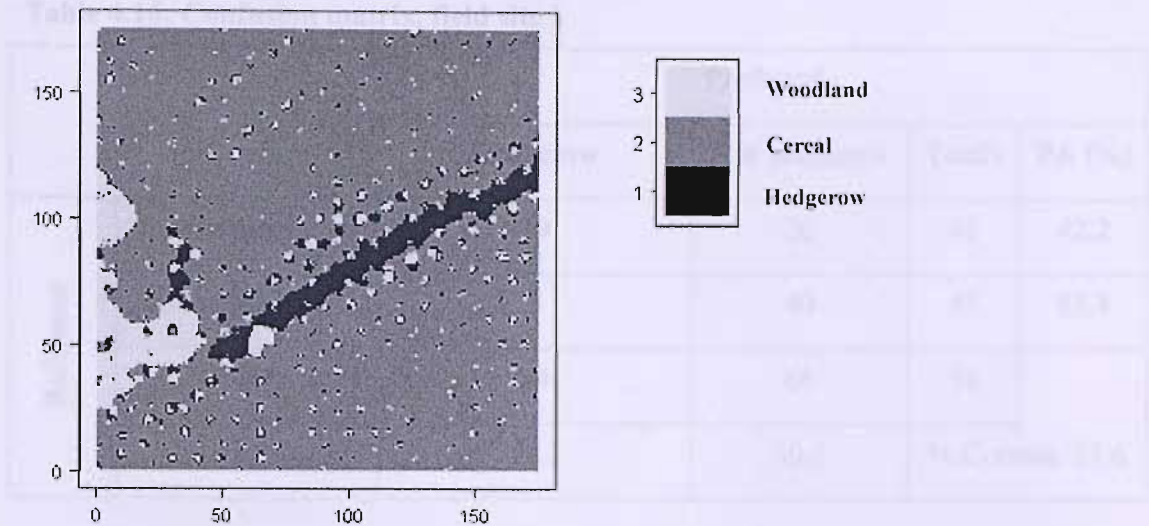


Figure 4.14: Super-resolved output, field site 1

A visual inspection reveals that the super-resolved output generally reflects the scene depicted in the original imagery. The hedgerow class was predicted as a continuous feature and the shape of the woodland class resembled the shape of the feature in the original imagery. However, misclassification was clearly evident throughout the image in the form of discrete blocks of error. Table 4.16 presents the accuracy assessment of the hedgerow class, where training data (reference) were used to produce a confusion matrix against the prediction (super-resolved output). Overall accuracy was predicted at 65.6%.


Table 4.16: Confusion matrix, field site 1

		Predicted			
		Hedgerow	Not hedgerow	Totals	PA (%)
Reference	Hedgerow	19	26	45	42.2
	Not hedgerow	5	40	45	88.8
	Totals	24	66	59	
	UA (%)	79.2	60.6	% Correct: 65.6	


4.3.3.2 Field site 2 (Site ref: B)

Field site 2 was soft classified with the FCM and MM. Figure 4.15 displays the predicted soft proportions for each class in field site 2.

(a)



(b)



(c)

(d)

Table 4.17: Accuracy assessment, 2004 data

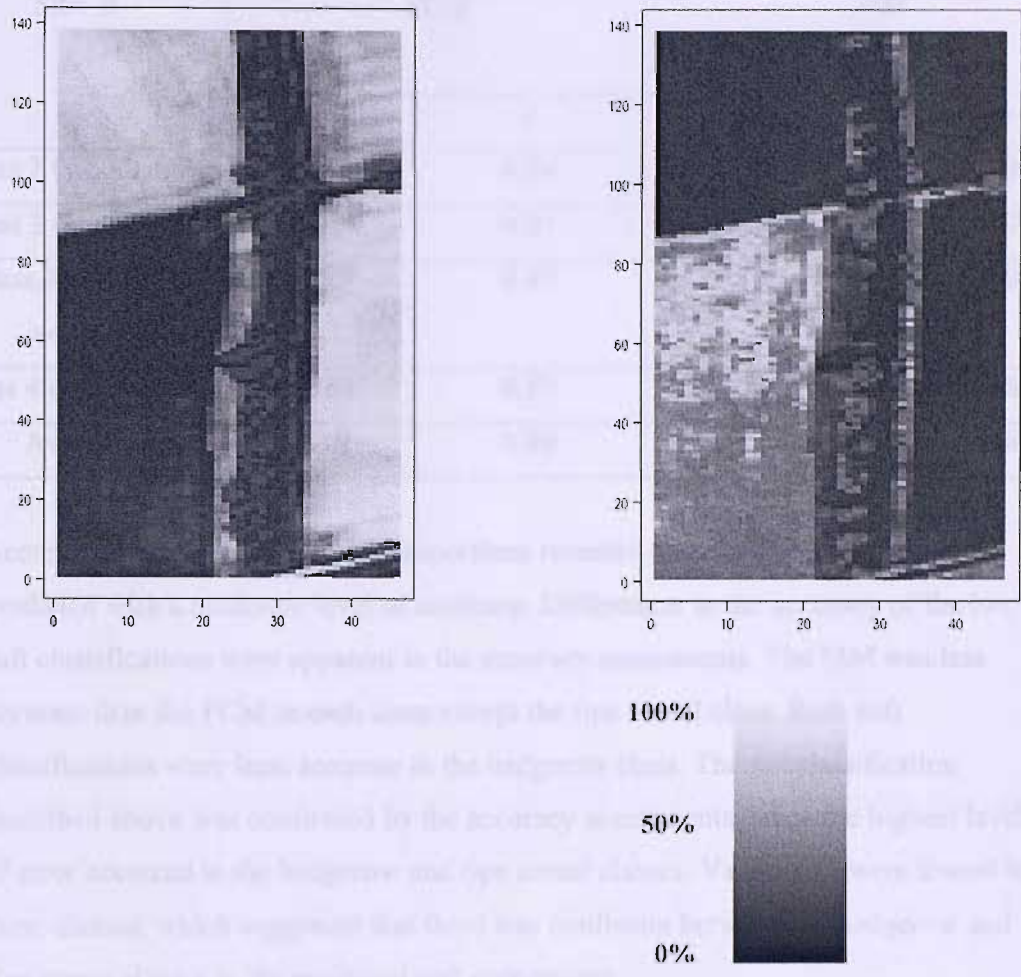


Figure 4.15: Proportion images for field site 2: (a) “woodland”, (b) “hedgerow”, (c) “non-ripe cereal”, (d) “ripe cereal”. Units: percentage of a pixel’s area.

In Figure 4.15, mixed pixels occurred within each land cover class, as depicted by grey pixels. Mixed pixels were most obvious within the ripe cereal class (Figure 4.15(d)). Misclassification was also apparent, particularly in the hedgerow class (Figure 4.15(b)), where light grey pixels in the ripe cereal field (lower-left corner of the image), suggest some confusion between the hedgerow and ripe cereal classes. Table 4.17 shows the predicted accuracy of the soft classification.

Table 4.17: Accuracy assessment, field site 2

Site: B	FCM		MM	
	RMSE %	<i>r</i>	RMSE %	<i>r</i>
Class 1 (woodland)	14.91	0.94	19.31	0.91
Class 2 (hedgerow)	25.95	0.81	28.08	0.76
Class 3 (non-ripe cereal)	6.56	0.99	15.81	0.96
Class 4 (ripe cereal)	24.64	0.83	23.36	0.85
Average	18.01	0.89	21.64	0.87

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a moderate level of accuracy. Differences in the accuracy of the two soft classifications were apparent in the accuracy assessments. The MM was less accurate than the FCM in each class except the ripe cereal class. Both soft classifications were least accurate in the hedgerow class. The misclassification described above was confirmed by the accuracy assessments, since the highest levels of error occurred in the hedgerow and ripe cereal classes. Values of *r* were lowest in these classes, which suggested that there was confusion between the hedgerow and ripe cereal classes in the predicted soft proportions.

The pixel swapping technique was then applied to the soft proportions from the FCM using a zoom factor of five. Figure 4.16 displays the super-resolved output.

Table 4.18 Confusion matrix, field site 2

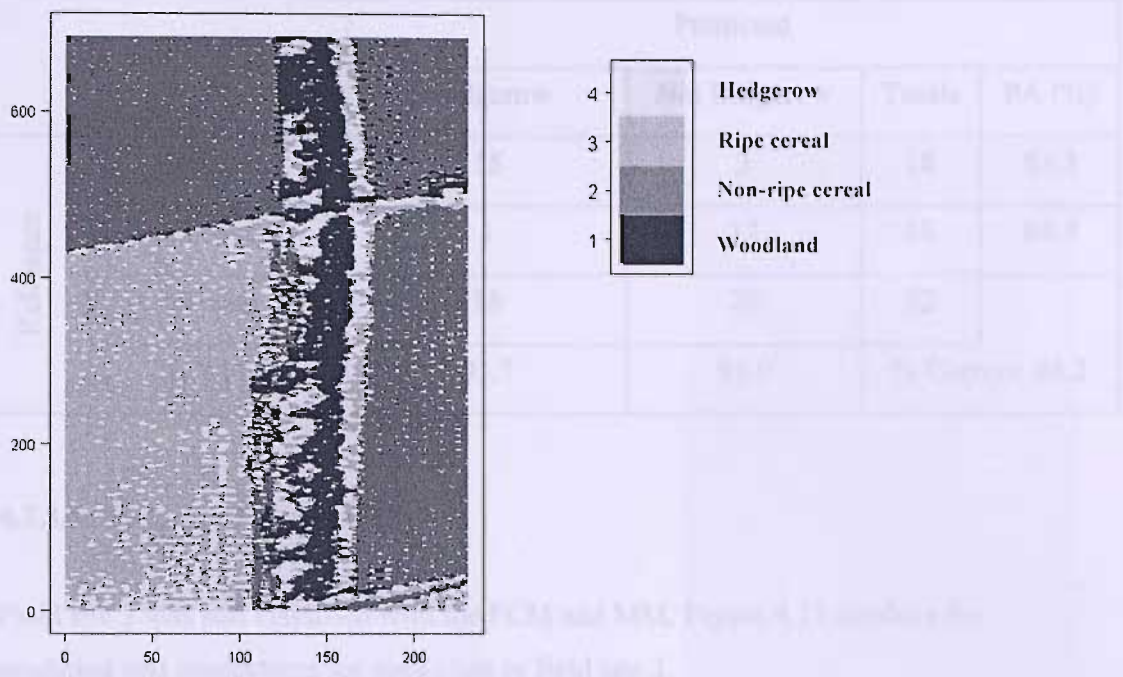


Figure 4.16: Super-resolved output, field site 2

A visual inspection reveals that the super-resolved output generally reflects the scene depicted in the original imagery. The hedgerow class was predicted as a continuous feature and the shape of the woodland class resembled the shape of the feature in the original imagery. Misclassification was clearly distinguishable throughout the image in the form of discrete blocks of error. Table 4.18 presents the accuracy assessment of the hedgerow class, where training data (reference) were used to produce a confusion matrix against the prediction (super-resolved output). Overall accuracy was estimated at 84.2%.

Table 4.18: Confusion matrix, field site 2


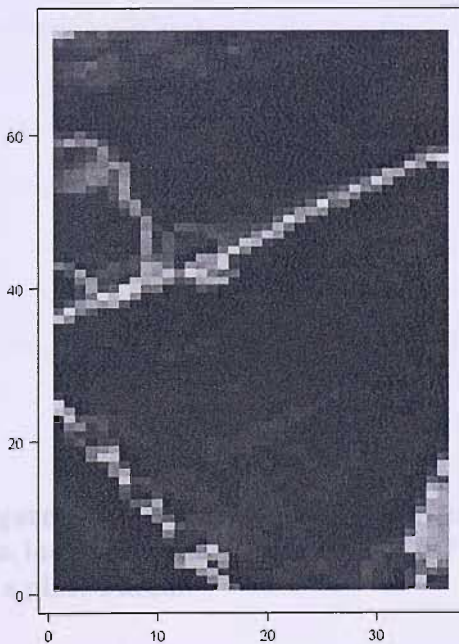
		Predicted			
		Hedgerow	Not hedgerow	Totals	PA (%)
Reference	Hedgerow	15	3	18	83.3
	Not hedgerow	1	17	18	94.4
	Totals	16	20	32	
	UA (%)	93.7	85.0	% Correct: 84.2	

4.3.3.3 Field site 3 (Site ref: C)


Field site 3 was soft classified with the FCM and MM. Figure 4.17 displays the predicted soft proportions for each class in field site 3.

(a)

(b)



(c)



(d)





Figure 4.17: Proportion images for field site 3: (a) “hedgerow”, (b) “woodland” (i.e. individual trees), (c) “ripe cereal”, (d) “non-ripe cereal”. Units: percentage of a pixel’s area.

In Figure 4.17, mixed pixels occurred within each land cover class, as depicted by grey pixels. Mixed pixels were most obvious within the ripe-cereal class (Figure 4.17(c)) where misclassification was apparent due to the existence of grey pixels in each of the other proportion images in the area of the image measured as non-ripe

cereal on the ground. Table 4.19 shows the predicted accuracy of the soft classification.

Table 4.19: Accuracy assessment, field site 3

Site: C	FCM		MM	
	RMSE %	<i>r</i>	RMSE %	<i>r</i>
Class 1 (hedgerow)	18.15	0.90	38.03	0.33
Class 2 (woodland)	20.03	0.88	18.84	0.89
Class 3 (ripe cereal)	12.21	0.95	12.71	0.95
Class 4 (non-ripe cereal)	10.63	0.97	16.80	0.93

Accuracy assessment of the soft proportions revealed that class proportions were predicted with a moderate level of accuracy. Differences in the accuracy of the two soft classifications were apparent in the accuracy assessments. In each class, the MM was less accurate than the FCM. The FCM was least accurate in the woodland class, whereas the MM was least accurate in the hedgerow class. With the exception of the non-ripe cereal class, in this field site, the MM was substantially less accurate than the FCM, ranging from 6% more error in the non-ripe cereal class to 20% more error in the hedgerow class.

The pixel swapping technique was then applied to the soft proportions from the FCM. Figure 4.18 displays the super-resolved output.

Table 4.20: Confusion matrix, field site 3

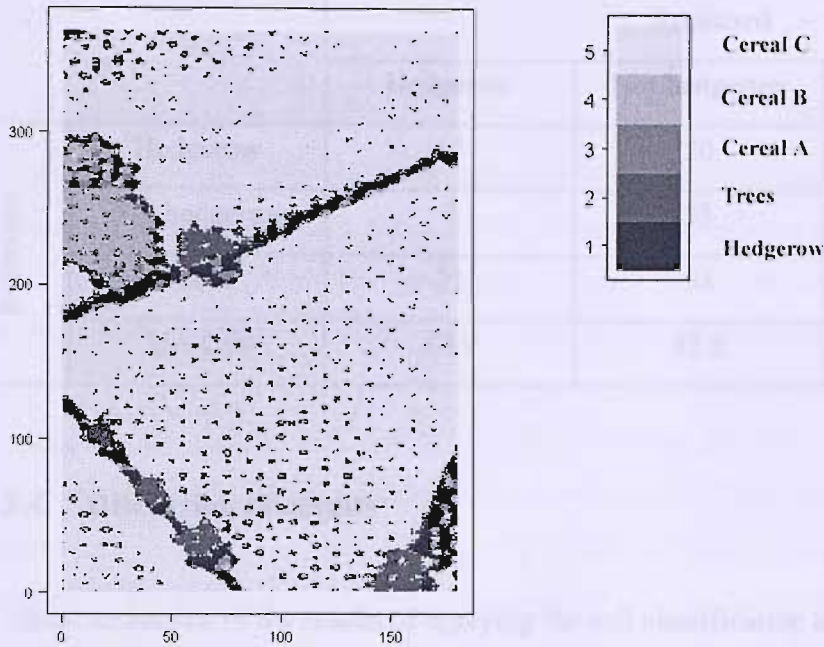


Figure 4.18: Super-resolved output, field site 3

A visual inspection reveals that the super-resolved output generally reflects the scene depicted in the original imagery. The hedgerows in the hedgerow class were predicted as continuous features and the shape of the woodland class resembled the shape of the feature in the original imagery. However, misclassification was clearly evident throughout the image in the form of discrete blocks of error. Table 4.20 presents accuracy assessment of the hedgerow class, where training data (reference) were used to produce a confusion matrix against the prediction (super-resolved output). Overall accuracy of the hedgerow class was estimated at 67.6%.

Table 4.20: Confusion matrix, field site 3.

		Predicted			
		Hedgerow	Not hedgerow	Totals	PA (%)
Reference	Hedgerow	17	20	37	45.9
	Not hedgerow	4	33	37	89.1
	Totals	21	53	50	
	UA (%)	80.9	62.2	% Correct: 67.6	

4.3.4 Discussion of results

A visual inspection of the results of applying the soft classification techniques to each of the three field sites indicated that the woodland class in each of the field sites was separated from surrounding classes adequately by soft classification techniques. In other classes, particularly the cereal classes, however, the existence of greys suggested mixing between classes. Accuracy measures, such as RMSE and the correlation coefficient were used to assess the accuracy of each of the soft classification techniques. The reported accuracy of the woodland class from each of the accuracy assessments supported the visual inspection of the imagery. However, accuracy assessments revealed more information about the soft classification than a visual inspection. For example, in field site 1, there was little difference in accuracy between the FCM and MM. In the hedgerow and cereal classes, the difference in RMSE was 1-2%. In the woodland class, the difference was 4%. However, in field sites 2 and 3, where there were both more land cover classes and individual classes which were less spectrally separable, accuracy assessments showed that the FCM was more accurate than the MM. Accordingly, only the output of the FCM was used as input to the sub-pixel mapping algorithm.

A visual inspection of the super-resolved output revealed that in most cases continuous features were predicted and hedgerows approximately reflected the size and geometry of the features in the original imagery. In some cases, though,

particularly field site 3, small gaps appeared within the hedgerows. Trees within hedgerows resembled their initial shape.

The super-resolved output from the pixel swapping algorithm contained detailed information about the rural land cover features of interest. Field site 1 represented the simplest scene containing only three classes (hedgerow, woodland and cereal) and these features were resolved with reasonable accuracy. Variation in width along the length of the hedgerow was mapped. Inspection of the original imagery confirmed the narrowing to the south-western and north-eastern ends of the hedgerow. In field sites 2 and 3, the classification task was more difficult. In field site 2, the primary mapping tasks for the algorithm were the area of woodland (around 20 m wide) and the adjacent hedgerow (1-2 m wide). The former feature was predicted relatively accurately. The hedgerow was predicted accurately on a per-pixel basis. However, visually the hedgerow was mapped as a series of discrete objects. The same was true of the mapping tasks in field site 3 (each of the hedgerows), although in this case, where the hedgerow was clearly distinguished from its background, greater contiguity was achieved. Remembering that the original pixel size was 2.6 m and the widths of the hedgerows were approximately equal to that, the results were reasonable. It was also demonstrated that the pixel swapping technique could be used to map land cover features within features. For example, in field site 3, in addition to mapping the hedgerows, individual trees within hedgerows were mapped. This represented a great potential for the pixel swapping algorithm and represented functionality that was not common within mainstream feature extraction techniques. With development, the pixel swapping algorithm could be used to provide geometric information on these features. Such information (e.g., size of tree crowns, number of tree crowns, relative density of tree crowns to hedgerow, width of hedgerow) could be of great utility to ecologists (e.g., in terms of characterising habitats for bird foraging, nesting and territories generally) for a variety of applications.

In the super-resolved outputs of each of the three field sites, there was evidence of error, which resulted in speckled patterns appearing in the output and affected the accuracy of the super-resolved output. The error was a result of misclassification in the soft classification stage. Visual inspection of the original imagery and the soft proportion images provided some possible reasons for this error. In the original

imagery (Chapter 3), variation within classes was apparent in each of the three field sites. The human eye was able to distinguish between the classes quite simply, that is, in field site 1, the hedgerow, the woodland and field classes were easily delineated. In field sites 2 and 3, it was also possible to delineate individual classes by eye, including more complex features such as trees within hedgerows. It was also apparent that within classes, for example, in each of the cereal classes, which the human eye delineates as “fields”, there was mixing within them, notably in field site B. On the ground, mixing within the field could have been caused by many things, for example, the stage of the growth of the cereals, where the spectral response of cereals change as they ripen, or areas of non-growth resulting in bare soil. In the case of the woodland class in field site B, the unusual shape of the feature in the super-resolved output, was a result of tree crowns and gaps within the canopy. Mixing within classes would have affected the accuracy of the soft classification, if the number of classes used in soft classification did not realistically describe the number of the classes on the ground.

To minimise the effect of the error on the super-resolved output a mathematical morphology technique was developed. This technique is described and presented in the next section.

4.4 Mathematical morphology

The accuracy of the pixel swapping technique was affected adversely by error in the soft classification input. Therefore, a two-step mathematical morphology (Heijmans, 1995; Serra, 1982) approach was used to remove small areas of error from the super-resolved output. Mathematical morphology techniques were developed to handle objects with a characteristic spatial structure, comprised of a specific arrangement of individual elements. Mathematical morphology techniques adjust and rearrange these individual elements to influence the state and appearance of the structure of objects. In this research, “closing and opening” operations were used. On a class-by-class basis, a filter (the “structuring element”), variable in both size and shape, was passed over the image, which initially *eroded* objects, by removing pixels around the objects’ edges. Then, the objects were *dilated*, where sub-pixels were added back in. This

completely removed small objects (assumed to be classification error) leaving larger objects (assumed to be actual objects) relatively untouched.

Mathematical morphology was used to remove error in the super-resolved output while leaving the feature of interest intact. In each case, the structuring element was disc-shaped and the size of the structuring element was varied on a class-by-class basis for each field site. Morphology was applied initially on all classes except the feature class (i.e. all non-feature error sub-pixels were removed first) to maximise the availability of feature class pixels in a subsequent inverse-distance weighting step. For example, in field site 1, a basic $[0, 1, 1]$ structuring element was used (i.e., no morphology operator was applied to the feature class, and the structuring element for the other two classes was of size 1 sub-pixel). In the remaining sites, where there were more classes and more significant error, larger elements were used, for example, $[0, 2, 2, 2]$ for field site 2 and $[0, 1, 2, 3, 3]$ for field site 3. Morphology was then applied a second time, acting only on the feature class (to remove feature class error pixels e.g. $[1, 0, 0]$ for site A, $[1, 0, 0, 0]$ for site B and $[1, 0, 0, 0, 0]$ for site C).

As a result of applying the mathematical morphology technique and removing small areas of error, some sub-pixels were no longer assigned to a class. Therefore, a simple inverse-distance weighting algorithm in a 3 by 3 window was applied, to assign a class to these sub-pixels, predicting the class based on the sub-pixels' neighbours.

The mathematical morphology technique described above was applied to the super-resolved output of each of the three Christchurch field sites. The result of applying the mathematical morphology technique to the super-resolved output of each of the three Christchurch field sites is presented in the following section, in the form of super-resolved land cover maps. The accuracy of the technique is assessed using confusion matrices.

4.4.1 Field Site 1

The mathematical morphology technique was applied to field site 1, using a disc-shaped structuring element, in two stages. The size (in sub-pixels) of the structuring

element was 0 on the hedgerow class, 1 on the woodland class and 1 on the cereal class (represented as [0,1,1]). The element was then reapplied in the form [1,0,0]. Inverse distance weighting was then used to assign classes to any sub-pixels without a class allocation. Figure 4.19: Super-resolved output with mathematical morphology applied, field site 1 shows the super-resolved output with morphology applied.




Figure 4.19: Super-resolved output with mathematical morphology applied, field site 1

A visual inspection of Figure 4.19 suggests that much of the error apparent in the standard super-resolved output was removed. Figure 4.19 was more visually appealing without the error and resembled more closely the type of output that was expected from the technique (three distinct classes). Noticeably, areas of error along the edges of the classes (for example, hedgerow pixels along the edge of the woodland class) were removed without affecting the shape of the woodland class. The single tree within the hedgerow remained, whilst other misclassified woodland pixels were removed. Table 4.21 presents the confusion matrices to assess the accuracy of the morphology technique.

Table 4.21: Confusion matrix, super-resolved output with mathematical morphology applied

		Predicted			
		Hedgerow	Not hedgerow	Totals	PA (%)
Reference	Hedgerow	21	24	45	46.7
	Not hedgerow	4	41	45	91.1
	Totals	25	65	62	
	UA (%)	84.0	63.1	% Correct: 68.8	

Accuracy assessment of the super-resolved output with morphology applied indicated that the accuracy of the prediction of the hedgerow class increased by 3%.

4.4.2 Field site 2

The mathematical morphology technique was applied to field site 2, using a disc-shaped structuring element, in two stages. The size (in sub-pixels) of the structuring element was 0 on the hedgerow class, and 2 for each of the other three classes (represented as [0,2,2,2]). The element was then reapplied in the form [1,0,0,0]. Inverse distance weighting was then used to assign classes to any sub-pixels without a class allocation. Figure 4.20 displays the super-resolved output with mathematical morphology applied.

Table 4.22: Confusion matrix, super-resolved output with mathematical morphology




Figure 4.20: Super-resolved output with mathematical morphology applied, field site 2

A visual inspection of Figure 4.20 suggests that much of the error apparent in the standard super-resolved output was removed, particularly in the non-ripe cereal class in the upper left corner and lower-right corner of the figure. Figure 4.20 was more visually appealing without the error and resembled more closely the type of output that was expected from the technique (four distinct classes). Error still existed, for example, ripe cereal sub-pixels were not removed from areas where hedgerow pixels were expected. Table 4.22 presents the confusion matrices to assess the accuracy of the morphology technique.

Table 4.22: Confusion matrix, super-resolved output with mathematical morphology applied

		Predicted			
		Feature	Not Feature	Totals	PA (%)
Reference	Feature	17	1	18	94.4
	Not Feature	1	17	18	94.4
	Totals	18	18	34	
	UA (%)	94.4	94.4	% Correct: 94.4	

Accuracy assessment of the super-resolved output with mathematical morphology applied estimated an increase in accuracy of 6%.

4.4.3 Field site 3

The mathematical morphology was applied to field site 2, using a disc-shaped structuring element, in stages. The size (in sub-pixels) of the structuring element was 0 on the hedgerow class, 1 on the tree class, 2 on the cereal class and 3 on the remaining cereal classes (represented as [0,1,2,3,3]). The element was then reapplied in the form [1,0,0,0,0]. Inverse distance weighting was then used to assign classes to any sub-pixels without a class allocation. Figure 4.21 displays the super-resolved output with mathematical morphology applied.

Table 4.23: Confusion matrix, super-resolved output with mathematical morphology applied, field site 3

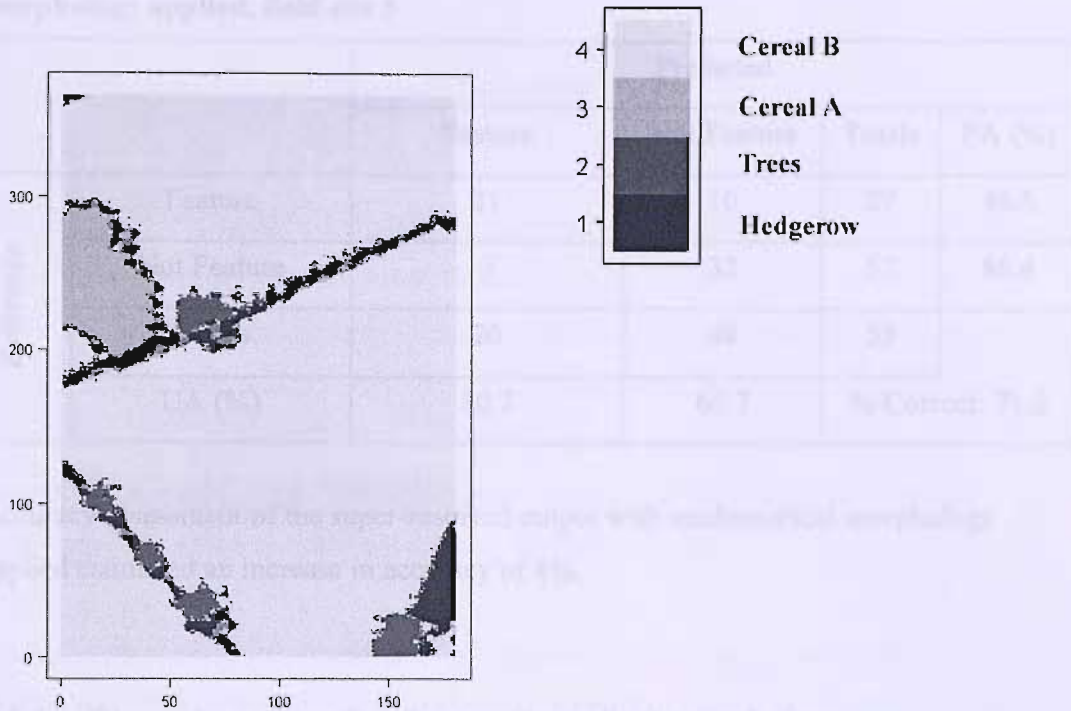


Figure 4.21: Super-resolved output with mathematical morphology applied, field site 3

A visual inspection of Figure 4.21 suggests that much of the error apparent in the standard super-resolved output was removed. Figure 4.21 was more visually appealing without the error and resembled more closely the type of output that was expected from the technique (four distinct classes). Table 4.23 presents the confusion matrices to assess the accuracy of the morphology technique.

Table 4.23: Confusion matrix, super-resolved output with mathematical morphology applied, field site 3

		Predicted			
		Feature	Not Feature	Totals	PA (%)
Reference	Feature	21	16	37	46.5
	Not Feature	5	32	37	86.4
	Totals	26	48	53	
	UA (%)	80.7	66.7	% Correct: 71.2	

Accuracy assessment of the super-resolved output with mathematical morphology applied estimated an increase in accuracy of 4%.

4.4.4 Discussion of results after mathematical morphology

The super-resolved output after mathematical morphology was applied showed some visual improvements. In most cases, mathematical morphology removed small areas of misclassification from the super-resolved output. When compared with confusion matrices for the standard super-resolved output, the confusion matrices for the super-resolved output after mathematical morphology had been applied, indicated an increase in overall accuracy. However, in field sites 2 and 3, where mixing between classes was complex, particularly in the ‘cereal’ classes, some obvious error was still apparent, even after application of the mathematical morphology.

The primary constraint of the pixel swapping algorithm is that class proportions (from a soft classification) are not modified, it only updates the spatial location of sub-pixels. The mathematical morphology technique violates this constraint since areas of misclassification are removed, and inverse distance weighting is used to interpolate unknown values. As previously discussed, the existence of error in the super-resolved output is solely a function of the input soft classification, and accordingly, a preferred solution to the problem would be to increase the accuracy of the soft classification, rather than post-process the imagery, as described above.

4.5 Discussion of chapter results

When applied to simulated imagery, the pixel swapping technique displayed results which were both comparable visually to the target and were accurate to 98% or greater. In each case, the super-resolution technique produced more accurate results than a standard hard classification and displayed the potential of the pixel swapping technique for land cover mapping. The simple shapes depicted in the simulated imagery were useful in identifying key characteristics of land cover features that would require consideration when developing the technique.

When the pixel swapping technique was applied to real imagery, using soft proportions derived from a soft classification technique, several problems occurred. The primary problem was the handling of error in the soft classification output. As the pixel swapping algorithm completed successive iterations, the algorithm swapped the misclassified sub-pixels as if they were correctly allocated. After several iterations, the misclassified sub-pixels clustered together and began to take a spatially structured form. This limited the accuracy of the pixel swapping technique. In order to understand this problem and develop the algorithm accordingly, it was important to investigate the causes of the error based on what was on the ground in the image.

There were several facets of the imagery, which affected the accuracy of the standard pixel swapping technique. In complex sites such as field site 2 and field site 3, class mixing existed *on the ground* (i.e., areas within the image exhibited complex mixing). In some fields, the percentage crop cover varied spatially and was mixed with areas of bare soil. This was problematic during both soft classification and subsequently during pixel swapping. The problem of mixing is a real life problem that exists as a challenge for any remote sensing research project, however, for the purposes of testing the algorithm, a different choice of field sites might have yielded more informative results. Had time permitted, an in-depth evaluation of the effect of mixing on the ground would have been carried out, initially by testing the algorithm on additional sources of real remotely sensed imagery.

Both classifiers displayed accuracies above 80%, yet both classifiers were less accurate where there was mixing within classes on the ground. This was most obvious in the more complex sites, such as field sites 2 and 3, where there was both bare soil and cereals in the agricultural fields, even though only 'cereal' was used in training. Further, as Foody (2000) points out, the existence of untrained classes in the image may have led to reduced accuracy of the soft classification. Consequently, the accuracy of the pixel swapping algorithm was affected, due to its reliance on accurate class proportion prediction. It is likely that more accurate soft classification data would have yielded more accurate output from the pixel swapping technique.

The trees and hedgerows in the imagery were commonly spectrally similar to the background, which made separating them from their background difficult. For example, the hedgerows on the ground were generally made up of green plants, such as short and long grasses (*Miscanthus*), ferns (*Athyrium*) and cow parsley (*Anthriscus sylvestris*). In some cases, there were also woody plants, such as sticks, brambles and young trees. The plants in the non-ripe cereal classes, were generally green, but changing to yellow as they ripened. Therefore, they often had similar spectral characteristics to the hedgerow classes.

Error in the ground data may have affected the assessment of the accuracy of the techniques. The phenological stage of the vegetation in the scene when the satellite sensor image was acquired was likely to differ slightly from that when the ground data were acquired. Additionally, land cover features may have changed: the width of the hedgerow might have differed to the width of the hedgerow in the satellite sensor image, due, for example, to hedgerow management practices. The GPS Position Dilution of Precision (PDOP) value, a function of satellite geometry, affects the precision with which measurements are taken. PDOPs of between 2.2 and 4.0 were recorded which related to an average precision of 68%. Furthermore, georectification of the imagery was accurate to approximately 2 m. Given that the average width of the hedgerows were 3 m, this accuracy would have had an effect on locating the hedgerow (measured in the field) in the actual image in accuracy assessment. These issues may have affected the accuracy of the assessment of the sub-pixel output.

4.6 Conclusion

The standard pixel swapping algorithm is a simple and efficient super-resolution technique. Initial testing of the pixel swapping algorithm on simulated imagery led to accurate predictions. Application to real satellite sensor imagery, however, yielded less promising results. During this stage of research, one key requirement for accurate sub-pixel mapping was highlighted, that is, accurate soft classification information are required as input. For features that were approximately half the width of a pixel, the algorithm failed to produce contiguous features, instead, mapping them as a series of discrete objects. It was evident, therefore, that the technique required some development in order to accurately predict linear features. Nevertheless, the super-resolution pixel swapping technique displayed obvious potential for mapping linear features.

Chapter 5 presents the next stage of the research – “linearisation”.

Chapter 5

Linearisation


Linearisation is a technique used to approximate a non-linear function by a linear function. It is based on the Taylor series expansion of the function around a point. The first-order Taylor series expansion is the linear approximation of the function at that point. The higher-order Taylor series expansions provide better approximations of the function, but they are more complex and computationally expensive. Linearisation is a useful tool for analyzing the behavior of non-linear systems near a point of interest. It is widely used in engineering, physics, and economics. In this chapter, we will discuss the theory of linearisation and its applications in various fields.

The first-order Taylor series expansion of a function $f(x)$ around a point x_0 is given by $f(x) \approx f(x_0) + f'(x_0)(x - x_0)$. This approximation is valid for small values of $x - x_0$. The second-order Taylor series expansion is given by $f(x) \approx f(x_0) + f'(x_0)(x - x_0) + \frac{1}{2}f''(x_0)(x - x_0)^2$. This approximation is valid for small values of $x - x_0$ and provides a better approximation of the function than the first-order expansion. The third-order Taylor series expansion is given by $f(x) \approx f(x_0) + f'(x_0)(x - x_0) + \frac{1}{2}f''(x_0)(x - x_0)^2 + \frac{1}{6}f'''(x_0)(x - x_0)^3$. This approximation is valid for small values of $x - x_0$ and provides a very good approximation of the function.

Linearisation is a powerful tool for analyzing the behavior of non-linear systems near a point of interest. It is widely used in engineering, physics, and economics. In this chapter, we will discuss the theory of linearisation and its applications in various fields. We will start by discussing the theory of linearisation and then move on to its applications in engineering, physics, and economics. We will also discuss the limitations of linearisation and how to overcome them. Finally, we will provide some examples of linearisation in various fields.

Chapter 5

5. Linearisation



5.1 Introduction

In Chapter 4 of this thesis, a super-resolution pixel swapping technique was evaluated for its potential to map land cover classes from remotely sensed imagery. When applied to simulated remotely sensed imagery, the technique mapped the features in the imagery accurately. When applied to real satellite imagery, however, the land cover classes in each of the three field sites were mapped less accurately as a result of error in the soft classification. After a morphology technique was applied to the super-resolved output to remove small areas of error, the locations of the classes in the output were more comparable with the target imagery. In particular, the locations and dimensions of the linear features within the super-resolved output were mapped with reasonable accuracy, displaying the obvious potential of the technique for mapping land cover. However, in some cases, the predicted features were not contiguous, instead, they were predicted as a series of discrete objects. Additionally, some feature-level characteristics, for example, variation in the width of a feature along its length, were not accurately represented. It was therefore necessary to develop the pixel swapping technique to accommodate these requirements.

In the following two chapters, research carried out to develop the pixel swapping technique is presented, where the objective was to specifically map *linear* features in remotely sensed imagery. Research into existing linear feature techniques (as described in Chapter 2 of this thesis) served as a useful background in defining some of the key characteristics of linear features (as described in Chapter 3). With this information, the pixel swapping technique was modified to map linear features. The research carried out to achieve this objective is presented in this chapter.

5.2 Linearisation

Linear features exhibit a wide range of geometric characteristics. These characteristics are described in Chapter 3 of this thesis. In order to map linear features within remotely sensed imagery with greater accuracy than the existing pixel swapping technique, the pixel swapping technique was developed to incorporate two of these characteristics: direction and width.

5.2.1 Introduction

In the standard pixel swapping technique, an exponential distance decay window is used to predict which sub-pixels to swap. A uniform isotropic exponential distance decay window is used on each pixel, regardless of the spatial and spectral arrangement of the sub-pixels. Anisotropy (the opposite of isotropy) is the property of being directionally dependent. A rural linear feature is anisotropic, that is, it may appear different, or have different characteristics in different directions. For example, we would expect the scale of spatial variation along the length of a hedgerow to be quite different compared to the scale of spatial variation across the width. Indeed, the existence of a linear feature could be predicted on the basis of measured anisotropy within pixels in a remotely sensed image.

In the new ‘linearised’ pixel swapping algorithm, in order to increase the likelihood of predicting linear features, a unique anisotropic exponential distance decay window was created for every pixel, on the basis of measured anisotropy within the class proportions. The new method is described below.

5.2.2 Linearising the pixel swapping technique

In order to develop the pixel swapping technique to increase the likelihood of predicting linear features, *a priori* information on the existence of linear features in the field site was required. One source of this information was soft classification. Soft classification provides information on the proportion of classes within individual pixels. Soft proportions do not provide any information on the location of features at the sub-pixel level. However, assumptions on the existence of linear features can be made using the arrangement of pixels which contain proportions of a particular class. Figure 5.1 illustrates this.

10	17	20	54	80
7	22	28	74	65
20	40	<i>i</i>	60	58
46	58	40	40	35
49	37	30	29	29

Figure 5.1: Predicting the direction of a linear feature from soft proportions. A constraint is imposed that a linear feature must pass through the central pixel of a moving window. In this example, a linear feature is predicted at 45° from north.

In a moving window applied to the class proportions of the feature class, with the centre pixel x_i as the 'target', the preferred direction is extracted as that in which the sum of two pixels in a straight line through the central pixel is greatest. Figure 5.2 illustrates this. The anisotropy ratio is also estimated, by dividing the sum of the pixels on the orthogonal of the estimated direction by the sum of the pixels in the estimated direction.

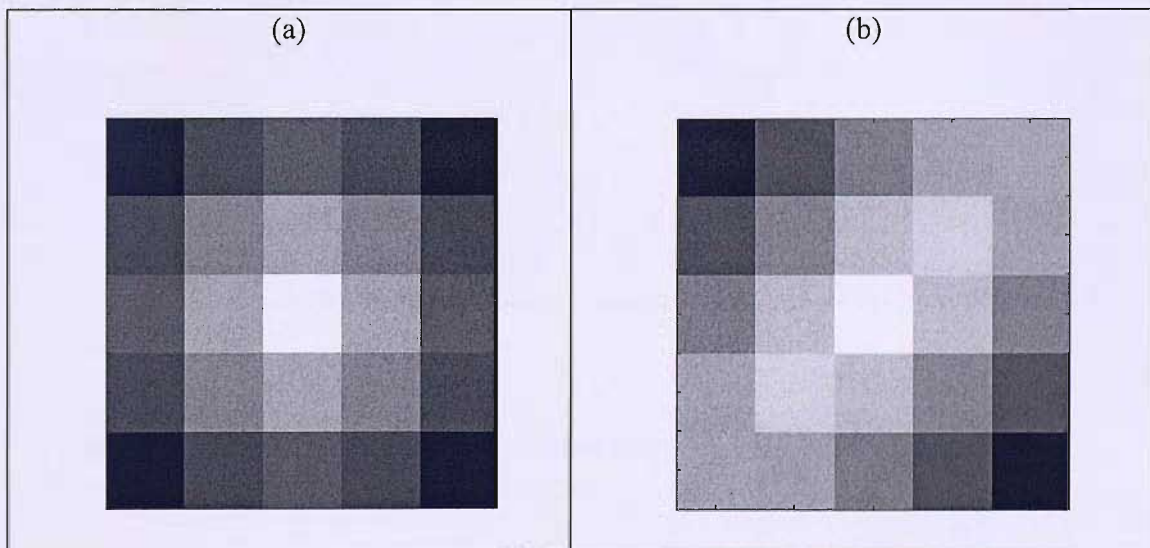
The direction estimated within the moving window applied to the class proportions image is then used to modify the exponential distance decay model using a standard anisotropic correction (see Goovaerts (1997)). The co-ordinates of the exponential

distance decay window, $\mathbf{h} = (h_x, h_y)^T$ (where T is the transpose), are transformed by the direction estimated from the class proportions and the anisotropy ratio, $\mathbf{h}' = (h'_\phi, h'_\theta)^T$, such that the transformed co-ordinates are isotropic:

$$g(\mathbf{h}) = g'(|\mathbf{h}'|) \text{ with } |\mathbf{h}'| = \sqrt{h'^2_\phi + h'^2_\theta} \quad (6)$$

where $g'(\cdot)$ is an isotropic model with a range equal to the minor range a_ϕ of anisotropy.

Figure 5.2(a) shows the 'standard' isotropic exponential distance decay model for one pixel. Figure 5.2(b) shows an anisotropic exponential distance decay model, with a direction of 45° and anisotropy ratio of 0.35. The shape of the model changes from circular to ellipsoid, with the anisotropy ratio defining the width of the ellipse.



**Figure 5.2: Exponential distance decay models, (a) Isotropic (b) Anisotropic.
Legend: White = 1, Black = 0**

In this investigation, two different sized moving windows were used: a 3 by 3 window and a 5 by 5 window. In the 3 by 3 window, four directions can be estimated, 0° , 45° , 90° and 135° (Figure 5.3(a)) and in the 5 by 5 window, eight directions can be estimated, between 0° and 157.5° at 22.5° intervals (Figure 5.3(b)). In the 5 by 5 window, where the line passes through two pixels adjacent to the central pixel (e.g.,

22.5° in Figure 5.3, pixels 8 and 9 and 17 and 18), the mean of the sum of the two pixels is used in estimating the proportion for that direction.

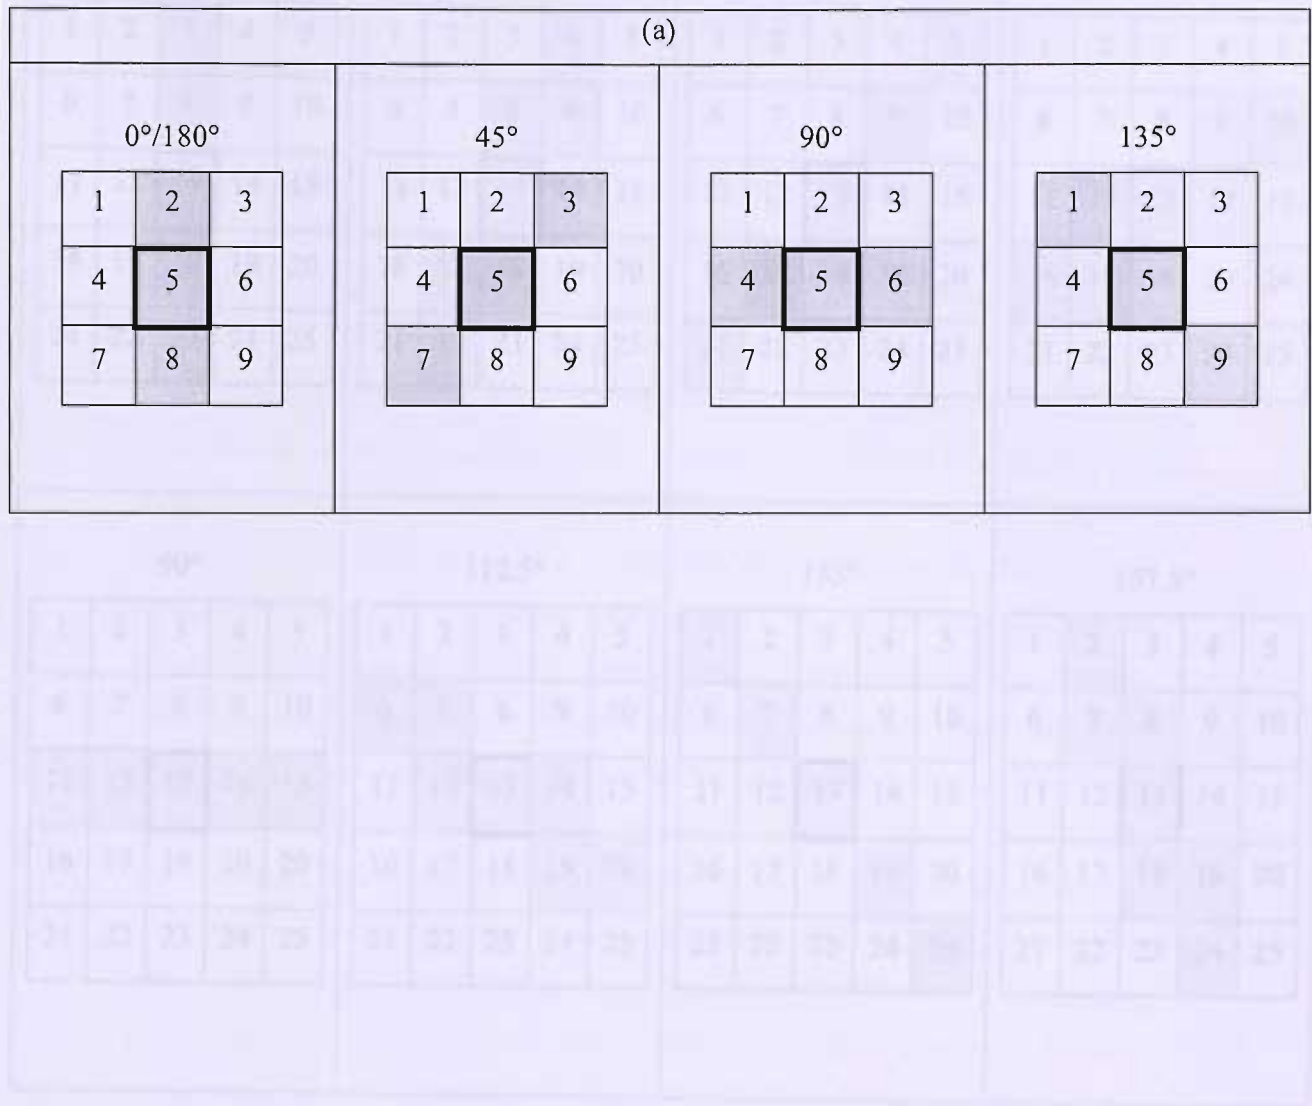


Figure 5.3: Directions, (a) The 4 directions in a 3 by 3 window, (b) The 8 directions in a 5 by 5 window

5.3.3 Multiple directions

In other cases, more than one direction is identified within the neighborhood. To allow for such occurrences, for each pixel, multiple directions are used to find the best class proportions when combined into the same histogram representation of the image window. For example, in a 3 by 3 window, directions at 90° and 135° might be

(b)

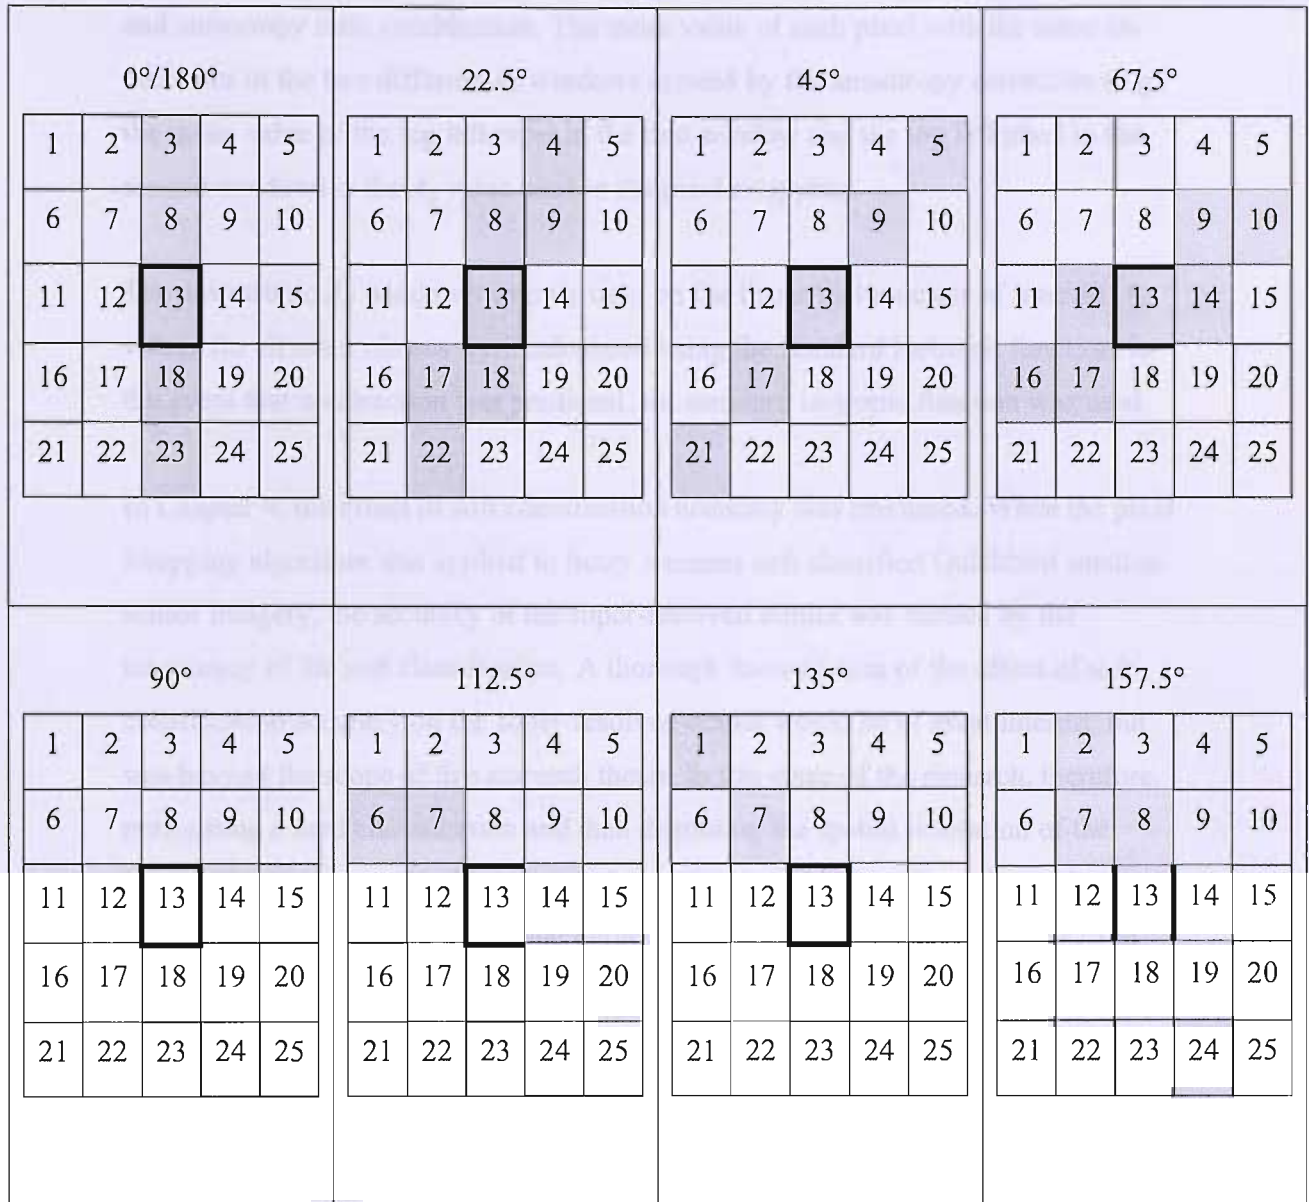


Figure 5.3: Directions, (a) The 4 directions in a 3 by 3 window, (b) The 8 directions in a 5 by 5 window

5.2.3 Multiple directions

In some cases, more than one obvious direction existed within the proportions. To allow for such occurrences, for each pixel, multiple directions estimated from the pixel class proportions were combined into the same anisotropic exponential distance decay window. For example, in a 3 by 3 window, directions of 90° and 135° might be

predicted. In this circumstance, two A_{ij} windows were created, one for each direction and anisotropy ratio combination. The mean value of each pixel with the same coordinates in the two different A_{ij} windows created by the anisotropy correction (e.g., the mean value of the top left pixel in the first window and the top left pixel in the second window) is the A_{ij} value used in the pixel swapping.

The anisotropic A_{ij} window was only used on the linear feature class of interest. A_{ij} values for all other classes were calculated using the standard isotropic function. In the event that no direction was predicted, the standard isotropic function was used.

In Chapter 4, the effect of soft classification accuracy was discussed. When the pixel swapping algorithm was applied to fuzzy c -means soft classified Quickbird satellite sensor imagery, the accuracy of the super-resolved output was limited by the inaccuracy of the soft classification. A thorough investigation of the effect of soft classification accuracy on the super-resolved output would be of great interest, but was beyond the scope of this research thesis. In this stage of the research, therefore, performing a hard classification and then degrading the spatial resolution of the imagery to produce soft proportions reduced the uncertainty attributed to soft classification error. The input aerial photography, at a spatial resolution of 0.25 m used in this section is described in Chapter 3.4.2. After hard classification had been performed, the spatial resolution of the classified imagery was degraded to create two sets of input imagery for use in this stage of research - by a factor of 10, simulating soft classified imagery at a spatial resolution of 2.5 m, and by a factor of 20, simulating soft classified imagery at a spatial resolution of 5 m.

In using hard classified aerial photography at a degraded spatial resolution as input to the pixel swapping technique, the super-resolved output was configured to be at the same spatial resolution as that of the original aerial photography, which enabled a full assessment of accuracy, including the creation of difference images.

5.3 Results

The following section presents the results of applying the pixel swapping technique to the input imagery. In each case, three versions of the pixel swapping algorithm were applied - original, linear (3 by 3) and linear (5 by 5). The size of the A_{ij} window was 9 pixels on each axis (81 pixels). This size was chosen through repeated testing and provided the best overall solution in terms of accuracy.

The results for each set of imagery are presented in the following formats: the input imagery (at the degraded spatial resolution), a hard classification, the initial random allocation, a direction map, confusion matrices and a difference image. The direction map is a simple representation of the shape and orientation of the local anisotropic A_{ij} window used in the linearised version of the pixel swapping.

5.3.1 Field site (spatial resolution 2.5 m)

Initially, the input imagery was hard classified using the maximum likelihood classifier. The hard classification is presented in Figure 5.4.




Figure 5.4: Hard classification, SR 2.5 m

The hard classification of the input imagery depicted the site poorly. The edges of the hedgerow and tree classes have taken on a “blocky” appearance. A confusion matrix of this classification is presented in Table 5.1.

Table 5.1: Confusion matrix, hard classification, 2.5 m SR

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6467	154	32	6653	97.2
	Hedgerow	236	2548	83	2867	88.8
	Tree	88	198	395	681	58.0
	Totals	6791	2900	510	9410	
	UA (%)	95.2	87.8	77.4	% Correct: 92.2	

Accuracy assessment of the hard classification revealed that the cereal class was mapped with reasonable accuracy, however, the other two classes, in particular, the tree class, were mapped with moderate accuracy.

The pixel swapping technique was applied to the 2.5 m spatial resolution input imagery, at a zoom factor of ten resulting in super-resolved output images at the same spatial resolution as the original aerial photography (0.25 m). Sub-pixels were allocated randomly and the pixel swapping technique iterated until the predicted image converged on a solution. Figure 5.5(a) shows the simulated input imagery and Figure 5.5(b) shows the random allocation.

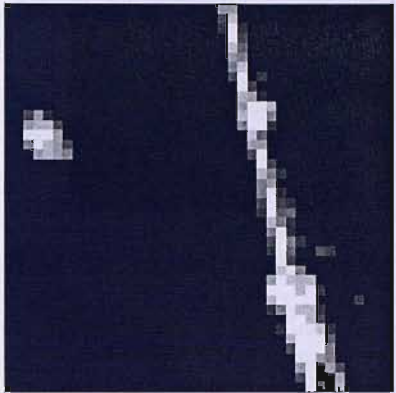
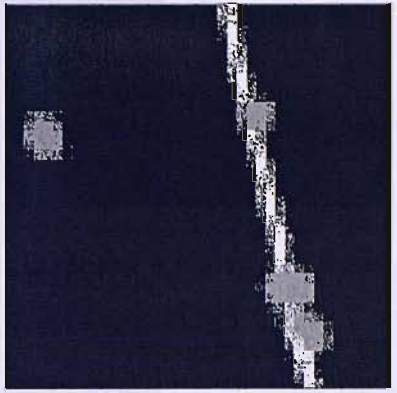
Spatial Resolution (SR)	Input imagery	Random allocation
2.5 m	<p style="text-align: center;">(a)</p> 	<p style="text-align: center;">(b)</p> 

Figure 5.5: Field site, spatial resolution 2.5 m, (a) Input imagery, (b) Random allocation

Anisotropic exponential distance decay windows from the soft proportions were calculated for the input imagery. Figure 5.6(a) shows the windows for a 3 by 3 window on the soft proportions and Figure 5.6(b) shows the windows for a 5 by 5 window on the soft proportions.

Figure 5.6: Direction maps (a) Anisotropic, 3 by 3 window, 2.5 m SR, (b) Anisotropic, 5 by 5 window, 2.5 m SR. Legend (A): Cyan = Low weight, Red = High weight.

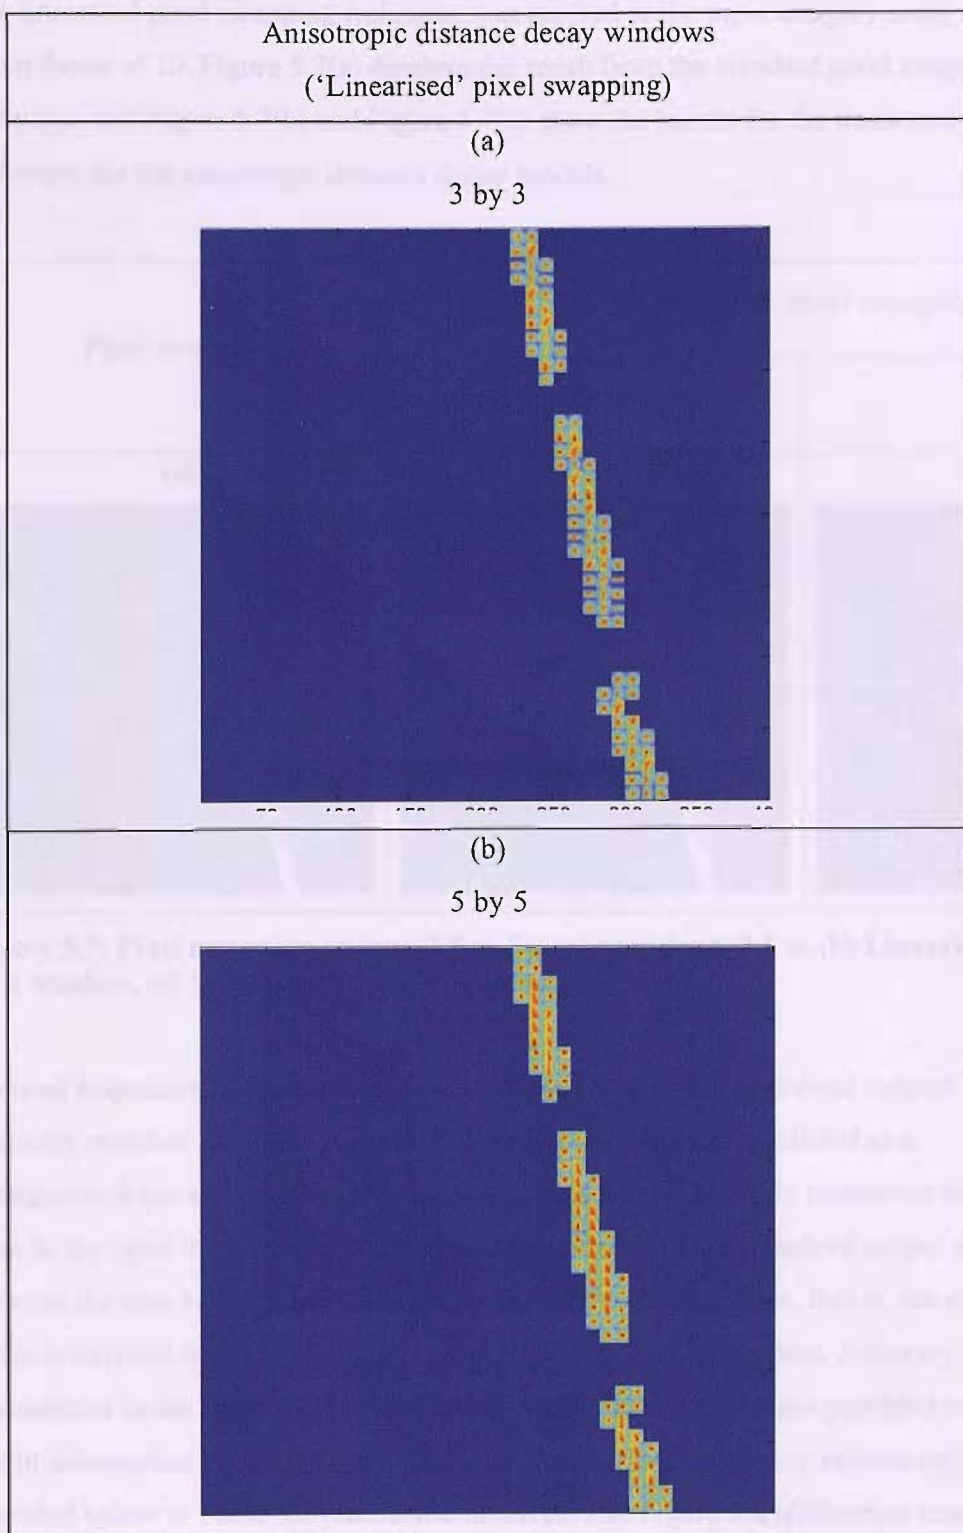


Figure 5.6: Direction maps (a) Anisotropic, 3 by 3 window, 2.5 m SR, (b) Anisotropic, 5 by 5 window, 2.5 m SR. Legend (A_{ij}): Cyan – Low weight, Red – High weight

The linearised pixel swapping technique was applied to the input imagery using a zoom factor of 10. Figure 5.7(a) displays the result from the standard pixel swapping technique and Figure 5.7(b) and Figure 5.7(c) show the results for the linearised technique for the anisotropic distance decay models.

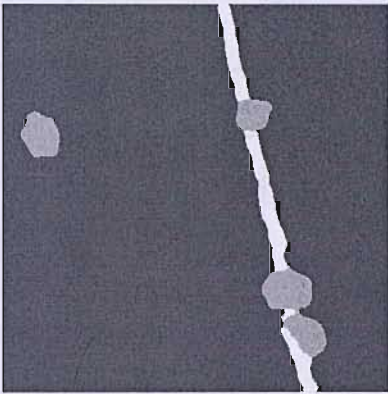
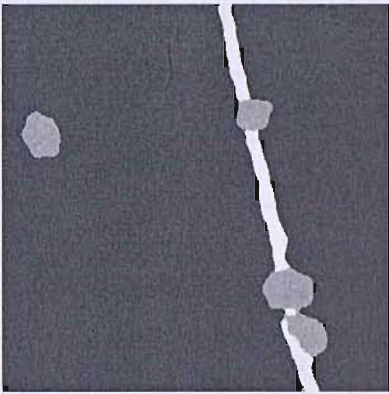
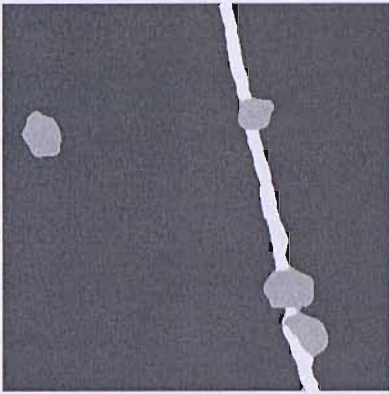
Resolution (SR)	Pixel swapping	'Linearised' pixel swapping	
		3 by 3	5 by 5
2.5 m	(a) 	(b) 	(c) 

Figure 5.7: Pixel swapping output, 2.5 m SR (a) Standard, 2.5 m (b) Linearised, 3 by 3 window, (c) 'Linearised', 5 by 5 window

A visual inspection of Figure 5.7 reveals that each of the super-resolved outputs generally matched the input imagery. The hedgerow class was predicted as a contiguous shape and the dimensions and shape of the trees closely resembled the trees in the input imagery. The main disagreement in the super-resolved output was between the area between the two trees embedded in the hedgerow, that is, the shape of the lowermost tree was predicted differently in each of the outputs. Accuracy assessments in the form of confusion matrices and difference images provided more useful information on the accuracy of the technique. These accuracy assessments are provided below in Table 5.2 (confusion matrices) and Figure 5.8 (difference images).

Table 5.2: Super-resolved output confusion matrices, 2.5 m spatial resolution, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window

(a)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6531	90	32	6653	98.1
	Hedgerow	74	2762	31	2867	96.3
	Tree	45	114	522	681	76.6
	Totals	6650	2966	585	9815	
	UA (%)	98.2	93.1	89.2	% Correct: 96.2	

(b)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6532	92	29	6653	98.1
	Hedgerow	74	2768	25	2867	96.5
	Tree	44	106	531	681	77.9
	Totals	6650	2966	585	9831	
	UA (%)	98.2	93.3	90.7	% Correct: 96.3	

(c)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6530	91	32	6653	98.1
	Hedgerow	75	2763	29	2867	96.3
	Tree	44	112	525	681	77.0
	Totals	6649	2966	586	9818	
	UA (%)	98.2	93.1	89.1	% Correct: 96.2	

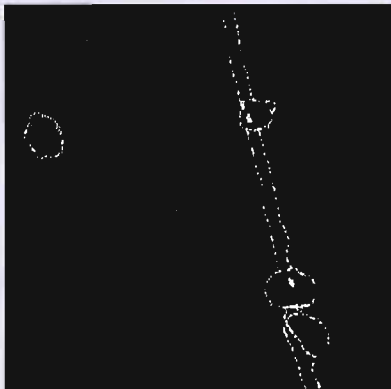
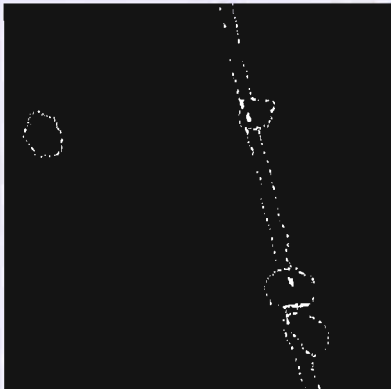
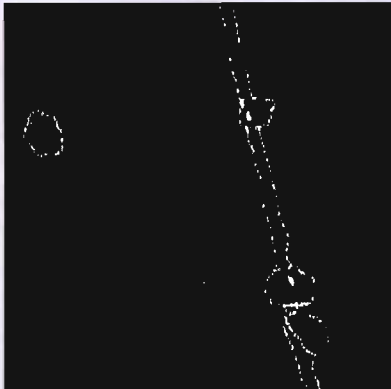
Spatial Resolution (SR)	Pixel Swapping	Linear Pixel Swapping	
		3 by 3	5 by 5
2.5 m	(a) 	(b) 	(c) 

Figure 5.8: Difference maps: spatial resolution 2.5 m, (a) pixel swapping, (b) 'linearised' pixel swapping, 3 by 3 window, (c) 'linearised' pixel swapping, 5 by 5 window. Legend: black – correct allocation, white – error of omission, grey - error of commission

Confusion matrices for the standard pixel swapping techniques estimated the overall accuracy as 96.2%, 4% greater than a standard hard classification. The overall accuracy for the linearised technique in the 3 by 3 window was estimated as 96.3% and for the 5 by 5 window as 96.2%. Closer inspection of the confusion matrices reveals that each of the techniques converged to virtually identical solutions,

separated only by disagreement between the hedgerow and tree classes. The difference images show the areas of error, which was most prevalent between the two trees embedded within the hedgerow. A significance test, as described in Chapter 2.6, was applied to evaluate the differences between the confusion matrices. In each case the linearised pixel swapping technique was compared with the standard pixel swapping technique.

Table 5.3: Significance tests, (a) PS & LPS, 3 by 3 window, (b) PS & LPS, 5 by 5 window

(a)

		Linearised Pixel Swapping		
		Correct	Incorrect	Total
Pixel Swapping	Correct	9815	0	9815
	Incorrect	16	370	386
	Total	9831	370	10201

$\chi^2 = 16$. Significant at 0.05%.

(b)

		Linearised Pixel Swapping		
		Correct	Incorrect	Total
Pixel Swapping	Correct	9814	1	9815
	Incorrect	4	382	386
	Total	9818	383	10201

$\chi^2 = 1.8$. Not significant at 0.05%.

Significance tests on the super-resolved output show that although the overall accuracy was very similar for each technique, the difference in accuracy between the standard pixel swapping technique and the linearised pixel swapping in a 3 by 3 window was significant. This result should be treated with caution, however, due to the large number of sub-pixels involved.

5.3.2 Field site (spatial resolution 5 m)

Initially, a hard classification of the input imagery was performed. This hard classification is presented in Figure 5.9.




Figure 5.9: Hard classification, 5 m SR

The hard classification of the field site at the 5 m spatial resolution delineated the classes in the site poorly. The hedgerow was depicted as a series of discrete objects and the intricate shape of the trees have been entirely lost. An accuracy assessment of this classification was performed and is presented in Table 5.4.

Table 5.4: Confusion matrix, hard classification, 5 m SR

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6190	361	102	6653	93.0
	Hedgerow	588	2172	107	2867	75.7
	Tree	203	267	211	681	30.9
	Totals	6981	2800	420	8573	
	UA (%)	88.6	77.5	50.2	% Correct: 84.0	

Accuracy assessment of the hard classification applied to the input imagery at a spatial resolution of 5 m, confirmed the analyses of the visual inspection, that is, the hedgerow and tree classes were predicted with poor accuracy.

The pixel swapping technique was applied to the 5 m spatial resolution input imagery, at a zoom factor of 20 resulting in super-resolved output images at the same spatial resolution as the original aerial photography (0.25 m). Sub-pixels were randomly allocated and the pixel swapping technique iterated until the predicted image converged on a solution. Figure 5.10(a) shows the simulated input imagery and Figure 5.10(b) shows the random allocation.

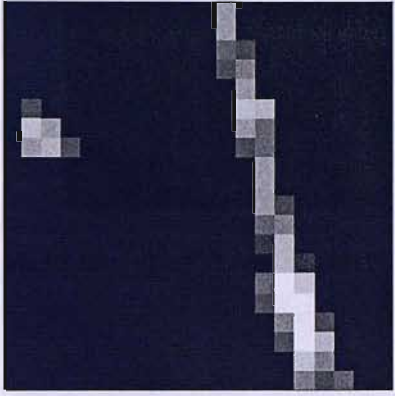
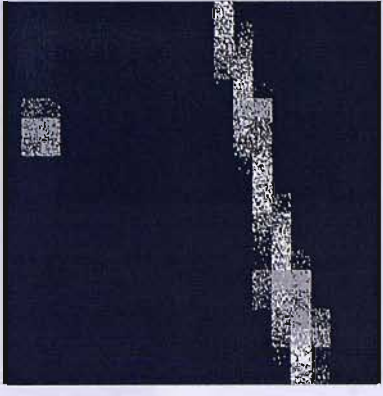
Spatial Resolution (SR)	Input imagery	Random allocation
5 m	<p style="text-align: center;">(a)</p> 	<p style="text-align: center;">(b)</p> 

Figure 5.10: Field site, spatial resolution 5 m, (a) Input imagery (b) Random allocation

Anisotropic exponential distance decay windows from the soft proportions were calculated for the input imagery. Figure 5.11(a) shows the windows for a 3 by 3 window on the soft proportions and Figure 5.11(b) shows the windows for a 5 by 5 window on the soft proportions.




Figure 5.11: Direction-weighted Anisotropic, 3 by 3 window, 5 m SR. (a) Anisotropic, 3 by 3 window, 5 m SR. Legend: High weight (red), Med weight (green), Low weight (blue).

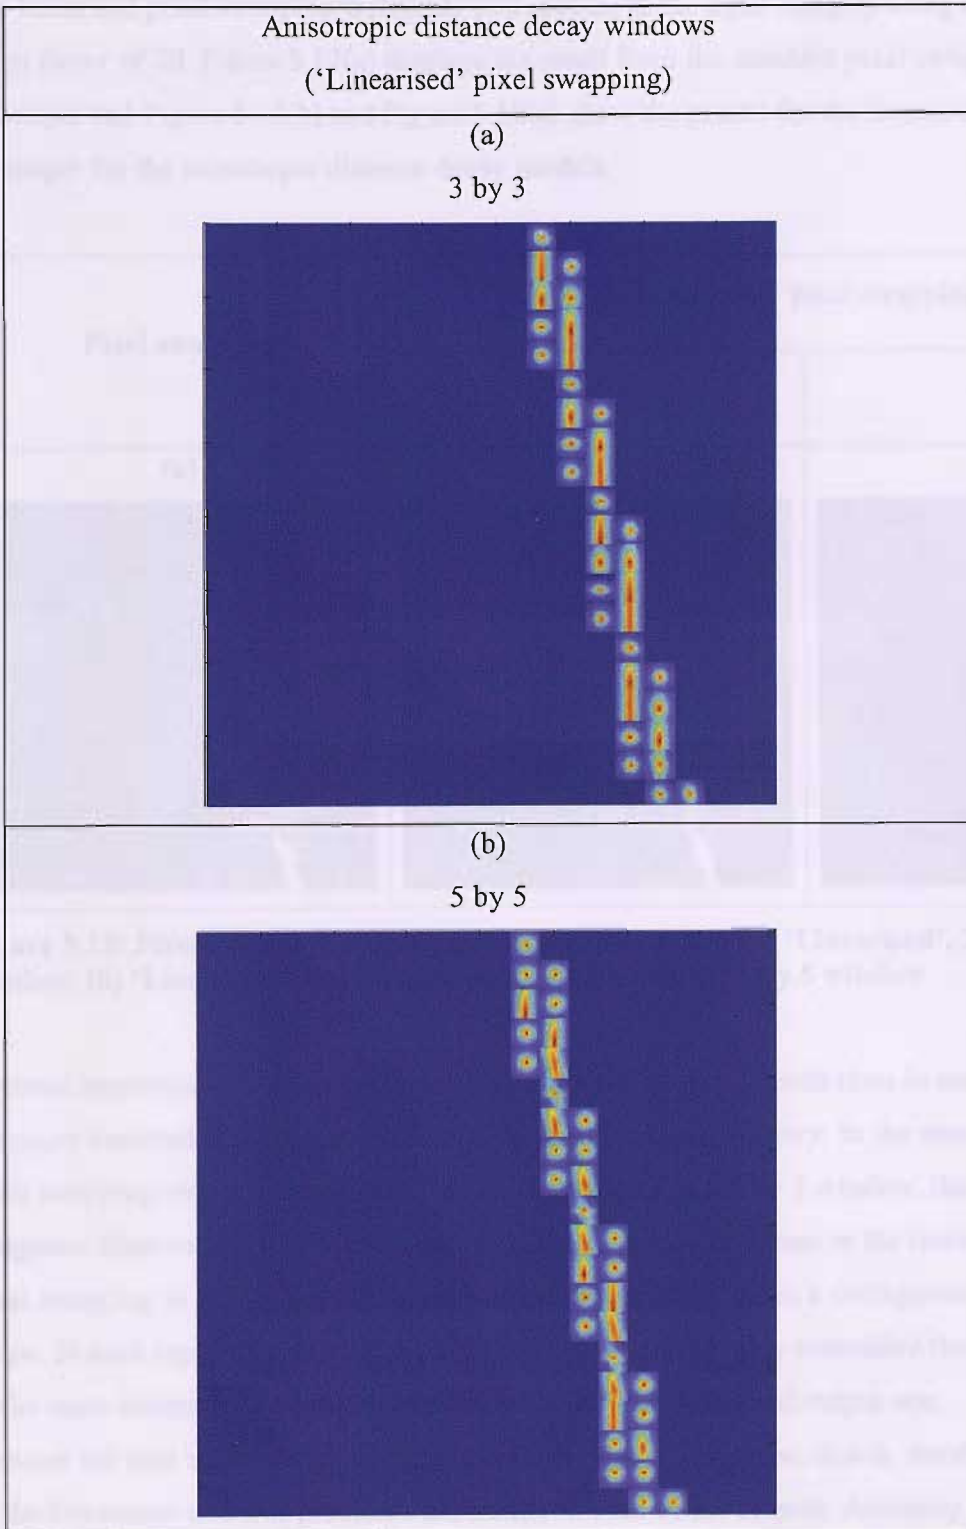


Figure 5.11: Direction maps (a) Anisotropic, 3 by 3 window, 5 m SR, (b) Anisotropic, 5 by 5 window, 5 m SR. Legend (A_{ij}): Cyan – Low weight, Red – High weight

The linearised pixel swapping technique was applied to the input imagery using a zoom factor of 20. Figure 5.12(a) displays the result from the standard pixel swapping technique and Figure 5.12(b) and Figure 5.12(c) show the results for the linearised technique for the anisotropic distance decay models.

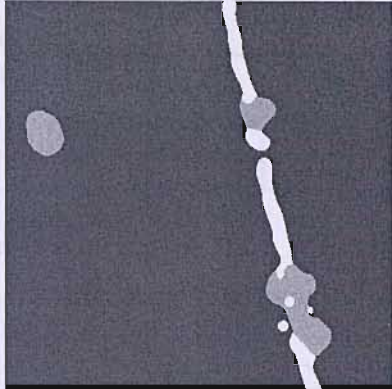
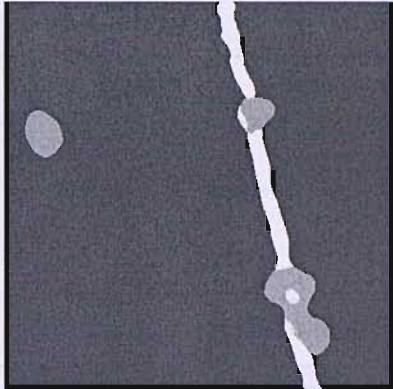
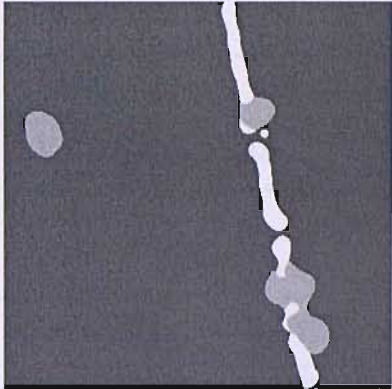
Spatial Resolution (SR)	Pixel swapping	'Linearised' pixel swapping	
		3 by 3	5 by 5
5 m	(a) 	(b) 	(c) 

Figure 5.12: Pixel swapping output, spatial resolution 5 m (a) 'Linearised', 3 by 3 window, (b) 'Linearised', 5 by 5 window, (c) 'Linearised', 5 by 5 window

A visual inspection of Figure 5.12 reveals that the orientation of each class in each of the super-resolved outputs approximately matched the input imagery. In the standard pixel swapping output and the linearised pixel swapping in a 5 by 5 window, the hedgerow class was predicted as a series of discrete objects, whereas in the linearised pixel swapping in a 3 by 3 window, the hedgerow was predicted as a contiguous shape. In each super-resolved output, the shape of the trees barely resembled the trees in the input imagery. The main disagreement in the super-resolved output was between the area between the two trees embedded in the hedgerow, that is, the shape of the lowermost tree was predicted differently in each of the outputs. Accuracy assessments in the form of confusion matrices and difference images provided more useful information on the accuracy of the technique. These accuracy assessments are provided below in Table 5.5 (confusion matrices) and Figure 5.13 (difference images).

Table 5.5: Super-resolved output confusion matrices, 5 m spatial resolution, (a) pixel swapping, (b) 'linearised' pixel swapping, 3 by 3 window, (c) 'linearised' pixel swapping, 5 by 5 window

(a)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6374	177	102	6653	95.8
	Hedgerow	177	2617	73	2867	91.2
	Tree	103	73	505	681	74.1
	Totals	6654	2867	680	9496	
	UA (%)	95.7	91.2	74.2	% Correct: 93.0	

(b)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6451	111	91	6653	96.9
	Hedgerow	131	2662	74	2867	92.8
	Tree	69	94	518	681	76.0
	Totals	6651	5867	683	9631	
	UA (%)	96.9	92.8	75.8	% Correct: 94.4	

(c)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6461	125	67	6653	97.1
	Hedgerow	125	2629	113	2867	91.6
	Tree	68	113	500	681	73.4
	Totals	6654	2867	680	9590	
	UA (%)	97.0	91.6	73.5	% Correct: 94.0	


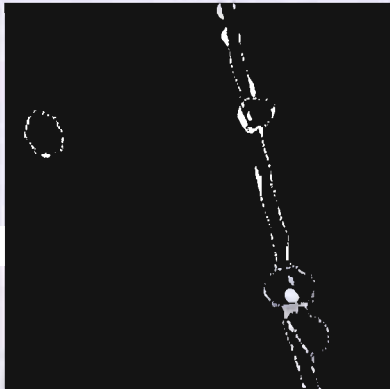
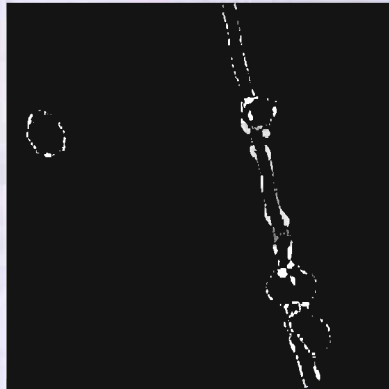
Spatial Resolution (SR)	Pixel Swapping	Linear Pixel Swapping	
		3 by 3	5 by 5
5 m	(a) 	(b) 	(c) 

Figure 5.13: Difference maps: spatial resolution 5 m, (a) pixel swapping, (b) ‘linearised’ pixel swapping, 3 by 3 window, (c) ‘linearised’ pixel swapping, 5 by 5 window. Legend: black – correct allocation, white – error of omission, grey - error of commission

Confusion matrices for each of the super-resolved outputs reported overall accuracy for the standard pixel swapping technique as 93.0%, in the linearised output in the 3 by 3 window as 94.4% and 94.0% in the linearised pixel swapping output in the 5 by 5 window. Closer inspection of the confusion matrices reveals that each of the techniques converged to virtually identical solutions, separated only by disagreement

between the hedgerow and tree classes. The difference images show the areas of error, which were most prevalent between the two trees embedded within the hedgerow. A significance test, as described in Chapter 2.6, was applied to evaluate the differences between the confusion matrices. In each case, the linearised pixel swapping technique was compared with the standard pixel swapping technique. These results are displayed in Table 5.6.

Table 5.6: Significance tests, (a) PS & LPS, 3 by 3 window, (b) PS & LPS, 5 by 5 window

(a)

		Linearised Pixel Swapping		
		Correct	Incorrect	Total
Pixel Swapping	Correct	9496	0	9496
	Incorrect	135	570	705
	Total	9631	570	10201

$x^2 = 133.0$. Significant at 0.05%.

(b)

		Linearised Pixel Swapping		
		Correct	Incorrect	Total
Pixel Swapping	Correct	9491	5	9496
	Incorrect	99	606	705
	Total	9590	611	10201

$x^2 = 86.7$. Significant at 0.05%.

It was evident from the significance tests that the super-resolved output of each version of the linearised pixel swapping technique were producing more accurate results than the standard pixel swapping technique.

5.3.3 Discussion of results

A hard classification was applied to each of the input images. Hard classification is a common application of remotely sensed imagery in the production of land cover maps. In this research, however, it was clear that hard classification was not suitable for application to fine spatial resolution remotely sensed imagery. The resultant maps were generally not comparable with the target imagery. In every case, the super-resolved output of the pixel swapping technique was both more accurate and created visually superior output images.

Figure 5.6 and Figure 5.11 show the exponential distance decay windows for each pixel across the entire image, i.e., the anisotropic windows used in the 'linearised' pixel swapping technique, for 3 by 3 pixel and 5 by 5 pixel windows respectively. The gaps in the direction maps within the hedgerow class were created by the trees evident within the hedgerow (in which case the default isotropic model was used).

The pixel swapping method operated efficiently when the input data were at a spatial resolution of 2.5 m. For each of the three versions of the technique, the accuracy of the map for the cereal class was 98.1%. The hedgerow class was mapped with an accuracy of 96% and the tree class was mapped with an accuracy of 80%. Using the 5 by 5 window was 3% more accurate than the standard method, whereas the 3 by 3 window produced results 5% more accurate than the standard method. When the input imagery was at a spatial resolution of 5 m, the standard pixel swapping algorithm did not map the hedgerow as a continuous feature. Instead, the feature was mapped as a series of discrete units. Each of the 'linearised' versions mapped a more contiguous feature, and through visual inspection, produced a result more comparable to the original 0.25 m spatial resolution imagery than the standard pixel swapping method. The accuracy of the maps in each case was > 93%. Both 'linearised' versions of the technique were more accurate than the standard pixel swapping technique. The confusion matrices indicated that the standard pixel swapping method mapped the tree class more accurately, in this case, and the significance tests confirmed these analyses.

'Linearised' pixel swapping with a 3 by 3 window was more accurate than with a 5 by 5 window at both spatial resolutions. Results at a spatial resolution of 2.5 m were comparable, but at a spatial resolution of 5 m the significance tests suggest that the 3 by 3 window achieved a better result in the hedgerow and tree classes, in spite of the overall accuracies being similar.

The difference maps show the areas in which error occurred. This was predominantly around the edges of the hedgerow and tree features, particularly around those trees embedded within the hedgerow.

5.4 Advanced linearisation – 12 directions

The super-resolved output from each of the three versions of the pixel swapping algorithm described in the previous section, mapped the predicted sub-pixel locations of land cover classes with reasonable accuracy. In the absence of fine spatial resolution remotely sensed imagery, the sub-pixel map, in most cases, will be of greater utility than a standard hard classification, particularly in the case of mapping features that are fine relative to their surroundings.

The linearised pixel swapping technique displayed an increase in mapping accuracy over the standard pixel swapping technique, through the use of anisotropically adjusted A_{ij} windows, based on an estimation of the direction of the linear feature within an image. In this stage of the research, the estimation of direction was made in a fixed number of directions, in either a 3 by 3 window (4 directions) or a 5 by 5 window (8 directions). Several possible means existed for increasing the number of directions and thus potentially increasing the precision of the estimation of the direction of the feature. Table 5.7 illustrates one possible approach in a 3 by 3 window, increasing the number of directions to 12.

The sequential grid cropping, with 17 directions in a 3 by 3 window, was applied to each of the input images, with the same parameters as in Section 5.3.4, above. Figure 5.24 shows the results.

Table 5.7: 12 possible directions in a 3 by 3 window, Scale: pixels

0°			15°			30°		
1	2	3	1	2	3	1	2	3
4	5	6	4	5	6	4	5	6
7	8	9	7	8	9	7	8	9
45°			60°			75°		
1	2	3	1	2	3	1	2	3
4	5	6	4	5	6	4	5	6
7	8	9	7	8	9	7	8	9
90°			105°			120°		
1	2	3	1	2	3	1	2	3
4	5	6	4	5	6	4	5	6
7	8	9	7	8	9	7	8	9
135°			150°			165°		
1	2	3	1	2	3	1	2	3
4	5	6	4	5	6	4	5	6
7	8	9	7	8	9	7	8	9

Such an approach would increase the likelihood of predicting a direction, especially on the edges of linear features where the spatial resolution is coarse relative to the width of the feature.

The linearised pixel swapping, using 12 directions in a 3 by 3 window, was applied to each of the input images, using the same parameters as in Section 5.3, above. Figure 5.14 shows the super-resolved output, Table 5.8 and Table 5.9 show the accuracy assessment (confusion matrices) and Figure 5.15 shows the difference images.

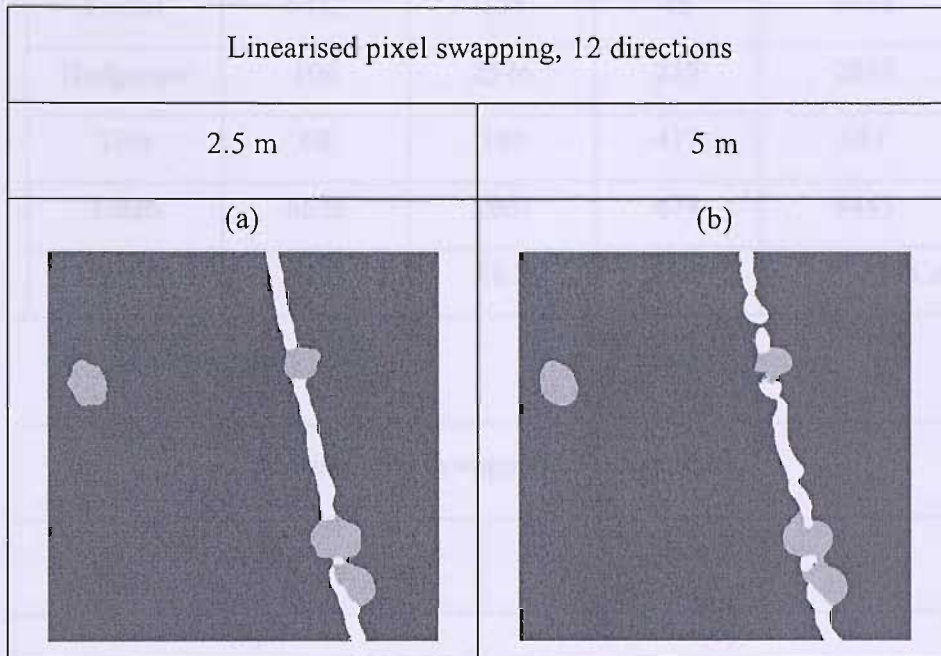


Figure 5.14: Super-resolved output. Linearised pixel swapping, 3 by 3 window, 12 directions (a) spatial resolution 2.5 m (b) spatial resolution 5 m

Table 5.8: Confusion matrices, linearised pixel swapping, 12 directions, spatial resolution: 2.5 m

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6549	91	13	6653	98.4
	Hedgerow	73	2768	26	2867	96.5
	Tree	28	107	546	681	80.1
	Totals	6650	2966	585	9863	
	UA (%)	98.4	93.3	93.3	% Correct: 96.7	

Table 5.9: Confusion matrices, linearised pixel swapping, 12 directions, spatial resolution: 5 m

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	6482	125	46	6653	97.4
	Hedgerow	106	2546	215	2867	88.8
	Tree	68	196	417	681	61.2
	Totals	6656	2867	678	9445	
	UA (%)	97.3	88.8	61.5	% Correct: 92.5	

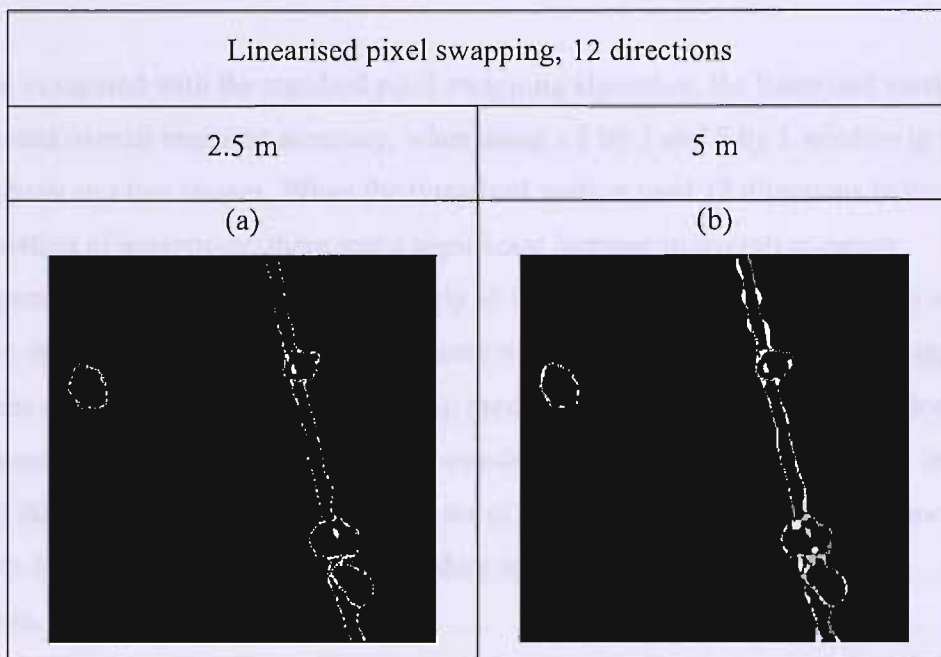


Figure 5.15: Difference maps. Linearised pixel swapping, 3 by 3 window, 12 directions (a) spatial resolution 2.5 m (b) spatial resolution 5 m. Legend: black – correct allocation, white – error of omission, grey - error of commission

The linearised pixel swapping technique using 12 directions was, at the 2.5 m spatial resolution more accurate than other versions of the linearised technique. However, the accuracy in the prediction of the hedgerow and tree classes at 5 m spatial resolution, in particular, decreased. The difference maps show the locations of the error, and it is

evident from these maps that the increase in the number of directions resulted in less error along the edges of the hedgerow and tree classes. Figure 5.15 illustrates this.

The difference images highlight the fact that complexity of the edges of the hedgerow class have not been characterised by the pixel swapping technique. In the super-resolved output in each site, the edges of the features have been over-simplified. When compared with the target image, the edges of the hedgerow and trees appear smoothed. The 12 directions was provided in this thesis as an example of developing the prediction of direction. Further discussion of avenues for developing the linearisation technique are discussed in Chapter 7.

5.5 Discussion of chapter

When compared with the standard pixel swapping algorithm, the linearised version increased overall mapping accuracy, when using a 3 by 3 and 5 by 5 window in the hedgerow and tree classes. When the linearised version used 12 directions in the calculation of anisotropy, there was a significant increase in overall accuracy compared with other versions, particularly in the 5 by 5 window. The accuracy with which the class containing the linear feature was predicted increased, which suggests that the direction of the feature was being predicted with greater precision. Indeed, confusion matrices for each of the super-resolved outputs suggest that either a large set of directions (12 or more) or a small set of directions (4) should be used, since results for 8 directions in the 5 by 5 window were least accurate in each of the outputs.

A visual inspection of the 'direction maps' (Figure 5.6 and Figure 5.11) provided information for consideration in development of the pixel swapping algorithm. Firstly, in each field site, measured anisotropy in a 3 by 3 window yielded more accurate results than in a 5 by 5 window. Anisotropy, used in determining the shape of the A_{ij} window, was measured from the soft proportions. The 3 by 3 window predicted one direction per pixel. As a result, anisotropic exponential function windows tended to be linear in a single direction and accuracy was increased. In the 5

by 5 window, multiple directions in each pixel were predicted more frequently. As more directions were predicted in each pixel, the exponential function window became increasingly isotropic (that is, similar to the exponential function window used in the standard pixel swapping technique) and accuracy was not increased above that of the standard pixel swapping technique. When 12 directions were used, however, multiple directions were predicted less frequently, since the prediction of direction was more precise, and accordingly the exponential function window remained anisotropic.

Secondly, the direction maps indicate that there were “gaps” in A_{ij} between pixels (less spatial dependence). This was most evident when applied to the 5 m spatial resolution input imagery in the 5 by 5 window. These gaps are one factor in accounting for the standard pixel swapping technique and the linearised pixel swapping technique failing to map a continuous feature (Figure 5.12). In the 3 by 3 window, there were fewer gaps between pixels and anisotropy in pixels was much more pronounced, which resulted in more accurate results and ultimately, the prediction of a contiguous feature. A further evaluation of the effect of measured anisotropy and the anisotropy ratio is presented in Chapter 7 of this thesis.

It should be noted that the input imagery to which the prediction task was applied, did not represent an “easy” prediction task for the algorithm. Since the simulated imagery used in development was derived from colour aerial photography, it retained many facets of real imagery. For example, the edges of hedgerows and trees were geometrically complex, displaying minute changes in direction, which affected overall accuracy. Despite these complex characteristics, the algorithm performed well. In Chapter 6, the second component of the linearised pixel swapping model is presented, which presents one solution for mapping such complex geometric characteristics.

5.6 Conclusion

In this chapter, the standard pixel swapping algorithm was developed to increase the likelihood of predicting linear features in simulated remotely sensed imagery. Using information from soft proportions on the position and direction of linear features, the standard exponential distance decay model was adjusted, using anisotropy, to influence the prediction of linear features. These modifications led to an increase in average accuracy over the standard pixel swapping technique. The effect of the size of the window from which the directions was estimated was evaluated, and in general, using a 3 by 3 window provided most accurate results.

The algorithm was developed such that the size of the window used in estimating directions and the number of directions that could be estimated could be adjusted. Potential for increases in accuracy using the linearised technique was demonstrated by increasing the number of predicted directions to 12. Increasing the number of directions was shown to increase the precision with which the directions in the soft proportions were predicted and accordingly some increases in the accuracy of the technique were observed.

Chapter 6

Image fusion

Image fusion is the process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application. In this chapter, we will explore some of the most common methods for image fusion, including linear, non-linear, and wavelet-based methods.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Image fusion is a process of combining two or more images into a single image. This can be done for a variety of reasons, such as to improve the quality of the image, to highlight certain features, or to create a new image that is more useful than the original images. There are many different methods for image fusion, and the choice of method depends on the specific application.

Chapter 6

6. Image fusion

6.1 Introduction

In the previous chapter, the standard pixel swapping technique was modified to increase the likelihood of resolving linear features within remotely sensed imagery. The output from the modified technique was on average 5% more accurate than the standard pixel swapping technique. In this chapter, the results of an additional development to the pixel swapping method are presented. Common sources of remotely sensed imagery, such as Ikonos or Quickbird (as used in this research), are available as both multispectral and panchromatic data. Panchromatic imagery often offers finer spatial resolution, enabling more accurate delineation of land cover classes. In this second stage of development of the pixel swapping algorithm, an image fusion framework was incorporated, enabling the use of fine spatial resolution panchromatic imagery when determining the quality of a swap.

6.2 Image Fusion

Image fusion is an image processing concept implemented to combine multiple sources of information. The resultant combination theoretically provides more information than would be available from each data source individually. A common approach to image fusion is the combination of two or more remotely sensed images, where the individual images contain information that the others do not, such as via a different spatial resolution or more spectral information through different spectral wavebands. For example, where the mapping of individual species of vegetation is

required, a fine spatial resolution panchromatic image could be combined with a hyperspectral image to create a fine spatial resolution hyperspectral image, i.e., an image with the spatial resolution of the panchromatic image and the spectral detail of the hyperspectral image, which could increase the accuracy with which the species are mapped at the panchromatic image spatial resolution.

A review of existing image fusion techniques was presented in Chapter 2 of this thesis. In this chapter, a new approach to image fusion is presented, whereby a panchromatic image was fused with multispectral imagery during pixel swapping. Simulated multispectral and panchromatic imagery were used to evaluate the fusion technique.

Raw data contained within a panchromatic image is commonly in the form of measured intensity values. The measured intensity value of each pixel represents a mixture of the land cover classes found within the pixel at the ground level – the intensity value is therefore assumed to be a linear combination of the spectral response of each of these land cover classes. Using information derived during soft classification on the location of spectral endmembers (pure pixels) in the multispectral imagery and assuming accurate geometric coregistration of the two images, it is possible to predict the location of pure pixels in the panchromatic imagery. The average intensity for each class can then be calculated, by averaging the panchromatic pixels that correspond to the pure pixels in the multispectral image. Figure 6.1 illustrates this. The location of a pure pixel (e.g., Class 1) in the multispectral imagery is used to locate the equivalent four pixels in the panchromatic image. The average value of the panchromatic pixels predicts an average intensity value of Class 1 of 36. This operation is performed for every pure pixel in each class.




Figure 6.1: Preliminary stage of fusion model

6.2.1 Methodology

The pixel swapping technique is initialised as per the standard technique (Chapter 4). For each of the two randomly selected sub-pixels (r_1, r_2), the binary values of the sub-pixels in a window with the selected sub-pixel in the centre are converted to intensity values using the average class intensities derived in stage 1 of the model. For example, during pixel swapping, pixel proportions are super-resolved to a binary hard classification at the sub-pixel scale, that is, each sub-pixel is assigned to be only one class (Class 1, 2 or 3). If the predicted average intensity for land cover class 1 was 37, then every sub-pixel assigned to land cover class 1 in the window around the selected sub-pixel would be assigned a value of 37. The intensity values in the window are averaged (P_{r_1}, P_{r_2}). The two selected sub-pixels are swapped, and the average intensity values for each sub-pixel are recalculated ($P_{r_1}^*, P_{r_2}^*$). The differences are calculated:

$$P_{r_1}d = P_{r_1} - P_{r_1}^* \quad (7)$$

$$P_{r2d} = P_{r2} - P_{r2}^*$$

and an overall P_{ij} value is calculated:

$$P_{ij} = P_{r1d} + P_{r2d} \tag{8}$$

Figure 6.2 illustrates the fusion model. The classified multispectral image is super-resolved into binary class values (1). Two sub-pixels are selected randomly within a pixel. Using the average class intensities derived from stage 1, an average intensity value for each relevant pan-pixel is calculated. Sub-pixels are swapped and the pan-pixel intensity value is recalculated. For each selected sub-pixel, the corresponding pixel in the panchromatic image is located (2). The value of this pixel is returned (3). The difference between the calculated average intensity value for each pan-pixel and the “actual” value before and after the swap is calculated.

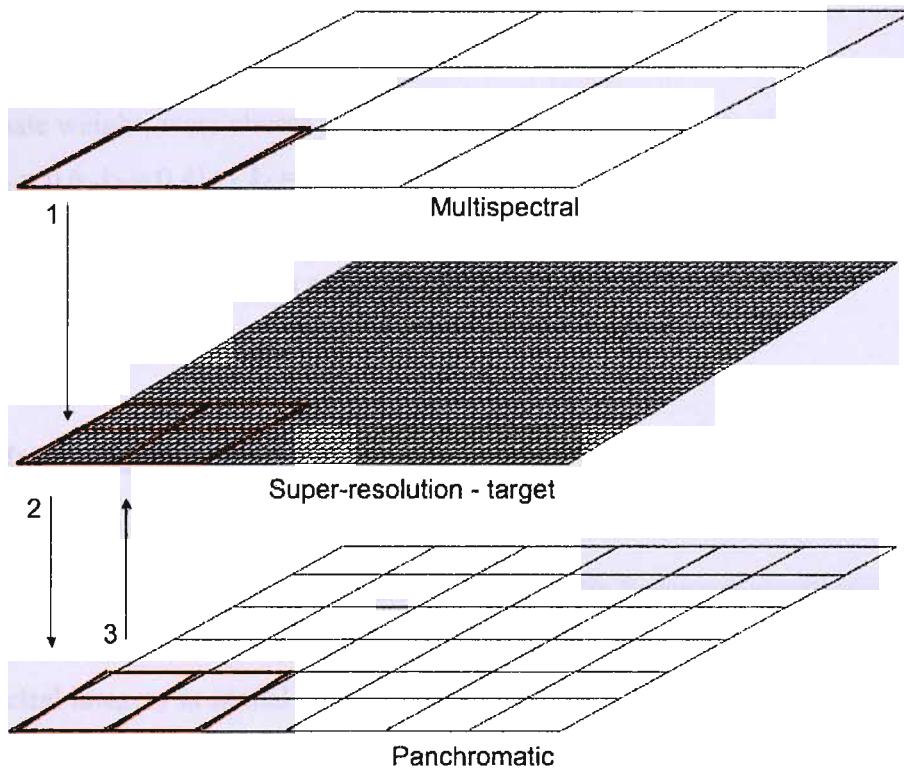


Figure 6.2: The fusion model

6.2.2 Defining the energy function

In the standard technique, a swap is retained on the basis of an attractiveness function (A_{ij}) using the exponential distance-decay function in a moving window. The linearised version of the pixel swapping technique uses an anisotropically adjusted exponential distance-decay window to calculate A_{ij} . The fusion process uses the anisotropically corrected exponential distance-decay window to calculate A_{ij} as well as P_{ij} . A linear combination of A_{ij} and P_{ij} determines whether a swap is accepted, thus:

$$E_{ij} - k_1 A_{ij} + k_2 P_{ij} \quad (9)$$

where k_1 and k_2 are weights and $k_2 = (1 - k_1)$. If the swap resulted in an overall increase in E_{ij} the swap was retained (as opposed to the simulated annealing approach where a swap was retained if the swap resulted in a decrease in the energy function O_{ij} : [$O_{ij} = E_{ij}^* - E_{ij}$]). Otherwise, an annealing schedule (Chapter 4) was used to determine whether the swap was retained or rejected.

Appropriate weights were chosen through repeated testing. Weights of $\{k_1 = 0.5, k_2 = 0.5\}$, $\{k_1 = 0.6, k_2 = 0.4\}$, $\{k_1 = 0.75, k_2 = 0.25\}$ and $\{k_1 = 0.9, k_2 = 0.1\}$ were evaluated. Weights of $\{k_1 = 0.75, k_2 = 0.25\}$ were found to be optimal. This evaluation is presented later in a subsequent section of this chapter.

6.3 Initialisation

The following sections present the results of applying the linearised pixel swapping technique with fusion to the input imagery described in Chapter 3.4.3. Simulated multispectral imagery at spatial resolutions of 2.5 m and 5 m were used as input. On the basis of the results in Chapter 5, the size of the linear anisotropic window was set at 3 by 3. A zoom factor of 10 was used to super-resolve the 2.5 m spatial resolution input imagery and a zoom factor of 20 was used to super-resolve the 5 m spatial resolution input imagery to the same spatial resolution as that of the input (25 cm).

Confusion matrices and difference images were constructed for each super-resolved output to evaluate the accuracy of the technique.

6.3.1 Simulating panchromatic imagery

A panchromatic image (Figure 6.3) at a spatial resolution of 1.25 m was simulated from the classified aerial photography described in Chapter 3. The process to create the panchromatic imagery was similar to that used to generate the simulated multispectral imagery, except that estimated class intensity values were used instead of binary class values (6.2.1). The advantage with simulating panchromatic imagery from the original multispectral imagery was that the panchromatic image and multispectral images used as input to the pixel swapping technique were automatically geometrically coregistered.




Figure 6.3: Simulated panchromatic image

6.4 Results

6.4.1 Field site (spatial resolution 2.5 m)

The linearised pixel swapping technique was applied to the 2.5 m spatial resolution input imagery and fused with the simulated panchromatic image, at a zoom factor of ten resulting in super-resolved output images at the same spatial resolution as the

original aerial photography (0.25 m). Sub-pixels were allocated randomly and the pixel swapping technique iterated until it converged. Weights of 0.75 were applied to A_{ij} and 0.25 applied to P_{ij} in the fusion process. Figure 6.4 displays the results.

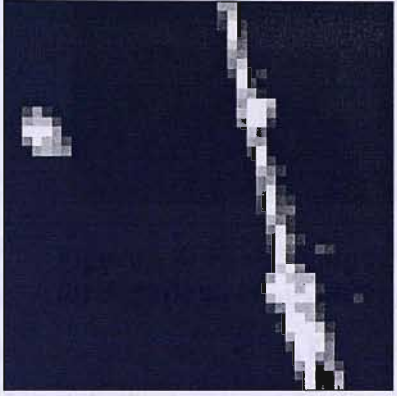
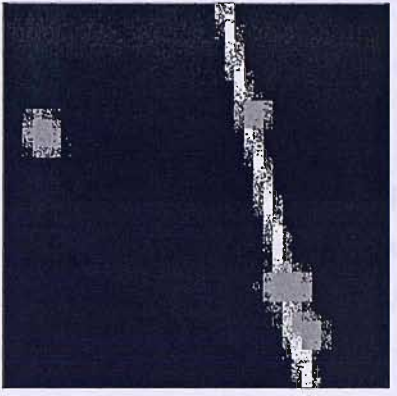
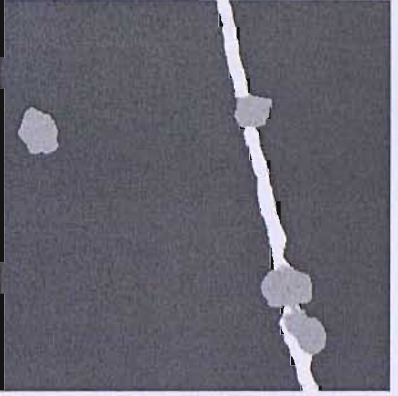
Spatial Resolution (SR)	Input imagery	Random allocation	Super-resolved output
2.5 m	(a) 	(b) 	(c) 

Figure 6.4: Fusion output, (a) Input imagery, 2.5 m SR, (b) Random allocation, (c) Super-resolved output

A visual inspection of the super-resolved output in Figure 6.4(c), revealed a close resemblance to the input imagery. The hedgerow was predicted as a continuous feature and the geometric characteristics of edges of the hedgerow and the trees, particularly the trees embedded within the hedgerow, bore a close resemblance to those in the input imagery. This was most evident on the lone tree, where the shape was a very close match to the input imagery. The two small areas of hedgerow that were visible in the input imagery in the topmost and middle trees embedded within the hedgerow, have been “removed”, i.e., they have been clustered with the main hedgerow, as opposed to remaining within the trees. Table 6.1 displays the accuracy assessments.

Table 6.1: Confusion matrix, super-resolved output (2.5 m input SR)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	9043	17	12	9072	99.6
	Hedgerow	15	709	3	727	97.5
	Tree	14	60	328	402	81.6
	Totals	9072	786	343	10080	
	UA (%)	99.7	90.2	95.6	% Correct: 98.8	

Overall accuracy was estimated at 98.8%. The cereal class was mapped with a high level of accuracy, as was the hedgerow class. The tree class was the least accurate, which the confusion matrix suggested was due to confusion between the tree and hedgerow classes. Figure 6.5 displays the difference image, which shows the location of error in the super-resolved output.




Figure 6.5: Difference image, super-resolved output, 2.5 m SR

Error is visible in the difference image along the edges of the hedgerow and tree classes. Misclassification occurred between the two trees embedded within the hedgerow. The small areas of hedgerow within the embedded trees, which were misclassified during the technique, are also visible in the difference, as white sections where the hedgerow class has been misclassified as the tree class. Particular attention should be paid to the areas of the difference image which were black, i.e., the

predicted image matched the target image, in the areas where error was expected to occur. For example, along the edge of the tree separate from the hedgerow, there are several areas of zero error, that is, the predicted shape of the tree matched the input imagery. The same was true of several sections along the edges of the hedgerow. A significance test was applied to evaluate the differences between the linearised pixel swapping technique in a 3 by 3 and the linearised pixel swapping technique with fusion in a 3 by 3 window. These results are displayed in Table 6.2.

Table 6.2: Significance test, linearised pixel swapping and linearised pixel swapping with fusion

		Linearised Pixel Swapping with fusion		
		Correct	Incorrect	Total
Linearised Pixel Swapping	Correct	10054	0	10054
	Incorrect	21	126	147
	Total	10080	126	10201

$\chi^2 = 21$. Significant at 0.05%.

Inspection of the significance test (Table 6.2) reveals that the fusion technique operated more accurately than the linearised pixel swapping technique.

6.4.2 Field site (spatial resolution 5 m)

The linearised pixel swapping technique was applied to the 5 m spatial resolution input imagery and fused with the simulated panchromatic image, at a zoom factor of 20 resulting in super-resolved output images at the same spatial resolution as the original aerial photography (0.25 m). Sub-pixels were allocated randomly and the pixel swapping technique iterated until it converged. Weights of 0.75 were applied to A_{ij} and 0.25 applied to P_{ij} in the fusion process. Figure 6.6 displays the results.

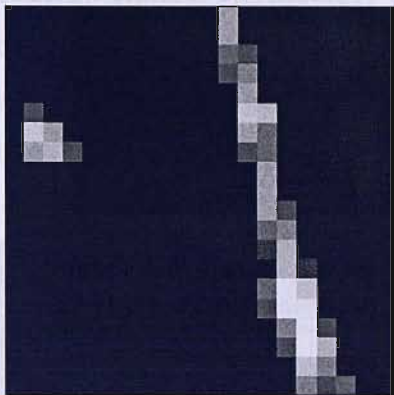
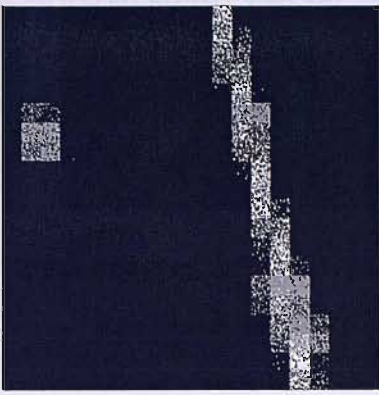
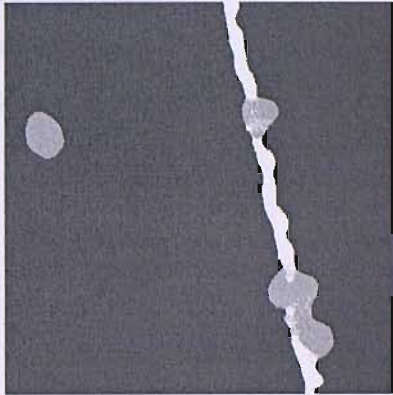
Spatial Resolution (SR)	Input imagery	Random allocation	Super-resolved output
5 m	(a) 	(b) 	(c) 

Figure 6.6: Fusion output, (a) Input imagery, 2.5 m SR, (b) Random allocation, (c) Super-resolved output

Figure 6.6(c) displays the super-resolved output. The hedgerow was predicted as a continuous feature, which shows an obvious improvement over the standard pixel swapping technique. Some of the geometric characteristics of the hedgerow and tree classes were lost and there was obvious confusion in the trees embedded within the hedgerow. In particular, the edges of the hedgerow show a greater degree of curvilinearity than was evident in the input image. Table 6.3 displays the accuracy assessment.

Table 6.3: Confusion matrix, super-resolved output (5 m input SR)

		Predicted				
		Cereal	Hedgerow	Tree	Totals	PA (%)
Reference	Cereal	8947	32	93	9072	98.6
	Hedgerow	43	680	4	727	93.5
	Tree	42	88	272	402	67.6
	Totals	9032	800	369	9899	
	UA (%)	99.0	85.0	73.7	% Correct: 97.0	

Overall accuracy was estimated at 97%. The cereal class was mapped with a high level of accuracy, as was the hedgerow class. The tree class was the least accurate, which the confusion matrix suggested was due to confusion between the tree and hedgerow classes. The accuracy was, however, greater than any of the equivalent super-resolved outputs from the previously described techniques (standard pixel swapping, linearised pixel swapping). Figure 6.7 displays the difference image, which shows the location of error in the super-resolved output.




Figure 6.7: Difference image, super-resolved output, 5 m SR

Error is visible in the difference image along the edges of the hedgerow and tree classes. Misclassification occurred between the two trees embedded within the hedgerow. The small areas of hedgerow within the embedded trees, which were misclassified during the technique, are also visible in the difference, as white sections where the hedgerow class has been misclassified as the tree class.

6.5 Discussion of results

Through visual inspection, the super-resolved output from the 2.5 m and 5 m spatial resolution multispectral inputs resembled the target image. At each spatial resolution, a contiguous feature was mapped. At both spatial resolutions, the accuracy of the super-resolved output in the cereal class was greater than 98%. Overall, the super-

resolved output was more accurate in each class when the input was at a spatial resolution of 2.5 m than at a spatial resolution of 5 m.

When compared with results from the linearised pixel swapping technique with the 3 by 3 window (Chapter 5), at each spatial resolution the super-resolved output from the fused method was more accurate in each class at both spatial resolutions. The difference images show less error along the edges of the images and less misallocation between the hedgerow and tree classes. In particular, through visual inspection, the edges of the target image are more realistically predicted in the output of the fusion method than the output of the standard pixel swapping method. Figure 6.8 illustrates this at the 2.5 m spatial resolution. Particular points of interest have been highlighted (black squares) in the target image for reference against each super-resolved output.




Figure 6.8: Comparison of super-resolution output with target image: (a) Target image, (b) Linearised pixel swapping output, (c) Linearised pixel swapping output with image fusion

In the linearised pixel swapping output, the edges and shapes of each of the tree features in the tree class have been smoothed. In the output of the fusion method, the shape and edges of each of the trees more closely resemble the original. For example, in Figure 6.8, consider the tree that is separate from the hedgerow. Figure 6.8(a) shows the target. In Figure 6.8(b), the tree was represented as a circular feature, with obvious smoothing on the edges. In Figure 6.8(c), the shape of the tree has been more realistically predicted, with sub-pixel level intricacy. Consider also the two trees

embedded within the hedgerow. In the linearised output (Figure 6.8(b), the two trees have been predicted as two separate objects, whereas in the target image it appeared that the two trees were “joined” (where partial canopies overlapped). In the linearised with fusion output (Figure 6.8(c)), the shapes of the trees were a closer match to the target than the linearised output and the canopy was more realistically predicted. The complexities of the edges of the hedgerow class are more realistically predicted in the fusion output than in the linearised pixel swapping method. The difference images support these points. This represents a key advantage to the fusion algorithm over the standard and linearised pixel swapping algorithms.

There are errors in the output. For example, in both the linearised pixel swapping output and the fusion output, the small area of hedgerow class that is visible within the tree within the hedgerow class at the upper end of the image is incorrectly mapped - the hedgerow sub-pixels are adjoined to other areas of hedgerow. The two trees in the lower area of the hedgerow are mapped as a contiguous feature in both super-resolved outputs, whereas in the original they are separate features. The fusion method predicts the locations of these two trees more accurately than the linearised pixel swapping method.

6.6 Discussion of chapter

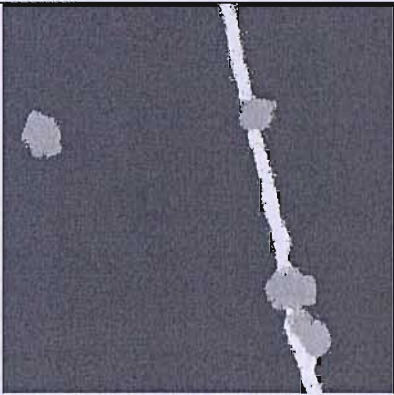
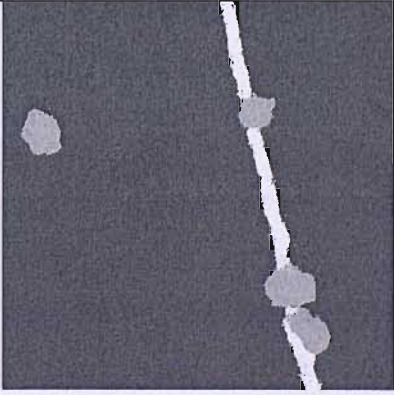
In this chapter, the linearised version of the pixel swapping technique was further developed to incorporate a fusion component. It has been previously noted, that the super-resolved outputs from the standard and linearised versions of the technique were statistically accurate but visually the edges of the classes appear smoothed. In the fusion output, however, the edges of land cover classes in the super-resolved output of the fusion technique when applied to simulated imagery are represented more realistically when compared with the target image. For example, in the standard and linearised pixel swapping technique, the individual tree in the left of the image was mapped as a more or less circular feature. In the output from the fused method, the individual tree and the edges of the hedgerow visually display more resemblance to the actual boundaries of these classes in the target image.

In developing the linearised pixel swapping algorithm to incorporate fusion, there were important considerations and decisions to be made. A discussion of these factors is now presented.

6.6.1 Parameterisation

6.6.1.1 Weightings

The fusion model introduced a further parameter to the pixel swapping technique – the weighting applied to $A_{ij}(k_1)$ and $P_{ij}(k_2)$. The weighting was used to determine the effect of each of these two values on the acceptance of a swap. The optimal value for 2.5 m spatial resolution input imagery was $k_1 = 0.75$, $k_2 = 0.25$. These values were derived through repeated testing. Figure 6.9 shows the effect of the weighting. A_{ij} was essential to clustering of sub-pixels whereas P_{ij} increased the precision of the clustering.

	Weights	Super-resolved output	% Correct
(a)	$k_1 = 0.5$ $k_2 = 0.5$		97.1
(b)	$k_1 = 0.6$ $k_2 = 0.4$		97.9

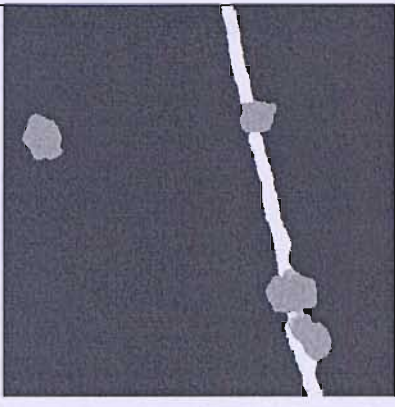
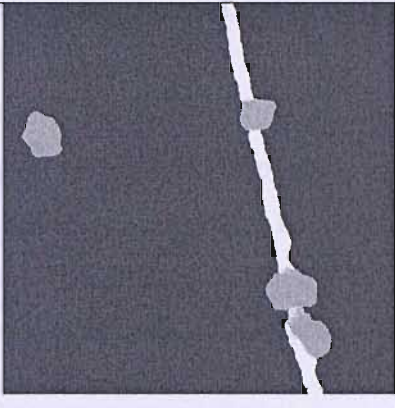
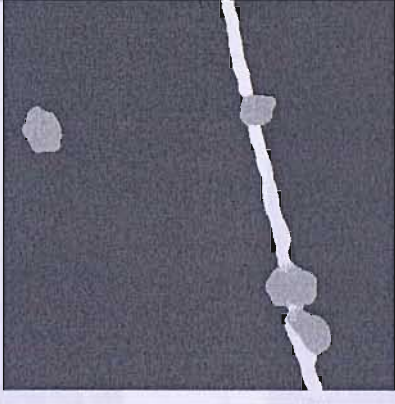
(c)	$k_1 = 0.75$ $k_2 = 0.25$		98.8
(d)	$k_1 = 0.9$ $k_2 = 0.1$		98.1
(e)	$k_1 = 1$ $k_2 = 0$		

Figure 6.9: Effect of weights on super-resolution output, (a) $k_1 = 0.5, k_2 = 0.5$, (b) $k_1 = 0.6, k_2 = 0.4$, (c) $k_1 = 0.75, k_2 = 0.25$, (d) $k_1 = 0.9, k_2 = 0.1$, (e) $k_1 = 1, k_2 = 0$,

6.6.1.2 Spatial resolution


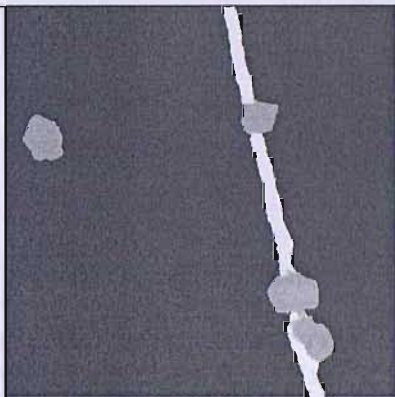

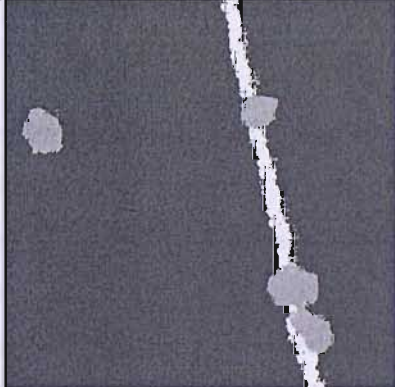
The extent to which fusion increased the accuracy of the prediction of the edges was, in this research, not dependent on the spatial resolution of the input imagery. 2.5 m spatial resolution and 5 m spatial resolution remotely sensed imagery was fused with 1.25 m spatial resolution panchromatic imagery. Therefore, the spatial resolution of the panchromatic image was twice or four times as fine as the multispectral imagery.

At both spatial resolutions, the effect of the fusion was an average increase in accuracy of 1-2% in each class.

6.6.1.3 Image coregistration

An important factor in any fusion technique is that all sources of data are accurately coregistered to each other. Coregistration of imagery to be used in the fusion is beyond the control of the pixel swapping algorithm, but nevertheless requires careful consideration.

In this research, simulated imagery were used and so perfect coregistration of the multispectral and panchromatic images was assured. However, the effect of inaccurately coregistered imagery was investigated. Panchromatic imagery was simulated that incorporated a horizontal shift, by displacing the panchromatic image to the left. Figure 6.10 shows the results.

Displacement: Pixels / Distance		Super-resolved output
(a)	<p>0 / 0</p> 	
(b)	<p>1 / 1.25 m</p> 	




Figure 6.10: Effect of inaccurate coregistration of input imagery, super-resolved output (a) Perfect registration, (b) 1 pixel displacement (1.25 m), (c) 2 pixel displacement (2.5 m), (d) 3 pixel displacement (3.75 m), (e) 4 pixel displacement (5 m)

The effect of inaccurate coregistration of the imagery on the accuracy of the super-resolved output is directly related to the spatial resolution of the imagery, the number of classes within the imagery, the spatial arrangement of land cover classes within the imagery and the size of the window used to measure A_{ij} and values of k_1 and k_2 . In this example, the deleterious effect of inaccurate coregistration was most prevalent when the displacement was 2.5 m, equivalent to the spatial resolution of the input imagery

and approximately equivalent to the width of the hedgerow class, as erroneous location of panchromatic pixels used in the fusion process will affect the acceptance of swaps. However, when the displacement was greater than the size of the features in the imagery, the effect of the displacement was to use only cereal class panchromatic pixels in the calculation of P_{ij} and, therefore, had a more limited effect. In real remotely sensed imagery where there may be more classes or the spatial arrangement of those classes may be more complex, the effect of inaccurate coregistration will likely be greater.

6.7 Conclusion

Introducing a fusion element to the linearised pixel swapping technique increased the accuracy with which the land cover classes were mapped. In particular, edges of classes (for example, the hedgerow and tree classes) were more comparable to the target imagery.

Accuracy assessment in the form of difference images, confirmed that along the edges of the features, the output from the fusion process more accurately resembled the target image than the output from the standard and linearised pixel swapping technique. This is of great utility in giving a realistic representation of the shape and geometry of the boundaries of these classes on the ground.

Further research

Chapter 7

Discussion

and a number of other studies. First, they are not
comparable to the present study because they are
not designed to investigate the relationship between
the variables of the study. In addition, the
methodology used in the present study is different
from the methodology used in the other studies.

References

Long, J. H. (1980). The relationship between
the variables of the study. In J. H. Long (Ed.),

Chapter 7

7. Discussion and further research

7.1 Introduction

The objective of the research presented in this thesis was to develop an algorithm to facilitate super-resolution mapping of linear land cover features. The super-resolution pixel swapping technique presented in this thesis, which included an estimation of anisotropy, has been demonstrated to be able to map linear land cover features from remotely sensed imagery with a spatial resolution that was around the same size as or coarser than the width of those features. An image fusion component of the model increased the ability of the algorithm to delineate the complexities of the edges of linear and other land cover features.

In this penultimate chapter, an in depth discussion of the key components of the model that make up the pixel swapping algorithm described in thesis is presented. Initially, the overall performance of the algorithm is examined along with the effect of the assumptions that the model makes on the accuracy of the super-resolved output, and some pertinent limitations of the model. Then the model, including the main parameters such as the zoom factor, is discussed. The chapter presents possibilities for future research and closes with a conclusion.

7.1.1 Overall performance

The pixel swapping technique is an iterative technique where one swap (two sub-pixels) is made in each pixel of the image per iteration. The standard technique was configured to swap sub-pixels after calculating an attractiveness value for every sub-

pixel in each pixel. The two sub-pixels that when swapped were most likely to result in an increase in spatial dependence (as indicated by the attractiveness measure) were swapped. The standard configuration was resource intensive as each iteration required substantial computation. In this research, the standard technique was modified to select sub-pixels randomly as per spatial simulated annealing approaches. This modification was used in each subsequent version of the technique (linearisation and fusion). This resulted in substantially less computation per iteration. The number of iterations required for the model to converge increased dramatically. However, the processing time for each iteration decreased. Table 7.1 shows an example of the performance.

Table 7.1: Example processing time for a 3-class classified image, Size: 100 pixels by 100 pixels, zoom factor: 5

Convergence criteria	Computed sub-pixel selection	Random sub-pixel selection
No. iterations	20	5000
Processing time (minutes)	30	< 1

The size of the image, the number of land cover classes and zoom factor each affected overall processing time. When increasing any of these factors when using random sub-pixel selection, the increase in processing time was approximately linear. However, when computing sub-pixel selection, the increase was approximately exponential. Random sub-pixel selection was occasionally affected by local minima (i.e., where different swaps produced the same result and were always accepted) and the algorithm would not converge. The use of an annealing schedule, however, reduced the likelihood of local minima, although increased the number of iterations required for the technique to converge.

7.1.2 Assumptions and considerations

In order to develop an algorithm for land cover mapping, certain assumptions about land cover were made. The underlying assumption of the pixel swapping model was that land cover classes within remotely sensed imagery were spatially dependent. The

goal of pixel swapping was to maximise spatial dependence. Spatial dependence was maximised by calculating a sub-pixel distance function using the exponential distance-decay model (A_{ij}). The distance function was modified by an anisotropic correction to increase the likelihood of predicting linear features using directions predicted from soft proportions. The correction facilitated the prediction of multiple directions within each pixel, which was assumed likely to increase the accuracy of the super-resolved output.

When applied to remotely sensed imagery, the impact of these assumptions could be evaluated. By assuming spatial dependence, for example, small groups of isolated sub-pixels of a particular class were occasionally incorrectly clustered with larger groups of sub-pixels of the same class. For example, in the field used in Chapter 5, a small part of the canopy of the lowermost tree within the hedgerow class was open, so that soft classification correctly predicted this small area as hedgerow. In the target image (in Section 3.4.3), this is shown as a small area of white. On most occasions in the super-resolved outputs, this area of hedgerow was incorrectly clustered with the main areas of the hedgerow class and did not remain in the middle of the tree. In addition, non-linear features will have the tendency to become elongated and linearised as a result of the anisotropic correction to the distance-decay window.

The pixel swapping algorithm accurately mapped features that were approximately the width of a pixel. In addition, characteristics of edges of features that were in adjoining pixels were also mapped accurately. However, features that were smaller than a pixel were often mapped inaccurately (for example, the small patch of hedgerow described previously). In general, as the size of the feature relative to the size of the pixel increased, the accuracy with which it was mapped increased.

7.1.2.1 Accuracy assessments

Accuracy assessment techniques were implemented to assess the accuracy of the super-resolved output of each technique and inform on the efficiency of the algorithm and the effect of the developments made to it. For example, confusion matrices were constructed to provide an overall accuracy statistic to provide a simple metric for

evaluating the effectiveness of the technique. However, in some cases, these accuracy assessments may not provide a realistic representation of the utility of the super-resolved output. In some cases, particularly when applied to input imagery at a spatial resolution of 5 m, overall accuracy was estimated at > 93%. Significance tests were used to evaluate the difference in accuracy between each of the techniques. However, visual inspection of the super-resolved output often suggested that the output was not comparable to the input imagery, for example, where a hedgerow was not predicted as a contiguous feature (e.g., Figure 5.14). Certain measures were implemented to evaluate this, for example, producing accuracy assessments of specific areas of a super-resolved output and the use of difference images. Indeed, the arbitrary values established in Chapter 2 to describe what results would be deemed “accurate”, may have been set too high and exceeded the limits of the model and, remembering that the objective of the research presented in this thesis was to

One of the primary benefits of the fusion technique was the increased precision with which the edges of features, such as the hedgerow, were predicted. In some cases, these increases were not reflected in the accuracy assessments, that is, the overall accuracy was not increased so as to be significant. That is, a small increase (e.g., 0.5%) in overall accuracy would not accurately reflect the improved delineation of feature boundaries at the sub-pixel level. However, such improvements were clearly evident in the super-resolved output and through visual inspection of the difference images, where error visible in the difference images along the edges of features decreased.

7.1.3 Limitations of the algorithm

The pixel swapping technique was dependent on the accuracy of the proportions information derived from the soft classification. As Chapter 4 demonstrated, where soft classification was unable to unmix land cover classes accurately, the accuracy of the super-resolved output was affected.

The spatial resolution of the input imagery was critical to super-resolving contiguous linear features. Where degrading the spatial resolution of the classified image (to simulate proportions) resulted in two or more contiguous pixels for which the

proportion of the hedgerow class were both less than 50%, the technique did not super-resolve a contiguous feature. Additionally, the random allocation of sub-pixels occasionally prevented the mapping of a contiguous feature if the allocation was such that sub-pixels allocated to the feature were further apart than the size of the window used to calculate A_{ij} . In this instance, spatial dependence between sub-pixels was not sufficient to cluster them as a contiguous feature, since spatial dependence between sub-pixels allocated to other classes was greater.

7.2 Discussion of the model

The pixel swapping algorithm was developed around a model which utilised a complex set of parameters. In this section, the parameters of the model are discussed. The effects of altering the parameters are evaluated and some considerations on determining optimal settings are presented.

The standard pixel swapping model has four parameters:

- Zoom factor
- Size of exponential function window
- Non-linear parameter (NLP) of distance-decay model
- Cooling rate of annealing schedule

In addition, the linearised pixel swapping model introduced two extra parameters:

- Size of direction window (and thereby number of predicted directions)
- Anisotropy ratio

The fusion component required one additional parameter:

- Weight applied to A_{ij} and P_{ij}

In total, therefore, the full model has seven parameters. Evaluating the effects of these parameters on the super-resolved output was performed to assess the operation of the

algorithm and to facilitate the prediction of optimal conditions. The majority of parameters were affected directly by one of the other parameters. For example, the size of the exponential function window and the rate of decay of the model (NLP) changed the scale of spatial dependence, since the rate of decay in the exponential function was calculated as a function of the size of the window and the NLP. Consider two windows (3 pixels by 3 pixels and 5 pixels by 5 pixels) together with a constant NLP. In this example, the rate of decay of the 3 by 3 window will be less than the rate of decay in a 5 by 5 window.

The value of each parameter was dependent on the characteristics of the input remotely sensed imagery, the frequency of spatial variation within the imagery and the intended use of the super-resolved output. Optimal parameter conditions will, therefore, vary considerably between datasets and need to be evaluated on a case-by-case basis.

Discussion of each of the parameters is now presented.

7.2.1 Zoom factor

Perhaps the most important parameter of the model was the zoom factor. The zoom factor was used to determine how many sub-pixels were created along each axis of a pixel. Increasing the zoom factor resulted in super-resolved output at a finer spatial resolution. A finer spatial resolution theoretically enabled more accurate identification of linear features. The zoom factor had a direct impact on the performance of the technique – as the zoom factor increased, the technique became more computationally intensive and required more iterations to converge.

The pixel swapping algorithm was developed such that the zoom factor could be easily varied. In order to evaluate the effect of the zoom factor, the linearised pixel swapping technique (using a 3 by 3 window to calculate directions) was applied to 2.5 m spatial resolution simulated imagery at different zoom factors and the super-resolved output and resultant accuracy assessments were evaluated. Figure 7.1 shows the super-resolved output at each zoom factor.

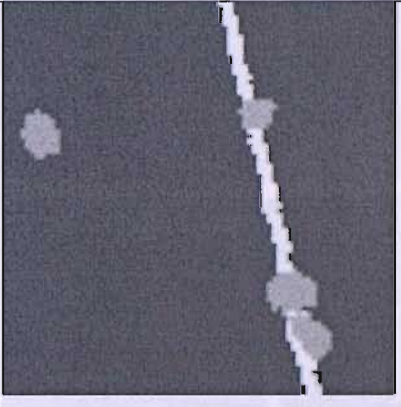
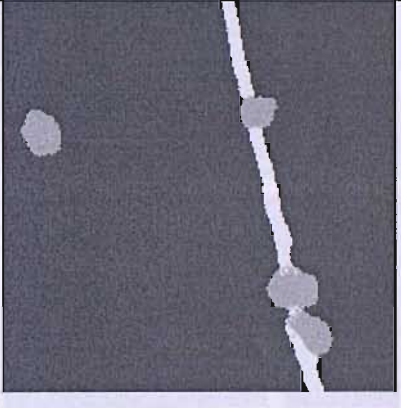
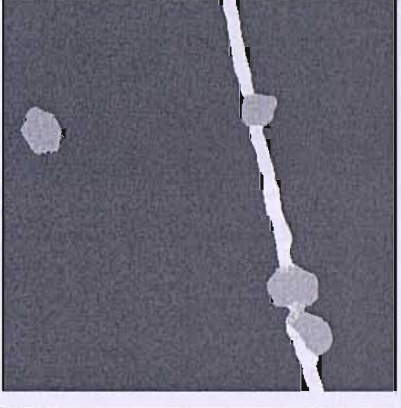
	Zoom factor / SR (m) / Dimensions (pixels)/	Overall accuracy (%)	Super-resolved output
(a)	2 / 1.25 / 80	95.1	
(b)	5 / 0.5 / 200	97.6	
(c)	8 / 0.32 / 320	97.9	

Figure 7.13 Effect of zoom factor on super-resolution accuracy. Zoom factor, SR, and dimensions are shown in the table.

As the zoom factor of 2 the super-resolved output of the

was fairly accurate. As the zoom factor increased, the super-resolution output became increasingly more accurate. However, as the zoom factor increased beyond a certain point (zoom factor = 70), the additional accuracy achieved by the super-resolution algorithm (Figure 7.13a-f).

In this context, the zoom factor of 80 is particularly interesting, as it super-resolved the input to match the spatial resolution of the output, which resulted in a 97.9% accuracy.

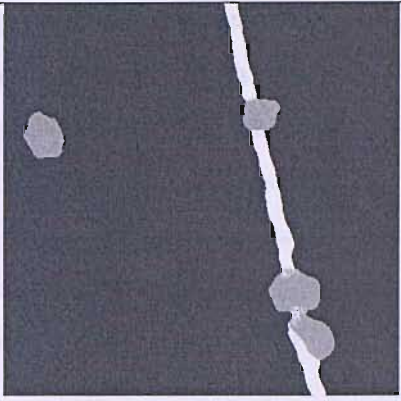
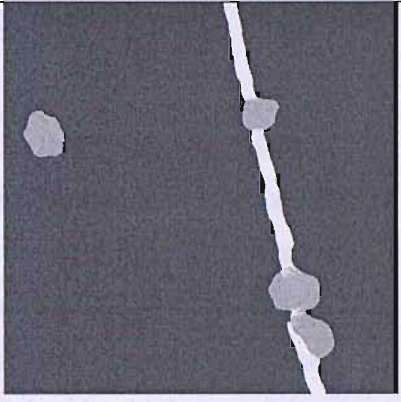
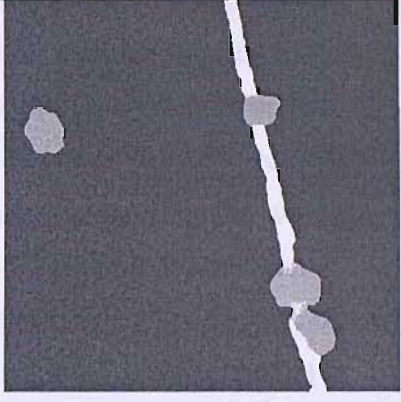
(d)	10 / 0.25 / 400	98.6	
(e)	12 / 0.20 / 480	98.6	
(f)	15 / 0.16 / 600	98.5	

Figure 7.1: Effect of zoom factor on super-resolved output (linearised pixel swapping, 3 by 3 window, 2.5 m SR)

At a zoom factor of 2 the super-resolved output of the hedgerow class and tree classes were poorly delineated. As the zoom factor increased, the super-resolved output started to resemble the target image. However, as the zoom factor increased beyond a certain point (zoom factor > 10), so the additional benefits relative to the target image diminished (Figure 7.1(e,f)).

In this research, the zoom factor on this particular dataset was 10, as it super-resolved the input to match the spatial resolution of the target, which enabled simple accuracy

assessment and the creation of difference images to represent visually areas of error. However, the results above suggest that there is little utility in using zoom factors greater than 10, since little or no extra information was discernible from the super-resolved output compared with the computational cost and the effect of smoothing on the edges of the classes, particularly in the example of the fusion technique, where the results displayed increased accuracy in predicting the edges of features, benefits that were less apparent if the input was super-resolved beyond a zoom factor of 10.

7.2.2 Exponential function window parameters

The exponential distance-decay model was used in a moving window to calculate the attractiveness function (A_{ij}) and was central to the swapping process. The size of the window determined how many sub-pixels surrounding the selected sub-pixels were used in calculating A_{ij} . The non-linear parameter defined the rate at which the distance model decayed spatially, that is, the rate at which values in the function window decreased from 1 to 0 and, thus, the value of A_{ij} . A model with a high rate of decay placed greater weights on sub-pixels that were close to the selected sub-pixel, whereas a model with a slower rate of decay performed in the opposite manner. A model with a high rate of decay accepted fewer swaps, converged faster and the super-resolved output was less prone to local minima. However, the super-resolved output was less accurate than if a lower rate of decay was used. Figure 7.2 illustrates this when applied to 2.5 m spatial resolution remotely sensed imagery in an exponential function window 9 pixels by 9 pixels with a zoom factor of 10, using the linearised pixel swapping algorithm with fusion.




Figure 7.2: Non-linear parameter of the exponential distance decay model, linearised pixel swapping with fusion: (a) low rate of decay, (b) increased rate of decay, (c) highest rate of decay

Figure 7.3 illustrates the effect of changing the size of the exponential function window. The pixel swapping technique was applied to 2.5 m spatial resolution imagery, with a non-linear parameter of 4.5 and a zoom factor of 10, using the linearised pixel swapping algorithm with fusion.

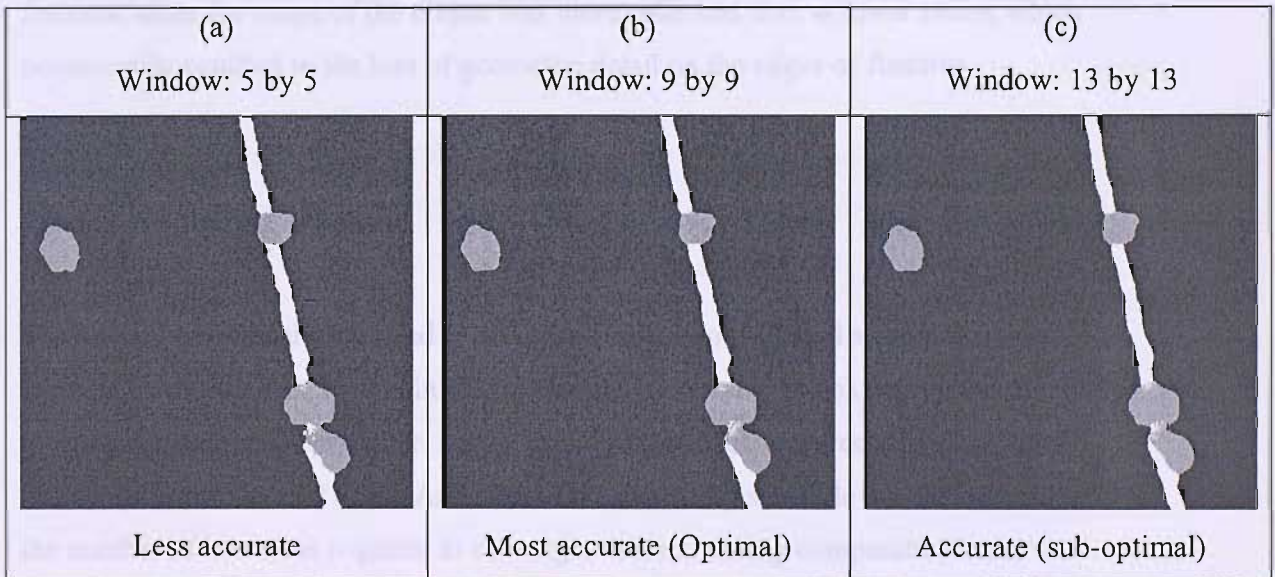


Figure 7.3: Size of exponential function window (a) 5 by 5, (b) 9 by 9, (c) 13 by 13. Scale: sub-pixels.

Increasing the size of the exponential function window increased the processing time substantially. The optimal size of the window was 9 by 9. In smaller windows, the edges of the hedgerow classes and tree classes are less accurately delineated. However, as the window size increased beyond 9 by 9, the additional benefits relative to the target diminished.

7.2.3 Anisotropy ratio

The anisotropy ratio was used during the prediction of direction in the linearised pixel swapping technique. The standard exponential function decayed isotropically in the exponential function window. In the linearised model, the function was modified to decay ellipsoidally in the direction predicted from the soft proportions. The length of the ellipse was determined by the actual soft proportion values and the width of the ellipse was determined by the anisotropy ratio. A small anisotropy ratio increased the prediction of linearity in the super-resolved output, since the shape of the ellipse was more linear than at larger ratios. As a result, non-linear features such as trees occasionally became linearised and accordingly, the pixel swapping algorithm was modified to only use the anisotropic exponential feature on classes that were known to contain linear features. A larger anisotropy ratio tended to produce more rounded features, since the shape of the ellipse was more spherical than at lower ratios, which occasionally resulted in the loss of geometric detail on the edges of features.

7.2.4 Annealing schedule

The annealing schedule was used to retain bad swaps during pixel swapping in order to avoid the problem of local minima. Accepting bad swaps “shook up” the system by introducing small amounts error. As the annealing schedule progressed and “cooled” bad swaps were retained often. An effect of the annealing schedule was to increase the number of iterations required to converge, thus increasing computation time.

7.3 Future research

The pixel swapping technique with linearization and image fusion components is a robust technique for super-resolution mapping of linear land cover features. However, there are parts of the technique that if investigated further could yield more accurate results and extend the applicability of the technique.

7.3.1 Prediction of direction

The prediction of linear features is a useful addition to the standard pixel swapping technique. Currently, directions are predicted using soft classification data and in a fixed number of directions (4, 8 or 12) in two window sizes (3 by 3 or 5 by 5). Further investigation into these parameters could enable more precise prediction of the direction of a linear feature, where the prediction of direction at any angle would likely be an optimal solution. Additionally, the techniques of other line detection methods such as snakes, could be incorporated into the pixel swapping method. For example, a snake could be applied to the panchromatic image and used to emphasise the edges of features. This could then be used to increase the precision with which the directions of the edges of features (and the feature itself) are predicted.

Alternatively, a method that predicts linear features by counting connected pixels, such as the Hough transform, could be used to inform on the existence of linear features within remotely sensed imagery and increase the precision of detecting directions during the linearisation stage.

7.3.2 Soft classification, noise, geometric error and real imagery

Error in soft classification is covered extensively in the literature (Congalton, 1991; Congalton, 1994; Bastin, 1997; Foody, 2002) and the causes of many of the problems experienced in the current research have been described previously. There are various factors influencing the results which, broadly speaking, are either methods-related

(limitations of the soft classification) or scene-related (nature and properties of the scene), with considerable overlap and inter-dependency between each. By understanding the nature of the error found within the super-resolved output, it is possible to recreate these errors in a controlled situation, using simulated imagery. Importantly, factors affecting the error specific to the scene (e.g., the number of land cover classes or the spatial arrangement of the classes) can be varied such that a characterisation of the error can be created and information on optimal parameters can be gained. A more thorough understanding of the causes of error can then be applied when soft classifying real satellite sensor imagery in order to produce more accurate soft proportions for use in the pixel swapping.

The fuzzy *c*-means classifier and linear mixture model have previously been shown to have certain limitations (Bastin, 1997). Specifically, the assumption of linear mixing between classes or the use of the mean to represent classes, may be unrealistic. Other soft classification techniques are available which could provide more accurate results, for example, a non-parametric soft classifier such as the *k*-nearest neighbour (McRoberts *et al.*, 2002). A non-parametric soft classifier is data-based (i.e. it uses all available information on each class) and could provide more accurate soft proportion predictions compared with trying to fit a model to complex data as (in parametric classifiers). Similarly, a multi-layer perceptron artificial neural network may in certain circumstances predict class proportions more accurately than the fuzzy *c*-means or linear mixture model.

An *absolute* classifier could also be investigated (Foody, 1990; Foody, 2000). As (Foody, 2002a) indicates, the existence of untrained classes within a soft classification can have deleterious effects on the overall accuracy of the classification. Accordingly, it is suggested that specific classes be excluded from training if they are of no interest to the investigation. Such a technique may be termed an “absolute” classifier, whereby only classes of interest are used in the training stages and all other classes are grouped together homogeneously. For example, in this investigation, training solely on the feature of interest, such as a hedgerow, would be possible, and could reduce error in the overall accuracy, by removing the potential for confusion between multiple classes that are of no interest to the overall classification. Sanchez-Hernandez *et al.* (2004) implemented a one class classification technique using

support vector machines for habitat monitoring, reporting increased classification accuracy over traditional classification methods.

The absolute classifier differs from traditional soft classification techniques in that class training is performed only on the feature of interest (e.g., the path or hedgerow, as opposed to all potential classes in the image) using a typicality measure, that is, the likeness of a pixel to each class independently of all classes. Then, any pixel, which does not contain any of that class, can be ignored in the super-resolution method. One method for achieving an absolute classification is by using a variation on the fuzzy *c*-means classifier, known as the *possibilistic c*-means, where the membership values for each class are predicted (distance measured from class centroids) independently of all other classes. This technique requires very careful training, so that the feature is characterised accurately to ensure none of the feature of interest is removed.

However, it has obvious potential benefits. Firstly, confusion between classes should be reduced as the soft classifier needs to distinguish only between two classes (is it the feature class or is it not the feature class?). Secondly, we can still train on additional classes (for example, trees, which are commonly found within hedgerows), maintaining class distinctions.

7.4 Conclusion

This chapter of the thesis has discussed further the parameters of the pixel swapping model and the effect that changing these parameters has on the output. The parameters of the model are in many ways directly dependent on the value of other parameters and the optimal configuration will vary depending on the type of imagery used as input.

Further avenues for developing the methodology were also discussed, such as refining the prediction of directions and a more rigorous assessment of the data used as input to the pixel swapping technique.

Chapter 8

Conclusions

The first part of the book is a general introduction to the theory of the firm, which is followed by a detailed treatment of the theory of the firm in a dynamic context.

References

Baron, R. (1985) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1986) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1987) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1988) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1989) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1990) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1991) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1992) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1993) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.
Baron, R. (1994) *The Art of Conjecture*. Cambridge, MA: Harvard University Press.

Chapter 8

8. Conclusions

8.1 Introduction

In this concluding chapter, the findings of the research presented in this thesis are drawn together. Initially the research, including the background, aims and objectives, and results are summarised. The implications of this research are then discussed. The chapter closes with concluding remarks.

8.2 Summary

The research in this thesis can be summarised in four parts: (i) background, (ii) aims and objectives, (iii) development and analysis, and (iv) results.

8.2.1 Background

Remote sensing represents one important source of land cover information, which can be an invaluable aid to mapping, monitoring and management of the natural and built environments. Remote sensing often represents the only feasible means of obtaining such information, particularly in cases where large study areas are required. However, the coarse spatial resolution provided by satellite sensor imagery has commonly been disadvantageous to the accuracy of such studies. For example, the pixel size, relative to the land cover features of interest, often leads to large numbers of mixed pixels, where there is more than one land cover feature within a pixel. Such mixed pixels

consequently produce error in land cover maps derived from traditional hard classification techniques, where pixels are allocated to a single class. The introduction of soft classification approaches, has, in recent years, helped to somewhat reduce these problems, producing more informative land cover maps. Nevertheless, while such classifiers predict the class composition of each pixel, there are few techniques for mapping the location of these classes within image pixels.

8.2.2 Research aims and objectives

The principal objective of this research was to develop an algorithm to perform super-resolution mapping of land cover features in remotely sensed imagery. These features exhibited many important geometric characteristics, perhaps most importantly of which, is the width of the feature. In most cases, however, the width of the feature is finer than the spatial resolution of the remotely sensed imagery. Soft classification techniques are able to predict the proportion of classes within each pixel, yet a technique was required to predict the location of these classes with sub-metre geometric accuracy. Super-resolution mapping was chosen as the tool to achieve this. An emphasis was placed on the identification of fine linear land cover features. For the purposes of evaluating the technique, hedgerows and other land cover objects such as trees, particularly as they are often found embedded within hedgerows, were used as target features.

8.2.3 Development and analysis

The super-resolution pixel swapping method was developed as a stochastic optimisation algorithm using the principles of simulated annealing as a framework. Image pixels were converted into sub-pixels converting the output from a soft classification into binary values, maintaining the predicted proportions, yet effectively giving the class proportions locations. Using the geostatistical phenomenon of spatial dependence, sub-pixels were selected randomly and the suitability of the swap was

analysed on the basis of estimated 'attractiveness'. The exponential distance decay model was used to estimate an attractiveness value as an indicator of spatial dependence. The model was developed to increase the likelihood of predicting linear features. The existence of linear features was estimated from soft class proportions; this information was used to apply an anisotropic correction to the shape of the exponential distance decay model, such that features exhibiting directional differences would be identified. Further, an image fusion component was developed to refine the prediction of the sub-pixel class, by merging information from a panchromatic image at a finer spatial resolution than the multispectral imagery and adjusting the criteria for accepting a swap.

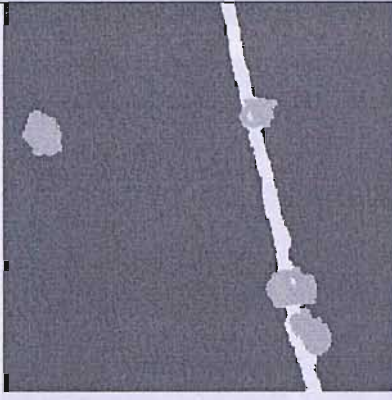
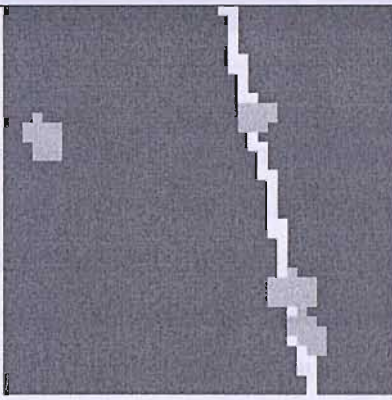
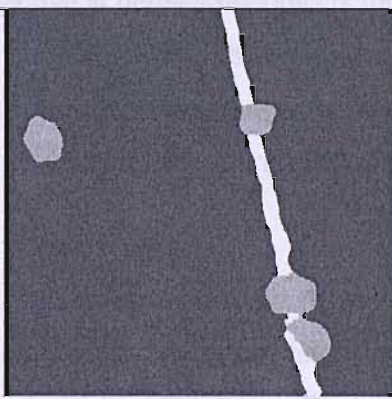
The pixel swapping technique was tested on multiple sources of data. Initially, the standard pixel swapping technique was applied to fully synthetic simulated imagery of basic shapes, allowing an evaluation of the function and suitability of the technique as well as ensuring a thorough understanding of the technique to enable its development. As the technique was developed, it was also applied to real satellite sensor imagery, aerial photography and pseudo-synthetic imagery (simulated imagery derived from aerial photography). The image fusion technique enabled the fusion of panchromatic imagery with the available multispectral imagery (satellite sensor and aerial photography). The scenes depicted in each set of imagery were semi-rural, where the land-use was mainly agricultural and the common land cover classes were cereal and grass fields, hedgerows, pathways, trees and forest.

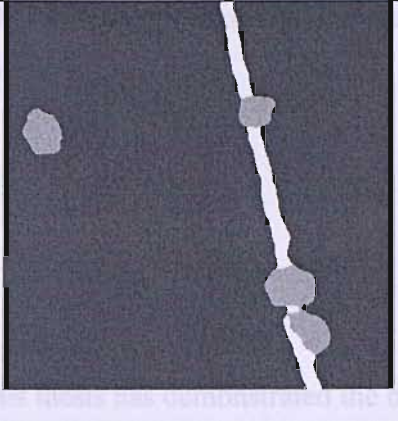
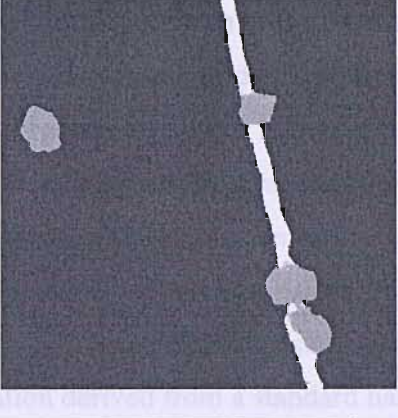
8.2.4 Results

In the majority of cases, the super-resolution pixel swapping technique produced accurate results. When applied to simple simulated imagery, in some cases, the accuracy of the super-resolved output was estimated as 99% when compared with the input imagery. In every case, the results were more accurate than was possible using a traditional hard classification technique. The development of the technique to identify linear features enabled the prediction of continuous features, where hard classification

techniques or, indeed, the standard pixel swapping technique, mapped the features as a series of discrete objects. The results of each technique are summarised in Table 8.1 for the simulated field site at a spatial resolution of 2.5 m.

Table 8.1: Summary of results, spatial resolution 2.5 m

Technique	Image output	% Correct
<p>Target</p>		<p>100</p>
<p>Hard classification</p>		<p>92.2</p>
<p>Pixel swapping</p>		<p>96.3</p>

<p>Linearised pixel swapping (3 by 3 window)</p>		<p>98.6</p>
<p>Linearised pixel swapping with fusion (3 by 3 window)</p>		<p>98.8</p>

The hard classified image is one of the most commonly produced land cover maps used in remote sensing applications. In many situations, a hard classification may be sufficient. However, in applications where a high-level of detail is required, it will not be sufficient. The pixel swapping technique has been demonstrated to provide more accurate land cover maps (6.6%) with a superior visual output. The linearised pixel swapping technique was effective at delineating objects found within other objects, for example, trees found within hedgerows. The inclusion of a fusion component, where simulated panchromatic imagery was merged with simulated aerial photography, resulted in more realistic characterisation of the edges of classes

There are factors that affect the accuracy of the pixel swapping technique presented in this thesis, such as, the accuracy of the input soft classification, the number of land cover classes and the frequency of mixing between these classes. Selection of a suitable zoom factor is important for accurate results; a higher zoom factor does not always produce more accurate results – a suitable zoom factor is dependent on many factors, for example, spatial resolution of the input imagery, the fuzziness of the soft

classification data, the number of land cover classes and the size of the window used in the exponential distance decay.

8.3 Implications

The research presented in this thesis has demonstrated the benefits of using super-resolution sub-pixel mapping techniques, such as the pixel swapping technique presented here, for the creation of accurate land cover maps. Certainly, the extension of existing techniques, to incorporate additional information (e.g., additional information from the original imagery, such as, estimated directions of linear features, or, alternatively the use of fusion to incorporate an additional source of data) has demonstrated the value and potential of such algorithms. Indeed, the information contained within such maps, for certain remote sensing applications, will be of far greater use than the information derived from a standard hard classification. The standard pixel swapping technique was efficient when applied to simple shapes. However, by introducing information on measured anisotropy, the technique predicted continuous linear features, whereas in the standard configuration, the pixel swapping technique predicted linear features as a series of discrete objects. As discussed, the ability to map fine linear features in remotely sensed imagery would be of great use to organisations involved in the monitoring and management of, for example, the countryside. Indeed, by merging panchromatic imagery with multispectral imagery, it was possible to further increase the accuracy of these maps.

8.4 Conclusion

The super-resolution pixel swapping technique presented in this thesis has been shown to be able to map linear land cover features from fine spatial remotely sensed imagery with greater accuracy than was previously possible. Further work is required to refine the technique for application to different remotely sensed imagery, however, the technique is a simple and accurate method, that could realistically be used for the

creation of land cover maps containing accurate information on the location of rural land cover features.

1. The first step in the process is to identify the rural land cover features that are to be mapped. This can be done by consulting with local stakeholders and using satellite imagery to identify the features.

2. The second step is to collect data on the location of the rural land cover features. This can be done using a variety of methods, including ground truthing, aerial photography, and satellite imagery.

3. The third step is to process the data and create a land cover map. This can be done using a variety of software packages, including ArcGIS, QGIS, and MapInfo.

4. The fourth step is to validate the land cover map. This can be done by comparing the map to ground truth data and using a variety of validation metrics, including accuracy, precision, and recall.

5. The fifth step is to disseminate the land cover map. This can be done using a variety of methods, including web-based maps, printed maps, and data exports.

6. The sixth step is to update the land cover map. This can be done using a variety of methods, including ground truthing, aerial photography, and satellite imagery.

7. The seventh step is to use the land cover map for a variety of purposes, including land use planning, resource management, and environmental monitoring.

8. The eighth step is to evaluate the land cover map. This can be done using a variety of methods, including stakeholder feedback and validation metrics.

9. The ninth step is to publish the land cover map. This can be done using a variety of methods, including web-based maps, printed maps, and data exports.

References

- Abd-Alamageed, W., C. E. Smith and S. Ramadan (2003). Kernel snakes: non-parametric active contour models. IEEE International Conference on Systems Man and Cybernetics, Washington DC.
- Acqua, F. and P. Gamba (2001). "Detection of urban structures in SAR images by robust fuzzy clustering algorithms: the example of street tracking." IEEE Transactions on Geoscience and Remote Sensing **39**(10): 2287-2297.
- Adams, J. B., M. O. Smith and P. E. Johnson (1985). "Spectral mixture modelling: a new analysis of rock and soil types at the Viking Lander 1 site." Journal of Geophysical Research **91**: 8098-8112.
- Agresti, A. (1996). An Introduction to Categorical Data Analysis. New York, Wiley.
- Aplin, P. and P. M. Atkinson (2001). "Sub-pixel land cover mapping for per-field classification." International Journal of Remote Sensing **22**(14): 2853-2858.
- Aplin, P., P. M. Atkinson and P. J. Curran (1997). "Fine spatial resolution satellite sensors for the next decade." International Journal of Remote Sensing **18**: 3873-3881.
- Ashton, E. A. (1998). "Detection of subpixel anomalies in multispectral infrared imagery using an adaptive Bayesian classifier." IEEE Transactions on Geoscience and Remote Sensing **36**(2): 506-517.
- Atkinson, P. M. (1997). Mapping sub-pixel boundaries from remotely sensed images. Innovations in GIS IV. London, Taylor and Francis: 166-180.
- Atkinson, P. M. (2004a). "Spatially weighted supervised classification." International Journal of Applied Earth Observation and Geoinformation **5**: 277-291.
- Atkinson, P. M. (2004b). "Super-resolution target mapping from soft-classified remotely sensed imagery." Photogrammetric Engineering and Remote Sensing **71**: 839-846.
- Atkinson, P. M. (2005). "Sub-pixel target mapping from soft-classified remotely sensed imagery." Photogrammetric Engineering and Remote Sensing **71**(7): 839-846.
- Atkinson, P. M. and P. Aplin (2004). "Spatial variation in land cover and choice of spatial resolution for remote sensing." International Journal of Remote Sensing **25**(18): 3687-3702.
- Atkinson, P. M. and P. Lewis (2000). "Geostatistical classification for remote sensing: an introduction." Computers & Geosciences **26**: 361-371.

- Atkinson, P. M. and A. R. L. Tatnall (1997). "Neural networks in remote sensing." International Journal of Remote Sensing **18**(4): 699-709.
- Baath, H., A. Gallerspang, G. Hallsby, A. Lundstrom, P. Lofgren, M. Nilsson and G. Stahl (2002). "Remote sensing, field survey, and long-term forecasting: an efficient combination for local assessments of forest fuels." Biomass and Bioenergy **22**(3): 145-157.
- Bastin, L. (1997). "Comparison of fuzzy c-means classification, linear mixture modelling and MLC probabilities as tools for unmixing coarse pixels." International Journal of Remote Sensing **18**(17): 3629-3648.
- Baudry, J., F. Burel, S. Aviron, M. Martin, A. Ouin, G. Pain and C. Thenail (2003). "Temporal variability of connectivity in agricultural landscapes: do farming activities help?" Landscape Ecology **18**(3): 303-314.
- Belward, A. S., J. E. Estes and K. D. Kilne (1999). "The IGBP-DIS global 1km land-cover data set DIScover: a project overview." Photogrammetric Engineering and Remote Sensing **65**: 1013-1020.
- Bernard, A. C., G. G. Wilkinson and I. Kanellopoulos (1997). "Training strategies for neural network soft classification of remotely-sensed imagery." International Journal of Remote Sensing **18**(8): 1851-1856.
- Bezdek, J. C., R. Ehrlich and W. Full (1984). "FCM: The fuzzy c-means clustering algorithm." Computers & Geosciences **10**(2): 191-203.
- Boucher, A. and P. C. Kyriakidis (2006). "Super-resolution land cover mapping with indicator geostatistics." Remote Sensing of Environment **104**: 264-282.
- Boutin, C., B. Jobin, L. Belanger and L. Choiniere (2002). "Plant diversity in three types of hedgerows adjacent to cropfields." Biodiversity and Conservation **11**(1): 1-25.
- Braswell, B. H., S. C. Hagen, S. E. Frolking and W. A. Salas (2003). "A multivariable approach for mapping sub-pixel land cover distributions using MISR and MODIS: Application in the Brazilian Amazon region." Remote Sensing of Environment **87**: 243-256.
- Butler, D. (2000). "Souped-up search engines." Nature **405**: 112-115.
- Campbell, J. B. (1996). Introduction to Remote Sensing. London, Taylor and Francis.
- Cannon, R. L., J. V. Dave, J. C. Bezdek and M. M. Trivedi (1986). "Segmentation of a thematic mapper image using the fuzzy c-means clustering algorithm." IEEE Transactions on Geoscience and Remote Sensing **24**: 400-408.
- Canny, J. (1986). "A computational approach to edge detection." IEEE Transactions on Pattern Analysis and Machine Intelligence **8**(6): 679-698.

Chavez, P. S., S. C. Sides and J. A. Anderson (1991). "Comparison of three different methods to merge multiresolution and multispectral data: Landsat TM and SPOT Panchromatic." Photogrammetric Engineering and Remote Sensing **57**(3): 295-303.

Chibani, Y. and A. Houacine (2002a). "The joint use of IHS transform and redundant wavelet decomposition for fusing multispectral and panchromatic images." International Journal of Remote Sensing **23**(18): 3821-3833.

Chibani, Y. and A. Houacine (2002b). "Redundant versus orthogonal wavelet decomposition for multisensor image fusion." Pattern Recognition **36**(4): 879-887.

Choi, E. and C. Lee (2001). "Optimizing feature extraction for multiclass problems." IEEE Transactions on Geoscience and Remote Sensing **39**(3): 521-528.

Cihlar, J. (2000). "Land cover mapping of large areas from satellites: status and research priorities." International Journal of Remote Sensing **21**: 1093-1114.

Congalton, R. G. (1991). "A review of assessing the accuracy of classifications of remotely sensed data." Remote Sensing of Environment **37**: 35-46.

Congalton, R. G., M. E. Martin, S. D. Newman and J. D. Aber (1998). "Determining forest species composition using high spectral resolution remotely sensed data." Remote Sensing of Environment **65**: 249-254.

Curran, P. J. and M. J. Atkinson (1998). "Geostatistics and remote sensing." Progress in Physical Geography **21**: 61-78.

Curran, P. J., E. J. Milton, P. M. Atkinson and G. M. Foody (1998). Remote sensing: from data to understanding. Geocomputation: A primer. P. A. Longley, S. M. Brooks, R. McDonnell and B. Macmillan, John Wiley & Sons Ltd.

Davies, E. R. (1987). "A new parameterisation of the straight line and its application for the optimal detection of objects with straight edges." Pattern Recognition Letters **6**: 9-14.

Deutsch, C. V. and A. G. Journel (1998). GSLIB Geostatistical Software Library and User's Guide. Oxford, Oxford University Press.

Duda, R. O. and P. E. Hart (1972). "Use of the Hough transformation to detect lines and curves in pictures." Communications of the ACM **15**(1): 11-15.

Editorial (2001). "Landsat - 30 years and counting." Remote Sensing of Environment **78**: 1-2.

Fisher, P. (1997). "The pixel: a snare and a delusion." International Journal of Remote Sensing **18**(3): 679-685.

- Foody, G. M. (1990). "Directed ground survey for improved maximum likelihood classification of remotely sensed data." International Journal of Remote Sensing **11**(10): 1935-1940.
- Foody, G. M. (1996). "Approaches for the production and evaluation of fuzzy land cover classifications from remotely-sensed data." International Journal of Remote Sensing **17**(7): 1317-1340.
- Foody, G. M. (1998). "Sharpening fuzzy classification output to refine the representation of sub-pixel land cover distribution." International Journal of Remote Sensing **19**(13): 2593-2599.
- Foody, G. M. (2002a). "Hard and soft classifications by a neural network with a non-exhaustively defined set of classes." International Journal of Remote Sensing **23**(18): 3853-3864.
- Foody, G. M. (2002b). "Status of land cover classification accuracy assessment." Remote Sensing of Environment **80**: 185-201.
- Foody, G. M. (2004). "Thematic map comparison: evaluating the statistical significance of differences in classification accuracy." Photogrammetric Engineering and Remote Sensing **70**(5): 627-633.
- Foody, G. M. and D. P. Cox (1994). "Sub-pixel land cover composition estimation using a linear mixture model and fuzzy membership functions." International Journal of Remote Sensing **15**: 619-631.
- Foody, G. M. and A. Mathur (2006). "The use of small training sets containing mixed pixels for accurate hard image classification: Training on mixed spectral responses for classification by a SVM." Remote Sensing of Environment **103**: 179-189.
- Forshaw, M. R. B., A. Haskell, P. F. Miller, D. J. Stanley and J. R. G. Townshend (1983). "Spatial resolution of remotely sensed imagery - A review paper." International Journal of Remote Sensing **4**: 497-520.
- Franco, D. (2002). "The scale and pattern influences on the hedgerow networks' effect on landscape processes: First considerations about the need to plan for landscape amelioration purposes." Environmental Management and Health **13**(3): 263-276.
- Fuller, R. J., S. A. Hinsley and R. D. Swetnam (2004). "The relevance of non-farmland habitats, uncropped areas and habitat diversity to the conservation of farmland birds." International Journal of Avian Science **146**: 22-28.
- Fuller, R. M., G. B. Groom and S. M. Wallis (1994). "The availability of Landsat TM images of Great Britain." International Journal of Remote Sensing **15**(6): 1357-1362.

Gebbinck, M. S. K. (1998). Decomposition of mixed pixels in remote sensing images to improve the area estimation of agricultural fields. Katholieke Universiteit, Nijmegen

Gerd, D., J. Deckers, J. Poesen, G. Govers, H. Sanchez, M. Ramirez, R. Vanegas, E. Tacuri and G. Loaiza (2006). "Spatial variability in crop response under contour hedgerow systems in the Andes region of Ecuador." Soil & Tillage Research **86**(1): 15-26.

Gillespie, A. R., A. B. Kahle and R. E. Walker (1986). "Colour enhancement of highly correlated images." Remote Sensing of Environment **20**: 209-235.

Gillings, S. and R. J. Fuller (1998). "Changes in bird populations on sample lowland English farms in relation to loss of hedgerows and other non-crop habitats." Oecologia **116**(120-127).

Goovaerts, P. (1997). Geostatistics for Natural Resources Evaluation. Oxford, Oxford University Press.

Gregoire, J. M., K. Tansey and J. M. N. Silva (2003). "The GBA2000 initiative: developing a global burnt area database from SPOT-VEGETATION imagery." International Journal of Remote Sensing **24**(6): 1369-1376.

Gunn, S. R. and M. S. Nixon (1994). A model based dual active contour. British Machine Vision Conference, Kingston University, London.

Gunn, S. R. and M. S. Nixon (1997). "A robust snake implementation: a dual active contour." IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(1): 63-68.

Heath, M. D., S. Sarkar and T. Sanocki (1997). "A robust visual method for assessing the relative performance of edge-detection algorithms." IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(12): 1338-1359.

Ho, T. K., J. J. Hull and S. N. Srihari (1994). "Decision combination in multiple classifier systems." IEEE Transactions on Pattern Analysis and Machine Intelligence **16**(1): 66-75.

Hochberg, E. J. and M. J. Atkinson (2003). "Capabilities of remote sensors to classify coral, algae, and sand as pure and mixed spectra." Remote Sensing of Environment **85**: 174-189.

Hough, P. (1962). Methods and means for recognizing complex patterns. United States.

Isaaks, E. H. and R. H. Srivastava (1989). Applied Geostatistics. New York, Oxford University Press.

Journel, A. G. and C. J. Huijbregts (1978). Mining Geostatistics. London, Academic Press.

Ju, J., S. Gopal and E. D. Kolaczyk (2005). "On the choice of spatial and categorical scale in remote sensing land cover classification." Remote Sensing of Environment **96**: 62-77.

Ju, J., E. D. Kolaczyk and S. Gopal (2003). "Gaussian mixture discriminant analysis and sub-pixel land cover characterisation in remote sensing." Remote Sensing of Environment **84**: 550-560.

Kass, M., A. Witkin and D. Terzopoulos (1988). "Snakes: active contour models." International Journal of Computer Vision **1**(4): 321-331.

Kerr, J. T. and M. Ostrovsky (2003). "From space to species: ecological applications for remote sensing." Trends in Ecology and Evolution **18**(6): 299-305.

Key, J. R., J. A. Maslanik and R. G. Barry (1989). "Cloud classification from satellite data using a fuzzy set algorithm: a polar example." International Journal of Remote Sensing **10**: 1823-1842.

Kovalev, V. and M. Petrou (1996). "Multidimensional co-occurrence matrices for object recognition and matching." Graphical Models and Image Processing **58**(3): 187-197.

Kristensen, S. P. (2001). "Hedgerow planting activities by Danish farmers: a case study from central Jutland." Danish Journal of Geography **101**: 101-114.

Kuo, B.-C. and D. A. Landgrebe (2002). "A robust classification procedure based on mixture classifiers and nonparametric weighted feature extraction." IEEE Transactions on Geoscience and Remote Sensing **40**(11): 2486-2494.

Leavers, V. F. (1993). "Which Hough transform?" CVGIP: Image Understanding **58**(2): 250-264.

Lee, C. and D. A. Landgrebe (1993). "Feature extraction based on decision boundaries." IEEE Transactions on Pattern Analysis and Machine Intelligence **15**(4): 388-400.

Li, H., B. S. Manjunath and S. K. Mitra (1995). "Multisensor image fusion using the wavelet transform." Graphical Models and Image Processing **57**(3): 235-245.

Li, S., J. T. Kwok and Y. Wang (2002). "Using the discrete wavelet transform to merge Landsat TM and SPOT panchromatic images." Information Fusion **3**: 17-23.

Lloyd, C. D. and P. M. Atkinson (2004). "Increased accuracy of geostatistical prediction of nitrogen dioxide in the United Kingdom with secondary data." International Journal of Applied Earth Observation and Geoinformation **5**: 293-305.

- Loveland, T. R. and R. S. DeFries (2004). "Observing and Monitoring land use and land cover change." Geophysical monograph **153**: 231-246.
- Manolakis, D., C. Siracusa and G. Shaw (2001). "Hyperspectral subpixel target detection using the linear mixing model." IEEE Transactions on Geoscience and Remote Sensing **39**(7): 1392-1409.
- Mather, P. M. (1999). Computer Processing of Remotely Sensed Images. Chichester, Wiley.
- Matheron, G. (1965). Les variables régionalisées et leur estimation. Masson, Paris.
- Matheron, G. (1971). The Theory of Regionalised Variables and Its Applications. Fontainebleau, Ecole Nationale Supérieure des Mines de Paris.
- Mayer, H. (1999). "Automatic extraction from aerial imagery - a survey focusing on buildings." Computer Vision and Image Understanding **74**(2): 138-149.
- McCollin, D. (2002). "Editorial: Hegerow policy and protection - changing paradigms and conservation ethic." Journal of Environmental Management **60**: 3-6.
- Mertens, K. C., B. De Baets, L. P. C. Verbeke and R. R. De Wulf (2004a). Direct sub-pixel mapping exploiting spatial dependence. Geoscience and Remote Sensing Symposium, IGARSS '04.
- Mertens, K. C., B. De Baets, L. P. C. Verbeke and R. R. De Wulf (2006). "A sub-pixel mapping algorithm based on sub-pixel/pixel spatial attraction models." International Journal of Remote Sensing **27**(15): 3293-3310.
- Mertens, K. C., L. P. C. Verbeke, E. I. Ducheyne and R. R. De Wulf (2003). "Using genetic algorithms in sub-pixel mapping." International Journal of Remote Sensing **24**(21): 4241-4247.
- Mertens, K. C., L. P. C. Verbeke, T. Westra and R. R. De Wulf (2004b). "Sub-pixel mapping and sub-pixel sharpening using neural network predicted wavelet coefficients." Remote Sensing of Environment **91**: 225-236.
- Mullerova, J., P. Pysek, V. Jarosik and J. Pergl (2005). "Aerial photographs as a tool for assessing the regional dynamics of the invasive plant species *Heracleum mantegazzianum*." Journal of Applied Ecology **42**(6): 1042-1053.
- Muslim, A. A., G. M. Foody and P. M. Atkinson (2003). Sub-Pixel analysis in Shoreline Mapping. RSPSoc 2003, Nottingham, UK.
- Muslim, A. A., G. M. Foody and P. M. Atkinson (2006). "Localized soft classification for super-resolution mapping of the shoreline." International Journal of Remote Sensing **27**: 2271-2285.

Nguyen, H. H. and P. Cohen (1993). "Gibbs random fields, fuzzy clustering, and the unsupervised segmentation of textured images." Graphical Models and Image Processing **55**(1): 1-19.

Nguyen, M. Q., P. M. Atkinson and H. G. Lewis (2005). "Superresolution mapping using a Hopfield neural network with lidar data." Geoscience and Remote Sensing Letters **2**(3): 366-370.

Nguyen, M. Q., P. M. Atkinson and H. G. Lewis (2006). "Superresolution mapping using a Hopfield neural network with fused images." IEEE Transactions on Geoscience and Remote Sensing **44**(3): 736-749.

Nixon, M. S. and A. Aguado (2002). Feature Extraction and Image Processing. Oxford, Newnes.

Nunez, J., X. Otazu, O. Fors, A. Prades and R. Arbiol (1999). "Multiresolution-based image fusion with additive wavelet decomposition." IEEE Transactions on Geoscience and Remote Sensing **37**(3): 1204-1211.

Paola, J. D. and R. A. Schowengerdt (1995). "A detailed comparison of backpropagation neural network and maximum-likelihood classifiers for urban land use classification." IEEE Transactions on Geoscience and Remote Sensing **33**(4): 981-996.

Parzen, E. (1962). "On estimation of a probability density function and mode." Annals Mathematical Statistics **33**: 1065-1076.

Piella, G. (2003). "A general framework for multiresolution image fusion: from pixels to regions." Information Fusion **4**: 259-280.

Pohl, C. (1999). "Tools and methods for fusion of images of different spatial resolution." International Archives of Photogrammetry and Remote Sensing **32**(7-4-3).

Pohl, C. and J. L. Van Genderen (1998). "Multisensor image fusion in remote sensing: concepts, methods and applications." International Journal of Remote Sensing **19**(5): 823-854.

Quackenbush, L. J. (2004). "A review of techniques for extracting linear features from imagery." Photogrammetric Engineering and Remote Sensing **70**(12): 1383-1392.

Rajan, D. and S. Chaudhuri (2002). "Data fusion techniques for super-resolution imaging." Information Fusion **3**(1): 25-38.

Richards, J. A. and X. Jia (1999). Remote sensing digital image analysis: an introduction. Berlin, Springer.

Richards, J. A. and X. Jia (2006). Remote sensing digital image analysis: an introduction. Berlin, Springer.

- Robinson, R. A. and W. J. Sutherland (2002). "Post-war changes in arable farming and biodiversity in Great Britain." Journal of Applied Ecology **39**: 157-176.
- Rosenqvist, A., A. Milne, R. Lucas, M. Imhoff and C. Dobson (2003). "A review of remote sensing technology in support of the Kyoto Protocol." Environmental Science & Policy **6**: 441-455.
- Sanchez-Hernandez, C., D. Boyd and G. M. Foody (2004). One class classification for EU priority habitats monitoring. RSPSoc, Aberdeen.
- Settle, J. J. (2002). "On constrained energy minimisation and the partial unmixing of multispectral images." IEEE Transactions on Geoscience and Remote Sensing **40**(3): 718-721.
- Settle, J. J. and N. A. Drake (1993). "Linear mixing and the estimation of ground cover proportions." International Journal of Remote Sensing **14**: 1159-1177.
- Shi, W. and C. Zhu (2002). "The line segment match method for extracting road network from high-resolution satellite images." IEEE Transactions on Geoscience and Remote Sensing **40**(2): 511-514.
- Shimazaki, Y. and R. Tateishi (2001). Land cover mapping using spectral and temporal linear mixing model at Lake Baikal region. 22nd Asian Conference on Remote Sensing, Singapore.
- Simone, G., A. Farina, F. C. Morabito, S. B. Serpico and L. Bruzzone (2002). "Image fusion techniques for remote sensing applications." Information Fusion **3**: 3-15.
- Skingley, J. and A. J. Rye (1987). "The Hough transform applied to SAR images for thin line detection." Pattern Recognition Letters **6**: 61-67.
- Stefanov, W. L., M. S. Ramsey and P. R. Christensen (2001). "Monitoring urban land cover change: an expert system approach to land cover classification of semiarid to arid urban centers." Remote Sensing of Environment **77**: 173-185.
- Stewart, C. V. (1999). "Robust parameter estimation in computer vision." SIAM review **41**: 513-537.
- Strahler, A. H., C. E. Woodcock and J. A. Smith (1986). "On the nature of models in remote sensing." Remote Sensing of Environment **20**: 121-139.
- Su, L., X. Deng, J. Wang and X. Li (June, 2003). Managing and distributing remote sensing images based on metadata and microimage. Image Processing and Pattern Recognition in Remote Sensing.
- Tatem, A. J., H. G. Lewis, P. M. Atkinson and M. S. Nixon (2001). "Multiple-class land-cover mapping at the sub-pixel scale using a Hopfield neural network." JAG **3**(2): 184-190.

Tatem, A. J., H. G. Lewis, P. M. Atkinson and M. S. Nixon (2002a). Super-resolution land cover mapping from remotely sensed imagery using a Hopfield neural network. Uncertainty in Remote Sensing and GIS. G. M. Foody and P. M. Atkinson, John Wiley & Sons, Ltd.

Tatem, A. J., H. G. Lewis, P. M. Atkinson and M. S. Nixon (2002b). "Super-resolution land cover pattern prediction using a Hopfield neural network." Remote Sensing of Environment **79**: 1-14.

Tatem, A. J., H. G. Lewis, P. M. Atkinson and M. S. Nixon (2003). "Increasing the spatial resolution of agricultural land cover maps using a Hopfield neural network." International Journal of Geographical Information Science **17**(7): 647-672.

Thenail, C. and J. Baudry (2004). "Variation of farm spatial land use pattern according to the structure of hedgerow networks in northeast Brittany." Agriculture, Ecosystems & Environment **101**(1): 53-72.

Tso, B. and P. M. Mather (2001). Classification Methods for Remotely Sensed Data. London and New York, Taylor and Francis.

Tu, T., S.-C. Su, H. Shyu and P. S. Huang (2001). "A new look at IHS-like fusion methods." Information Fusion **2**: 177-186.

Veitch, N., J. R. Treweek and R. M. Fuller (1995). The land cover map of Great Britain – a new data source for environmental planning and management. Advances in Remote Sensing. F. M. Danson and S. E. Plummer. Chichester, John Wiley: 157-170.

Verhoeve, J. and R. R. De Wulf (2002). "Land cover mapping at sub-pixel scales using linear optimisation techniques." Remote Sensing of Environment **79**: 96-104.

Vijayaraj, V., N. H. Younan and C. G. O'Hara (2006). "Quantitative analysis of pansharpened images." Optical Engineering **45**(4).

Wald, L. (1999). "Some terms of reference in data fusion." IEEE Transactions on Geoscience and Remote Sensing **37**(3): 1190-1193.

Wang, F. (1990). "Fuzzy supervised classification of remotely sensed images." IEEE Transactions on Geoscience and Remote Sensing **28**: 194-201.

Woodcock, C. E. and A. H. Strahler (1987). "The factor of scale in remote sensing." Remote Sensing of Environment **21**: 311-332.

Wu, J., H. Honglin, J. Tian and J. Liu (2005). Remote sensing image data fusion based on local deviation of wavelet packet transform. Autonomous Decentralized Systems, ISADS, Chengdu, China.

Wu, S., M. J. Er and Y. Gao (2001). "A fast approach for automatic generation of fuzzy rules for generalised dynamic fuzzy neural networks." IEEE Transactions on Fuzzy Systems **9**(4): 578-594.

Wyatt, B. K., N. G. Davies, R. G. H. Bunce, R. M. Fuller and M. O. Hill (1993). "Comparisons of land cover definitions." Report to the Department of the Environment.

Zhan, Q., M. Molenaar and K. Tempfli (2002). Finding spatial units for land use classification based on hierarchical image objects. Symposium on Geospatial Theory, Processing and Applications, Ottawa.

Zomet, A. and S. Peleg (2002). Multi-sensor super-resolution. Applications of Computer Vision, Orlando, Florida, USA.

Appendix 1

The principal output of the research presented in this thesis was a fully functional computer application for super-resolution pixel swapping. The operational development and key components of this software are described below.

Technical Background

The standard pixel swapping algorithm was originally written by Peter Atkinson as a simple iterative subroutine in the S programming language. The subroutine was executed in S-PLUS. S-PLUS is statistics software that provides basic and advanced statistical calculations in an integrated environment. In particular, S-PLUS provides the functionality for a user to write custom scripts. However, the software has a large application overhead, which makes such scripts unsuitable for processing very large datasets such as remotely sensed imagery (as described in Chapter 7) and, therefore, unsuitable for the purposes of this research. Accordingly, it was decided to port the subroutine to an alternative programming language.

Three programming languages were considered as part of the process to “translate” the standard pixel swapping algorithm. These languages were: Fortran, Perl and C++.

Fortran was considered due to the availability of many pre-existing image processing routines, such as the NAG libraries. Perl was considered due to its reputation of excellence in data mining and processing of large datasets and the authors’ previous exposure to the language. However, C++ was ultimately chosen as the programming language for a number of reasons. Foremost, the author had prior exposure to and experience of the language and its constructs, which meant the learning curve was much lower. In addition, C++ is an industry standard programming language and the wide-reaching and common use of C++ in every day programming meant that there was a wealth of help and support available in books and on the Internet. As a compiled language, C++ was guaranteed to provide substantial improvements over the procedural approach of S-PLUS (or Perl) and was considered to be a more

straightforward (and up to date) approach than Fortran, especially with the availability of an excellent open source integrated development environment (DevC++, <http://www.bloodshed.net/devcpp.html>) that provided invaluable development and debugging tools. Finally, applications written in C++ can be executed in virtually any operating system that can run the compiler. The pixel swapping algorithm was primarily developed under both the Win32 and Linux environments, but was also tested under two of the distributed computing environments offered by the University of Southampton (Linuxcompute and IRIDIS). The application could therefore process very large datasets efficiently.

Programming tasks

Converting the pixel swapping subroutine code from the S programming language to C++ involved a number of key tasks. C++ does not natively provide the same file management (e.g., opening input files, creating data structures), memory management and visualisation functionality that S-PLUS does, and so, each of these areas needed consideration. A generic file reading routine was developed to handle all input and output file related tasks, such as reading the output from a soft classification or simulated soft proportions as well as writing results (e.g., super-resolved output and accuracy assessments) to files. In addition, the portable graymap image format (<http://netpbm.sourceforge.net/doc/pgm.html>) was adopted to enable visualisation of all image output. A set of functions to manage all memory allocation tasks was also developed, which simplified the process of creating multi-dimensional arrays, which contained image and proportions data.

In developing the application, the pixel swapping algorithm was initially converted from S to C++. During this process, various optimisations were made to increase the efficiency of the algorithm, such as, using delegates and functions for performing repetitive tasks and changing the datatype of the arrays from real number arrays to bit arrays (to represent the binary hard classification of the input proportions in a true Boolean format), saving approximately two bytes of memory overhead per element of each array.

In addition to optimising the original algorithm, the pixel swapping algorithm was primarily developed in three ways: simulated annealing, linearisation and image fusion. These developments were programmed, however, in such a way that the application could be executed in a specific mode, such that differences in the accuracy of the super-resolved output between each of the methods could be compared easily. For example, the software can be executed in any of the following ‘modes’:

- Pixel swapping
- Linearised pixel swapping
- Pixel swapping with fusion
- Linearised pixel swapping with fusion

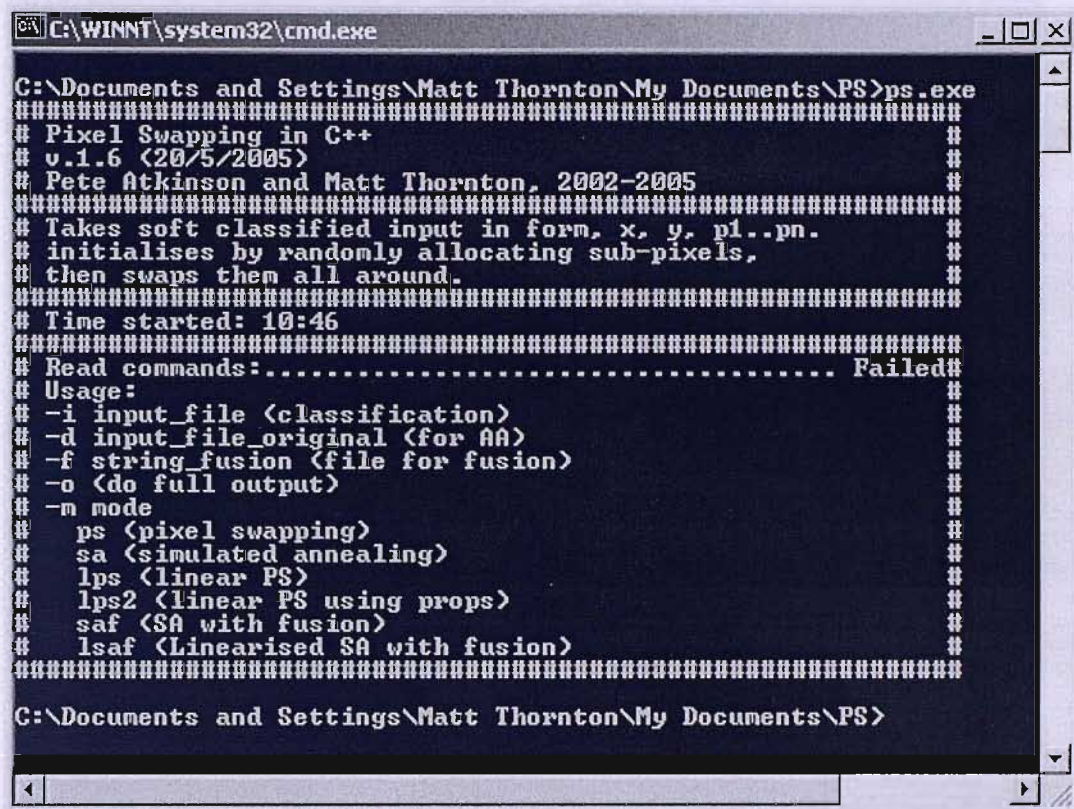
In each mode, sub-pixels to be swapped can be selected using either the method described in the original pixel swapping technique or randomly (simulated annealing). It was observed, however, that the algorithm converged to the same result through either method. Screenshots of the application are presented in Appendix 2.

In its output, the pixel swapping application also provides a complete set of accuracy assessments, depending on the mode in which it is to be run. As required, it can output maps of estimated direction, difference images and confusion matrices as well as complete operation logs and debugging information.

Additional programming

The pixel swapping application used information from a soft classification as input. These data were provided in different formats from many different sources. Soft classification of real remotely sensed imagery was performed using the fuzzy *c*-means and linear mixture model techniques. These algorithms were acquired in Fortran from the openly available Numerical Algorithms Group (NAG) Library. In some cases, data were reformatted in Matlab. Non-standard data visualisation was carried out in either S-PLUS or Matlab.

Appendix 2



```
C:\WINNT\system32\cmd.exe
C:\Documents and Settings\Matt Thornton\My Documents\PS>ps.exe
#####
# Pixel Swapping in C++                                     #
# v.1.6 (20/5/2005)                                         #
# Pete Atkinson and Matt Thornton, 2002-2005               #
#####
# Takes soft classified input in form, x, y, pl..pn.      #
# initialises by randomly allocating sub-pixels,         #
# then swaps them all around.                             #
#####
# Time started: 10:46                                       #
#####
# Read commands:..... Failed#
# Usage:                                                    #
# -i input_file <classification>                          #
# -d input_file_original <for AA>                         #
# -f string_fusion <file for fusion>                      #
# -o <do full output>                                     #
# -m mode                                                  #
#   ps <pixel swapping>                                   #
#   sa <simulated annealing>                              #
#   lps <linear PS>                                        #
#   lps2 <linear PS using props>                          #
#   saf <SA with fusion>                                   #
#   lsaf <Linearised SA with fusion>                      #
#####
C:\Documents and Settings\Matt Thornton\My Documents\PS>
```

Screenshot 1: Application with no command-line arguments

```

C:\WINNT\system32\cmd.exe
C:\Documents and Settings\Matt Thornton\My Documents\PS>ps.exe -m ps -i aerial2/
aerial_nh_40.txt
#####
# Pixel Swapping in C++
# v.2 (20/6/2006)
# Pete Atkinson and Matt Thornton, 2002-2006
#####
# Takes soft classified input in form, x, y, pl..pn.
# initialises by randomly allocating sub-pixels.
# then swaps them all around.
# Now with Linearisation and Fusion.
#####
# Time started: 10:50
#####
# Set mode:..... Pixel Swapping#
# Read commands:..... Done#
#####
# Setup
# ..
# Opening file..... Done#
# Import soft classification data..... Done#
# Create class proportions arrays..... Done#
# Create random values array..... Done#
# Delete random values array..... Done#
# Initialised sub-pixels..... Done#
#####
# Pixel Swapping
# ..
# Create exponential function..... Done#
# Iteration: 1
# Iteration: 2
# Iteration: 3
# Iteration: 4
# Iteration: 5
# Iteration: 6
# Iteration: 7
# Iteration: 8
# Iteration: 9
# Iteration: 10
# Swap pixels..... Done#
#####
# Output
# ..
# Output final image..... Done#
#####
# Clean up
# ..
# Release memory..... Done#
# And breathe..... Yummy#
#####
# Finished, at: 10:50
#####
C:\Documents and Settings\Matt Thornton\My Documents\PS>

```

Screenshot 2: Application in pixel swapping mode

Screenshot 3: Field work data capture sheet

Site Name:				Sub-site No:									Quick description:
Date:				Sheet No:									
No.	Position	Width		Height	B1	B2	B3	B4	B5	B6	B7	B8	Notes (note any anomalies, points of interest etc.)
		B	T										