

University of Southampton Research Repository ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS

School of Engineering Sciences

**A MODEL FOR GENERATING OBJECT-BASED CHANGE
INFORMATION FROM MULTI-TEMPORAL
REMOTELY-SENSED IMAGERY**

by

John S Bevington

Thesis for the degree of Doctor of Philosophy

December 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS

SCHOOL OF ENGINEERING SCIENCES

Doctor of Philosophy

A MODEL FOR GENERATING OBJECT-BASED CHANGE
INFORMATION FROM MULTI-TEMPORAL REMOTELY-SENSED
IMAGERY

by John Scott Bevington

As world populations increasingly are clustered in urban areas, so there is a tangible need for accurate mapping of these regions by national mapping agencies. A consequential impact of growing cities is that greater numbers of people across the globe are vulnerable to the effects of natural disasters or anthropogenic catastrophes. Tools such as remote sensing have been widely used by researchers to monitor urban areas for applications such as land use and land cover changes and population distribution to name a few. Air- and space-borne sensors with fine spatio-temporal resolutions have facilitated these analyses, offering an effective and efficient data source for multi-temporal analysis of urban areas.

Alongside the increased data availability from remote sensors is a demand for efficient algorithms for interpretation of these images. This thesis describes the development of a conceptual framework for the iterative processing of fine spatial resolution optical images. It consists of two central components, object detection and object comparison.

In the object detection phase, buildings are identified in the image and extracted as objects stored in a scene model. Object attributes describing the location, geometric, spectral and textural characteristics of each object are stored in a database, allowing the on-demand display as vector or raster entities. The thesis implements the model through exemplars for the detection of circular and cylindrical features on several remote sensing and simulated datasets.

The object comparison phase allows automated change information to be generated describing per-object and intra-object brightness variability over time, hence, allowing change to be quantified for each detected feature. These descriptors facilitate the manual use of qualitative scales for damage assessment. A detailed discussion is presented on the merit of the conceptual model, its limitations and describes how future expansion of the model to full implementation could be achieved.

TABLE OF CONTENTS

1 INTRODUCTION	2
1.1 BACKGROUND	2
1.1.1 Nomenclature and Definitions	4
1.2 REMOTE SENSING	6
1.3 APPLICATIONS OF REMOTE SENSING	8
1.3.1 Introduction	8
1.3.2 Damage Detection	8
1.3.3 Monitoring Non-disaster Urban Change	10
1.4 ALGORITHM DEVELOPMENT	11
1.5 RATIONALE FOR RESEARCH	12
1.5.1 Aims & Objectives of the Thesis	13
1.5.2 Scope of this Research	15
1.6 THESIS FORMAT	16
2 LITERATURE REVIEW	22
2.1 INTRODUCTION	22
2.2 CHANGE DETECTION	22
2.2.1 Change detection methods	22
2.2.2 Application of Change Detection: (i) Monitoring of Urban Areas	28
2.2.3 Application of Change Detection: (ii) Damage Detection	30
2.3 OBJECT DETECTION AND FEATURE EXTRACTION	32
2.3.1 Background	32
2.3.2 Edge Detection Methods	33
2.3.3 Shape-matching Methods	34
2.3.4 The Hough Transform	35
2.3.5 Active Contours	43
2.4 SCENE MODELS	44
2.4.1 Background	44
2.4.2 Image Primitives	46
2.4.3 Optimising Primitive Matches	48
2.5 UNDERSTANDING “CHANGE”	53
2.5.1 Background	53
2.5.2 Quantitative Change Descriptors	53
2.5.3 Qualitative Change Descriptors	55
2.6 SUMMARY	58
3 THE CONCEPTUAL MODEL	61
3.1 BACKGROUND	61
3.1.1 Aims of the Conceptual Model	61

3.2	MODEL INPUTS	64
3.2.1	Image Characteristics	64
3.2.2	Pre-processing	65
3.3	OBJECT DETECTION MODULE	65
3.3.1	Segmentation Algorithm	67
3.3.2	Object Validation	68
3.3.3	Validation Iteration	70
3.3.4	Primitive Matching	71
3.4	OBJECT COMPARISON MODULE	72
3.4.1	Updating the Scene Model	72
3.4.2	Iteration	73
3.5	OUTPUTS	73
3.5.1	Scene Models	73
3.5.2	Change Descriptors	74
3.6	SUMMARY	74
3.6.1	Implementation	75
4	STUDY SITES & DATA	77
4.1	INTRODUCTION	77
4.2	SITE 1: SIMULATED IMAGES	77
4.3	SITE 2: BUNCEFIELD OIL DEPOT, UK	79
4.3.1	Background	79
4.3.2	Location of the Site	79
4.3.3	Imagery Details	79
4.4	SITE 3: DESERT IRRIGATION FACILITIES, SAUDI ARABIA	82
4.4.1	Background	82
4.4.2	Location and Imagery Details	83
5	TWO-DIMENSIONAL OBJECT EXTRACTION	86
5.1	INTRODUCTION	86
5.1.1	Background	86
5.1.2	The Application Development Environment	87
5.1.3	Method – The Circular Hough Transform (HT-)	87
5.2	ALGORITHMIC TESTING: THE HT- APPLIED TO SITE 1: SIMULATED IMAGERY	90
5.2.1	Input Parameters	91
5.2.2	Results	92
5.2.3	Discussion	96
5.3	APPLICATION TO SITE 2: BUNCEFIELD OIL DEPOT	98
5.3.1	Introduction	98
5.3.2	Method: The HT+ extension	99
5.3.3	Application to the Buncefield Imagery	102
5.3.4	Results	103
5.3.5	Discussion	106
5.3.6	Summary of the HT+ Buncefield Exemplar	107

5.4 APPLICATION TO SITE 3: CIRCULAR IRRIGATION FACILITIES, SAUDI ARABIA	108
5.4.1 Introduction	108
5.4.2 Study Site and Data	109
5.4.3 Method	109
5.4.4 Results	112
5.4.5 Discussion	115
5.4.6 Summary of the HT+ Saudi Arabia Exemplar	116
5.5 SUMMARY	116
6 THREE-DIMENSIONAL OBJECT EXTRACTION	119
6.1 INTRODUCTION	119
6.1.1 Background	119
6.1.2 Image Considerations	119
6.2 STUDY SITE AND DATA	123
6.2.1 Choice of Input Imagery	123
6.2.2 Image Metadata Preparation	123
6.3 METHOD	124
6.3.1 Object Detection	124
6.3.2 The Hough Transform	125
6.3.3 Automatic Circle Removal	126
6.3.4 Similarity Matrices & Filters	127
6.3.5 The Score Matrix	132
6.3.6 Circle Matching	134
6.3.7 Optimisation by Simulated Annealing	134
6.3.8 Automatic Object Completion & Cylinder Orientation	137
6.3.9 Automatic Determination of ' η '	141
6.3.10 Repeated SA Using New Value of η	141
6.3.11 Monte Carlo Simulations	142
6.4 RESULTS	142
6.5 DISCUSSION	149
6.6 MODELLING THE THREE-DIMENSIONAL SCENE	153
6.6.1 POV-Ray 3-D Models	153
6.6.2 Google Earth Models	155
6.7 SUMMARY	158
7 DETECTION OF OBJECT CHANGE	160
7.1 INTRODUCTION	160
7.1.1 The Input Scene Model	160
7.2 METHOD	161
7.2.1 Image Preparation and Pre-processing	162
7.2.2 Grounding the Object Bases	163
7.2.3 Projection of Top Circle	163
7.2.4 Forward Modelling	166
7.2.5 Quantification of Change	167
7.2.6 Describing Change	168
7.3 RESULTS	169

7.4 DISCUSSION	174
7.4.1 Object Changes due to Differences in Image Geometry	174
7.4.2 Object Differences Described in Quantitative and Qualitative Terms	174
7.5 SUMMARY	177
8 DISCUSSION	180
8.1 BACKGROUND	180
8.2 THE CONCEPTUAL MODEL	181
8.2.1 Object Detection	181
8.2.2 Object Comparison	184
8.3 AIMS AND OBJECTIVES	184
8.3.1 Objective 1: Provide a framework or guidelines for processing of RS images	185
8.3.2 Objective 2: Provide methods that identify objects of interest and provide pre-event information as a basis for change detection	186
8.3.3 Objective 3: Provide a model-based platform for per-building analysis	187
8.3.4 Objective 4: Provide semi-automated and iterative methods	188
8.3.5 Objective 5: Provide a change detection method for capturing urban morphological change	189
8.3.6 Objective 6: Provide results in a form that facilitates the generation of qualitative and quantitative change descriptors	190
8.4 POTENTIAL MODEL ENHANCEMENTS	192
8.4.1 Enhancement of the Object Detection Module	193
8.4.2 Enhancement of the Object Comparison Module	194
8.5 SUMMARY	196
9 CONCLUSIONS	198

LIST OF FIGURES, TABLES & EQUATIONS

Figure 1.1. Global economic loss risks from natural disasters as a proportion of GDP per unit area.....	3
Figure 1.2. Global urban and rural population change 1950-2050 (UN, 2007).....	3
Figure 1.3 Flowchart of thesis sections.....	18
Figure 2.1 Workflow of a circular Hough Transform algorithm	36
Figure 2.2. Stages of the circular Hough Transform algorithm.	37
Figure 2.3 Generation of an Accumulator Space for a HT circle detector (from Sonka <i>et al.</i> , 1999)	38
Figure 2.4 The five-parameter dimensions of a cylinder	40
Figure 2.5 Traditional gradient descent optimisation.	49
Figure 2.6 An example of a simulated annealing cooling schedule.....	52
Figure 2.7 Classification of damage to masonry buildings.....	56
Figure 3.1 The Conceptual Model showing four phases: 1. Inputs; 2. Object Detection; 3. Object Comparison; 4. Outputs. ($n =$ iteration number).....	63
Figure 3.2 The effect of different view angles on object footprint delineation	66
Figure 4.1. Location of the Buncefield Oil Depot, Hemel Hempstead, UK.	80
Figure 4.2. Detailed layout of the Buncefield Oil Depot Site (BMIIB, 2006).....	80
Figure 4.3. Pre-event aerial image of the Buncefield site.	81
Figure 4.4. Post-event aerial image of the Buncefield site	81
Figure 4.5. Location of Site 3 in Saudi Arabia (marked by red cross).	83
Figure 4.6. Image showing circular irrigation facilities in Saudi Arabia.....	84
Figure 5.1. Greyscale simulated image generated for image 2b.	91
Figure 5.2. Error information generated for image 2b.	92
Figure 5.3. Accumulator space with detected circles overlaid for synthetic image-set 1.....	93
Figure 5.4. Accumulator space with detected circles overlaid for synthetic image-set 2.....	94
Figure 5.5. Accumulator space with detected circles overlaid for synthetic image-set 3.....	95
Figure 5.6. An exaggerated example of the discretisation problem.....	96
Figure 5.7. Subset of the Buncefield image used as input to the model	101
Figure 5.8. Results of 200 iterations of HT+ applied to Buncefield imagery.....	104
Figure 5.9. RMSE percentage & number of circles detected at each iteration	105
Figure 5.10. Correlation coefficient & no. circles detected at each iteration.....	105
Figure 5.11. Subset of the Saudi Arabia image used as input to the model.....	109
Figure 5.12. Results from the Canny edge detection in the x (a) and y (b) directions	111
Figure 5.13. Circles detected from iteration number 1 (a), 25 (b), 50 (c), 75 (d) and 100 (e)	113
Figure 5.14. RMSE (grey-levels) & number of circles detected at each iteration ..	114
Figure 5.15. Correlation coefficient & number of circles detected at each iteration	114
Figure 6.1. Image, sensor and illumination geometry.....	120
Figure 6.2. Comparative geometry of a map (orthographic projection: a), and vertical aerial photograph (perspective projection: b).	121

Figure 6.3. Examples of how common shapes differ in appearance when displayed in perspective projection (a) and orthogonal projection (b).....	122
Figure 6.4. Two methods of automatic cylinder detection from perspective imagery.	125
Figure 6.5. Workflow for automatic cylinder detection.....	126
Figure 6.6. Example of the criteria for automatic removal of circles.	127
Figure 6.7. Creation of the direction filter for images where $90^\circ < D_s < 180^\circ$	130
Figure 6.8. Example of the circle swapping process.....	135
Figure 6.9. Potential erroneous matching of top circles. These potential errors can be reduced by reducing the distance threshold	137
Figure 6.10. Orientation of cylinders with respect to sensor location (perspective projection)	138
Figure 6.11. Generation of top and base node information.....	139
Figure 6.12. Circles detected by the Hough Transform algorithm	143
Figure 6.13. Circle matching and simulating cylinders	144
Figure 6.14. Local energy fluctuations for each iteration of SA.....	146
Figure 6.15. Temperature and lowest energy reduction for 1 run of SA	146
Figure 6.16. Top & base circle matches for 50 iterations of three-cylinder creation	147
Figure 6.17. Total RMSE and number of correct cylinder detections for 50 iterations of three-cylinder creation. Red, orange and green columns show one, two and three correct matches respectively	147
Figure 6.18. Grouped RMSE of best F matched cylinders.....	147
Figure 6.19. The modal combination of cylinders for $\eta = 3$	148
Figure 6.20. Correct cylinder detection from Monte Carlo iteration 47	148
Figure 6.21. POV-Ray rendering of the three correct cylinders	155
Figure 6.22. GIS-produced map of top (red) & base (blue) circles in geographic coordinates.	156
Figure 6.23. Oblique-angle view of Google Earth-generated 3-D objects.	157
Figure 6.24. Nadir view of post-event Buncefield image with detected cylinders overlaid.....	157
Figure 7.1. Workflow of change detection methodology	162
Figure 7.2. Subset of the resampled post-explosion image of the Buncefield oil depot (2006).....	163
Figure 7.3. Cylinder geometry showing displacement (d) of top and base circles .	164
Figure 7.4. Calculation of the magnitude of projected displacement (d) between top, C_t and base, C_b of an object. Also shown is the object height, h	165
Figure 7.5. Calculation of the direction of displacement	166
Figure 7.6. Calculation of per-pixel texture information	168
Figure 7.7. Locations (nadir points) of the <i>before</i> (green) and <i>after</i> (red) image sensors in relation to the study area (blue).....	169
Figure 7.8. Pre-event (a) and post-event (b) Scene Model objects projected in greyscale with 1999 and 2006 image geometry respectively, and including intra-object spectral properties.....	171
Figure 7.9. Display of the six change measures used for description of intra-object pixel brightness changes	172
Figure 7.10. Intra-object textural comparison.....	173
Figure 7.11. Damage to the Buncefield Oil Depot as classified by the European Scale of Industrial Accidents (ARIA, 2009).....	176

LIST OF TABLES

Table 1.1. Rationale and Aims of the Thesis	17
Table 1.2. Thesis outline and context of chapters	19
Table 4.1. Comparison of data used in the application of the conceptual model	77
Table 4.2. Simulated images used in chapter 5	78
Table 5.1. Input parameters used to test the HT algorithm on the synthetic datasets.	91
Table 5.2. Statistical results for image-set 1	93
Table 5.3. Statistical results for image-set 2	94
Table 5.4. Statistical results for image-set 3	95
Table 5.5. Input parameters used for application on the Buncefield image.....	102
Table 5.6. Additional statistical results from the 200 <i>HT+</i> iterations.....	106
Table 5.7. Input parameters used for the Saudi image application	110
Table 5.8. Additional statistical results from the 100 <i>HT+</i> iterations	112
Table 6.1. Parameters used in the <i>HT-</i> implementation	126
Table 6.2. Sensor and illumination geometry for the Buncefield image subset	142
Table 6.3. Most frequent circle matches and Score for 50 MC implementations of SA where $\eta=3$	145
Table 7.1. Object location data from the pre- (a) and post-event (b) Scene Models	170
Table 7.2. Per-object spectral data from SM_1 and SM_2 . Units are greyscale brightness values.	170
Table 7.3. Qualitative descriptors of change from a manual interpretation using the EMS-98 and the WFS Damage Scales.....	172

DECLARATION OF AUTHORSHIP

I, John Bevington, declare that the thesis entitled *A Model for Generating Object-based Change Information from Multi-temporal Remotely-Sensed Imagery* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:

Bevington, J.S., Lewis, H.G., Atkinson, P.M. (2007) Quantitative modelling of unexpected change using remotely sensed imagery. In *Proceedings of the 32nd International Symposium on Remote Sensing of Environment (ISRSE07), San Jose, Costa Rica, 25-29 June 2007*.

Signed:

Date: 10 December 2009

ACKNOWLEDGEMENTS

I would like to foremost acknowledge my supervisory team of Dr. Hugh Lewis in the School of Engineering Sciences, and Prof. Peter Atkinson, Head of the School of Geography, for their diligent and unwavering support of my academic and personal development through the course of this Ph.D. research. They have offered guidance when necessary from the initial stages of this research through to the final stages of completion, and, with their support, this thesis has been possible.

I would also like to acknowledge the input and advice offered to me by Dr. David Holland and Dr. Isabel Sargent from the Research team at Ordnance Survey, Southampton. They have lent their varied experiences in remote sensing, photogrammetry and object detection to keep this project grounded and applicable beyond the academic realm.

I would like to recognise the advice offered to me from the other members of the Astronautics Research Group, most notably my internal examiner, Dr. Adrian Tatnall, who proved an important sounding board at the mid-point of this research. And a fond word or two of thanks should go to the numerous researchers who have become friends whilst passing through the Ph.D. offices of the Astronautics Research Group. They would be embittered if they looked at this and didn't see a mention by name, so a big thanks to Jay, François, Nick B, Nick L, Piers, Rebecca, Sam, Stefano and Tarek, amongst others. I would mention all the Aero. guys as well, but they were better at football than us...

The continued support from my current employers should be recognised, as they have offered me the platform to build on the research carried out during this Ph.D, and to further my research career in a challenging and stimulating environment. I wish to thank Dr. Beverley Adams, Managing Director of ImageCat Ltd, for her continued support and for enabling me to accomplish many professional goals.

Many housemates and friends, some now living far and wide, have shared this journey, and should be thanked for their friendship and patience, whether in Southampton or Surrey. I would also like to acknowledge the fellow members of the 2007 UK GRAD programme, Exeter, for the camaraderie of shared experiences.

The most praise and thanks is reserved for my family, especially my Mum and Dad, who have selflessly given their unconditional support, not just in the duration of this research, but also throughout my life, and the choices I have made. Sadly, Dad didn't get to see the final bound version of this thesis, or indeed graduation, but he knew I had completed this journey, and it filled him with enormous pride. He will forever be missed.

Lastly, I dedicate this thesis to Rachel, who has lived through all the ups and downs of Ph.D. research by my side. I am indebted to her for her unfailing patience during the writing of this thesis, and she is the best partner I could ever wish for.

ABBREVIATIONS AND ACRONYMS

AHT	Adaptive Hough Transform
AI	Artificial Intelligence
ANN	Artificial Neural Network
ATOMI	Automated reconstruction of Topographic Objects using aerial images by means of vectorised Map Information
AVHRR	Advanced Very High Resolution Radiometer
BMIIB	Buncefield Major Incident Investigation Board
BUS	Building Unit Shape
CNES	Centre National pour les Etudes Spatiales (French National Space Agency)
CAT	Computed Tomography
CVA	Change Vector Analysis
DEM	Digital Elevation Model
DMC(ii)	Disaster Monitoring Constellation (International Imaging Ltd.)
DMSP / OLS	Defense and Meteorological Satellite Program / Operational Linescan System
DN	Digital Number
DRHT	Dynamic Randomised Hough Transform
DSM	Digital Surface Model
EC	European Commission
EMS	European Macroseismic Scale
EO	Earth Observation
EPA	Environmental Protection Agency (USA)
ESA	European Space Agency
ESIA	European Scale of Industrial Accidents
FEMA	Federal Emergency Management Agency
FHT	Fast Hough Transform
GA	Genetic Algorithm
GEO	Group on Earth Observations
GEO-CAN	Global Earth Observation Catastrophe Assessment Network
GEOSS	Global Earth Observation System of Systems
GIS	Geographic Information Systems
GMES	Global Monitoring for Environment and Security
GPS	Global Positioning System
GSD	Ground Sample Distance
HHT	Hierarchical Hough Transform
HT	Hough Transform
<i>HT- / HT+</i>	The unmodified/modified Hough Transform (as described in Chapter 5)
HRV	High Resolution Visible
HRVIR	High Resolution Visible & Near Infrared
IKONOS	Not an acronym. A commercial Earth observation satellite
IFOV	Instantaneous Field Of View
JPEG	Joint Photographic Experts Group (industry image standard)

KML	Keyhole Markup Language
LBS	Location-Based Services
MATLAB	Matrix Laboratory
MODIS	Moderate-resolution Imaging Spectroradiometer
MC	Monte Carlo iteration
MRI	Magnetic Resonance Imaging
MSS	Multispectral Scanning System (Landsat satellite sensor)
NaN	Not a number
NDVI	Normalised Difference Vegetation Index
NGD	National Geographic Database
NOAA	National Oceanic and Atmospheric Administration
OWL	Web Ontology Language ¹
PC	Personal Computer
PDA	Personal Digital Assistant
PDNA	Post-disaster Needs Assessment
POV-Ray	Persistence of Vision Raytracer
PPHT	Progressive Probabilistic Hough Transform
RESPOND	Not an acronym. Defined as follows: <i>“RESPOND is an alliance of European and International organisations working with the humanitarian community to improve access to maps, satellite imagery and geographic information.”</i> (RESPOND, 2008)
RGB	Red, Green, Blue (Spectral Band combination)
RMSE	Root Mean Square Error
RS	Remote Sensing
SA	Simulated Annealing
SAFER	Services and Applications for Emergency Response
SAR	Synthetic Aperture Radar
SHT	Standard Hough Transform
SI	Simulated Image
SM	Scene Model
SPOT	Système Pour l’Observation de la Terre
SSTL	Surrey Satellite Technology Limited
TLM	Topographic Landscape Model
TM	Thematic Mapper (Landsat satellite sensor)
TSP	Travelling Salesman Problem
UAV	Unmanned Aerial Vehicle
UN	United Nations
UNOSAT	Defined as follows: <i>“UNOSAT is a United Nations programme created to provide the international community and developing countries with enhanced access to satellite imagery and Geographic Information System services”</i> (2007)
USAID	United States Agency for International Development
WFS	Wind and Flood Scale

¹ “Although the character Owl from Winnie the Pooh wrote his name WOL, the acronym OWL was proposed without reference to that character... and, to quote Guus Schreiber, “Why not be inconsistent in at least one aspect of a language which is all about consistency?” “. (Wikipedia entry for Web Ontology Language)

CHAPTER 1

INTRODUCTION

1 INTRODUCTION

1.1 BACKGROUND

Recent years have seen an increasing trend for human populations across the globe to live in urban areas (UN, 2007: Figure 1.1). Significantly, much of this increase has taken place in less-developed countries, with 93% of the increase before 2020 associated with Third World cities (Baudot, 2001). The Chinese city of Shanghai, for example, has been one of the fastest growing cities in the world; its population increasing from 13.5 million to 21.5 million in the past 15 years and the physical size of the city increasing over six times from 100 km² to 680 km² (BBC, 2007). A decade ago, it was estimated that 70% of the population of western countries were living in urban areas (Henderson and Xia, 1997). These densely populated regions, especially in developing nations, can expand at a rapid rate, making the stored geographic data of the regions obsolete in a short time. However, significant problems lie beyond acquiring accurate geographic data of expanding metropolises.

Many cities experiencing rapid population growth are at risk of severe population and economic losses due to natural disasters (Dilley *et al.*, 2005: Figure 1.2), whether as a result of a single or a combination of factors, for example, geologic (earthquake, volcanic eruption, landslide etc.), hydrologic (flood, tsunami, etc.) and climatic (drought, windstorms, wildfires, etc.) factors (Figure 1.2). The frequency of occurrence and intensity of natural disasters are expected to intensify as sea surface temperatures rise due to global warming, causing climatic extremes and an increase in the incidence and magnitude of severe storm events (Salinger, 2005; Jiang and Perrie, 2007).

Large numbers of people and infrastructure are also at risk from anthropogenic or human-induced disasters. Examples of such unpredictable incidents have been witnessed around the globe in recent years, ranging from the 11 September 2001 terrorist attacks in the United States, to the 2005 Buncefield Oil Depot fire in the United Kingdom. There is a need for rapid reaction to both natural and anthropogenic disasters to provide data and methods from which to assess the

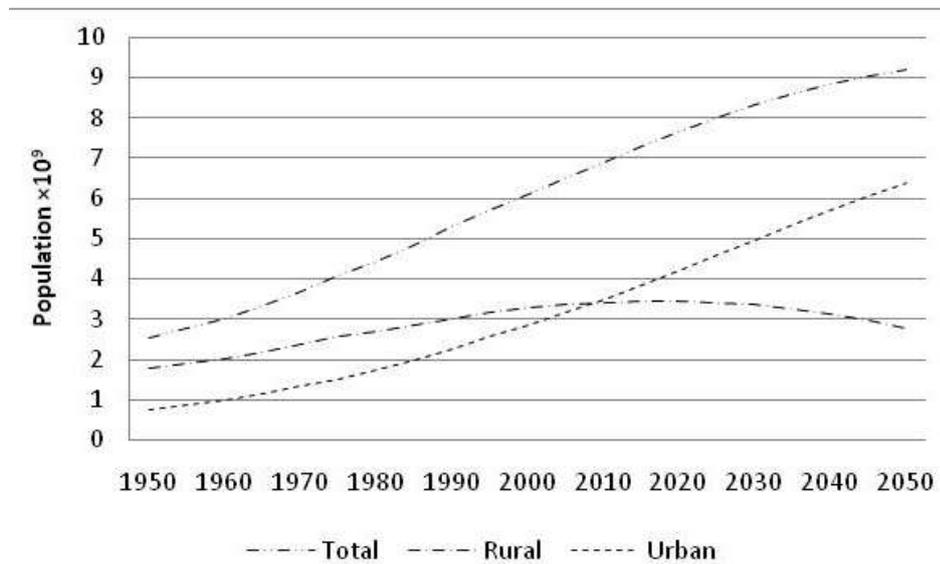


Figure 1.1. Global urban and rural population change 1950-2050 (UN, 2007)

scale of the event, thereby limiting the casualty and fatality numbers that may occur after such events.

Monitoring unpredictable changes in urban areas, whether as a result of natural or anthropogenic disasters, is one of several important reasons for maintaining an accurate inventory of building and infrastructure stock in burgeoning cities across the globe. The demand for geographic information is increasing in tandem with the advancement of technologies such as location-based services (LBS) through the internet, global positioning system (GPS), handheld personal digital assistants (PDAs) and mobile phones. The need for rapid retrieval of geographic

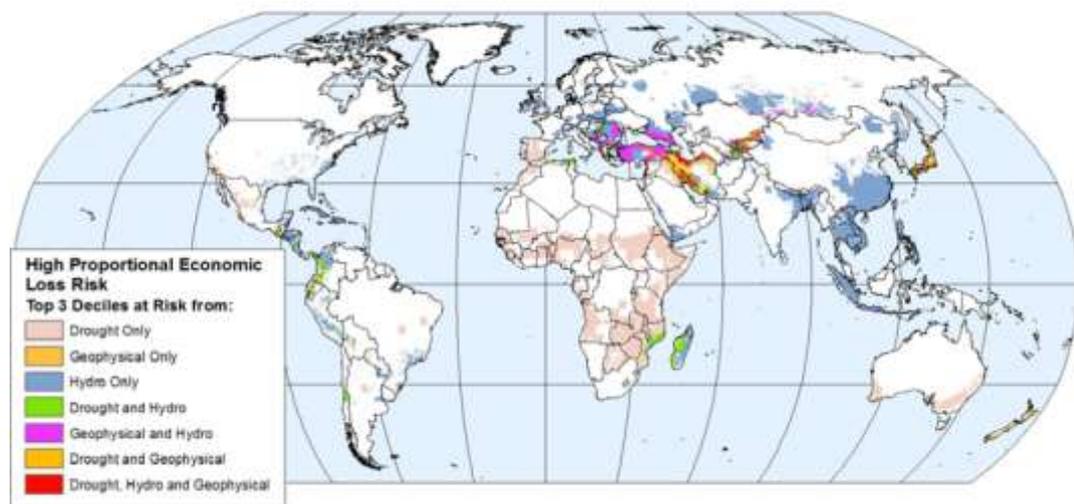


Figure 1.2. Global economic loss risks from natural disasters as a proportion of GDP per unit area (from (Dilley *et al.*, 2005).

data, especially in newly industrialised and developing nations, increases as the use of the internet and mobile LBS continues to grow.

Monitoring the evolution and morphology of the urban environment is a major remit for national mapping agencies such as the Ordnance Survey in Great Britain, who are tasked with the upkeep of the National Geographic Database (NGD). Ordnance Survey is constantly refining techniques for data capture and semantic retrieval of data; in the past, they have relied on ground-based surveys by (human) surveyors as a method of data collection for updating the NGD. This traditional method is increasingly complemented by data collected from remotely-sensed imagery (described in section 1.2). Imagery can be processed through a combination of techniques, with the aim of extracting information to compare with Ordnance Survey's flagship digital vector dataset OS MasterMap[®] (Ordnance Survey, 2009a).

Working in collaboration with Ordnance Survey Research, the work described in this thesis focuses on the dual applications of monitoring both urban morphology and damage caused by natural or anthropogenic disasters. It describes the development of a new change detection and change description method for multi-temporal imagery captured by remote sensors.

1.1.1 Nomenclature and Definitions

At this stage, it is important to introduce several terms, and explain their usage in this thesis. Terms are used subjectively to differentiate between changes witnessed as a result of natural and anthropogenic disasters, and changes resulting from dynamic urban morphology.

Change

It can be said that natural disasters or anthropogenic incidents can occur suddenly and without warning, yet these phenomena are not immediately recognisable as the cause of change seen in remotely-sensed imagery. The term *sudden* can normally be used to describe these events, but from the viewpoint of the analyst, the suddenness of change has more to do with the image sampling rate. Change could be said to be

“sudden” if the analyst was comparing images with a temporal difference of 10 years. Therefore, throughout this thesis, change caused by natural disasters, industrial, or human-induced accidents will be classed as *unexpected* change events.

Conversely, much of the evolution of urban areas is foreseen and deliberate (such as building demolition, extensions or new developments that may have proceeded through a planning process). When monitored regularly, these appear to be subtle, gradual changes. However, as the sampling rate of image collection is not at regular incremental time periods, the use of “gradual” can be misleading. To provide contrast and differentiation to *unexpected* change, the change attributed to the normal growth and evolution of urban areas will be classed as *planned* change. Because of these factors, temporal descriptions of any phenomena are misleading if there is a discrepancy between change in the phenomena of interest and the sampling rate.

Spatial Resolution

A note should be made at this point on the spatial resolution of sensors and imagery. The spatial resolution of a remote sensor refers to the size of the smallest feature that can be detected within an image. It largely depends on the ground sample distance (GSD) or instantaneous field of view (IFOV) – the area of the ground visible to the sensor at a given altitude at a point in time. The size of this area is calculated by multiplying the IFOV by the distance from the ground to the sensor, and is referred to as the resolution cell. This determines the maximum spatial resolution of a sensor (Driggers, 2003; CCRS, 2008). The resolution cell can also be referred to as a pixel (a cell in the matrix that makes up an image), which is normally square and represents a certain area on an image. For a single feature to be detected, it should be equal to, or larger than, the pixel size. Below this size, the feature may not be detected, unless its reflectance dominates the resolution cell. For example, an image with a GSD of 1 m has pixels representing a ground area measuring 1 m x 1 m. Features with dimensions less than this may not be detectable by traditional feature extraction techniques, and may rely on sub-pixel detection techniques.

For the purposes of this thesis, spatial resolution and pixel size will be used interchangeably. If a spatial resolution of 2 metres is stated, this refers to an image displayed at full resolution, with each pixel representing an area of 2 m x 2 m, as

observed on the ground. However, it is acknowledged that images may be displayed with a pixel size different than the resolution.

1.2 REMOTE SENSING

One of the most accessible and widely utilised techniques for the detection and monitoring of planned or unexpected change is the use of remotely-sensed imagery captured before and after the change event, combined with algorithms designed to interpret multi-temporal change. The definition of *remote sensing* varies across applications, ranging from medical imagery (e.g. magnetic resonance imaging (MRI), or computed tomography (CAT) scans), to meteorology (e.g. Radar). Remote sensing for Earth observation (EO) is effectively summarised by Campbell (2002) as:

“the practice of deriving information about the Earth’s land and water surfaces using images acquired from an overhead perspective, using electromagnetic radiation in one or more regions of the electromagnetic spectrum, reflected or emitted from the earth’s surface”. (Campbell, 2002)

The use of remote sensing (RS) for EO has evolved from the use of cameras in hot air balloons in 19th century France (Campbell, 2002) to modern-day constellations of EO satellites. RS has become an indispensable technology for monitoring change in terrestrial phenomena; from sea-surface temperatures to global land cover change.

The evolution of modern-day optical sensors (those that sense the visible radiometric wavelengths of the electro-magnetic spectrum) has occurred in three stages, or generations (Donnay *et al.*, 2001), from the first generation, including the Landsat Multispectral Scanning System (MSS) sensor in 1973 (79 metre spatial resolution), the Landsat Thematic Mapper (TM) sensor onboard Landsat IV, V and VII satellites (30 m spatial resolution), to second generation moderate spatial resolution sensors, such as the high resolution visible (HRV) sensor onboard Système Pour l’Observation de la Terre I-IV (SPOT) satellites (10 m spatial resolution). The third generation saw licences granted to commercial companies in the late 1990s to capture very fine spatial resolution images (sub-10 metres)

following the declassification of large amounts of military reconnaissance data (Aplin, 2003).

The third generation, optical commercial satellites include IKONOS, Quickbird, and GeoEye-1 (41 cm spatial resolution in panchromatic mode). These commercially available sensors, with fine spatial, temporal and spectral resolutions, now produce imagery of isolated areas of the world not previously studied in such fine detail, and can be commercially programmed on an *ad hoc* basis to capture data on a specific region in a matter of hours.

EO sensors have evolved further since Donnay's description, with a parallel advancement in low-cost mini- and micro-satellites that have the capability to carry fine spatial resolution sensors (Sweeting and Underwood, 2003). These satellites operate independently, or as part of a constellation, offering reduced revisit rates and fine spatial resolution. The TopSat micro-satellite, developed by Surrey Satellite Technology Ltd. (SSTL), the Rutherford Appleton Laboratory and QinetiQ, and launched in 2005, has captured panchromatic optical imagery at 2.8m spatial resolution (QinetiQ, 2007). The Disaster Monitoring Constellation (DMC), also developed by SSTL and administered by DMC International Imaging (DMCii), consists of seven nationally owned micro-satellites from UK, Algeria, China, Nigeria, Spain and Turkey. These fly in Sun-synchronous orbits, allowing daily revisit rates and maximum spatial resolutions of 4 m (DMCii, 2009).

Optical data are also collected through the use of airborne sensors mounted in aircraft, or unmanned aerial vehicles (UAVs). These forms of aerial imagery have several advantages over optical imagery captured by satellite sensors. The spatial resolution of aerial sensors can be finer than that achieved by spaceborne sensors, although the spatial coverage of each image is smaller than those captured by satellite sensors due to the operational altitude of the airborne sensor. The low altitude at which aerial imagery is captured (1-2 km) reduces the effect of occlusion by clouds and haze on ground features than compared to imagery captured by satellite sensors. The revisit rates of aerial sensors are also not constrained by orbital characteristics, and so are much more flexible.

1.3 APPLICATIONS OF REMOTE SENSING

1.3.1 Introduction

There is a vast array of tangible and important applications for imagery captured by airborne or spaceborne sensors. These range from meteorological uses (Rao *et al.*, 1989), to military purposes (Deer, 1999), or management of natural resources, such as forests (Hall *et al.*, 1991; Foody, 2003). However, this thesis primarily focuses on the use of fine spatial resolution optical remotely sensed imagery for the purposes of post-disaster damage assessment for urban areas, as well as monitoring of other, less dramatic, urban change - significant for urban mapping applications. Several of the image processing methods described in chapter 2 may be transferable across multiple remote sensing applications, however the method and exemplars described in later chapters have been developed for the dual applications of urban change assessment, both in post-disaster and non-disaster situations.

1.3.2 Damage Detection

Increasingly, researchers are using optical imagery captured from spaceborne and airborne sensors to detect and assess damage to built-up areas after natural or anthropogenic disasters (Chiroiu and Andre, 2001; Liu *et al.*, 2004; Gusella *et al.*, 2005b). As the use of RS technology for damage detection has evolved, the need for international cooperation has become more apparent. Consequently, the use of RS data in this context has improved the coordination of agencies and organisations offering humanitarian aid, and search and rescue services by focussing response efforts (Miura *et al.*, 2005). International agreements and policy initiatives, such as the International Charter on Space and Major Disasters², have been created to oversee the collection and distribution of geographic data. The Charter was formed by the European Space Agency (ESA) and the French Space Agency (Centre National pour les Etudes Spatiales; CNES), and is a legally non-binding intention to share knowledge and data free of charge to any country affected by natural or

² Charter on Cooperation to Achieve the Coordinated Use of Space Facilities in the Event of Natural or Technological Disasters: hereafter: "The Charter" (International Charter, 2007).

technological disasters, excluding war, armed conflicts, ice monitoring, drought, and routine epidemiological outbreaks (Ito, 2005).

Other methods of coordination exist globally, including the RESPOND and Services and Applications for Emergency Response (SAFER) initiatives, backed by the European Commission's and ESA's Global Monitoring for Environment and Security (GMES) initiative. These European schemes evolved from ESA's UNOSAT project to bring together value added companies and public agencies to provide vital geo-information to coordinate humanitarian aid (RESPOND, 2004; ESA, 2007). RESPOND and SAFER have delivered map products following a wide array of natural and human-induced incidents, most notably, in the ongoing monitoring of the Darfur crisis (Bally *et al.*, 2005), and its partners include the DMC. In 2002, the G8 nations of industrialised countries, and the World Summit recognised the need for coordinated international efforts to provide a global network for environmental monitoring using spaceborne sensors. This led the Group on Earth Observations (GEO) to initiate the Global Earth Observation System of Systems (GEOSS), with "Societal Benefit Areas" including disasters and climate (GEO, 2009). Whilst not specifically intended for post-disaster response, GEOSS has a remit for monitoring ongoing environmental restoration following natural or technological disasters.

At a local level, policies have been created for the integration of remote sensing technology and geographic data for post-disaster response. Laben (2002), describes the example of the Pacific Disaster Center and uses it as an organisational and technological model for local, national and global initiatives for disaster management. Specifically, Laben uses the model of disaster management initially devised by the US Federal Emergency Management Agency (FEMA). This model addresses the four phases of a disaster scenario: Mitigation, Preparation, Response, and Recovery (Montoya, 2003).

1. *"MITIGATION relates to activities before an event that reduce or eliminate the chance or effects of a disaster. Risk assessments and the reduction of potential effects of disasters are examples of mitigation.*
 2. *PREPARATION involves planning response if a disaster occurs and working to increase resources to respond effectively.*
-

3. *RESPONSE covers activities and measures implemented immediately after a disaster that provide aid to victims and reduce the chance of secondary damage.*
4. *RECOVERY lasts until all systems and infrastructure return to a normal state after an event. This is the last phase of the disaster management model.”*

The magnitude of a natural or anthropogenic disaster is not the only parameter that affects the scale of response and the resulting number of agencies involved in each of FEMA’s stages. The disaster location also plays a vital role in response and recovery. Less developed regions of the world are often poorly advanced in the mitigation and preparation stages due to smaller budgets for such activities and inferior infrastructure, geographic data, technology and knowledge, as was apparent in the aftermath of the 2004 Indian Ocean Tsunami. Remote sensing technologies provide a demonstrable alternative data source to fill to this data gap (Laben, 2002; Vu *et al.*, 2007; Yang *et al.*, 2007). There is clearly a need for coordinated, international approaches to the dissemination of geographic data after disasters and the primary mechanism for this is through The Charter, or through FEMA in the USA (Laben, 2002; ESA, 2007).

RS data captured before a change event can be used to provide baseline information for areas of the world where it is lacking. This can then be compared to the post-event image information to generate change information that describes the extent of damage. Remote sensors already provide a readily available global data archive and reliable data source on areas affected by natural or anthropogenic incidents globally, regardless of the availability of a current geographic building inventory.

1.3.3 Monitoring Non-disaster Urban Change

A major use of optical remotely-sensed imagery has been for the monitoring and updating of geographic datasets (Baillioeul *et al.*, 2003). Increasingly, national mapping agencies, such as the Ordnance Survey are using fine spatial resolution images to capture data at the scale of individual buildings to increase the accuracy and aid the updating of their data products; essentially, collecting richer data more

rapidly (Holland and Tompkinson, 2003). As analysis methods, such as feature extraction and change detection techniques, become further refined and automated, the goal of rapid data collection becomes closer to reality. Tompkinson *et al.* (2003) acknowledged the lack of existing automated methods of change detection, stating that if a technique relies on input from a human operator, including the interpretation of classified change maps, then it cannot be described as automated.

Due to the extensive number of datasets held in the National Geographic Database, it is imperative to assign meaning to geographic features so that data can be retrieved efficiently. Part of Ordnance Survey's ongoing research focuses on developing machine-readable semantic retrieval systems. By assigning meaning to data, it becomes easier to integrate into users' systems, and, thus, allow the efficient retrieval of data using Web Ontology Languages (OWL), and the Semantic Web (Ordnance Survey, 2009).

1.4 ALGORITHM DEVELOPMENT

It is apparent that remote sensing is a readily-available data source for gathering information on both planned and unexpected change events. Consequently, there is a parallel need for powerful algorithms for processing of RS data to describe change efficiently and effectively after a planned or unexpected event in both quantitative and qualitative terms.

The progression of data collection technologies in RS has occurred alongside a growth in the use of personal computers and the subsequent increase in processing power available to users. This has furthered the development of automated analytical techniques that focus on image interpretation, often for images captured by remote sensors. The use of these techniques, from the fields of computer science, engineering and mathematics, has reduced the dependency on human-supervised methods. Feature extraction techniques designed in the field of machine vision to detect objects such as faces in photographs have been shown to be generalisable in their application to EO imagery (Cross, 1988; Nixon and Aguado, 2002). Examples of these techniques are presented in detail in chapter 2.

The development of algorithms for use with RS data is necessary for several reasons. Firstly, the multitude of features in the imagery makes analysis of individual features complicated. There are also complex physical processes to be monitored periodically using imagery. Finally, the large amount of data availability captured from a variety of sensors demands algorithms that are flexible and generalisable across imagery types. EO data of urban areas can incorporate a large variety of land covers and land uses, and contain many more features adding to the complexity of the imagery. Cloud, haze and shadows are some of the phenomena that can occlude image features targeted in the analysis. One method of reducing the complexity of the image is by identifying and extracting only the image features of interest. This segmentation of the image can allow delineation of thematic classes such as vegetation, water, or individual buildings (depending on the spatial resolution of the image), which can be stored in a model – a simplified extraction of real world features of interest. For regions of the world lacking comprehensive and accurate geographic data coverage, this modelling approach allows individual buildings to be extracted from a pre-event image, using one of many existing feature extraction techniques, and providing baseline information for change detection.

There are existing networks for the provision of RS data after unexpected events, as described previously, but there are no standardised methods for processing these data to determine the event's extent which has implications for the response and recovery stages of the disaster cycle. Having automated or semi-automated methods that are generalisable across change events and locations would be an effective method of building resilience to unexpected change events in the preparedness stage of the disaster cycle.

1.5 RATIONALE FOR RESEARCH

Monitoring urban change and damage detection are important priorities for national and international governments and agencies. The ability to respond effectively after natural or anthropogenic disasters relies on a timely and accurate assessment of what has changed. Remote sensors can provide a rapid source of fine spatial resolution data, fine enough to detect the characteristics of individual buildings in an urban

area. With knowledge gained from a multi-temporal comparison of fine spatial resolution imagery, the disaster response can be coordinated and recovery efforts monitored. Similarly, with knowledge of subtle urban changes gained from analysis of multi-temporal imagery, national mapping agencies are capable to provide data and mapping services to the population for use in geo-spatial services, ranging from satellite navigation to three-dimensional city models.

This research hypothesises that both urban morphology and damage from unexpected events can be monitored and assessed by adopting similar methods. Using pre- and post-event RS imagery as datasets, techniques can be created by modifying and amalgamating existing algorithms from the fields of remote sensing and computer vision. Chapter 2 describes some of the methods utilised previously by researchers in these fields for a variety of applications, from object detection and feature extraction, to the detection and description of change. Several key requirements have been identified and form the rationale for this work.

1.5.1 Aims & Objectives of the Thesis

The thesis describes a novel conceptual model has been created that has two central components – object detection and object comparison. Briefly, existing methods of detecting primitive features are extended to extract both 2-D and 3-D objects of interest from fine spatial resolution imagery captured before a change event. The objects are then modelled using the image geometry of the post-event image and statistics generated to describe the changes in objects between the time periods. This thesis uses exemplars to describe the development of this conceptual model, and investigates its potential as a valid method for the joint applications of identifying change for post-disaster damage assessment, and monitoring of urban change. Chapters 5-7 describe the development of the method and testing on aerial imagery. The conceptual model was developed as a method for integrating and addressing the following strategic objectives (Table 1.1):

- 1.** Processing techniques following natural disaster events are mainly chosen on an *ad hoc* basis. There are no set of standard procedures or guidelines for measuring post-disaster damage, and providing a standardised framework for damage

assessment using multi-temporal data would go some way to speeding up time-critical damage assessments.

AIM: provide a framework or guidelines for processing of RS images following significant change events.

2. There is a necessity to provide adequate change detection methods that are generalisable for the dual applications of urban monitoring and post-disaster damage detection. Currently, there is an abundance of optical data collected by existing remote sensors. There are also several national and international initiatives for data sharing in the aftermath of unexpected events. However, there are no common, object-based image processing methods for fast and efficient processing of these data.

AIM: Provide methods that identify objects of interest in imagery and provide pre-change information as a baseline dataset from which to detect change.

3. Images contain information on a large number of complex phenomena. Image features targeted by extraction techniques can be modelled allowing analysis of urban building stock, and retaining geometric knowledge of the target objects irrespective of occlusions or noise.

AIM: Provide a model-based platform for per-building detection from imagery, and subsequent analysis of these data.

4. There are time pressures when processing post-disaster data that is vital for prioritising response initiatives. Thus, there is a need to provide automated methods that reduce processing time and user-interaction to speed the process. Many methods require optimal parameterisation, which can be time consuming, and do not generalise well. However, the importance of precise parameterisation can be reduced using an iterative processing method.

AIM: Provide an automated or semi-automated method of image processing, designed in an iterative fashion to intelligently reduce the incidence of poor parameterisation of algorithms, reducing the need for time-consuming precise initial parameterisation.

5. National mapping agencies are investigating the use of remote sensing technologies to aid the updating of the National Geographic Database. Algorithms and models for processing imagery are therefore of great interest to Ordnance Survey in Great Britain, and change detection from multi-temporal imagery is one way of providing new vector data for comparison to their existing datasets

AIM: Provide a robust change detection method for capturing and mapping of regular urban changes.

6. Modelling objects captured from imagery allows quantitative statistical analysis of multi-temporal change, highlighting trends and facilitating object-based comparisons. Assigning meaning to the detected changes can allow the development of change descriptors that may be correlated with the cause of the change. Relationships between objects can be highlighted, and change can be classified using existing formalised damage or change classes.

AIM: Provide quantification and qualitative descriptors of per-building change captured from multi-temporal imagery.

1.5.2 Scope of this Research

The main crux of the research described in this thesis is to develop a model-based framework for the detection of object-based change, captured from a pair of remotely-sensed images of the same geographic area at two time periods. These images display a range of phenomena, or features, e.g. buildings, roads or vegetation. Each feature is projected on the image plane and can be delineated by the feature footprint (i.e. the areas of the image that are obscured by each building) which is extracted using semi-automated algorithms, thus identifying the features of interest. In the exemplars described, the image features that are extracted and stored as *objects* relate to, and only represent, buildings in the imagery. Once identified and extracted, objects are stored in a scene model. The conceptual model provides the modular framework for this approach, and is described in detail in chapter 3. Initial testing of the model has focused on the extraction of basic geometric features, both in two and three dimensions. These primitives have been amalgamated using logic

and evidence of similarity between primitive features. Qualitative and quantitative change descriptors have been developed to describe the change between two image dates.

1.6 THESIS FORMAT

The thesis is presented in a manner that describes the development and evolution of the conceptual model for change detection. The thesis is split into three main sections and sub-divided into nine chapters, organised as follows, and described in Figure 1.3 and Table 1.2:

1) Project Background (chapters 1-2)

A basic outline of the fundamental issues key to this research has been discussed in this chapter, and is extended in more detail in chapter 2, where an overview of current and previous research is presented. It is sub-divided into the main subject areas portrayed in the thesis and provides an introduction to the technique described later in the thesis, in parallel with a critique and discussion of previous related research.

2) Model Development and Application (chapters 3-7)

The full methodological approach and novel conceptual model framework is described in detail in chapter 3. The object detection phase of the conceptual model is applied in a real-world context to detect two-dimensional (chapter 5) and three-dimensional (chapter 6) objects, preceded by a detailed overview of the datasets used (chapter 4). Preliminary investigations into and early results of the object comparison module of the conceptual model are outlined in chapter 7.

Table 1.1. Rationale and Aims of the Thesis

	<i>Current situation</i>	<i>Requirements & Rationale</i>	<i>Aims</i>
1	<i>Ad-hoc</i> processing techniques exist for change detection of unexpected incidents	Need for generalisable methods and guidelines for post-disaster data processing	Provide a framework or guidelines for processing of RS images
2	Abundant data availability from remotely-sensed imagery	Need for suitable image processing techniques to identify and extract objects from fine spatial resolution imagery	Provide methods that identify objects of interest and provide pre-event information as a basis for change detection
3	Complex image scenes can constrain analysis efforts	Need to model the urban environment to exclude unimportant information and reduce complexity of study area	Provide a model-based platform for per-building analysis
4	Damage assessment is necessary in the response phase of the disaster cycle to prioritise post-disaster relief and response efforts	Need for semi-automated methods to detect extent of change Reduce the dependency on time-consuming supervised parameterisation of algorithms	(a) Provide semi-automation to speed up processing time and provide damage maps in a timely manner (b) Reduce the necessity for optimal parameterisation of the system through intelligent iteration of methods
5	National mapping agencies require continuous urban monitoring to provide up-to-date geodata and map products.	Ongoing requirement to improve the efficiency of data collection methods	Provide a robust change detection method for capturing normal urban morphology
6	Basic quantification of urban features as objects enables use of statistical measures for change description. Attributing descriptive meaning to geographic data allows semantic retrieval of geodata	Need to quantify and describe the amount of detected change, both for urban morphology and disaster applications	(a) Provide methods for quantification of per-building change (b) Provide qualitative descriptors for change in buildings using existing categories

3. Discussion & Future Advancement (chapters 8-9)

A discussion of the methods and results of the object detection phase is presented in chapter 8, with recommendations for further work. A summary of the thesis and conclusions of this research are found in chapter 9.

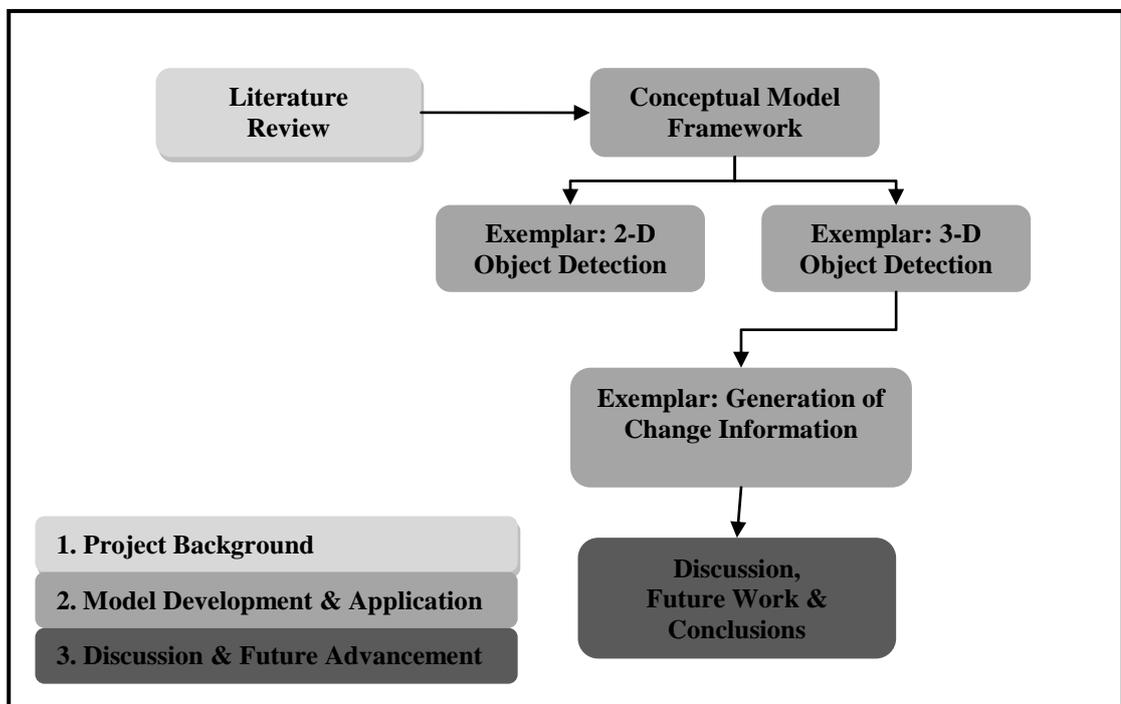


Figure 1.3 Flowchart of thesis sections

Table 1.2. Thesis outline and context of chapters

Project Aim: Develop an object-based model for change detection through an iterative, object detection and object comparison process.

	Chapter 1: Introduction	Chapter 2: Methods	Chapter 3: Conceptual Model	Chapter 4: Study Sites & Data
Aim of chapter	Introduce the project. Provide high-level background descriptions of key project concepts	Examine existing techniques and methods used in previous research, and provide a critique of these methods	Introduce the conceptual model built and tested in this research. Provide a detailed description of each component of the model	Describe the study areas and datasets used in the testing of the conceptual model
Chapter in context of overall project	Ch. 1 provides scientific context for the project and shows how the concept is grounded in tangible academic and real-world applications	Ch. 2 provides solid theoretical foundation for the development of a conceptual model.	Ch. 3 outlines the model framework, and provides context for future chapters that test individual model components	Ch. 4 provides suitable data for which to test the model
Methods used	Review the current political mechanisms for, and institutional applications of, change detection methods	Review existing scientific literature in fields of remote sensing, and machine vision, amongst others.	Amalgamate existing methods and new concepts into a novel conceptual model	N/A
Data used	Background Literature	Background Literature	N/A	Fine spatial resolution aerial and satellite imagery
Study site	N/A	N/A	N/A	Site 1. Simulated Imagery Site 2. Buncefield, UK Site 3. Saudi Arabia

Table 1.2 (continued) Thesis outline & context of chapters

Project Aim: Develop an object-based model for change detection through an iterative, object detection and object comparison process.

	Chapter 5: 2-D Object Detection	Chapter 6: 3-D Object Detection	Chapter 7: Object Change	Chapter 8: Discussion
Aim of chapter	Tests the object detection phase of the conceptual model on two diverse datasets	Extends the testing from Ch. 5 of the object detection phase and describes the generation of 3-dimensional scene models	Show advancement of object comparison phase of the model – towards detecting change	Discuss the research findings of previous chapters both in the development of a conceptual model and the described testing of the object detection phase
Chapter in context of overall project	Develops and tests the first half of the conceptual model – object detection	Extends the capability of the object detection phase of the model	Provides a preliminary example of how the object comparison phase can be utilised to detect and describe change	Reviews the overall project and offers insight into how the model can be advanced and applied holistically as a change detection method
Methods used	Hough Transform Forward and Inverse modelling	Hough Transform Forward and Inverse modelling Simulated Annealing	Summary statistics Spatial metrics	Discussion of research findings
Data used	Aerial imagery Satellite imagery	Aerial imagery	Aerial imagery Statistics Change descriptors	N/A
Study site	Site 2: Buncefield, UK Site 3: Saudi Arabia	Site 2: Buncefield, UK	Site 2: Buncefield, UK	N/A

CHAPTER 2

LITERATURE REVIEW

2 LITERATURE REVIEW

2.1 INTRODUCTION

Following the rationale outlined in chapter 1, this chapter gives a detailed insight into methods that have been used in the fields of change detection, damage assessment, and several branches of computer science and mathematics. Chapter 1 provided a brief description of how remotely-sensed data can be used to monitor changes in terrestrial phenomena from two standpoints; planned urban morphology and unexpected, post-disaster damage assessment. The literature reviewed in this chapter is discussed in relation to both applications. In addition, the chapter is subdivided into sections relating to object detection, modelling, and change description, providing background information for much of the rationale articulated in section 1.5.

2.2 CHANGE DETECTION

2.2.1 Change detection methods

There has been a long history of detecting and monitoring change in global phenomena through remotely-sensed observations, and there is a consequent abundance of image comparison techniques. The concept of “change detection” when applied to the field of remote sensing has a variety of meanings, dependent on what the specific algorithms or techniques are designed to measure. Coppin *et al.* (2004), in a review of change detection methods for ecological applications, defined digital change detection as quantifying temporal phenomena from multi-date imagery commonly acquired from satellite sensors. Despite the large number of applications, change detection is based on the same elemental assumption: remote

sensing technology detects a certain portion of the electro-magnetic spectrum, so changes in land cover must result in radiance change that is greater than radiance changes due to other factors (Ingram *et al.*, 1981). These factors include atmospheric conditions, Sun angle or soil moisture (Singh, 1989; Deer, 1999). Change detection is reliant on the assumption that the phenomena of interest must be visible in the imagery, and changes in the phenomena are identifiable by spectral or textural differences between multi-temporal images (Sui *et al.*, 2008).

Change detection for the analysis of land cover has been achieved for numerous ecological, environmental or sociological applications. A landmark review of the most commonly used methods for detecting change was presented by Singh (1989), providing two major classes of techniques – classification comparison, and direct comparison. A breakdown of 10 of the most common application areas is described by Lu *et al.* (2004), with additional reviews of change detection methods extending these two categories (Deer, 1999; Lu *et al.*, 2004). A comprehensive description of change detection stratified change detection strategies into 7 classes – direct comparison, classification, object-oriented method, model method, time-series analysis, visual analysis, and hybrid methods (Sui *et al.*, 2008). The remainder of this section provides a brief description of the methods that fall into these strata.

a) Direct comparison

Change detection by direct comparison can include several techniques reliant on the relationship between digital numbers (DN) of two images of the same geographic area. The most commonly used examples of this method are image differencing and image ratioing. Models have also been produced that relate digital numbers to describe physical properties, such as reflectance models - as reviewed by Lu *et al.* (2004). Direct comparison methods have been widely used to describe land-cover changes (Singh, 1989; Muchoney and Haack, 1994; Sohl, 1999), and for generating post-disaster damage information (Zhang *et al.*, 2003; Gusella *et al.*, 2005a) however the setting of thresholds above which “change” has occurred is widely regarded as the most important parameter in these types of techniques (Sohl, 1999; Rogerson, 2002; Sui *et al.*, 2008). Sui *et al.* (2008) provides a succinct explanation of the main techniques for setting thresholds: trial-and-error, statistical estimation,

and alternative, more complex techniques, not often practically used, such as Bayes' theory, Markov models or artificial neural networks (ANN).

Alternative methods exist for describing bi-temporal change (i.e. change between two image dates). Change vector analysis (CVA) is an example of one such technique, as reviewed by Singh (1989), and Johnson and Kasischke (1998). They described how CVA estimates the direction and magnitude of change between images, based on the identification of a radiometric threshold, above which change was deemed to have occurred. Lambin and Strahler (1994), described CVA as the multi-temporal observation of the coordinates of any pixel in multi-dimensional space. Change vectors describe the movement of each pixel from one time period to the next in this space. In addition, Seto *et al.* (2002) used the CVA technique to detect land-use change in China from bi-temporal Landsat TM data, achieving 93.5% classification accuracy when compared to field measurements.

b) Classification techniques

Many of the techniques reviewed by Singh (1989) involved the classification of two or more images of the same geographic area and the comparison of these classified images, resulting in a matrix of change information. A major advantage of the classification method is the reduced reliance on calibrated imagery, as the classification is conducted on separate images, and does not require raw DN's (Sui *et al.*, 2008). Classification techniques have been shown to perform better than more sophisticated methods (Singh, 1989), yet they are sensitive to image misregistration, and require images with similar viewing angles. Issues also arise during the image differencing technique because of variance in spectral responses within the same land cover class. Other problems can arise from these methods including the perpetuation and multiplication of classification errors in both images (Seto *et al.*, 2002; Coppin *et al.*, 2004). Singh (1989) gave examples of this method that resulted in 67% change detection accuracy with the use of the Landsat Multispectral Scanning System (MSS) sensor. It has also been applied to finer spatial resolution sensors, such as a fusion of the 10 metre panchromatic imagery collected by the SPOT HRV sensor with multi-spectral imagery from Landsat TM to achieve a greater change detection accuracy of 86% for buildings with widths of 10-20 m (Zhang, 2001). There are numerous examples of how multi-temporal classifications

of Landsat TM imagery can be used for land cover change to urban areas (Song *et al.*, 2001), and forest monitoring over a nearly-20 year period (Hall *et al.*, 1991).

c) Object oriented methods

The third method for review is an object-to-object comparison, where individual objects are extracted from multiple images. The comparison is between the objects only, therefore, much of the background information from the image is excluded from the analysis. As was seen in (b), this method is also insensitive to precise radiometric calibration between multi-temporal images, which is a positive aspect of the method. However, precise delineation and extraction of the objects in the first place is more of a challenge, and often additional datasets (such as GIS data) are required to target the analysis (Lo and Shipman, 1990; Agouris *et al.*, 2000). Object-based methods have been used in post-disaster response applications for identifying building damage (Gusella *et al.*, 2005a) and has been compared to pixel based methods in the study described by Bitelli *et al.* (2004).

d) Model methods

Sui *et al.* (2008) describe the model-based methods of change detection as “application oriented” change detection methods. They describe these methods in terms of approach-based and processing-based models, and recognise the main limitation is the initial conceptualisation and set-up of the model. Examples of model-based approaches include Herold *et al.* (2002), who used fractal dimension and other landscape metrics for detecting changes in urban areas; Gopal and Woodcock (1996), and Hese and Schmillius (2003), who both incorporated an ANN in their method of change detection to forested areas.

e) Time series analysis

Techniques that fall into this category of change detection specifically focus on the detection of change using multi-temporal imagery, rather than solely a bi-temporal analysis. These often require the use of imagery from coarse spatial resolution sensors such as the Moderate Resolution Imaging Spectroradiometer (MODIS; Zhan *et al.*, 2000), or the Advanced Very High Resolution Radiometer (AVHRR) onboard the National Oceanic and Atmospheric Administration’s (NOAA) NOAA-9 and

NOAA-10 platforms (Lambin and Strahler, 1994). The latter technique used change vectors to describe annual land-cover change processes in West Africa using the normalised difference vegetation index (NDVI; Lambin and Strahler, 1994). However, there are many examples that use finer spatial resolution sensors such as Landsat TM (Hall *et al.*, 1991; Song *et al.*, 2001), Landsat MSS (Lyon *et al.*, 1998), and even IKONOS and Quickbird imagery (Saito *et al.*, 2009). In the latter study, sub-metre satellite imagery was used as a data source for automated and semi-automated techniques for information extraction to monitor long-term post-disaster recovery in Pakistan and Thailand. (Sui *et al.*, 2008) also introduce additional methods of change detection using video as a data source.

f) Visual methods

Visual methods of change detection employ the human understanding of the operator visually identify changes between two images. These methods can benefit from the users' knowledge of an area, and allow the use of contextual information to identify features of interest. Methods that aid this analysis include layering of image bands to visually interpret changes over time (Lu *et al.*, 2004). Visual analysis is still irreplaceable though (Sui *et al.* 2008), and this holds true in the field of post-disaster response and damage assessment using one image from before and one after the disaster (Belward *et al.*, 2007; Wikantika *et al.*, 2007). Using such an approach, Yamazaki *et al.* (2005), and Miura *et al.* (2006) digitised individual buildings manually from imagery captured before and after the 2003 Bam, Iran and the 2006 Yogyakarta, Indonesia earthquakes respectively, to classify buildings according to well established, formalised damage classes. This was an extremely time-consuming exercise, although the authors, in both cases, were confident of a high accuracy in the resulting damage maps.

Accuracy was one of the reasons that a novel method of visual interpretation of post-earthquake damage was employed in the aftermath of the January 2010 Haiti earthquake (Bevington *et al.*, 2010). Over 600 engineers and image analysts from around the globe volunteered to analyse fine spatial resolution aerial imagery as part of the Global Earth Observation Catastrophe Assessment Network (GEO-CAN). In this event, analysts manually identified collapsed and heavily damaged buildings from a comparison of pre- and post-event imagery. This 'crowdsourcing' initiative

combined the results from volunteers in 23 different countries and detected nearly 30,000 damaged or collapsed structures within a month of the earthquake for use in the building damage post-disaster needs assessment (PDNA) involving the UN, EC, World Bank and other international donors.

Adams *et al.* (2004), in their visual study of damage after the 2003 Boumerdes, Algeria and Bam earthquakes, contend that devising universally applicable, automated techniques is difficult due to the varying spectral signatures of damaged buildings worldwide. This problem was exacerbated in the aftermath of the Bam earthquake because of the use of a widely-available local sandstone as a construction material (Huyck *et al.*, 2004b). Here, a measurement of the relative damage (compared with the original building state) was produced and validated against damage maps created by the United States Agency for International Development (USAID). However, Huyck *et al.* (2004b) recognised that changes needed to be made to make the method more robust, including connecting the degree of change to formal damage classes before the method could be effective in other future earthquake locations. This initiative is investigated in this thesis, as demonstrated in objective number 6 (chapter 1).

g) Hybrid methods

Hybrid change detection relates to methods that amalgamate two or more of the techniques from sections a - f. One of the benefits of adopting a hybrid method is to utilise the performance of certain algorithms at different sectors of the change detection process. Examples of hybrid methods include those that integrate GIS data with remote sensing imagery, of which there are many such examples (Lo and Shipman, 1990; Olsen *et al.*, 2002). A study monitoring post-disaster recovery (Saito *et al.*, 2009), utilised GIS and satellite imagery to produce comprehensive datasets showing the reconstruction of areas damaged by tsunami and earthquake. Other techniques have utilised spatio-temporal geographic gazetteer information to prioritise target areas for change detection from imagery (Agouris *et al.*, 2000). An alternative hybrid method generated three-dimensional object-based information from stereo pairs of aerial and IKONOS imagery before a comparison of the 3-D objects was completed (Heller *et al.*, 2001).

A large number of change detection methods for remote sensing data have been described in this section. Providing a robust change detection method for the dual applications of urban monitoring and post-disaster damage detection are key to the main objects of this thesis, as described in section 1.5 (Objectives **1, 2, 5 & 6**).

2.2.2 Application of Change Detection: (i) Monitoring of Urban Areas

Change detection using data captured from remote sensors has many tangible, real-world applications. For example, Ordnance Survey, and other national mapping agencies are actively seeking to develop cost-effective and efficient solutions for updating their geographic information stock held in the National Geographic Database. The provision of information describing populated areas is invaluable for the description of scenes before a change event, and forms a basis from which to measure *change* after the event (objective number **5**, chapter 1). This is a responsibility of mapping agencies, and is highlighted in the main remit of Ordnance Survey:

“to be the content provider of choice for location-based information in the new information economy”. (Ordnance Survey, 2007)

Ordnance Survey is constantly seeking to improve its data capture techniques to enhance the accuracy of height, vector and imagery primitives (such as spectral, contextual or shape characteristics; Parker and Stileman, 2005). Terrestrial features (especially buildings) are updated when significant changes are detected, with their flagship vector dataset - OS MasterMap[®] - holding over 450 million geographic vector features (Ordnance Survey, 2009b). Eventually it is anticipated that a fourth dimension of spatio-temporal information will exist, vastly improving data provision for dynamic events such as disasters (Parker and Stileman, 2005), where Ordnance Survey are the key information provider through their *Mapping for Emergencies* unit. Developing change detection methodologies that include quantitative and qualitative descriptions of urban areas therefore corresponds directly to the Ordnance Survey research strategy. For example, Holland and Tompkinson (2003),

provided a human operator with *cues* of change detected probabilistically from an object-to-object change detection. These cues were formed from changes in building shapes or from the introduction of new shadows in the second image, which served as a proxy for potential new buildings. Additional applications of change detection methods for use in updating official map products are evident in the literature.

Bailloeuil *et al.* (2003) proposed two methods for updating urban map products using fine spatial resolution imagery combined with ancillary data from a geographic information system (GIS). Motivated by commonly experienced problems with fine spatial resolution imagery of urban areas (shadows and building occlusions), they used urban information from the GIS datasets to implement a deformable active contour which delineates objects in multi-temporal images. The segmented objects from both images were directly compared, informing the authors of areas of change. Limitations arose in the research due to different geographic extents of building footprints and constraints on the accuracy of the GIS-to-image registration.

Another national mapping agency, the Swiss Federal Office of Topography, aimed to generate a 1 m spatial resolution national topographic landscape model (TLM), showing the urban landscape as three-dimensional objects in a GIS (Käser, 2001), through the ATOMI project (Automated reconstruction of Topographic Objects using aerial images by means of vectorised Map Information). The author acknowledged that one of the major roadblocks in this implementation, which impacts the profitability of such a method, is the lack of suitable automated methods for feature extraction and update.

Remote sensing technologies have also been utilised beyond the traditional urban realm for mapping of informal settlements (Mason and Baltsavias, 1997; Ruther *et al.*, 2002; Mayunga *et al.*, 2007). Also known as *shanty towns* or *squatter settlements*, these informal areas can appear very rapidly, often after the onset of unexpected events, such as civil disturbance, and natural disasters, and have been defined as

“dense settlements comprising communities housed in self-constructed shelters under conditions of informal or traditional land tenure”

(Hindson and McCarthy, 1994; Mason and Baltsavias, 1997)

Generating information about informal settlements is imperative for the multi-temporal, spatial description of how these settlements evolve. Typically, they are found in developing countries where there the mapping infrastructure available is poor and unable to provide the necessary formalised information about these areas. Hence, basic understanding of these areas, such as population, is missing. Developing an inventory of buildings has been used in several approaches in the past as a proxy for these data. For example, Souza *et al.* (2002) inferred population statistics of urban areas directly from a combination of pansharpened IKONOS imagery and census information for five census neighbourhoods in a Brazilian city. Population statistics were inferred by multiplying the number of buildings within each census area by the average per-household occupancy ratio (extracted from the census information). The results were promising; with 3% mean error for the pansharpened imagery. However, the authors noted that overestimation of population numbers was evident due to unoccupied housing, and variations in building materials and occupancy size (Souza *et al.*, 2002). It is apparent from the literature described in this section that remote sensing is a flexible and appropriate tool for estimating intra-census populations for informal settlements, or cities with burgeoning populations. To be effective in this regard, suitable methods should exist to monitor change over time. The following section describes several change detection methods for remotely-sensed imagery.

2.2.3 Application of Change Detection: (ii) Damage Detection

Many techniques adopted for change detection can be applied to the more time-critical scenarios encountered when detecting the extent of damage arising from a natural disaster. However, some existing change detection methods such as correlational analysis, have limited usage for damage detection after natural disasters due to structural damage and illumination effects (Shinozuka and Rejaie, 2000). Many methods utilise image differencing and classification techniques to detect damage after several types of natural hazard (Tralli, 2001; Tralli *et al.*, 2005).

Although urban areas in the developed world have archives of images spanning much of the last decade (since the introduction of commercial sensors), it

is common that pre-event imagery is not available when gauging post-event damage in less developed nations. One damage assessment method that aimed to address this problem used a manual, visual assessment of post-earthquake damage (2001 Bhuj, India earthquake), but was limited in its ability to quantify object change beyond visual interpretation of the post-event imagery (Saito and Spence, 2004). However, their technique can be used to prioritise damaged locations needing further analysis. In contrast, remotely-sensed pre- and post-event data from optical and non-optical sources were used effectively to provide a three-dimensional representation for ongoing monitoring and change assessment of debris volume at the site of the World Trade Center after the 11 September 2001 terrorist attacks (Huyck and Adams, 2002). Such contrasting incidents emphasise the wide spectrum of attainability of remotely-sensed data for regions susceptible to both natural and human-induced disasters.

Various methods of change detection have involved the fusion of optical and non-optical remotely-sensed data. SPOT HRVIR optical imagery and synthetic aperture radar (SAR) data were combined to estimate damage in Eguchi *et al.*'s (2003) study of the 1999 Marmara, Turkey earthquake. The short revisit time, combined with the ability of the sensor to penetrate occlusions make SAR data valuable when an affected area is not illuminated, under cloud or smoke cover. However, scene-wide variations in the radar and spectral returns from the two sensors suggested large radiometric changes between the two time periods. SAR data were also used with some success to quantify damage after the 1995 Kobe, Japan earthquake (Aoki *et al.*, 1998). However, the significant randomness in backscatter led to misclassification, which limits the operational use of SAR for damage detection in urban areas. Finally, the measurement of night-time lighting levels using Defense Meteorological Satellite Program / Operational Linescan System (DMSP/OLS) imagery in urban areas as a proxy for damage detection was studied by Hayashi *et al.* (2000). This work relied on the assumption that much of the infrastructure of the studied cities had been affected by earthquakes to the extent that electricity was limited and road usage had been reduced. In addition, the level of detail that could be achieved was constrained by the coarse spatial resolution of the sensor (2.7 km resolution). The use of DMSP/OLS data is therefore limited to coarse resolution detection of damage in large cities, and may be only useful for focussing

finer spatial resolution damage assessment efforts on areas which appear changed in the night-time imagery.

The use of optical remotely-sensed imagery for post-disaster damage detection is not without problems. For example, methods rely on the identification of a variety of parameters, including angle of incidence (Chiroiu and Andre, 2001; Liu *et al.*, 2004; Gusella *et al.*, 2005b), and spatial and temporal resolutions (Chiroiu and Andre, 2001; Liu *et al.*, 2004; Rejaie and Shinozuka, 2004) which can be difficult. These problems can be exacerbated by the focus on per-pixel change measures, whereas object-oriented damage detection methods have been shown to reduce the effects of some of these limitations: variability arising from occlusions and sensor angle in multi-temporal and multi-sensor scenarios has been reduced using the technique described by Gusella *et al.* (2005b), whereby statistical classification of buildings as objects was used to estimate the number of collapsed buildings following the Bam earthquake, with preliminary accuracy of 70.5%. The detection of buildings as objects in imagery is part of the rationale of this thesis (Objective 2). The following section focuses on techniques that facilitate the individual extraction of buildings as objects from imagery.

2.3 OBJECT DETECTION AND FEATURE EXTRACTION

2.3.1 Background

Section 2.2 focused on two important applications of change detection: monitoring urban change, and damage detection. These applications rely on the ability to determine the state of objects at two points in time – before and after a change event - and require multi-temporal imagery of the same geographic location. National mapping agencies (such as Ordnance Survey) maintain building inventories in the form of the NGD, or other digital repositories. It is often the case, however, that such information is outdated, or is incomplete, so RS data provide a supplementary resource that can provide information about the study site before and after change events.

RS data has evolved substantially since the emergence of satellite based sensors (as described in chapter 1). However, complementary analytical techniques to interpret these data have not always been developed accordingly (Barnsley *et al.*, 1993). This is alluded to in objective number **2** of Table 1.1, and provides a key motivation for the research in this thesis. There are numerous feature detection and extraction methods for applications for computer science, machine vision and even medical imaging that have potential application in the field of remote sensing. The form of these methods depends on the geometric or spectral nature of the required objects, which range from low-level line and edge detectors to shape matching and shape deforming methods for features with complex geometries. This section outlines several of these feature extraction methods and describes their application to remotely-sensed imagery.

2.3.2 Edge Detection Methods

Edge detection is one of the most elementary methods of feature extraction. Described as a *low-level* feature extraction method (i.e. features can be extracted automatically from imagery without shape or spatial relationship information; Nixon and Aguado, 2002), the basic premise of edge detection is the identification of a series of interconnected pixels that lie on the boundary between two segments of an image (Gonzalez *et al.*, 2004). Edge detection works by looking for steep gradients in brightness between adjacent pixels using a differentiation process. The resulting edge is described by a vector showing the magnitude and direction of the detected gradient (Sonka *et al.*, 1999). The most common types of edge detectors are reviewed by many sources (Sonka *et al.*, 1999; Nixon and Aguado, 2002; Gonzalez *et al.*, 2004), with the most widely-used algorithms including the Roberts, Sobel, Prewitt and Canny operators.

In a review of edge detector methods, Petrou and Kittler (1988) outline the constituents of an “optimal” edge operator: good signal-to-noise ratio, good locality and maximum suppression of false responses. For example, Roberts’ method (Roberts, 1965), as one of the first edge detectors, was disadvantaged by being over-sensitive to noise. The Canny operator is one of the most widely used edge detection

algorithms, and represents an improvement on Roberts' method because it is insensitive to noise (Sonka *et al.*, 1999).

The Canny method of edge detection uses a Gaussian convolution filter to reduce noise in the input image. A second order derivative method is then applied to the smoothed image to produce edge information in the form of gradient magnitude images. This combination of a symmetric two-dimensional Gaussian filter and subsequent differentiation perpendicular to the edge direction, forms a simple and effective operator (Sonka *et al.*, 1999). The Canny method then uses two thresholds to identify the edges in the edge strength images. Firstly, an upper threshold is applied to detect significant edges, before a second, lower threshold is applied to include weaker edges that are connected to pixels with strong gradients.

2.3.3 Shape-matching Methods

Several feature extraction methods employ algorithms that search for specific geometric shapes in imagery. One such method is template matching, whereby features in an image are matched to a pre-defined shape outline. The template shape is centred on points throughout the image and the number of points along the shape edge that match features in the image is measured (Nixon and Aguado, 2002). Scores are accumulated within a parameter space so the centre point with the highest score has the highest probability of matching the target shape. The technique is one of the most fundamental methods of *high-level* feature extraction, which differs from low-level detectors through the introduction of shape intricacies.

Shape matching methods can be generalised to accommodate variable target shapes and orientations, but this involves the manual creation of all expected sizes and rotations, and is not an automated method. Nixon and Aguado (2002) describe several drawbacks of the method, including discretisation issues, where gaps can appear at the edges of the template shapes due to the choice of coordinate system of the image. Direct extraction of features is often preferred because of this (Nixon and Aguado, 2002).

2.3.4 The Hough Transform

A common method of feature extraction from imagery through shape matching is the Hough Transform (HT); a well-established method that utilises evidence gathering in a parameter space for the extraction of lines, circles, ellipses and arbitrary shapes (Hough, 1962; Duda and Hart, 1972). The standard HT is applied in three steps (Kälviäinen *et al.*, 1995): (i) calculating and accumulating parameter values; (ii) finding local maxima in the parameter space, which correspond to object segments in the image space; and (iii) extracting the objects based on the local maxima information. A strength of the HT, when compared to template matching, is that it is faster and more efficient at detecting objects (Nixon and Aguado, 2002).

The HT has been used to detect a variety of primitive geometric shapes such as lines and curves (Duda and Hart, 1972; Princen *et al.*, 1989; Kesidis and Papamarkos, 1999), circles (Chan and Siu, 1991; Guil and Zapata, 1997; Atiquzzaman, 2005), and even to detect 3-dimensional shapes (Vosselman *et al.*, 2004; Rabbani and van den Heuvel, 2005; Khoshelham, 2007).

Hough Transform for circle detection

The HT algorithm comprises of several key processes, shown in Figure 2.1. Firstly, an edge detection algorithm is applied to an input image (of dimension, x_{max} , y_{max}), resulting in an intermediate image showing the graded strength of edges in both the x and y directions across the image. A suitable threshold, M^t , must be identified as an input to this process based on acceptable edge strengths for the particular implementation of the HT. Secondly, each pixel from this edge image is processed in turn by the HT. When the edge strength is greater than M^t (e.g. a dark pixel adjacent to a bright one, in either the x , or y direction), the pixel is mapped into a parameter space, or *accumulator* space, of equal dimensions to the image space (a_{max} , b_{max}). This pixel, is defined as a circle, with radius r , and centred at (a, b) , thus:

$$(x - a)^2 + (y - b)^2 = r^2$$

(2-1)

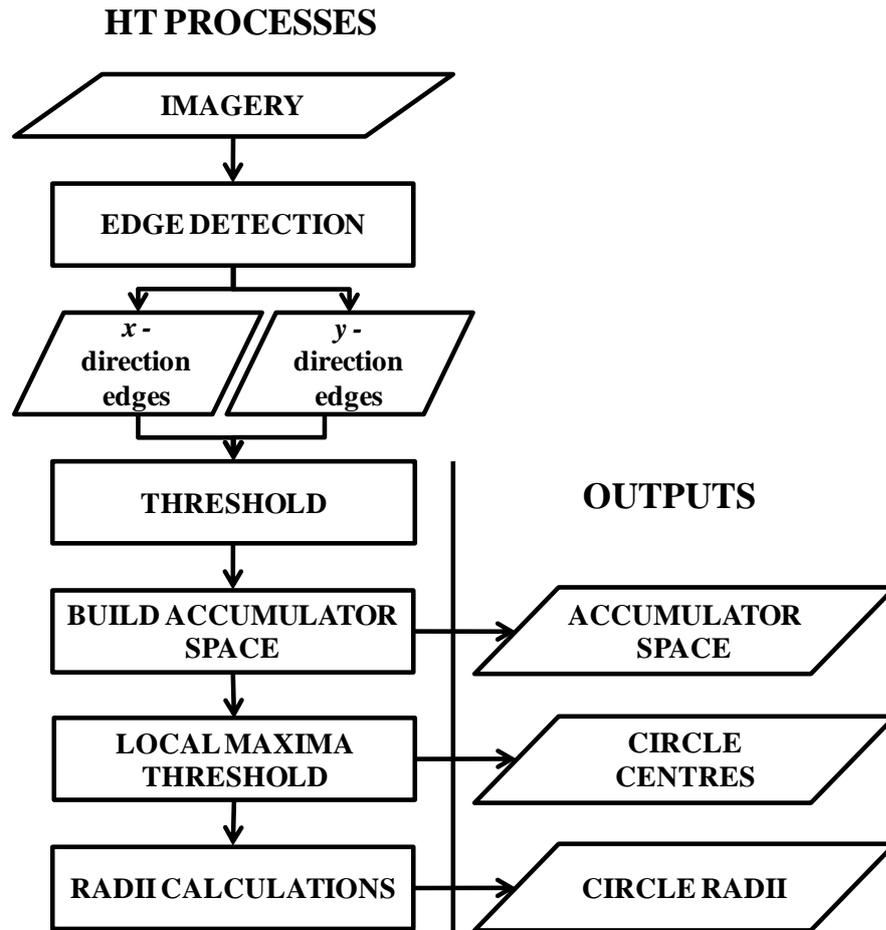


Figure 2.1 Workflow of a circular Hough Transform algorithm

The HT algorithm is designed to detect circles within a range of radii, and so the value of r is variable. Because of this, the parameter space can be thought of as three-dimensional with each pixel in the image space mapping to equivalent cells in the parameter space (Nixon and Aguado, 2002; Figure 2.2). This parameter space acts as a repository for evidence gathered for circle detection, collected incrementally for every possible value of r , based on the x, y location of the pixel in the edge image Figure 2.2c.

Thirdly, for each pixel in the image space, evidence in the form of votes are stored in the parameter space where each cell accumulates votes for all possible circles passing through that pixel in the image space. As a result, peaks form in the parameter space for cells with the greatest number of votes. These peaks correspond to probabilistic evidence for the locations of circle centres (Sonka *et al.*, 1999; Figure 2.3). This is described in stages:

- The figure shows a dark circle on a bright background (a).

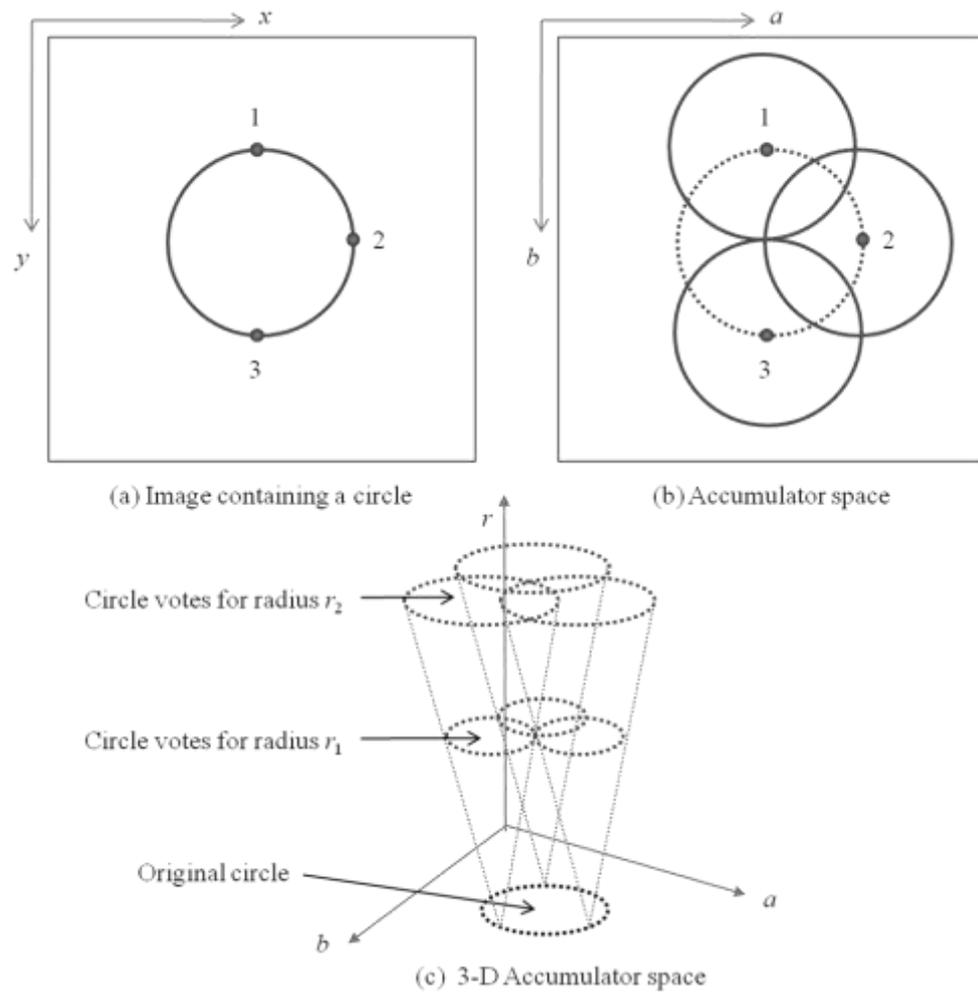


Figure 2.2. Stages of the circular Hough Transform algorithm.
(from Nixon and Aguado, 2002),

- The HT defines a potential location of a circle centred on each dark pixel (b), repeated for all dark pixels.
- The frequency of dark pixels within each candidate circle is assigned a parameter space, resulting in a greater frequency for true centre circles, marked ● in (c & d).
- The correct centre circles can also be detected when circles are incomplete, or are occluded (d). Due to the process of vote accumulation, the parameter space is often known as the *accumulator space*, which will be the preferred term for the remainder of this thesis.

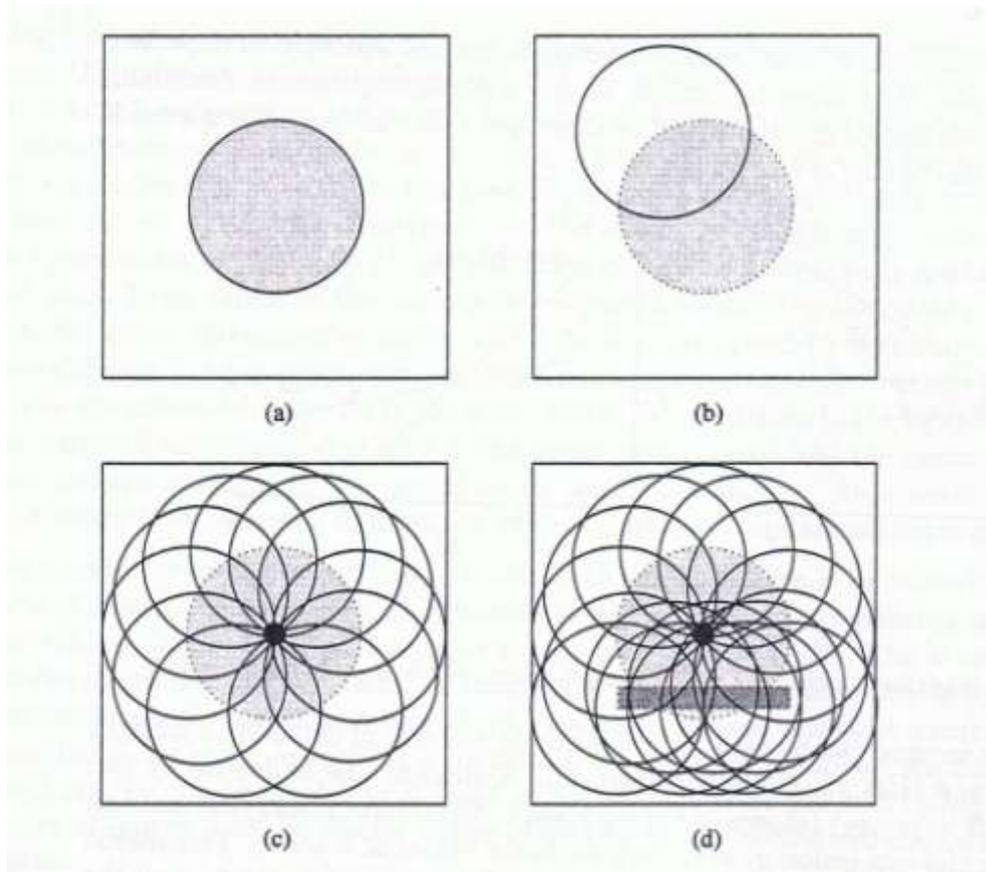


Figure 2.3 Generation of an Accumulator Space for a HT circle detector (from Sonka *et al.*, 1999)

(a) A dark circle on a light background. (b) A circle with radius r is centred on every dark pixel. (c) The frequency of pixels occurring in the centre-circle loci is determined – the most frequent represents the centre of the dark circle (marked by ●). (d) When occlusions occur, the HT can still directly detect the centre of the dark circle despite incomplete circle information.

Princen *et al.* (1989) described the accumulator space as

“a discrete representation of the continuous multi-dimensional space which spans the possible parameter values that describes the shape”.

Princen *et al.* (1989)

The accumulator space, whilst visualised as a three-dimensional phenomenon in Figure 2.2, is generated in two dimensions (a, b ; Figure 2.2b), where the radius is parameterised before the algorithm is run as a range of values. For each edge pixel, votes are generated for every radius where,

$$r_{min} < r < r_{max}.$$

(2-2)

A location (x, y) in the image space is mapped to a locus of points (a, b) in the accumulator space (i.e. from Fig 2.2a to Fig2.2b) according to

$$a = x - r \cos\theta \quad (2-3)$$

$$b = y - r \sin\theta \quad (2-4)$$

θ describes the trace of the perimeter of a circle (Nixon and Aguado, 2002), and forms a circle of votes when all values of θ (from 0° to 360°) are considered. The formulae are repeated for each possible value of r , consistent with (2-2).

Once the voting process is completed, a local maxima filter is applied to the accumulator space to identify peaks, which correspond to the circles with high probability of existing in the input image. Parameters that control the local maxima filter, L^t , are provided by the user when the HT is initialised. The peaks identified are used to estimate radius values, r , for each candidate circle (centred at the image x, y coordinates corresponding to the cells containing the maxima peaks in the accumulator space). Radii are calculated for all peaks in the accumulator space above the L^t threshold.

Previous applications of the circular Hough Transform

The circular HT has been used as a feature extraction tool for a wide variety of applications, from quality control in biscuit manufacturing (Davies, 1984), industry (Princen *et al.*, 1989), to the detection of lunar craters (Sawabe *et al.*, 2006). In Earth observation applications, the HT has been used to detect circular geologic structures from Landsat MSS (Cross, 1988) and SPOT HRV imagery (Taud and Parrot, 1992). Other implementations of the HT have aimed to maximise the computational efficiency of the algorithm (Chan and Siu, 1991; Guil and Zapata, 1997).

A fusion of SAR data and SPOT HRV optical images provided the basis for the implementation of a circular HT for the detection of oil drums (Wang *et al.*, 2004). SAR data were utilised to narrow the search area so drums appeared bright in comparison to the background. The HT was then applied to the optical images in the localised areas of interest. The method was found to be insensitive to the common problems of object shape distortion, shadows and occlusions found in optical

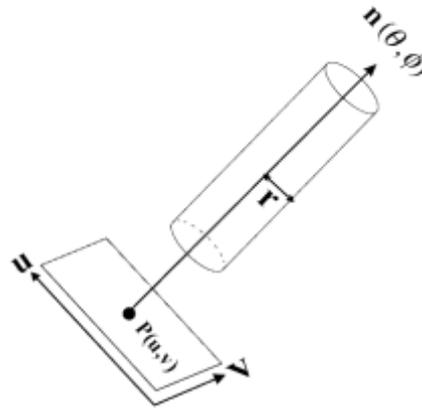


Figure 2.4 The five-parameter dimensions of a cylinder
(Rabbani and van den Heuvel, 2005)

imagery. The use of non-optical additional data in Wang *et al.*'s technique would be redundant in the technique described in this thesis as the areas of interest are known previously due to the nature of the change events where algorithms can be localised quickly without the added computational constraints of using ancillary data.

In other remote sensing applications, the Hough Transform technique was used to detect seepages from hydrocarbon reservoirs as measured from the spectral properties of aerial photographs (van der Werff, 2006). Liu *et al.* (2005) demonstrated a method of building extraction from pan-sharpened Quickbird imagery using geometric primitives identified by a HT method. In their study, a Progressive Probabilistic Hough Transform (PPHT) was used to identify linear features representing roof edges. Constraints were added to possible object shapes, thus, creating a model of building roofs, under the assumption that roof shapes were rectilinear, with four internal 90-degree angles. The dominant lines identified by the PPHT were amalgamated into objects representing building outlines, through colour and texture similarity measures. The authors were confident their experimental results were a success, despite the limited application of the technique.

Hough Transform for 3-dimensional object detection

In addition to the applications outlined above, for detecting two-dimensional (i.e. planar) objects, other applications have considered using HT methods designed to detect three-dimensional objects. These approaches are important for industrial applications, and in the field of reverse engineering; Vosselman *et al.* (2004),

Rabbani and van den Heuvel (2005), and Khoshelham (2007) describe methods for detecting three-dimensional objects from point data clouds and laser range data.

Rabbani and van den Heuvel (2005) detected three-dimensional cylinders from laser scanned point clouds for reverse engineering applications where cylindrical shapes are abundant. They opted not to increase the complexity of the HT algorithm by trying to detect whole cylinders in a single approach (which would require a five-parameter model; Figure 2.4), but rather split the cylinder shape into a set of two-dimensional subsets to calculate the orientation and radius of a cylinder. The method reduced the inherent dimensionality problem of the HT by splitting objects into geometric primitives, and was successfully applied to data from industrial sites. Vosselman *et al.* (2004) also showed an effective method for identifying planes, cylinders and spheres through the detection of geometric primitives, such as points, lines and circles. Hence, this method reduced the complexity of the generated parameter space. In this instance, the potential applications of their approach include the detection of objects for industrial applications, generation of city landscapes and the creation of digital elevation models (DEMs). This work highlighted the ability of a sequential methodological process to reduce computational and time constraints, and to make automatic detection of three-dimensional objects both feasible and manageable.

Extensions to the Hough Transform

There have been several previous attempts to extend the HT algorithm, both to improve the method, and also for specific applications. One method that improved the HT method described an inverse HT approach, *backmapping*, that extends the HT so that individual peaks in the accumulator space can be traced back to specific edge points, allowing modification of the accumulator space by removal of edge points (Gerig and Klein, 1986; Gerig, 1987). As each point in their image space adopts the associated parameters from the equivalent point in the accumulator space, one potential limitation of Gerig and Klein's approach would be that a given point should only ever be on one curve in the image space.

Cross (1988) used a circular HT to detect geologic features and found that simple thresholding, to segment the peaks in the accumulator space, was an ineffective method of extracting objects, due to circle irregularity and the presence

of peaks formed by chance edge pixel alignment. Using Gerig's backmapping method (Gerig, 1987), and under the assumption that true positive peaks would always be stronger than false positive peaks, the cells with the most votes in the initial accumulator array were extracted, whereas other accumulator cell values were reset to zero. This was an effective method of eliminating spurious peaks, and for reducing the number of false-positive detections.

The Dynamic Randomised Hough Transform (DRHT), an extension of a linear HT (Kälviäinen *et al.*, 1995), uses a pre-set maximum accumulator threshold to identify one cell in the accumulator space. A line associated with this cell was created in a new image, and a second iteration of the HT was applied to this image. For the second iteration, the parameter space was reset, and accumulated for the line at a finer resolution of the accumulator cells and at a higher maximum accumulator threshold. The addition of a second iteration improved the precision of line detection. In comparisons with traditional HT methods, there was an increase in precision, but at the cost of storage space and computational time (Kälviäinen *et al.*, 1995).

Hough Transform for line detection

Four common HT methods for line detection are: the Standard (SHT), Adaptive (AHT), Fast (FHT), and Hierarchical Hough Transforms (HHT) (Princen *et al.*, 1989). When applied to detect common industrial shapes, the AHT and FHT are computationally more efficient than the SHT. However, they identify only the most prominent peaks in the accumulator space, due to noise associated with *phantom peaks* (Princen *et al.*, 1989). The HHT is similar to the SHT, but offers computational advantages. The method described by Kesidis and Papamarkos (1999) used the sinusoidal curves generated in the parameter space of a linear HT algorithm. They described an inverse approach that recreated the original image by mapping individual peaks in the accumulator space to their original points as edges in the image space.

2.3.5 Active Contours

Active contours (also known as snakes) were first described by Kass *et al.* (1988) as a method of delineation of image features by evolving through growing or shrinking a polygon (a set of interconnected points) from an initial starting point to completely enclose the external or internal boundary of a target feature. The snakes approach can be thought of as an energy minimisation process where the target feature is a minimum of a suitably formulated energy functional, (E_{snake}) where,

$$E_{snake} = \int_{s=0}^1 E_{int}(\mathbf{v}(s)) + E_{image}(\mathbf{v}(s)) + E_{con}(\mathbf{v}(s)) ds \quad (2-5)$$

(Nixon and Aguado, 2002) when E_{int} represents the internal energy which controls the natural behaviour of the snake ($\mathbf{v}(s)$) and attributed points, E_{image} represents the image energy that attracts the snake to low-level image features such as edge points, and E_{con} represents the constraint energy, allowing additional information to control the snake's evolution (Nixon and Aguado, 2002).

The active contour approach is a well tried and tested technique, one of many deformable methods that have been applied for image segmentation in applications such as RS (Neuenschwander *et al.*, 1997), media post-production (Pitié *et al.*, 2004), medical imaging (Bamford and Lovell, 1998; Hamarneh *et al.*, 2000; Boscolo *et al.*, 2002) and head boundary segmentation (Neuenschwander *et al.*, 1997; Gunn and Nixon, 1998).

Several limitations of the snakes approach have been identified. Defining the starting point of the active contour is a key problem described by Gunn and Nixon (1997), Neuenschwander *et al.* (1997), and Guopu *et al.* (2006). One approach sought to reduce the dependency on user input for snake initialisation through the development of inbuilt iterative correction techniques with very high effectiveness for accurate object segmentation (> 99%) (Bamford and Lovell, 1998). An additional method was introduced by Gunn (1996), was implemented in further work by Gunn and Nixon (1997, 1998), and utilised a dual snakes approach. This method combines two contours that evolve simultaneously, both internally (expanding contour) and externally (contracting contour) to the target object. A convergence between the two contours was reached along a target boundary.

Bamford and Lovell (1998) applied a dual active contour, combined with a global search method, to 20,000 medical images to detect cancerous cell nuclei with over 99% accuracy. The method applied a dynamic Viterbi algorithm to these to find a polygon of best fit, thus determining the global energy minimum. Bamford and Lovell also employed a simple correction technique based on feature centroids to ensure the snakes were initialised correctly every time. Guopu *et al.* (2006) applied a dual snakes method on imagery of a hand and a cup with a corresponding Chan-Vese geometric contour method (Chan and Vese, 2001). Their results suggest the dual approach is especially robust to poor contour initialisation. In addition, they also cite the potential of prior knowledge to further reduce the dependency on precise initialisation of the snakes.

Snakes have been applied to remotely-sensed imagery for the mapping of informal settlements at a per-building scale. R  ther *et al.* (2002) described a method of using active contours initialised from ‘blobs’ on a digital surface model (DSM) to delineate building features from aerial imagery. Their method had a detection accuracy of 62%, which they attributed to the extreme density of the studied settlement. They also found that manual delineation of the buildings was quicker than the utilised snake technique. In contrast, Mayunga *et al.* (2007) were able to reduce the detection time for individual buildings by 32% when compared to manual human digitisation of buildings. They adopted a three-process approach using (i) preprocessing, (ii) snakes and dynamic programming, and (iii) manual fine tuning of building outlines. The authors found that using Quickbird imagery was less effective than finer spatial resolution aerial imagery, due to the small nature of some informal buildings. The snakes were also limited by occlusion arising from high levels of noise and shadow in the Quickbird imagery.

2.4 SCENE MODELS

2.4.1 Background

The extraction of image features from model-based approaches, such as the HT, allows the generation of models as a method of simplifying the image into a series

of simple geometric shapes. This enables analysis focusing solely on objects of interest and therefore reduces the complexity of the image space. Complex geometric shapes can be modelled through a process of splitting into simple geometric primitives. The subsequent amalgamation of these elementary shapes allows the creation of models of objects, known as *scene models*. The concept of a scene model is simply storing the knowledge of where objects are located in an image, and the image characteristics (image spatial resolutions, pixel spectral values etc).

The scene model approach has several advantages over image-based object comparison. The model holds knowledge of the shape and potential size of the objects, meaning it can be robust to object occlusion (Khoshelham, 2007), a problem described in previous object detection applications (Wang *et al.*, 2004). Modelling also allows direct comparisons with only the objects of interest in the image, and models can be manipulated to display user-defined characteristics to aid this comparison.

The delineation of building footprints has been a major focus of several studies (Mason and Baltsavias, 1997; Brenner *et al.*, 2001; Gruen and Wang, 2001; Sohn and Dowman, 2001; Sarabandi *et al.*, 2008a). Creation of a scene model that represents a composite of these building or object footprints has informed decision making during damage detection and assessment. Scene models can be vector representations of the ground footprints of objects, rather than image-based entities. Gerke *et al.* (2001) described a method of creating scene models at a range of scales and from an array of remote sensors to build a simulation system for training emergency responders. The scene model for buildings was created by identifying roof parts and shadow information within aerial imagery and amalgamating these data within a model of a settlement. A scene model was also utilised by Shinozuka and Rejaie (2000) to assess damage at a range of scales from regional scale to individual buildings. This method manually assigned buildings within aerial imagery to damage classes and overlaid this on vector building data in a GIS. Correlation analysis was used to detect buildings that have changed, and object edges were used to reconstruct a 3-D representation of the scene.

The use of scene models, rather than per-pixel approaches, in damage detection or urban monitoring is a method more akin to human perception as people

perceive objects in context with other objects and continua. Objects can be presented at a variety of scales (e.g. from building to block, to town) to inform the decision making process. A method trialled by Vu *et al.* (2005), created building inventory databases from Quickbird sensor imagery in preparation for object-based damage assessment after earthquake events. The authors emphasised the benefits of ignoring discrete pixel details in preference to an object-based scene model approach, and provided results that suggested that object-based image analysis has potential as a feature extraction method.

In this thesis, models of objects extracted from imagery are referred to as scene models, and their creation to describe objects in a remotely-sensed image requires the development and use of algorithms that can detect features of interest and extract them from images (Objective 1 & 3; section 1.5). A series of these were described in the previous section. The following section describes several methods developed in previous research that divided complex geometric shapes into object primitives and subsequently re-amalgamated them in a model space environment.

2.4.2 Image Primitives

Primitive geometric shapes include lines, arcs, circles, rectangles, and other two-dimensional shapes. The detection of these primitives from imagery is fundamentally a simpler task than detecting complex geometric shapes directly. These primitives, once identified and extracted from imagery, can be amalgamated to create models of all of the objects of interest in the scene. In their detection of primitives, Rabbani and van den Heuvel (2005) have stated that of the objects found in industrial scenes, 85% can be approximated by planes, spheres, cones and cylinders. Khoshelham (2007) stresses the importance of modelling detected objects in his description of the detection of 3-D objects from laser range data, described previously in relation to the HT. He describes the progressive process of primitive detection and model building to form a 3-D abstraction, and visualisation of objects. The author also stresses the importance that precise object detection has for accurate object modelling.

Several extraction techniques focus on the disaggregation of buildings in RS imagery. Once detected, the usual approach is to re-aggregate primitives to form an outline of the building footprint. Brenner *et al.* (2001) described an approach that automatically generated three-dimensional building models using a combination of two-dimensional ground plans and a laser-detected DSM. The building ground plans were divided into a set of rectangles and planes, representing the major linear features of a roof as viewed from directly overhead. Building height information was inferred from the DSM, and details such as overhanging eaves were introduced. The primitives were then amalgamated into a simple building “skeleton”, and textured using aerial and ground-level imagery. The approach was limited by the availability of accurate ground plans, a suitable DSM, and the difficulty of splitting some buildings into primitives due to their complex geometric shape. Niederöst (2001) described the use of a similar DSM and 2-D ground plan data technique used in the Swiss ATOMI project (described in section 2.2.2). Here, the quality of the image and height information constrained the method.

Another method for semi-automatic building extraction from RS imagery was presented by Gülch and Müller (2001). The inJECT software they proposed, unlike the methods mentioned previously, utilised digital aerial imagery as the sole input dataset, and did not require ancillary vector information. The system relied on the input of an operator to select building type, gable location, roof shape, and ground height. The inJECT software had a pre-defined set of building primitives, covering the most frequent building types and roof shapes, and matched the user input to one of these using a probabilistic approach, with a high degree of accuracy when tested on both fine spatial resolution aerial and IKONOS imagery. Liu *et al.* (2005) also focused on the delineation of roof outlines using the PPHT, as described in the discussion of HT applications in section 2.3.3.

Sohn and Dowman (2001) described a method of building extraction from panchromatic IKONOS imagery at 1 m spatial resolution, which used the Burns algorithm (an extension of a Sobel edge detector) to identify linear features within an intensely urban area. A clustering method was applied to a histogram of the orientation angles of the detected lines to group dominant features into building unit shapes (BUS's). This method offers an effective alternative image segmentation

technique which detects lines individually as opposed to grouping homogeneous pixels based on their spectral or textural values.

2.4.3 Optimising Primitive Matches

The key challenge to an automatic approach for amalgamating geometric primitives for scene model creation is described: Each primitive has a possibility of being matched with any other primitive, producing many possible matches for each feature, thus, producing a potentially huge search space of possible combinations from which to detect the optimum combination and arrangement of features. If the choice of feature matches is to be random, the computing effort needed to find the optimum choice of matches rises dramatically. Hence, there is a requirement to reduce the size of the search space to find the optimal result within a reasonable time. This form of optimisation problem is known as the *Travelling Salesman Problem* (TSP) (origin unknown). The basic premise of the TSP is based on a salesman who has to visit a set number of cities in the least mileage possible. From a random starting point, the salesman can make a change in city pairs and accept the updated route if the mileage is shorter (Carr, 2008). This method of optimising the TSP will often secure local minima in little time, but struggles to find a global solution and, therefore, is a suitable heuristic for time-limited optimisation.

Global minima in TSP's are difficult to achieve because of the huge potential for matching cities, or variables, for which Kirkpatrick *et al.* (1983), and de Vicente *et al.* (2003), suggest that an increase in n (number of cities) will require computing effort that increases exponentially with n . Schmitt (2004) uses the following analogy to effectively highlight the factorial nature of the number of possible solutions:

*“The age of the universe is:
 $1.5 \times 10^{10} [\text{yr}] \cdot 365 [\text{d/yr}] \cdot 24 [\text{hr/d}] \cdot 3600 [\text{sec/hr}] = 4.73 \times 10^{17} [\text{sec}]$
 An exhaustive search of all possible paths through only 36 points would
 require about $36! = 3.72 \times 10^{41}$ comparisons to find the shortest path. This
 would take more time than the age of the universe.” (Schmitt, 2004)*

It is, therefore, necessary to use an optimisation method that can calculate a suitable solution within the time constraints applied to the algorithm. Hence, a heuristic

method is often preferred to a traditional algorithmic approach. A heuristic is not guaranteed to find an optimal solution, but rather an acceptable time-constrained solution (Rutenbar, 1989). What follows is a brief description of two widely used optimisation techniques that go beyond standard gradient descent methods, where the minimum of a smooth search space can be found, irrespective of its potential to be a potentially erroneous local minimum, or an optimal global minimum (Figure 2.5)

Genetic Algorithms

A common optimisation method is the Genetic Algorithm (GA) approach. Introduced by Holland (1975), and extended in works such as Goldberg (1988), the GA heuristic is a technique for finding the maximum or minimum of a function through a gradient descent method. As a heuristic, the global maximum or minimum is not always found, rather a suitable local solution is found, which is often the preferable solution due to time limitations of detecting the global solution. Designed to imitate the process of natural selection, GAs iteratively and randomly vary an amalgamation of a set of samples (*chromosomes*) to create a *population* with an assigned energy function, or *fitness function* from which it reproduces, creating the next population (Srinivas and Patnaik, 1994; Mitchell, 1998; Rowland and Weisstein, 2009).

Genetic algorithms differ from traditional gradient descent techniques (Figure 2.5), as their stochastic nature allows fluctuations across the search space, reducing the need for a smooth gradient function. They have been implemented in a variety of applications including ecology and machine learning (Mitchell, 1998).

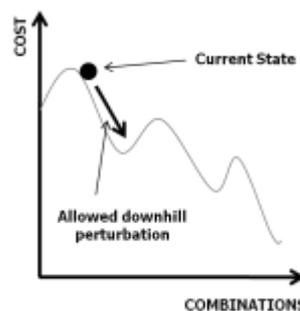


Figure 2.5 Traditional gradient descent optimisation.
Several local minima are shown, which are not equal to the overall, global minimum

The acceptance of a heuristic “optimal” solution is based on the combination of matches from the sampling pool – the *population* – and is not limited to individual combinations of matches.

Simulated Annealing

Another heuristic method of optimisation is simulated annealing (SA); an algorithm developed in the field of statistical mechanics to establish a solution to the global optimisation problem through the process of random searching. It is suitable for locating a minimum in a large search space that may have many parameters, and, therefore, the possibility of many local minima. Initially proposed by Kirkpatrick *et al.* (1983), SA is analogous to the process of annealing metals through a gradual and controlled cooling from a high temperature at molten state. The cooling rate has a critical effect on the form the final metal takes: cooling too quickly will not allow the elements to form as cohesively compared with a slower cooling schedule, where cooling slowly will often produce a more perfect result. In optimisation, this means that if a heuristic result is achieved in a short space of time, it may correspond to a local minimum in the search space, but it may not be the global minimum solution sought after. However, if the algorithm is allowed to run for longer, and local minima can be escaped from (Figure 2.5), then there is a greater chance the final, converged solution will be the global minimum.

According to Kirkpatrick *et al.* (1983), the two basic strategies of optimising large search spaces are: (i) divide and conquer, which combines the results of several small problems; and (ii) iterative improvement, where a known configuration of n elements is incrementally improved; a strategy prone to local minima effects. The SA method aims to improve on these methods by introducing a search space that is stochastically sampled and attributed with probabilistic performance measures, allowing perturbations to move uphill in a controlled manner (Rutenbar, 1989). The SA process searches for an optimal solution by randomly selecting a potential solution to the problem, and iteratively swapping this configuration of variables for a nearby configuration in the search space. Each potential configuration has an energy function, and the new *swap* is retained in preference to the old if the energy has been reduced. This swapping process continues, each time comparing the internal energy of the new configuration with

the previous best solution. The process can accept bad swaps (i.e. where the energy is not reduced), if a certain probability function is met.

The simulated annealing algorithm has few input parameters to define. However, an appropriate choice of cooling schedule is important. As SA progresses and the observed energy is reduced iteratively, a synthetic “temperature” (T) is lowered incrementally at a constant rate α , in accordance with the *annealing* analogy. The lower the value of T , the less probable it is that the algorithm will accept a bad swap, accepting only those swaps that lower the energy converging on a heuristic local minimum. The relationship between the temperature and probability of accepting a bad swap is shown in an example of a simulated cooling schedule in Figure 2.6.

A comprehensive comparison of the main types of cooling schedules was provided by Nourani and Andresen (1998). In summary, there are two recommended procedures: (i) exponential

$$T(t) = T_0 k^t \tag{2-6}$$

and (ii) linear

$$T(t) = T_0 - At, \tag{2-7}$$

where “time”, t , is the iteration count. An additional logarithmic cooling schedule is described in Nourani and Andresen’s work (1998), but dismissed as “utterly impractical” as its extremely slow temperature decrease reduces the effect of the constant value included in the exponential schedule (2-6). The value of k is retained as a constant in the range $[0,1]$, slowly reducing the temperature, T , so the system evolves into a Boltzmann distribution, with an ever-decreasing temperature (Kirkpatrick *et al.*, 1983).

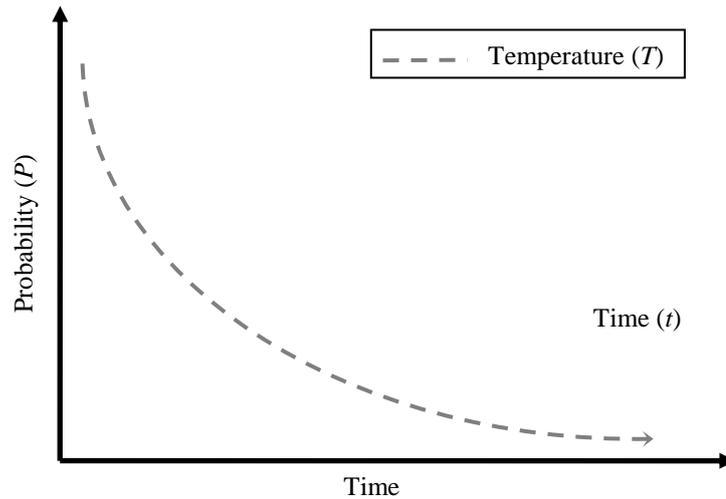


Figure 2.6 An example of a simulated annealing cooling schedule

In their method to improve the TSP, Kirkpatrick *et al.* (1983) used the Metropolis procedure (Metropolis *et al.*, 1953), where small random displacements of variables are made during an iterative improvement method. These displacements, or swaps, are accepted and retained when the energy measured (E) is reduced by the swap, or if a probability function (P) exceeds a threshold value. The probability function

$$P = \exp(-\Delta E / kT) \quad (2-8)$$

applies to non-zero temperatures. ΔE is the decrease in energy and k is a constant, slowing the temperature decrease. As such, the acceptance of bad swaps (where the energy is not reduced, but where a probability criterion is met) prevents the algorithm from falling into local minima, thus, reducing the problem of iterative improvement by allowing the energy to temporarily move uphill in a controlled manner (Rutenbar, 1989). This characteristic also allows the SA heuristic to explore more of the search space and return a more probable final result than that achieved by using a traditional gradient descent technique.

SA methods have been applied to a variety of applications. Buckham and Lambert (1999) found a SA strategy for the TSP that could optimise 6000 cities, improving on the previous largest solution of 318 cities. They did stress the difficult

nature of choosing the most suitable cooling strategy, describing it as “*a black art*” (Buckham and Lambert, 1999). Their paper also described applications of SA to image reconstruction, circuit design and robotic path generation. Kirkpatrick *et al.* (1983), Kirkpatrick (1984) and Rutenbar (1989), also described SA applications to the traditional optimisation problems of TSP and circuit design.

2.5 UNDERSTANDING “CHANGE”

2.5.1 Background

Until now, this chapter has focused on methods that detect and extract objects of interest from imagery. These objects can be detected directly, or be amalgamated from primitives into objects in scene models that describe a simplified reality of a particular part, or key features, of the image. The extraction of image features into models can be used as an information input in a process of change detection involving two phases of data processing. The first stage is to identify and capture the extent of the change occurrence. As described in the previous sections of this chapter, this can be done in several ways, either through direct image-to-image pixel comparison of *before* and *after* images, or by detecting and extracting image features into a scene model and modelling the change. The second phase of data processing involves generating a combination of quantitative measures and qualitative descriptors of the difference arising from the changes observed over a period of time. This section describes previous research into both of these phases, and introduces techniques that can be utilised for these stages of data processing.

2.5.2 Quantitative Change Descriptors

Statistical methods

There are several methods of quantifying change in scene model objects. The most elementary form of quantifying change uses first order statistical descriptions. Measures such as *area* for two-dimensional shapes, or *volume* for three-dimensional

objects, as well as perimeter or height are only some of the basic statistics that can be generated easily, and provide useful, elementary descriptors of object change. In addition, contextual information can be added on how the object has changed in the image by quantifying how far an object has moved in both image and geographic space. Imagery can also provide informative data on the radiometric properties of the image objects at both the pre- and post-event time periods. These can be per-pixel or per-object spectral characteristics, which can be used to derive qualitative descriptors of the change mechanism. Holland and Tompkinson (2003) utilised first-order descriptors of change in their discussion of a new change detection framework. They proposed that basic detected changes in both size and compactness at an object level can guide a professional operator to areas of potential change, allowing a supervised method of change description.

Second-order description also provides a valuable platform from which to quantify what has changed between images of the same location. Comparing the spatial relationships between features in the scene model can create a greater understanding of neighbourhood changes. Methods have been used previously for these types of descriptors for applications of explaining economic patterns in a city (Parker *et al.*, 2001), or second-order neighbourhood information to describe urban area complexity (Herold *et al.*, 2002). Example metrics from the Fragstats program that Herold *et al.* implement are patch density, edge density and fractal dimension.

Geometric Change

Modelling actual object change from pre- and post-event imagery can provide a valuable method for gauging the magnitude of per-building change. Assessing geometric changes in objects through a modelling approach can be a twofold process. Objects can be detected from both the before and after images, and stored in two models. A model-to-model comparison can be made, and the object differences recorded and stored. Alternatively, the geometric boundaries of the pre-change model can be modified and compared to the after image through the generation of a simulated image. In this instance, shape-morphing algorithms can be used for the modification of object boundaries. One particular method that has been investigated for this purpose is the use of active contours. Described earlier in the chapter, these “snakes” can be initialised on the extracted objects from the pre-event feature

extraction, and provide information on the direction and magnitude of geometric change to be stored.

2.5.3 Qualitative Change Descriptors

Previous studies have described several methods for generating qualitative change descriptors. This section is divided into those classes developed for the dual applications that are tackled in this thesis; damage detection and urban monitoring.

Damage Classes

Generating quantitative change statistics on a per-object basis facilitates the use of existing qualitative scales to describe change. For example, manually assigned damage classes from the European Macroseismic Scale (EMS, 1998) were used to classify high-spatial resolution pre- and post-earthquake imagery (Yamazaki *et al.*, 2005; Miura *et al.*, 2006). This scale, which differentiates damage to masonry buildings, is described in

Figure 2.7. This was a hugely labour-intensive task, and introducing some form of automation to this is desirable, as stated in objective number 4. Huyck *et al.* (2004a) found that it was difficult to differentiate all of the EMS-98 scale, and recommended that damage states 1, 4 and 5 were easily detectable from RS data. Applying a combination of quantitative descriptors to these damage classes could allow more effective differentiation between damage states. Eguchi *et al.* (2005) advocated the development of standardised damage scales to ensure consistent methodologies for damage detection and classification.

Formalised Classification – Urban Change

Formalised classifications of change also exist in the realm of urban monitoring. Ordnance Survey currently uses a proprietary classification system to describe the change specification important to them. Coppin *et al.* (2004) described four types of change for spatial entities that could be applied to objects: (i) objects can change class category; (ii) objects can expand, shrink, or change shape; (iii) objects can change position; and (iv) objects can fragment or unite. Knowledge of the potential

Classification of damage to masonry buildings	
	<p>Grade 1: Negligible to slight damage (no structural damage, slight non-structural damage)</p> <p>Hair-line cracks in very few walls.</p> <p>Fall of small pieces of plaster only.</p> <p>Fall of loose stones from upper parts of buildings in very few cases.</p>
	<p>Grade 2: Moderate damage (slight structural damage, moderate non-structural damage)</p> <p>Cracks in many walls.</p> <p>Fall of fairly large pieces of plaster.</p> <p>Partial collapse of chimneys.</p>
	<p>Grade 3: Substantial to heavy damage (moderate structural damage, heavy non-structural damage)</p> <p>Large and extensive cracks in most walls.</p> <p>Roof tiles detach. Chimneys fracture at the roof line; failure of individual non-structural elements (partitions, gable walls).</p>
	<p>Grade 4: Very heavy damage (heavy structural damage, very heavy non-structural damage)</p> <p>Serious failure of walls; partial structural failure of roofs and floors.</p>
	<p>Grade 5: Destruction (very heavy structural damage)</p> <p>Total or near total collapse.</p>

Figure 2.7 Classification of damage to masonry buildings
(EMS, 1998)

changes that can occur in image features facilitates the selection of statistical metrics used to quantify change.

Ontologies

Change detection in damage detection studies mainly focus on visual detection or basic quantitative measures of change, from which damage assessment is established. None of the literature studied to date has included a systematic description of what has actually happened to objects in a scene. To describe objects

and classify their meaning automatically would involve making the computer behave in a way much more akin to human behaviour than discrete pixel analysis seen in traditional change detection methods. Computers analyse images in a much more methodical way based mainly on per-pixel analyses of images rather than use the contextual understanding that is an inherent aspect of human perception. The ability to describe objects from images qualitatively would also make semantic retrieval of data more efficient.

One method of describing objects is through the design of an ontology of features. Ontology, in literal terms, is “*the science or study of being*” (OED, 2004), but applied to artificial intelligence (AI) the meaning of an *ontology* has been refined to “a formal specification of a shared conceptualization” (Borst, 1997). Based on a structured database of machine-readable primitive descriptors, it is proposed that ontologies can be created for a variety of scales ranging from individual building features to conglomerates of these, represented by a scene model. Ontological primitive descriptors can be created for the scenes captured in both the before and after images. Such granularity of scale allows scene description at several levels; object, neighbourhood, up to image level.

It is apparent from the available literature that there is a lack of semantic descriptive methods for remotely-sensed imagery. Much early ontology research focused on indexing images in the arts, to aid semantic retrieval and suffered from computational restraints, especially with attributional queries (Gudivada and Raghavan, 1995; Cibelli *et al.*, 2004). A prototype method of retrieving remotely-sensed images was proposed by Sun *et al.* (2005): an ontology was defined for a number of images based on an existing framework providing a pre-defined set of descriptive terms. This approach was tested against a grid retrieval system and was shown to be more efficient in terms of precision and speed. The authors acknowledged the importance of ontology engineers for precise definitions, a requirement alluded to by Hart and Greenwood (2003). A differing approach is suggested with an ontology to describe interrelated concepts as “building blocks” for an underlying geodata model. It is recognised that such a model has the danger of becoming overly complex, however. Such an approach is one of the main areas of interest for the Ordnance Survey, who are currently investigating machine-readable

ways of representing the meaning of Ordnance Survey's data and geographical knowledge using semantic web technologies (Ordnance Survey, 2009).

Durbha and King (2004) proposed an ontological framework for semantic retrieval of Landsat and Moderate-resolution Imaging Spectroradiometer (MODIS) data using Web Ontology Language (OWL), a commonly used framework. Their results were at various levels of granularity and relied on support vector machines and other machine learning algorithms for information retrieval. Applications that use efficient semantic systems to source imagery are increasing, especially with advanced mobile telephone technology. Searching for information by capturing objects digitally on mobile phones is being investigated by advertisers (Biever, 2006). For example, some mobile recognition software uses techniques that match and reference identical objects between two images (Biever, 2006).

It is also possible to further the change ontology concept and create an ontology for different scenarios or phenomena that may have contributed to change, a concept that would formalise the previously mentioned need for connecting levels of change to damage classes (Huyck *et al.*, 2004b). Objects or scenes may have distinct characteristics after specific events such as fire, earthquake, storm surge, flood etc. and researchers are actively seeking to develop scenario-specific damage scales for use after such natural disasters (Eguchi *et al.*, 2005; Womble, 2005; Friedland *et al.*, 2006). Ontologies can be created based on expert knowledge from remote sensing operators. A hierarchy of qualitative object data enables more effective management of changing objects and provides the user with a base to prioritise variables to aid decision making after a change event.

2.6 SUMMARY

This chapter has provided a wide ranging overview of some of the key methods of object detection from remotely-sensed images. It has provided background information from which to create the rationale for this thesis described in section 1.5. Several issues are apparent from this review. There is a lack of semi, or fully automated methods for damage assessment and urban monitoring, as intimated by

Tompkinson *et al.* (2003). Many existing techniques are dependent on human input during the data analysis, and there are no systems that can extract and reconstruct buildings from images to the level required by practical production (Zhang *et al.*, 2001). Despite this, it has been shown that there is an abundance of existing methods for the extraction of primitive geometric features from imagery, and several of these have been applied to RS data. However, some are often overly reliant on operator delineation of parametric inputs.

Considerable potential has been shown for the generation of object-oriented scene models as a method of detecting change from fine spatial resolution data. Primitive geometric shapes can be identified using one-or two-dimensional object description, therefore reducing the complexity of identifying building whole features using a single algorithm. These primitives can be amalgamated to reform the objects within a scene model, allowing comparisons over time at the object level. Scene models allow statistical, textural and geometric measures of change at both per-object and image levels. Modelling the objects of interest has the added advantage of being less sensitive to image occlusions, or differences in observational geometries between pre- and post-event images, and therefore addresses some of the problems associated with per-pixel change detection methods (Singh, 1989). Describing change using qualitative and quantitative descriptors provides additional understanding of object morphology through a combination of basic statistical measures, or more formalised change-specific classifications.

The subsequent chapters of this thesis describe the development of a conceptual model framework for change detection that aims to confront several of the problems highlighted in this chapter and section 1.5. Chapter 3 describes this model, with preliminary applications described in chapters 5-7.

CHAPTER 3

THE CONCEPTUAL MODEL

3 THE CONCEPTUAL MODEL

3.1 BACKGROUND

An analysis of the current state-of-the-art in the broad field of change detection has identified several needs, described in the aims and objectives of this thesis (section 1.5). These have been further described in the literature review in chapter 2, highlighting several omissions and limitations of the literature. A method for tackling several of these needs has been devised and is articulated in this chapter. A model-based framework has been created and is described in relation to the thesis rationale. The conceptual model (Figure 3.1) consists of two key phases of *object detection* and *object comparison*, bookended by *Input*, and *Output* phases. These distinctions between model phases are seen in Figure 3.1, and retained throughout the chapter. The individual components of the model are also described, detailing the key benefits and premises of the model. The conceptual model is implemented on several image sets as exemplars for its development (described in chapters 5-7).

3.1.1 Aims of the Conceptual Model

The conceptual model is a novel, modular, model-based framework for the detection of change arising from the dual applications of post-disaster damage and urban development. It aims to reduce the amount of user-input and provide automated extraction of objects from imagery, and model the changes in each object that have occurred between pre- and post-change time periods. In doing so, the algorithm has been designed to provide a rapid approach to detect changes resulting from disaster events and urban expansion. The conceptual model utilises existing techniques from the fields of remote sensing, computer science and mathematics (described in chapter 2) as modular components designed to be customisable to fit

differing damage detection scenarios. The review of the literature, and the thesis rationale (section 1.5) provide context for the following aims of the model

The model aims to

- i. adopt existing algorithms in a modular environment
- ii. identify and extract objects as geometric primitives from RS imagery
- iii. use iterative optimisation techniques to amalgamate primitives for the delineation of objects of interest
- iv. be reconfigurable for differing objects of interest or event scenarios
- v. generate a scene model of objects from the pre-event image which can be manipulated to ‘fit’ the post-event image
- vi. generate quantitative and qualitative statistics and descriptors of changes

Briefly, a pair of co-registered, fine spatial resolution, pre- and post-event images acts as inputs to the model, where an object detection algorithm is applied to the pre-event image to define building outlines. These outlines are stored as objects extracted from image space into vector space – the scene model. A forward model then creates a simulated image based on the objects in the scene by assigning brightness values to them in several wavebands. The brightness values of this simulated image are compared to real-world data from the pre-event remotely-sensed image and a brightness error image generated. An iterative optimisation process, where each iteration is informed by the outputs of its predecessor, is applied to update the scene model and effectively fit the spectral values of the model to the pre-event image.

In the second half of the model, the brightness properties of the pre-event scene model are compared iteratively with the brightness values of the image captured after the change event and a second error image produced based on these differences. A second forward model is used to update the spectral and geometric properties of the objects in the scene model, in an iterative manner with the error image informing each iteration. This results in the reduction of error in the scene, as the scene model gradually morphs to “fit” the post-event image. All spectral and geometric change vectors are stored automatically at each iteration, with object-based statistics quantifying the changes in objects between the images. The individual components of the model are described in detail in the following sections.

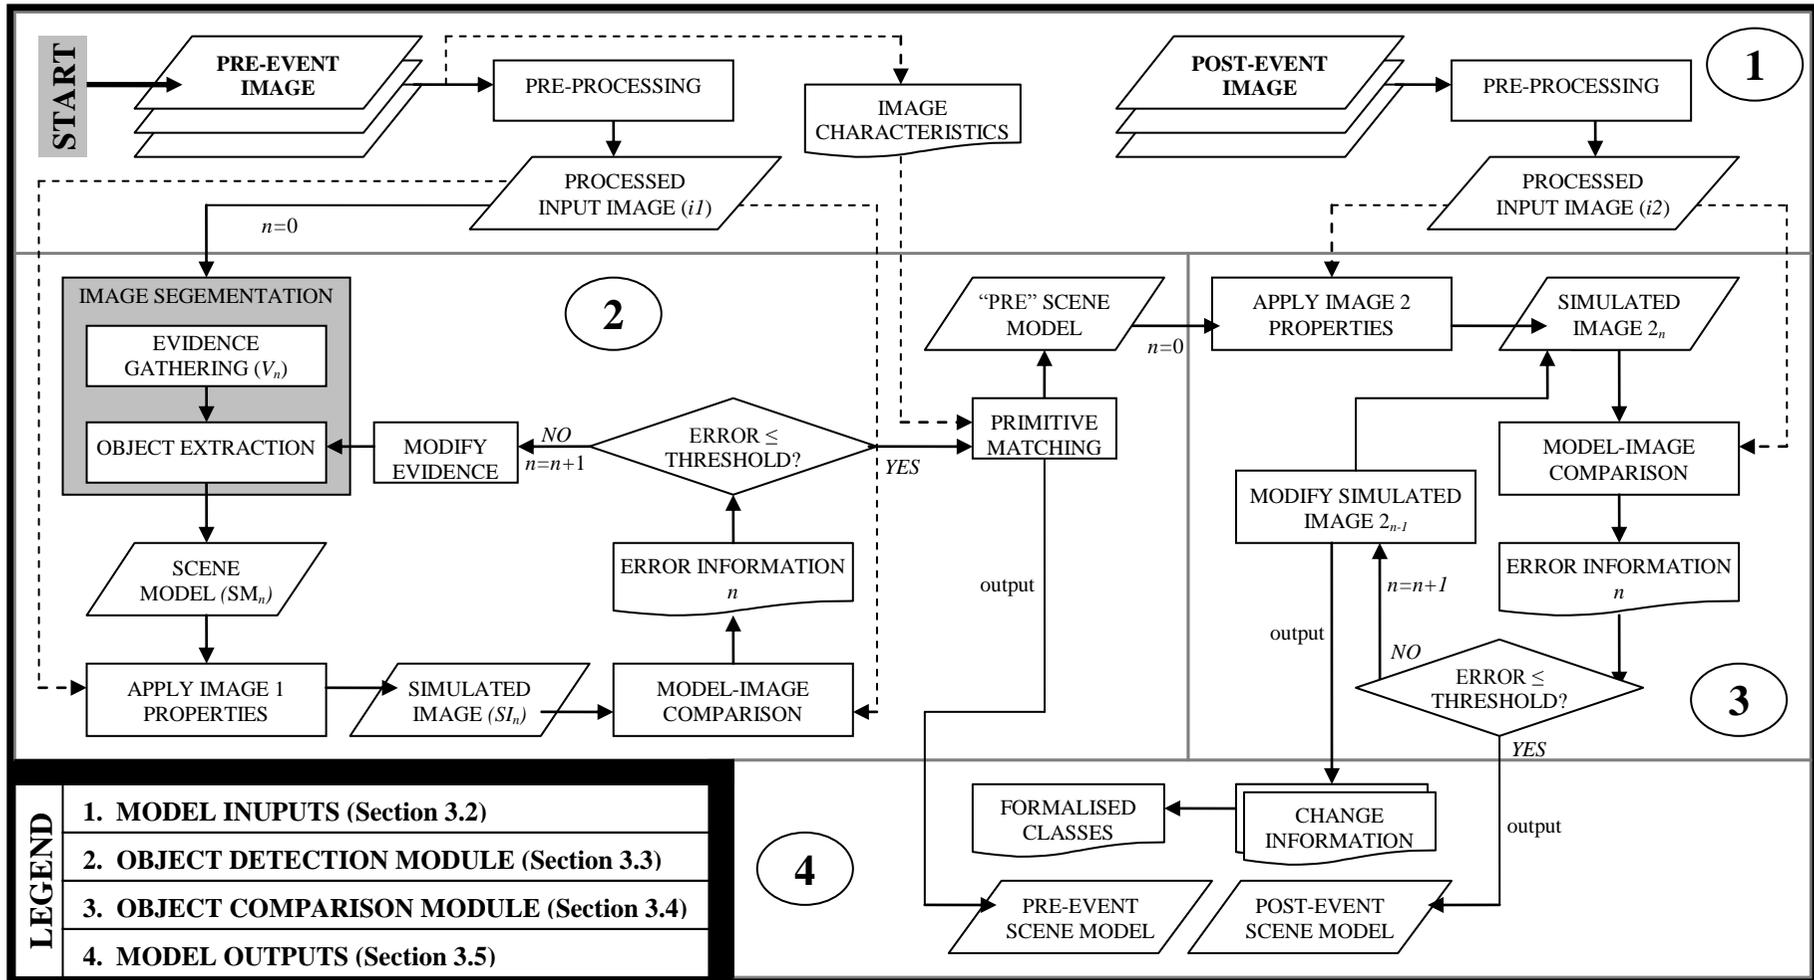


Figure 3.1 The Conceptual Model showing four phases: 1. Inputs; 2. Object Detection; 3. Object Comparison; 4. Outputs. (n = iteration number)

3.2 MODEL INPUTS

3.2.1 *Image Characteristics*

A pair of fine spatial resolution, remotely-sensed, optical images (captured before and after a change event) is used as the input to the model. The model is designed to accommodate both images from satellite and aerial sensors, and a combination of these data sources can be used with an additional pre-processing stage. The model is designed to work at a per-building scale, but imagery with a coarse spatial resolution could be utilised depending on the purpose of the application of the model.

The initial stage of the model involves the pre-processing of both input images. To facilitate precise detection of change to image objects, both the pre- and post-event images should be co-registered. This ensures the bases of all image features existing in both images exist in the same image location. This is a fundamental aspect of image-based change detection, without which would result in an increase in spurious detections and false changes detected, as is widely acknowledged in the literature (Dai and Khorram, 1998; Lu *et al.*, 2004; Sui *et al.*, 2008). There is not a necessity to orthorectify for the imagery, however, there is a need for both images to be at least co-registered, so if overlaid on top of each other, co-existing image features would be in the same image space.

Contextual information is required to initialise the model. Acquisition characteristics of both the sensor and the images are utilised by the model and so information on the date, time and location of the sensor at the time of capture is required. Obtaining these data allows the calculation of the solar location relating to the time and date of data capture. This information can be calculated using readily-available web resources such as Sun angle calculators (Conrad, 2009). This solar information is also added to the model.

Knowledge of image characteristics provides necessary ancillary information. Information about the geographic location, combined with the spatial and spectral resolution of the image should be known when ordering data, and provides contextual information about the scene captured in the images. Preliminary implementations of the model (described in chapters 5-7) have utilised eight-bit imagery. However, the model is not limited by constraints on radiometric precision.

There are also several pre-processing stages that need to be computed depending on the choice of input data and the required geographic extent of the implementation.

3.2.2 Pre-processing

The conceptual model was implemented in the form of a modular framework, allowing the substitution of components depending on the specific application. Several of the algorithms used in the object detection phase stipulate specific data formats to be used. Pre-processing steps account for these variable formats, with examples including the conversion of 16-bit to 8-bit imagery, resampling the spatial resolution, and conversion from red-green-blue (RGB) colour bands to greyscale images. The pre-processing stage of the model essentially consists of several object detection algorithm-dependent routines that were utilised only for manipulation of the data into the required format. The resulting processed image is used as the input to the object detection module.

3.3 OBJECT DETECTION MODULE

The object detection module is one of two key phases within the conceptual model, with the overriding objective to delineate objects of interest in the input pre-event image. These target features are identified within the imagery, and extracted into a binary, single-class scene model of the pre-event scene. A fundamental objective at the design phase of this module was to devise a method with high levels of automation, which could delineate precise target boundaries from a variety of image types. Retaining a level of generalisability of this module was also a key aim necessary if the model was to be used for multiple applications.

Several characteristics of this module were devised to ensure that the object detection was as automated as possible, whilst retaining generalisability of the model to a variety of object, or event scenarios.

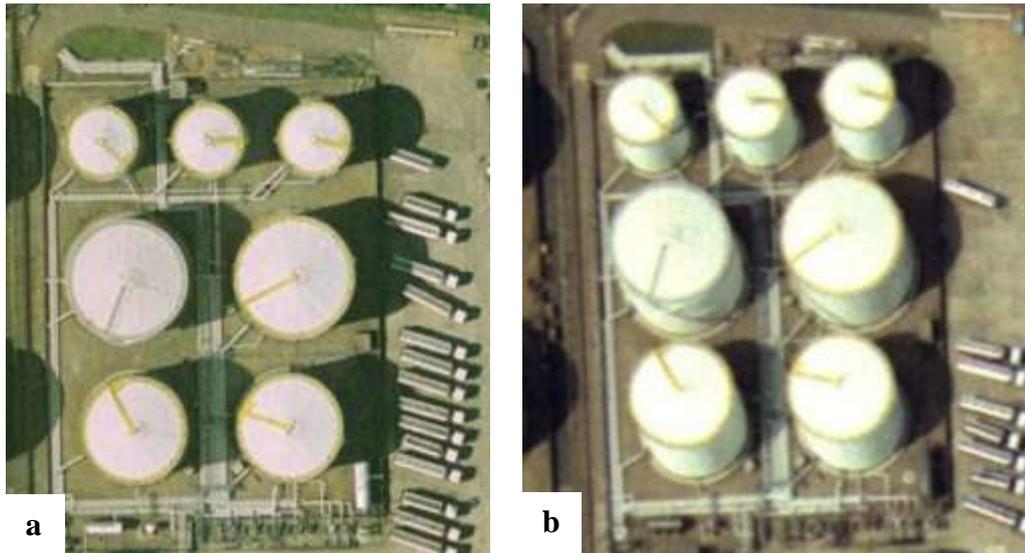


Figure 3.2 The effect of different view angles on object footprint delineation
Images copyright (a) Google Earth (b) GetMapping Ltd.

1. objects are detected using a well-established feature extraction method (segmentation)
2. candidate objects from 1 are compared to the observed image, and adjusted through an iterative forward and inverse modelling process (validation).

The conceptual model has been designed to incorporate object detection in two or three dimensions, depending on the geometric nature of the object as seen in the image, resulting from differing view angles between sensor and object. Figure 3.2 shows an example of the same objects (oil storage drums) captured in different images with unequal observational angles. It is clear that, in this instance, the footprint of each object in the image varies considerably depending on the view angle. Simply detecting the tops of the objects in this example would not allow accurate conversion from image space to geographic space, as the tops of the drums would not be in the same geographic location as the bases. It is therefore possible, needs dictating, to detect three-dimensional objects from oblique-angle imagery, depending on the view angle and the presence of occlusions in the image.

3.3.1 Segmentation Algorithm

The central component of the object detection stage is a suitable algorithm to detect primitive geometric shapes in images. Several methods were described in the previous chapter designed to detect specific geometric shapes in images (ranging from edge or line detectors, to shape-matching methods). The choice of feature extraction algorithm at this stage of the framework is dependent on the features of interest in the pre-event input image. The choice of using a probabilistic method allows the generation of evidence on the strength of fit of an object, and enables assessment of the relative importance of each object to each other. In these algorithms, such as the Hough Transform, the objects are extracted on the basis of this evidence gathering phase. The object detection module has been designed with the aim of reducing the importance of precise parameterisation of the segmentation algorithm (objective 4, section 1.5.1); a time consuming process. An iterative method is therefore introduced to update the evidence for a candidate feature, with the intention of allowing the relaxation of a rigid, single-iteration approach. Therefore, the algorithm is allowed to detect many objects, irrespective of them being true positive or false positive detections. The main goal at this stage is to limit the number of false-negative detections (where target objects in the input image are missed by the extraction algorithm).

The resulting objects from this algorithm are “candidates”, as their validity has not been confirmed. All of these objects are combined in a *scene model* – a set of object descriptors describing, for example, geometric, location or textural characteristics of each object. This scene model facilitates the generation of synthetic binary images of equal dimensions to the input image showing the spatial footprints of candidate objects. The verification of candidate objects involves the generation of error information for the detected objects, followed by an iterative error minimisation process, described in the following section.

3.3.2 Object Validation

The candidate objects detected by the segmentation algorithm are passed into an iterative forward and inverse modelling loop, with the intention of progressively increasing the number of correctly delineated objects, including both those that exist in the input image, but were missed by the segmentation algorithm, and the removal of erroneous objects (candidates that are not present in the input image). This is achieved through an iterative error minimisation process, where the error is generated by comparing the detected objects (displayed in a simulated image) with the input remotely-sensed image. Through this comparison, per-pixel, and scene-wide error information can be calculated, with high levels of error shown when:

- i. objects exist in the scene model where no objects exist in the input image,
- ii. objects exist in the input image, but have not been detected during segmentation, and therefore don't exist in the scene model,
- iii. objects exist in the scene model but do not precisely match the locations of the corresponding objects in the input image.

For the initial extraction of objects, $n = 0$ (where n is the current number of iterations through the validation loop), SM_0 is the primary scene model of extracted features. The validation loop consists of several steps, starting with scene model n (SM_n). Using the vector information in the scene model, the spectral characteristics of the pixels within the detected objects in the input image are derived and assigned to the corresponding pixels in SM_n , creating a simulated representation of the candidate objects – a *simulated image*, s_n .

Simulated Image

The simulated image is effectively a binary mask of *objects* and *background*, and is created by averaging the spectral values of every non-object pixel and attributing this to the equivalent pixels in s_n . In addition, the modal value of all *object* pixels within each feature is calculated, giving a single most common value found in the detected features in the image space. The mode is chosen over the mean brightness value of the pixels to avoid using a skewed mean value arising from spectral anomalies, caused by several scenarios: misalignment of object boundaries can significantly lower the mean value of pixels within a detected circle if the circle is

bright and the surrounding pixels are dark. With false-positive candidate object detections, the mean and mode values can vary substantially if the majority of the “object” is dark, with a few light pixels. Testing has shown that these differences can be greater than 20% of the range of possible values for a greyscale image (0-255). Hence, the modal value is assigned to all pixels within the boundaries of the candidate objects within s_n .

Model-image comparison & error generation

The next stage of the validation process involves the automatic generation of error information through a comparison of the binary simulated image, s_n , with the observed, input image, \mathbf{i} . The processing at this stage is simple, with a per-pixel differencing of the two images of the same dimensions so that

$$\mathbf{e}_n = \mathbf{s}_n - \mathbf{i}_1, \quad (3-1)$$

where \mathbf{e}_n is the resulting error image, and \mathbf{i}_1 corresponds to the pre-event input image. \mathbf{e}_n is a matrix of per-pixel error information, and is used to highlight spurious detections of objects, whether false positives, or false negatives. New error information is generated for each iteration of the process (described in section 3.3.3).

The objective of the iterative validation process is to reduce the error, and, hence, delineate all objects of interest in the input image. This is achieved through the generation of image-wide statistical measures of error, such as (i) the root mean square error (RMSE: 3-2):

$$\rho_n = \sqrt{\frac{1}{N} \frac{1}{M} \sum_{i=1}^N \sum_{j=1}^M (e_{ij})_n^2}, \quad (3-2)$$

where N is the number of pixels in each row of the image, and M is the number of pixels in each column. e_{ij} refers to the error at pixel (i,j) of the image array, for iteration n .

(ii) Mean absolute error, τ , (3-3) of the error image (\mathbf{e}_n) is calculated by

$$\tau_n = \frac{1}{N} \frac{1}{M} \sum_{i=1}^N \sum_{j=1}^M |e_{ij}|_n, \quad (3-3)$$

where the same notation from (3-2) applies.

(iii) correlation coefficient, c , (3-4) of the simulated pixel values, A , and the observed pixel values, B :

$$c = \frac{\sum_m \sum_n (A_{mn} - \bar{A})(B_{mn} - \bar{B})}{\sqrt{\left(\sum_m \sum_n (A_{mn} - \bar{A})^2 \right) \left(\sum_m \sum_n (B_{mn} - \bar{B})^2 \right)}}. \quad (3-4)$$

3.3.3 Validation Iteration

Validation is an iterative process that continues until either the error generated in equation (3-1) converges on a minimum value, or a preset error threshold is met. This is depicted in Figure 3.1 by a decision point in the workflow. If there is a reduction in the error, or the error increases, the aforementioned validation steps described in section 3.3.2 are repeated.

Each additional iteration of the model beyond $n = 0$, does not use the segmentation algorithm because the input RS image is unchanged, and so the results of another object identification would be the same as the initial candidate object set. A modified accumulator space is used instead for iterations where $n > 0$. The evidence in the accumulator space is modified, based on normalised error information, \mathbf{e}_n^1 . The accumulator votes, \mathbf{A} , are combined with the normalised error by

$$\mathbf{A}_n = \mathbf{e}_{n-1}^1 (\mathbf{A}_{n-1}) \quad (3-5)$$

It is apparent from (3-5) and Figure 3.1 that the iteration number, n , increases before the evidence is modified, allowing the new “votes” to feed the subsequent iteration. Essentially, the evidence is deliberately modified to the advantage of the system. Votes corresponding to false positive detections will be reduced, while those corresponding to false negative detections will be “boosted” through the use of the error information. The boundaries of correctly detected objects may also be modified through this per-pixel approach, based on large positive or negative error values in neighbouring pixels.

The model subsequently uses the new evidence, \mathbf{A}_n , without recourse to the original image segmentation method, to update the set of detected objects. A new binary simulated image is created as before, with newly calculated *object*, and *background* values generated, and s_n is compared with the observed image using equation (3-1). The resulting error information is used to determine a stopping point, based on two possible outcomes: convergence, or a supervised termination criterion. Convergence is achieved when there is no further decrease in energy for 5 or more iterations of the algorithm. The supervised termination criterion can be manually set by the operator to finish the processing after a set duration. Once the algorithm has run its course, the current set of candidate objects become the definitive pre-event scene model, otherwise a new iteration begins.

3.3.4 Primitive Matching

Images with oblique view angles can provide three-dimensional object information as discussed earlier, and shown in Figure 3.2. In these cases, information about the locations of the base and top of the object should be extracted through the segmentation and validation processes. The final stage of object detection is then the amalgamation of these primitive geometric shapes that constitute the target object. Matching the primitives of a complex object is a combinatorial optimisation problem similar to the travelling salesman problem, described in chapter 2. Potentially, the search space for possible matches is large, thus, a heuristic method is used to find a robust match within realistic time constraints. The image

characteristics identified in section 3.2.1 can be used as ancillary data to aid the matching of primitive shapes.

3.4 OBJECT COMPARISON MODULE

The object comparison module is the second of the two key stages core to the conceptual model. This module is designed to compare the objects, detected by the object detection module, with the post-event image. This stage differs from previous change detection methodologies, as it does not rely on direct pixel-to-pixel comparisons of two images. Rather, its main function is the gradual modification of the pre-event scene model to fit the post-event image. With the magnitude of change stored at each iteration, statistics can be generated showing how each object has changed in the interim time period between pre- and post-event, whether by a sudden, unexpected event, or gradual urban changes.

3.4.1 *Updating the Scene Model*

In the initial processing stage of object comparison, objects in the pre-event scene model are assigned spectral properties taken from the post-event input image, creating a new simulated image of the target objects (\mathbf{s}_2). A second validation loop is used to compare the simulated image with the observed, post-event image, \mathbf{i}_2 . This loop generates error information showing the strength of similarity between the geometric and spectral characteristics of the model and the post-event images,

$$\mathbf{e}_{2n} = \mathbf{s}_{2n} - \mathbf{i}_2. \tag{3-6}$$

If the error in the model is below a preset threshold, or has converged to a minimum, the candidate objects are confirmed as the objects for the post-event scene model. If the error remains large, the iterative process of updating the geometric and radiometric characteristics of the model continues.

3.4.2 Iteration

The error information generated by (3-6) is gradually reduced through a process of updating the simulated image s_{2n} . This process is two-fold, as the model can be updated spectrally and geometrically. Spectral updating of the model is described in chapters 5-6 in this thesis, and potential geometric updating of models is discussed section 8.4.2.

3.5 OUTPUTS

There are several outputs from the conceptual model, depicted in Figure 3.1. Combined, these form a record of the location and characteristics of the extracted objects. In addition, the iterative updating of the forward and inverse modelling loops allows the gradual generation of change information, culminating in a holistic view of the geometric and radiometric magnitude of per-object change.

3.5.1 Scene Models

Two scene models are output from the model, and store information on the location as well as spectral, geometric and texture characteristics of each detected object :

- 1 Pre-event - representing the objects detected in the pre-event image
- 2 Post-event – representing objects of interest in the post-event image

Both scene models are representations of objects of interest defined by spectrally homogeneous objects extracted from input imagery. They can be converted to vector format (e.g. shapefile) for inclusion in a GIS for additional analysis, inclusion in existing data products, or disseminated using tools such as Google Earth (section 6.6.2).

3.5.2 *Change Descriptors*

Change characteristics are output at every iteration of the object comparison verification loop. These characteristics describe the magnitude and direction of changes made to objects in the scene model at each iteration. This change information can be amalgamated for each object to show the total extent of change, quantifying the differences between features located at the same geographic position at two time periods and facilitating the use of formalised classification systems appropriate for the dual applications of urban monitoring, or post-disaster damage detection.

3.6 SUMMARY

The development of the conceptual model is the most fundamental aim of this thesis (objective number **1**, section 1.5). It provides a novel framework for the detection of objects in a model-based environment (objective **3**). It aims to provide semi-automated algorithms for the detection and amalgamation of image primitives, whilst not being reliant on the optimal, initial parameterisation of these algorithms (objective **4**). This allows *loose* input parameterisation, and *tightening* around correct objects. Khoshelham (2007) highlighted this importance, stating that the modelling of objects is dependent on accurate object detection.

Atiquzzaman (2005) stated that traditional methods of addressing the high computational requirements of various techniques can be addressed using *a priori* information or to design efficient algorithms. Both of these problems are addressed by the described conceptual model as the post-event scene model is modified based on pre-existing information from the pre-event model, rather than extracted directly from the post-event image. The output object models, containing pre- and post-event scene information, are also complemented by both qualitative and quantitative descriptors of change (objective **6**).

3.6.1 Implementation

The conceptual model described in this chapter has evolved throughout the project period, and represents potential advancement of methods for object extraction and comparison by reducing the amount of user-interaction in post-disaster change detection, and provide a semi-automated algorithmic framework. The method also is designed to model individual features, and qualitatively and quantitatively assess changes in objects, and the image scene. The remainder of this thesis focuses on the application of the first stages of the model to various datasets: The object detection module has been applied on two images, both for the automatic detection of 2-dimensional and 3-dimensional objects as described in chapters 5 and 6, respectively. Preliminary testing of the object comparison is described in chapter 7, where change information has been generated between pre- and post-event image objects, with both qualitative and quantitative descriptors assigned to describe the associated change. Chapter 8 and 9 focus on an open discussion of these applications, and describe future advancement of the conceptual model.

CHAPTER 4

STUDY SITES & DATA

4 STUDY SITES & DATA

4.1 INTRODUCTION

The development of a system that enables object detection and comparison relies on the use of appropriate input datasets. The model described in the previous chapter is designed to be generalisable across datasets. This work has focused on the implementation of the model on, and its application to, fine spatial resolution optical imagery, captured either from airborne or spaceborne sensors. This chapter describes the data used in the application of the model, as described in the subsequent chapters of this thesis. Table 4.1 compares the data used for the implementation of the model testing in the ensuing chapters.

Table 4.1. Comparison of data used in the application of the conceptual model

Study Site	Objects of interest	Geometric primitives	Data type	Spatial Resolution	Source
1. Simulated images	Circles	Circles	Simulated	n/a	Customised for purpose
2. Buncefield, UK	Oil storage drums	Circles, lines	Aerial imagery	25 cm 6.25 cm	GetMapping Cities Revealed
3. Saudi Arabia	Irrigation facilities	Circles	Screenshots of satellite sensor imagery	9 m	Google Earth

4.2 SITE 1: SIMULATED IMAGES

For testing of the circular Hough Transform algorithm (chapter 5), a series of synthetic images were customised for this purpose. The images contained features of a controlled size and shape.

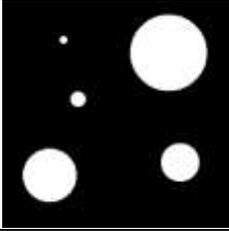
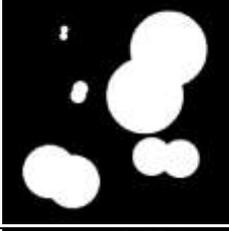
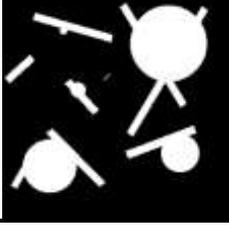
Twelve images were generated and are described in Table 4.2 based on three main image formats and containing a variety of circular features and additional obscuring elements. All images were 300×300 pixels in row \times column dimensions

and binary in nature with the pixels within each of the simulated circles assigned a value of 1 and a background value of 0. The images were converted to greyscale and were designed to vary in complexity to demonstrate the capacity of the algorithm to deal with a range of occlusions and image effects. A description of these synthetic images follows:

- Image 1 consists of a black background with five white circles distributed across the image. None of the circles are overlapping and circles have radii of 5, 10, 25, 35 and 50 pixels.
- Image 2 is of equal dimensions to image 1 and retains the centre locations and radii of the five circles. In addition, five further circles of equal radial values are added to overlap each of the original circles.
- Image 3 again retains the five original circles in addition to several straight lines at varying angles, partially obscuring the footprints of the original circles.

A range of image conditions was applied to the three image configurations, based on varying levels of smoothing applied to the images. A Gaussian smoothing algorithm was applied to each image, with a radius of one, three and five pixels for sub-image

Table 4.2. Simulated images used in chapter 5

Image-set	Preview (image <i>a</i>)	Description	Dimensions (pixels)	Circle Radii	Image conditions
1		5 circles only	300 × 300	5, 10, 25, 35, 50	1a. None
					1b. 1 pixel Gaussian
					1c. 3 pixel Gaussian
					1d. 5 pixel Gaussian
2		5 circles with 5 overlapping circles	300 × 300	5, 10, 25, 35, 50	2a. None
					2b. 1 pixel Gaussian
					2c. 3 pixel Gaussian
					2d. 5 pixel Gaussian
3		5 circles with added linear features	300 × 300	5, 10, 25, 35, 50	3a. None
					3b. 1 pixel Gaussian
					3c. 3 pixel Gaussian
					3d. 5 pixel Gaussian

b, *c* and *d* respectively, where sub-image *a* describes the unsmoothed image condition (Table 4.2). The Gaussian filter was used with varying radii to simulate progressively degraded edge information as the radius, *r*, is increased. However, it should be noted that techniques exist that specialise in smoothing images whilst retaining edge information. Many of these have been widely used to process Synthetic Aperture Radar images to reduce a phenomenon known as *speckle* (Lee, 1980; Frost *et al.*, 1982; Lopes *et al.*, 1990; Xiao *et al.*, 2003).

4.3 SITE 2: BUNCEFIELD OIL DEPOT, UK

4.3.1 Background

On 11th December 2005 an overflowing fuel tank at the Buncefield Oil Storage Depot, Hemel Hempstead, UK triggered an explosion that registered 2.4 on the Richter scale and caused the largest fire ever seen in peacetime Europe, devastating many of the depot's fuel tanks and buildings. The Buncefield site was chosen as a site on which to apply the conceptual model due to the simple circular shape of the oil tanks and the magnitude of spectral and geometric changes of the drums resulting from the incident.

4.3.2 Location of the Site

Buncefield is located at 51° 45' 55" N, 0° 25' 36" W. Figure 4.1a shows the location of the site, with (b) showing the extent of the smoke plume from the fire across much of southern England. Figure 4.2 details the layout of the site, showing the extent of burned area.

4.3.3 Imagery Details

Two images of the Buncefield site were purchased showing the scene before and after the incident, and are shown in Figure 4.3 and Figure 4.4 respectively.



Figure 4.1. Location of the Buncefield Oil Depot, Hemel Hempstead, UK.
 (a) Adapted from BMIIB (2008) (b) Extent of smoke plume in south of England (Wikipedia, 2009)

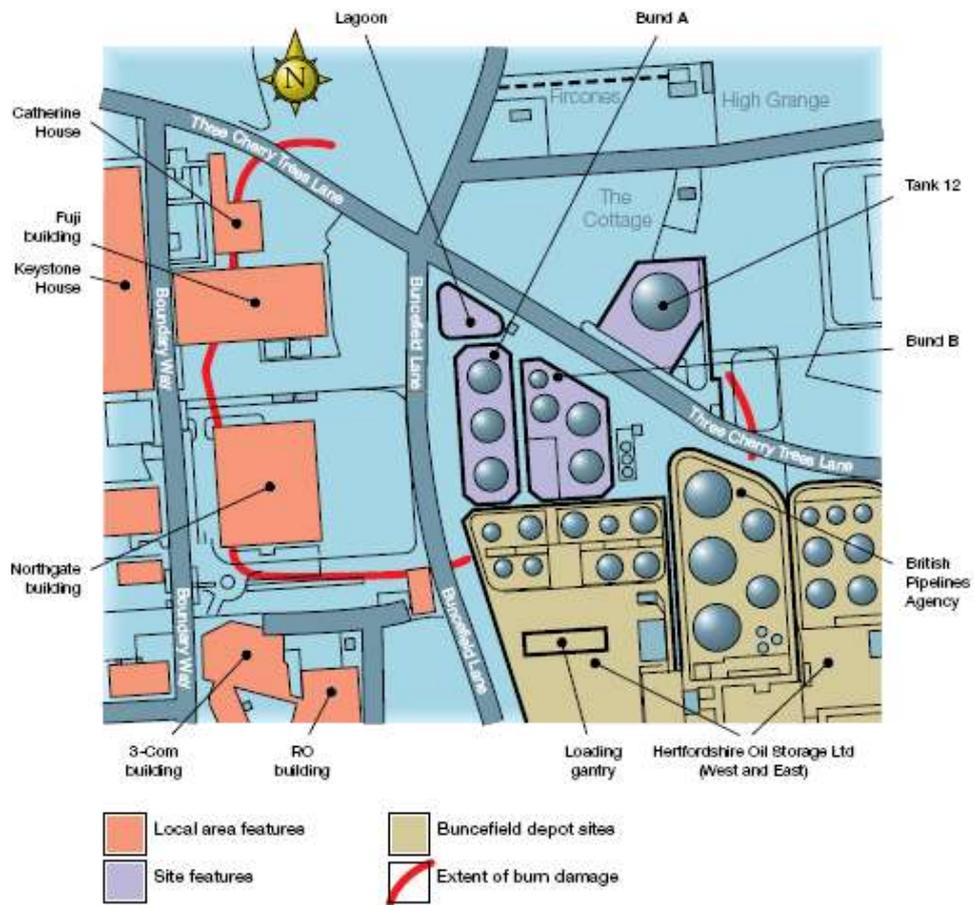


Figure 4.2. Detailed layout of the Buncefield Oil Depot Site (BMIIB, 2006)



Figure 4.3. Pre-event aerial image of the Buncefield site.
© 1999 GetMapping



Figure 4.4. Post-event aerial image of the Buncefield site
Cities Revealed aerial photography. © 2005 The GeoInformation Group

Pre-event Image

The pre-event image was sourced from GetMapping Ltd, a UK supplier of aerial photography. It was captured on 9th May 1999, six years and seven months before the explosion. It shows a region of 0.5 km × 0.5 km, at a spatial resolution of 25 cm. It was delivered in an 8-bit JPEG format, and provides RGB information only.

Post-event Image

The post-event image was sourced from Cities Revealed, at the GeoInformation Group, a UK supplier of aerial photography. It was captured on 28th January, 2006, seven weeks after the explosion. It shows an area of 0.7 km × 0.4 km, at a spatial resolution of 6.25 cm. It was delivered in an 8-bit JPEG format, and provides RGB information only.

Much of the Buncefield site suffered high levels of damage in the fire, and as a result, the objects of interest, the oil drums, have experienced significant geometric and spectral changes, as is apparent from a visual inspection of Figure 4.3 and Figure 4.4. The objects of interest in this image-set can be segmented into primitive linear and circular geometric shapes using appropriate algorithms, and the sensor properties of these images facilitate an investigation of the extraction of three-dimensional features.

4.4 SITE 3: DESERT IRRIGATION FACILITIES, SAUDI ARABIA

4.4.1 Background

To contrast the study area chosen for Site 2, two-dimensional image features were identified to test the Object Extraction module of the conceptual model. The study site chosen is an area of the Saudi Arabian desert incorporating several irrigation facilities, which form circular patterns when viewed from above.



Figure 4.5. Location of Site 3 in Saudi Arabia (marked by red cross).
©2009 Google Maps. Data ©2009 Europa Technologies

4.4.2 Location and Imagery Details

The imagery sourced for this study site consists of publically-available fine spatial resolution imagery from Google Earth of a small facility in the chosen region. Google Earth provides a cost-effective library of high-spatial resolution images, appropriate in this instance for algorithmic testing on multiple areas of interest. There are limitations on the use of Google Earth for commercial use, however, but for research purposes the use is allowed.

The chosen area is located at $25^{\circ} 02' 60''$ N, $45^{\circ} 34' 49''$ E, and the size of the image is $11 \text{ km} \times 7.8 \text{ km}$. The spatial resolution of the image has been estimated to be approximately 9 m per pixel, giving the image a similar resolution to SPOT HRV panchromatic imagery or imagery captured by other moderate spatial resolution sensors. The image is shown in Figure 4.6.

The predominant background of the image is sandy desert, while the objects of interest are scattered non-uniformly throughout the image. These irrigated areas contrast sharply with the surrounding desert land-cover, making the image an appropriate choice to demonstrate the method. The data sourced is of a different data type and format to Site 2 data, allowing testing of the model on a variety of datasets.

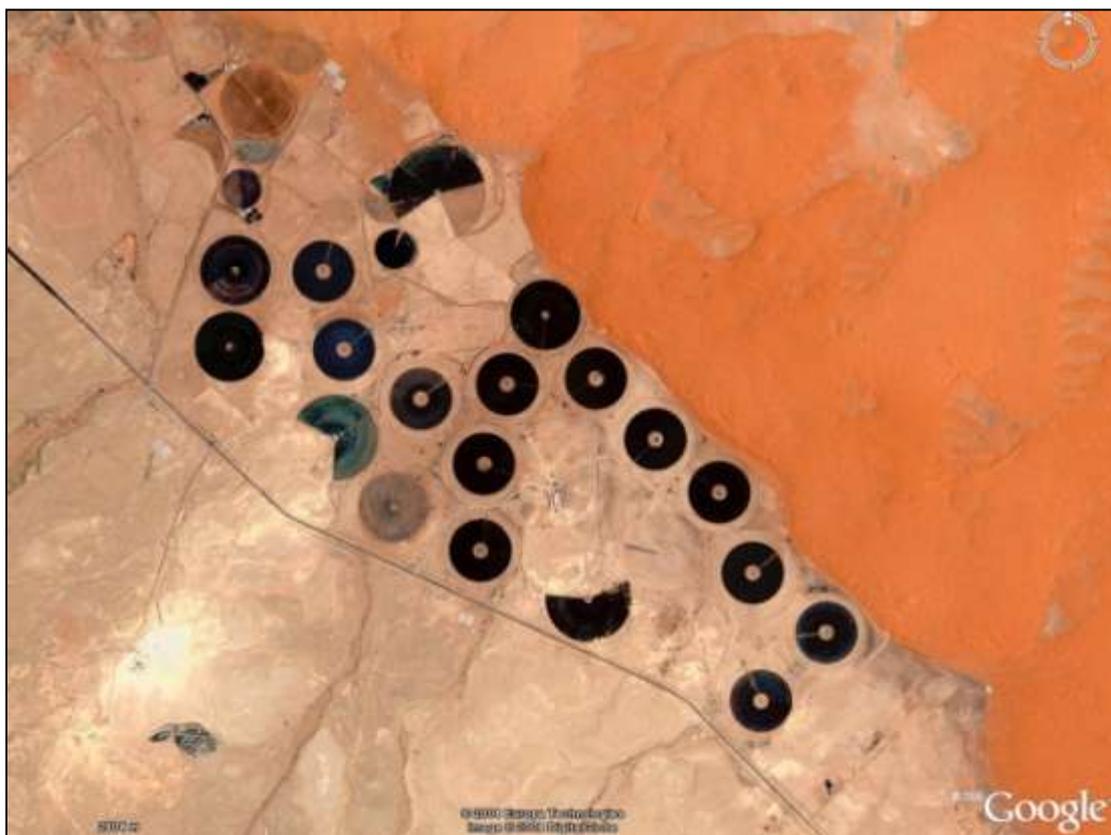


Figure 4.6. Image showing circular irrigation facilities in Saudi Arabia.
©2009 Google. Image ©2009 Digital Globe

CHAPTER 5
TWO-DIMENSIONAL OBJECT
EXTRACTION

5 TWO-DIMENSIONAL OBJECT EXTRACTION

5.1 INTRODUCTION

5.1.1 Background

Whilst providing a holistic vision for change detection, the conceptual model, outlined in chapter 3, was designed for implementation in several stages. The first major stage (object detection) focuses on the automatic identification and extraction of features from imagery into a scene model; an abstracted representation of objects of interest that can be used to qualitatively and quantitatively describe features within imagery. The development of this module required the integration and customisation of existing feature extraction techniques with the aim of detecting fundamental elements from an image; geometric primitives that can be combined to form and describe more complex shapes.

One type of two-dimensional primitive shape is the circle, which can be described by a single parameter: its radius. To describe these in spatial terms, their geographic x , y location would also be required. Because of its primitive nature, circle detection was chosen as the primary method used in the development and testing of the object detection module of the conceptual model. To achieve this, imagery containing prominent circular features were either created or obtained accordingly, as described in chapter 4. To comply with the ideals of the conceptual model, a widely used feature extraction method for detecting circular features in images was sourced; a circular Hough Transform (HT) algorithm.

This chapter describes the application of the HT algorithm to a variety of images and so describes the chronological development of the object detection module. Three main sections describe the advancement of the method and its application to a combination of custom images and images captured by remote sensors. The imagery used is described as follows:

-
- (i) Synthetic images containing controlled circular features
 - (ii) Remotely-sensed imagery of the Buncefield Oil Depot containing cylindrical oil storage tanks viewed from an oblique angle
 - (iii) Remotely-sensed imagery of circular irrigation sites in the Middle East viewed from near-nadir

The initial application of the model using the synthetic imagery was designed to provide a quality assessment of the algorithm in unaltered form (hereafter known as the *HT-*), applied to a variety of image conditions (section 5.2). The second phase of this work (5.3) introduced an extension to the HT (known as the *HT+*), which was applied to remotely-sensed aerial imagery. Finally, the *HT+* was further refined and applied to an additional remotely-sensed image containing circular features (5.4).

5.1.2 The Application Development Environment

The object detection module described in this chapter was developed using new algorithms, and extensions of existing code. MATLAB was used as the development environment and also chosen as the programming language, to allow the integration of the necessary datasets and algorithmic customisations for the development of the described method. MATLAB was chosen for its ability to combine matrix datasets and multi-dimensional arrays that can be used to represent multi-dimensional images, and also for its ability to visualise and interrogate results in an interactive manner.

5.1.3 Method – The Circular Hough Transform (HT-)

A HT algorithm was sourced from the MATLAB Central file exchange website (Mathworks, 2009), a MATLAB-hosted open repository for sharing user-created content and algorithms. The algorithmic methodology comprises five processes, described in Figure 2.1. Firstly, an edge detection algorithm is applied to a greyscale input image (of dimensions, x_{\max} , y_{\max}), resulting in an image showing the strength of edges in both the x and y directions across the image. The value of the gradient threshold (M^b) must be identified, based on acceptable edge strengths for the particular implementation of the HT. The remainder of the HT is initialised on a

thresholded edge strength image created from the combination of the x and y direction information. Each pixel from this edge image is processed in turn by the HT.

When the edge strength is greater than the gradient threshold, M^t , (e.g. when a dark pixel is adjacent to light one, in either the x , or y direction), it is mapped into a parameter space of the same dimensions as the image space (a_{\max}, b_{\max} , where $a_{\max} = x_{\max}$ and $b_{\max} = y_{\max}$). This pixel, which is located at x, y in the image space, is used to define a circle through a process of vote gathering, thresholding and radii estimation for each circle with radius r , that is located at cell (a, b) in the parameter space, so that:

$$(x - a)^2 + (y - b)^2 = r^2 \quad (5-1)$$

The HT algorithm is designed to detect non-concentric circles of multiple radii, and so the value of r is a variable in the range r_{\min} to r_{\max} . The parameter space can be thought of as 3-dimensional, with each pixel in the image space mapping to equivalent cells in the parameter space (Figure 2.1). This parameter space gathers evidence for circles incrementally for every possible value of r , based on the x, y location of the pixel in the edge image.

Accumulator Space

For each pixel in the image space, evidence is gathered in the form of votes stored in an array (the accumulator space) whereby each cell in the array accumulates votes for all possible circles passing through that particular point in the image space. As all of the edge strength pixels are processed, votes accumulate, and peaks form in the parameter space for those cells with the greatest number of votes. These peaks correspond to the image pixels with a high probability of being circle centres (Figure 2.3). Due to this process of vote accumulation, the parameter space is often known as the *accumulator space*, which will be the preferred term for the remainder of this thesis.

The accumulator space stores votes in cells in two dimensions (a, b) , for a range of radii of detected circles initialised as a range of values. For each edge pixel, votes are generated for every radius, r , where

$$r_{\min} < r < r_{\max} \quad (5-2)$$

The accumulation of HT votes is based on mapping a point location from an image to a locus of points in the accumulator space, where votes are cast in all cells within that circle. This process is described in Figure 2.2 and on page 38-39.

Once the voting process is completed for every image pixel, a local maxima filter is applied to the accumulator space to identify peaks corresponding to the candidate objects with the greatest probability of existing in the input image. This local maxima filter, L^1 , is parameterised by the user on initialisation of the HT. The identified peaks are used to calculate radius values for each candidate circle (centred at the image x, y coordinates corresponding to the cells (a, b) containing the maxima peaks in the accumulator space). Radii are calculated for all peaks in the accumulator space above this L^1 threshold.

Performance assessment

To assess the functionality of the algorithm, the detected circles can be compared to the original image and the resulting error information can be used as a performance indicator. The objective of this stage is to identify erroneous candidate circles; either *false positive* (where a circle is identified by the HT, that doesn't exist in the image in that location), or *false negative* (where a circle that exists in the input image is not identified by the HT) detections. To achieve this, a forward model can be used to create a difference image using a simulated image (\mathbf{s}) that shows the footprints of detected circles (Figure 5.1). The pixels in \mathbf{s} within the detected circles are assigned brightness values from the objects of interest from the input image (\mathbf{i}) and the whole image is compared to \mathbf{i} , so that

$$\mathbf{e} = \mathbf{s} - \mathbf{i} \quad (5-3)$$

The resulting error array (\mathbf{e}) contains values between -255 and 255 (for an 8-bit, greyscale image), where large negative values equate to false negative detections and large positive error values correspond to false positive detections. This error array can also be used to calculate per pixel errors, p , accumulated to describe the

overall performance of the algorithm. One such metric is the root mean square error, ρ (equation 3.2, p69).

Outputs

Radius values (r) for each detected circle are one of several outputs generated by the HT. Also produced are the edge strength properties of the image, data containing the centre points of candidate circles, the number of circles detected, and the accumulator space of voting results. All of these data are generated automatically by the HT code chosen and modified for this project. The circle centres and radii information can be used to reproduce a simulated scene model of these circles by plotting the detected circles and attributing all pixels within the circles with matching image characteristics from \mathbf{i} . This scene model can be subsequently compared to the input image and a variety of metrics calculated to quantify the performance of the algorithm. These calculations and additional data are generated in a series of MATLAB commands generated specifically for this project.

5.2 ALGORITHMIC TESTING: THE *HT*- APPLIED TO SITE 1: SIMULATED IMAGERY

To gain an understanding of the HT algorithm chosen for use in the object detection phase of the conceptual model, it was applied to a series of twelve synthetic images, customised for this purpose (described in section 4.2). The images contained features of a controlled size and shape, with a range of smoothing extents applied to the images. These bespoke images were used to assess the algorithm quality to help further development through the modification of individual components. This section describes the application of the unmodified *HT*- on the series of synthetic images.

5.2.1 Input Parameters

The following parameters were used to initialise into the *HT*- algorithm. The same parameters were retained for all applications using the synthetic images described in Table 4.2.

The range of radii was set using basic *a priori* knowledge of the fundamental size of objects in the image, and was set to between one and 50 pixels, in order to encompass all possible circle sizes in the image-sets. The M^t and L^t parameters, set to determine the gradient threshold and local maxima filter respectively, remained at their default values of $M^t = 5\%$ of the maximum brightness value, and $L^t = 8$ pixels.

Table 5.1. Input parameters used to test the HT algorithm on the synthetic datasets. The default parameters for the HT algorithm were chosen.

INPUT PARAMETERS
$1 < r < 50$ pixels
$M^t = 5\%$
$L^t = 8$

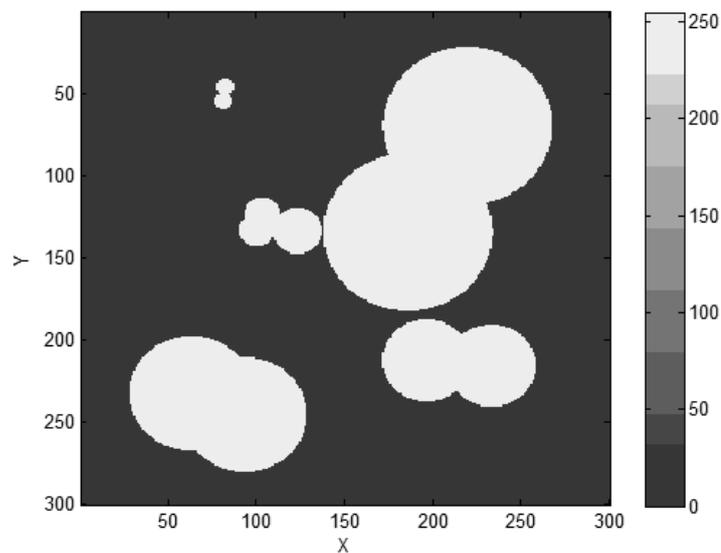


Figure 5.1. Greyscale simulated image generated for image 2b. Units are shown in greyscale brightness values

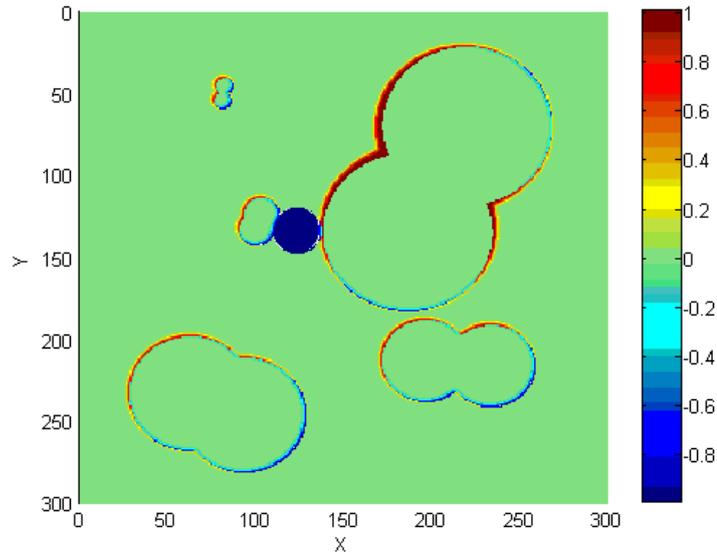


Figure 5.2. Error information generated for image 2b.

High positive values show the presence of false negative detections (i.e. not detecting a circle where one exists), and values <1 represent false positive detections (i.e. where a circle is detected that doesn't exist in the image).

5.2.2 Results

Several metrics were applied during the application of *HT-* to the three synthetic image-sets to measure the performance of the algorithm for each image condition. The number of detected circles, the number of correctly identified circles and a per-pixel RMSE for the whole image are described in Table 5.2 to Table 5.4.

An example of the resulting error array from (5-3) can be seen in Figure 5.2. The results from the *HT-* application on the twelve synthetic images are displayed in several formats in Figure 5.3 to Figure 5.5. Figure 5.3 shows the accumulator space with voting results for each of the sub images in image-set 1, with the statistical results for image-set 1 shown in Table 5.2. The results for image-set 2 take the same format, and are shown in Figure 5.4 and Table 5.3, and images-set 3 shown in Figure 5.5 and Table 5.4.

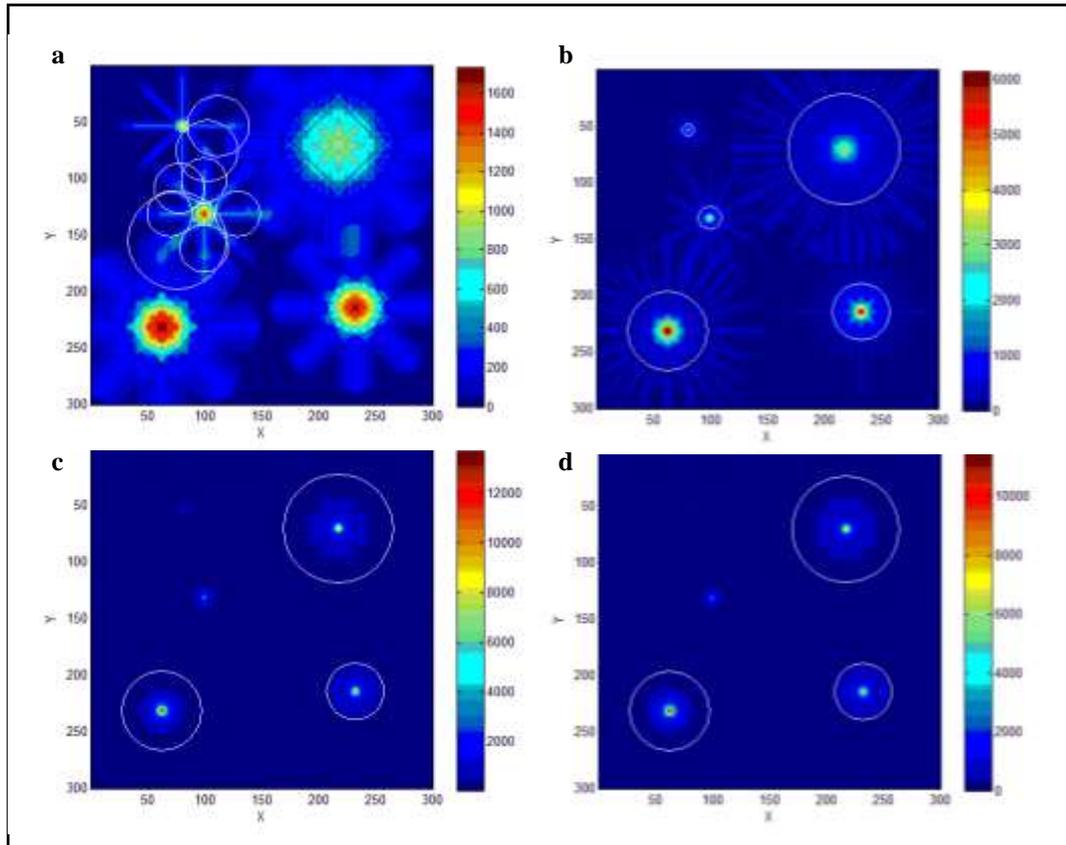


Figure 5.3. Accumulator space with detected circles overlaid for synthetic image-set 1.

With units in number of votes, dark blue pixels in these results represent low levels of scoring in the accumulator space, with dark red showing the highest levels. Scoring was more evenly spread around pixels in the unfiltered image (1a), resulting in many false positive circle detections. The filtered images saw votes increasingly cast in the locations of the circle centres, with the highest peak totalling close to 11,000 votes. The smoothing in 1c and 1d resulted in false negative detections of the smallest two circles, due to reduced definition of the object edges.

Table 5.2. Statistical results for image-set 1

IMAGE	NUMBER OF CIRCLES DETECTED		RMSE
	# DETECTED	CORRECT (IMAGE TOTAL)	
1a	10	2 (5)	123.76
1b	5	5 (5)	19.98
1c	3	3 (5)	26.49
1d	3	3 (5)	29.25

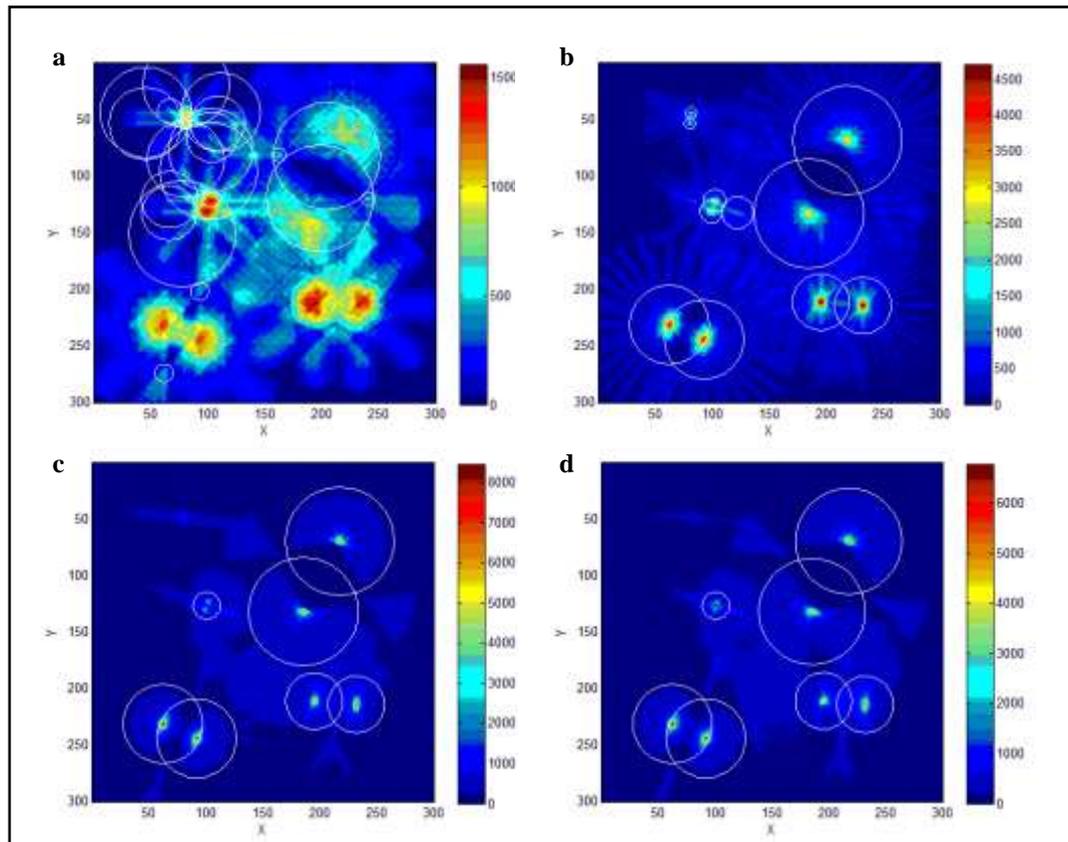


Figure 5.4. Accumulator space with detected circles overlaid for synthetic image-set 2

With units in number of votes, dark blue pixels represent low amounts of scoring in the accumulator space, with dark red showing the highest levels. Scoring was widespread across the unfiltered image (2a), and the accumulation of these votes meant there were many false positive circle detections, although the highest peaks occurred in the centres of the circles. The filtered images saw votes more precisely cast in the locations of the circle centres. The smallest circles (5 pixel radii) could not be detected in c or d, which also saw the two 10-pixel circles detected as a single object as the edge detail was degraded by the smoothing. However, circles with the largest radii were detected correctly.

Table 5.3. Statistical results for image-set 2

IMAGE	NUMBER OF CIRCLES DETECTED		RMSE	
	# DETECTED	CORRECT (IMAGE TOTAL)		INCORRECT
2a	22	2 (10)	20	148.37
2b	11	10 (10)	1	34.00
2c	7	6 (10)	1	29.22
2d	7	6 (10)	1	34.08

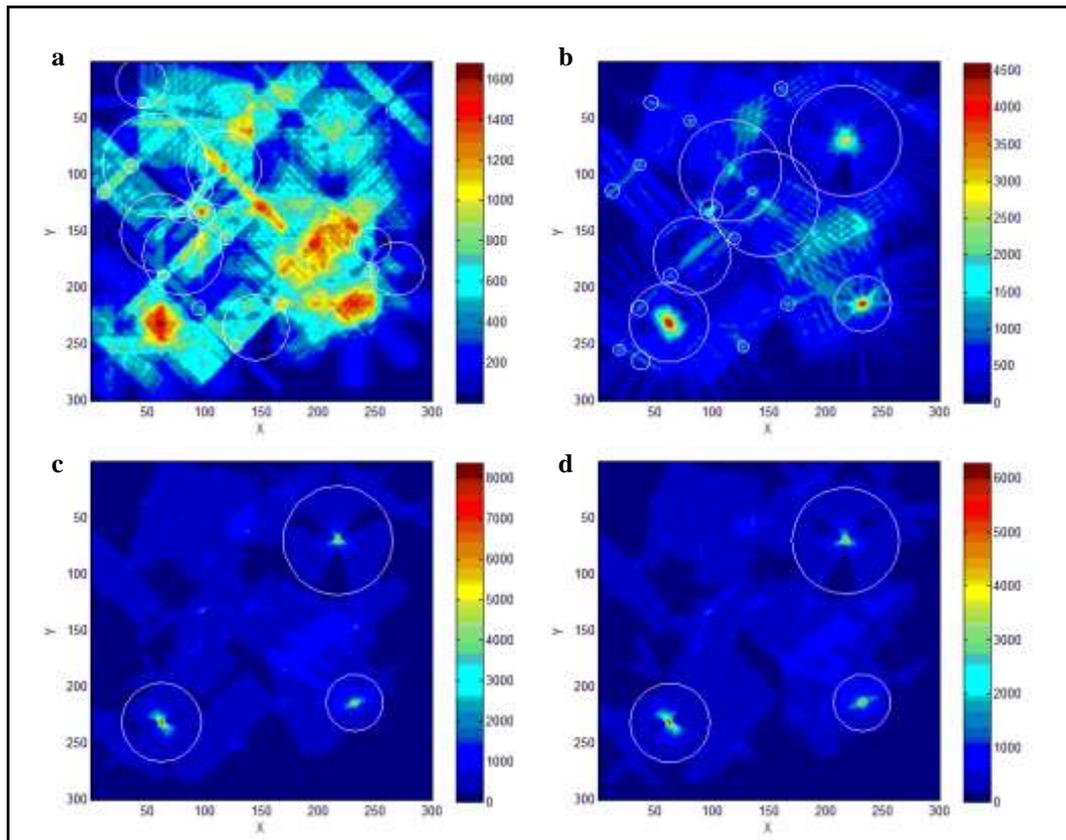


Figure 5.5. Accumulator space with detected circles overlaid for synthetic image-set 3

With units in number of votes, dark blue pixels represent low amounts of scoring in the accumulator space, with dark red showing the highest levels. The unfiltered image saw 21 incorrect circles detected as the edges of the linear occlusions were mistaken as circle edges. In 1b, the ends of the lines were falsely detected as circles due to the large amount of edge information in these regions of the image. The number of false positives was reduced in the filtered images, although the smallest features were progressively lost due to the reduction in edge definition.

Table 5.4. Statistical results for image-set 3

IMAGE	NUMBER OF CIRCLES DETECTED		RMSE	
	# DETECTED	CORRECT (IMAGE TOTAL)		INCORRECT
3a	21	0 (5)	21	153.60
3b	20	5 (5)	15	107.79
3c	3	3 (5)	0	57.20
3d	3	3 (5)	0	53.70

5.2.3 Discussion

Several trends are apparent from a visual comparison of Figure 5.3 to Figure 5.5. The complementary Table 5.2 to Table 5.4 describe a trend of greater numbers of false positive detections occurring when the *HT*- was applied to unmodified images rather than their respective 1, 3 or 5-pixel smoothed images. One probable explanation for this observed phenomenon is the issue of discretisation. Essentially, the *HT*- creates a series of circles represented as outputs as circle polygons. There is therefore an error inherent in fitting these outputs to a series of raster pixels from the input image. The spatial resolution of the raster image means that irrespective of how fine the spatial resolution (resulting in smaller and smaller pixels) it will never be able to fit exactly alongside a vector representation of a polygon feature. Figure 5.6 shows an exaggerated example of this. Evidence of this phenomenon can also be seen in the error information created for image 2b (Figure 5.4). Large positive error is evident in the pixels to the top left of the circle positions, with large negative error in the pixels on the diametric opposite sides of the circles.

One method of reducing the effect of discretisation is to smooth the images, resulting in reduced clarity of edges. When the *HT*- was applied to the smoothed synthetic images, the precision of the circle detection was increased. Lower RMSE values consistently resulted from images with at least 1-pixel smoothing when compared to the unmodified synthetic images. Application of the *HT*- to these images also resulted in much higher accumulator peaks (a four to five-fold increase), due to the increase in votes generated by the greater number of edge pixels that exist in the smoothed images.

When applied to a series of images with overlapping circles in image-set 2, the smoothed *HT*- was able to identify every circle within the image (Figure 5.4b).

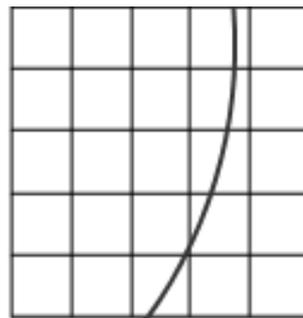


Figure 5.6. An exaggerated example of the discretisation problem. The curved line (vector) does not match with the underlying (raster) grid cells

This ability was reduced as the Gaussian smoothing window applied to the image increased in size (image 2c and 2d). Figure 5.4c showed a decreased RMSE when compared to Figure 5.4b. However, the circle detection in Figure 5.4c missed the two smallest circles (5 pixels radius), and one of the circles with a 10 pixel radius. The smoothing used in images 2c and 2d reduced the clarity of the features to such an extent that the circles with radii of 10 pixels suffered reduced edge definition and were detected by a single wide peak in the respective accumulator spaces. In contrast, results of the 1 pixel smoothing (image 2b; Figure 5.4), showed two thin peaks created for these circles.

The results of the *HT*- applied to image-set 3 (Figure 5.5) show that the amount of occlusion of the 5 and 10 pixel radii circles is detrimental to the detection of these circles in the smoothed images. The signature circular shape of the small circles is reduced by the addition of lines in the image, and eroded further by the smoothing. RMSE values were relatively low for images 3c and 3d, compared to images 3a and 3b. Yet this does not equate to correctly identified circles in these images. Rather, the algorithm failed to detect the two smallest circles that were severely obscured by the lines that were added to image set 3.

There is a strong positive correlation of $c = 0.938$ between the number of incorrectly identified circles in each image and the resultant RMSE. This highlights the strength of using RMSE as a performance metric. The false positives that were detected by the method for images 1a, 2a and 3a were reflected in RMSE values of up to six times greater than when there were zero false positive detections. Conversely, when false negative detections were seen when *HT*- was applied to images 1c, 1d, 3c and 3d, the RMSE remained low, due to the small radii of the missed circles. This is one reason why a smoothing technique cannot be recommended as a simple method of optimising the *HT*-, as it cannot assure the detection of all circular features within an image. A smoothing algorithm, while seemingly reducing some of the problems associated with the discretisation issue, reduces the clarity of the boundary signature of a feature, and has been seen to miss smaller or obscured features.

The application of the *HT*- to a series of synthetic images has exposed the algorithm's sensitivity to detecting false positive and false negative features within images. The remainder of this chapter focuses on the development of an extension to the *HT*- with the aim of improving the technique through a process of self-

regulation: an iterative extension is added to the algorithm to allow convergence on correctly identified objects in more complex imagery than the synthetic images. The extended algorithm is described as the “*HT+*” to distinguish the algorithm from the previous unmodified algorithm described in section 5.2.

5.3 APPLICATION TO SITE 2: BUNCEFIELD OIL DEPOT

5.3.1 Introduction

Objective number 4 described in Table 1.1 described the need to provide a level of automation in object detection methods to reduce the emphasis on human input during the feature extraction process. 4b suggested this could be achieved using an iterative updating method. The previous assessment of the *HT-* algorithm applied to the synthetic image-sets highlighted limitations in its application to unmodified images. The algorithm used spectral gradients to locate the edges of circles, with evidence for each candidate circle gathered in an accumulator space. The outline and centre of a circle were determined solely from evidence stored in the accumulator space. The extension of the algorithm to include a feedback loop designed to interact with the evidence gathered from one application of the *HT-* was aimed at reducing the incidence of false positive and false negative circle detections.

This section describes the extensions that were developed to advance the *HT-* algorithm by interacting automatically with the accumulator space to synthetically reduce the evidence of falsely detected circles, whilst adding additional evidence by increasing votes in the accumulator space for false negative circle detections (i.e. circles that have been missed during the detection process). The first application focuses on study site 2 described in detail in chapter 4. The preliminary focus is on an image subset of the Buncefield Oil Refinery, UK, with the aim to detect the tops of the oil drums in the image using the extended *HT+* algorithm.

The objective of this section is to describe the object detection phase of the conceptual model to demonstrate a novel algorithmic extension for object detection through the following processes:

- 1) Automatically detect the tops of the circular oil drums in the image
- 2) Update the accumulator space based on per-object error
- 3) Quantify the error associated with each circle using appropriate performance metrics calculated during each iteration of the algorithm

5.3.2 Method: The *HT+* extension

The implementation of the object detection module of the conceptual model follows the stages set out in section 2 of Figure 3.1. The *HT* circle detection algorithm (described earlier in this chapter; section 5.2) was used and further extended: The extended Hough Transform (*HT+*) algorithm builds on its predecessor by incorporating an error correction process, whereby each set of candidate circles detected by the *HT-* is compared to the observed image (using simulated brightness values). The resulting error information is used to update the accumulator space defined in the previous iteration and to generate a new set of candidate circles. This process is described in detail in the following stages:

Stage 1: Initial *HT-* run

The first stage involves one application of the unmodified *HT-* algorithm to the chosen input image. The outputs were the same as in previous applications of the method (section 5.2): circle radii, circle centre locations and the accumulator votes. These outputs form a scene model for candidate detected objects and are the basis for updating the accumulator space.

Stage 2: The Simulated Image

The initial phase of the extended *HT+* algorithm is the generation of a simulated image of the circles detected by the *HT-*. This simulated image, s , has dimensions equal to the input image and consists of two brightness values applied to within, and outside, the detected circles.

Stage 3: The Error Image

An error image, e_n , (where n = the iteration number), was used at each iteration to update the previous accumulator space by adding synthetic evidence to the original accumulator space. This was added in the form of simulated ‘voting’, where the

evidence in each cell of the accumulator space was automatically modified where error was high in the related cells of the simulated image, so that

$$\mathbf{e}_n = 1 + \left(\zeta \times \left(\frac{\mathbf{s}_n - \mathbf{i}}{i_{\max}} \right) \right) \quad (5-4)$$

where \mathbf{e}_n is generated from the comparison of the simulated image to the input image (Step 2) by subtracting brightness values of the observed image, \mathbf{i} , from the simulated brightness values, \mathbf{s} , and normalised by dividing by the maximum possible brightness value, i_{\max} . This results in a per-pixel error array of equal dimensions to the input image. The error image was modified further by adding a minor smoothing parameter, ζ , as a control to reduce the impact the error information has on the accumulator space at each iteration, avoiding excessive single changes and smoothing the transition that occurs iteratively between accumulator spaces. ζ is the proportion of the normalised error used to modify the accumulator space, and typically lies in the range [0:1]. Large values of ζ result in significant changes to the accumulator space, which will greatly affect the next iteration of the object detection. Smaller values, such as 0.2 were found to provide the appropriate scaling of the generated error information, resulting in subtle changes to the accumulator space which, over time, provides progressive modification of accumulated evidence.

Stage 4: Modification of the Accumulator Space

\mathbf{e}_n is used to modify the initial accumulator array, \mathbf{A}_0 , generated from the HT, creating a modified accumulator space, \mathbf{A}_n , where positive influence is given to correctly identified circles (defined by the error image), reducing the influence of false positive detections:

$$\mathbf{A}_n = \mathbf{e}_n (\mathbf{A}_{n-1}). \quad (5-5)$$

Each pixel within a false positive object has an error of $e < 0$. When combined with the original accumulator space, \mathbf{A}_{n-1} , the evidence within the accumulator space is reduced for these pixels.



Figure 5.7. Subset of the Buncefield image used as input to the model

Stage 5: Iteration of $HT+$

Using the modified accumulator space, \mathbf{A}_n , as an input, the local maxima filtering of the accumulator space is repeated iteratively, each time producing an updated accumulator score, \mathbf{A}_{n+1} , from which a new set of candidate circles can be identified. In short, the accumulator space from iteration n drives the circle definition for the following iteration ($n + 1$).

Stage 6: Generation of Statistics

At each iteration, it is necessary to assess the performance of the algorithm in terms of the new set of candidate circles that has been detected. Quantitative measures were generated at each stage to gauge (i) the RMSE (presented as RMSE, and also as a percentage of the range of brightness values), (ii) mean absolute error (MAE; equation (3-3) of the error image, \mathbf{e}_n , and (iii) correlation coefficient (c ; equation (3-4) of the simulated pixel values and the observed pixel values. RMSE was provided both in terms of grey-level values, and also expressed as a percentage of the range of possible brightness values. This percentage measure is used throughout the thesis, and was adopted for better understanding of the relative changes in RMS error. It should therefore be noted that all RMSE values were scaled equally, and so expressing as a percentage does not alter the error in any way.

5.3.3 Application to the Buncefield Imagery

In the Buncefield application, the *HT+* was applied to a RGB three-band image of dimensions 292×292 pixels (Figure 5.7). This image was generated from a subset of the 1999 Buncefield image (captured before the explosion and fire; Figure 4.3) containing three large white oil tanks as target features. It was converted to greyscale, and the *HT+* was applied for 200 iterations.

The forward model assigned a uniform brightness value of 240 to every pixel in the simulated image that was located within each circle detected by the *HT-*. This value was obtained from a manual calculation of the mean brightness value of the pixels representing the tops of an oil drum. Similarly, pixels found outside the circles were given a generic mean ‘background’ spectral brightness of 160 (calculated from the mean brightness value of all non-circle pixels).

The error information and accumulator values were created and modified using the parameters described in Table 5.5. A prior sensitivity analysis had shown that having a control of $\zeta = 0.2$ was a suitable rate of slowing the error reduction. The i_{\max} parameter was set to 255 as the 8-bit greyscale image, had 256 possible brightness values (0-255). A-priori information on the approximate size of the objects was used to set the range of r . The default HT parameters for M^t and L^t were retained. 200 iterations provided the algorithm with an appropriate length of time for error reduction.

Table 5.5. Input parameters used for application on the Buncefield image.

INPUT PARAMETERS

$$\zeta = 0.2$$

$$i_{\max} = 255$$

$$1 < r < 50 \text{ pixels}$$

$$M^t = 5\%$$

$$L^t = 8$$

$$\text{Iterations} = 200$$

5.3.4 Results

Figure 5.8 shows the input image with the circles detected and the accumulator space at iteration numbers 1, 50, 100 and 150. The numbers of detected circles for each iteration are plotted alongside the RMSE and correlation coefficients in Figure 5.9 and Figure 5.10 respectively. Additional statistical results for the aforementioned sampled iterations are shown in Table 5.5.

The first iteration of *HT*- identified seven circles (Figure 5.8a). Three circles roughly matched the tops of the oil tanks with a further four circles detected where no tanks existed (false positive detections). Figure 5.8c and Figure 5.9e show that results from later iterations displayed an oscillation of circle numbers, but the location of the detected circles became increasingly centred on the tops or sides of the drums in the initial image.

Secondly, from a quantitative analysis of the method, the initial accumulator space shows an abundance of votes (> 800) cast in areas of the accumulator space outside the boundary of detected circles, as represented by peaks of less than 800 votes in areas where no circle exists. A comparison with the modified accumulator spaces seen in later iterations (Figure 5.8f and h) reveals that the votes at these small peaks have been set to 0 whilst correct peaks have been made increasingly prominent where tanks exist. The largest peaks (> 800 votes) remained at the centres of the oil tanks.

Figure 5.9 shows the RMSE values and the number of circles detected at each iteration. It is apparent that RMSE values were reduced from $\rho = 20\%$ to $\rho = 10\%$ in the first 84 iterations. The error increased, however, after this point as the number of circles detected declined. There was an increase of 2% in RMSE at iteration 118, followed by a fluctuating increase and decrease of 1% as the number of circles changed by ± 1 for iterations > 125 .

The correlation coefficient (Figure 5.10) increased to 0.63 at 20 iterations despite the presence of more than three detected circles. The correlation coefficient between the simulated and input images oscillated with values remaining at 0.63 and 0.68 until iteration 118. Many iterations > 118 and < 200 detected no circles. This had a noticeable effect on the correlation coefficient values, resulting in $c = 0$. As no simulated images were created, due to the lack of detected circles, the correlation

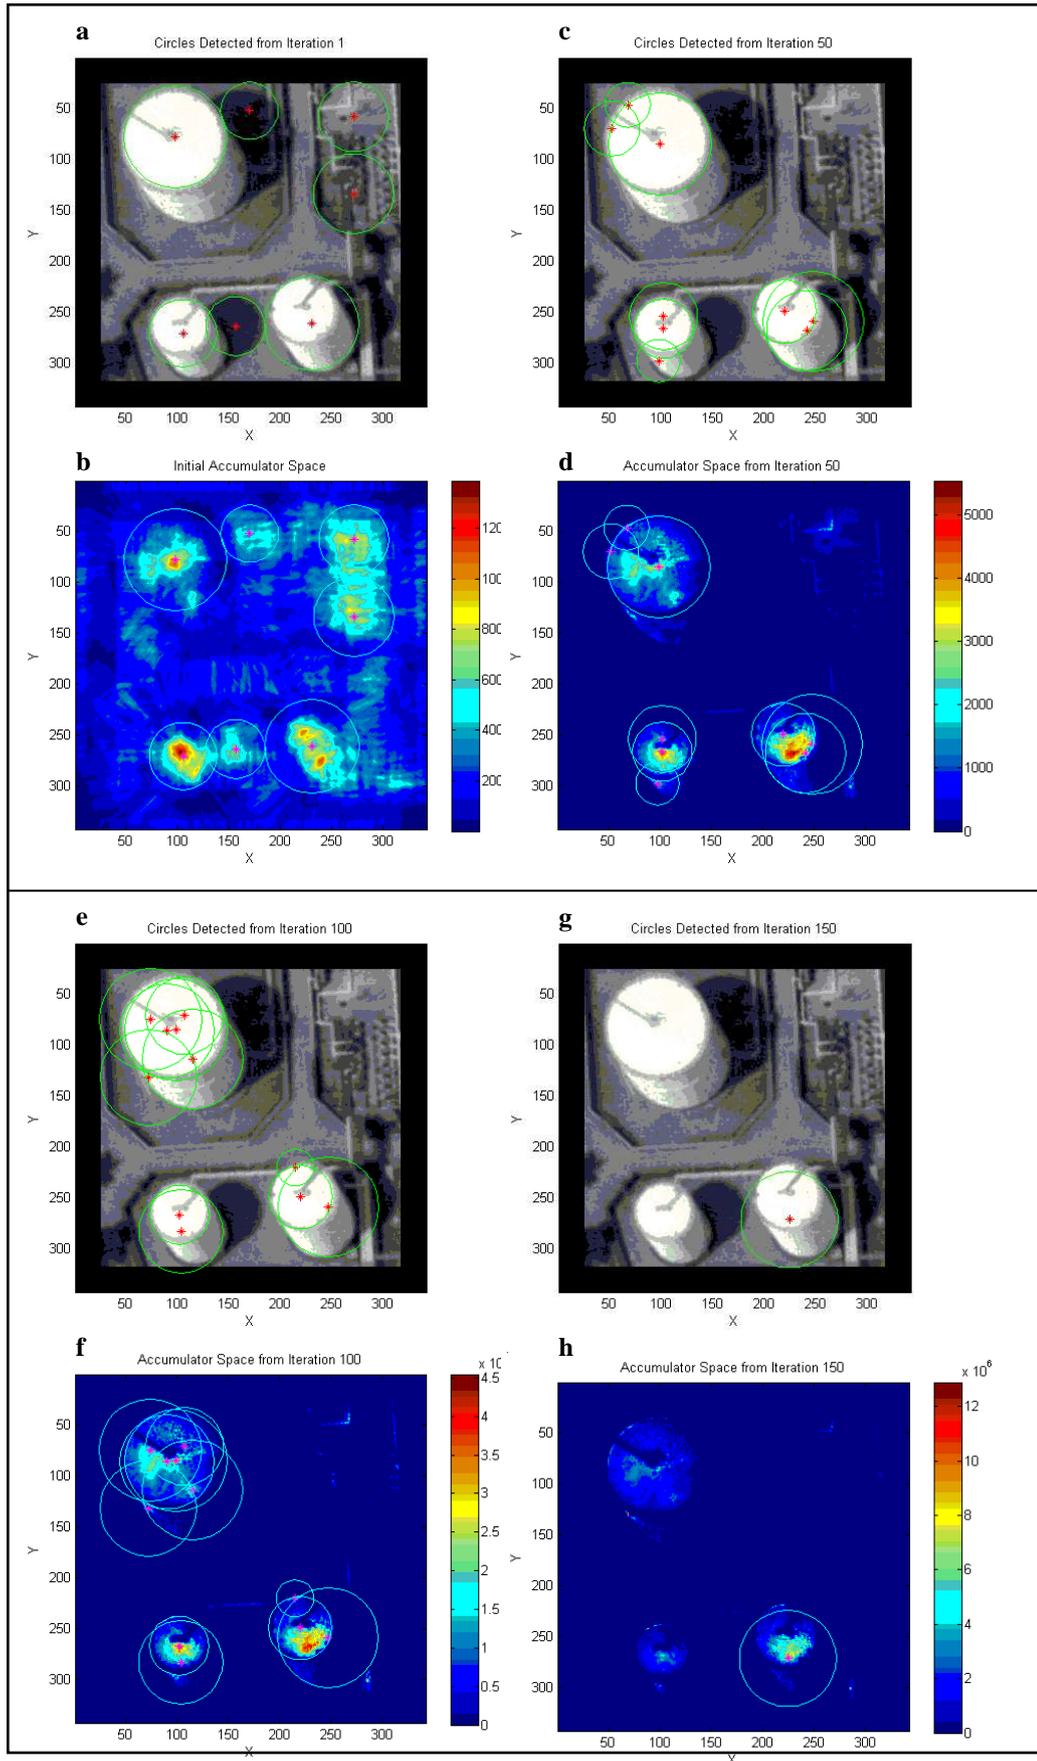


Figure 5.8. Results of 200 iterations of HT+ applied to Buncefield imagery. Results after 1 iteration (a&b), 50 (c&d), 100 (e&f) and 150 (g&h) iterations

between simulated and observed image could not be performed, hence the jagged fluctuations seen in Figure 5.10. Additional statistical results are displayed in Table 5.6, which shows a monotonic decrease in mean error.

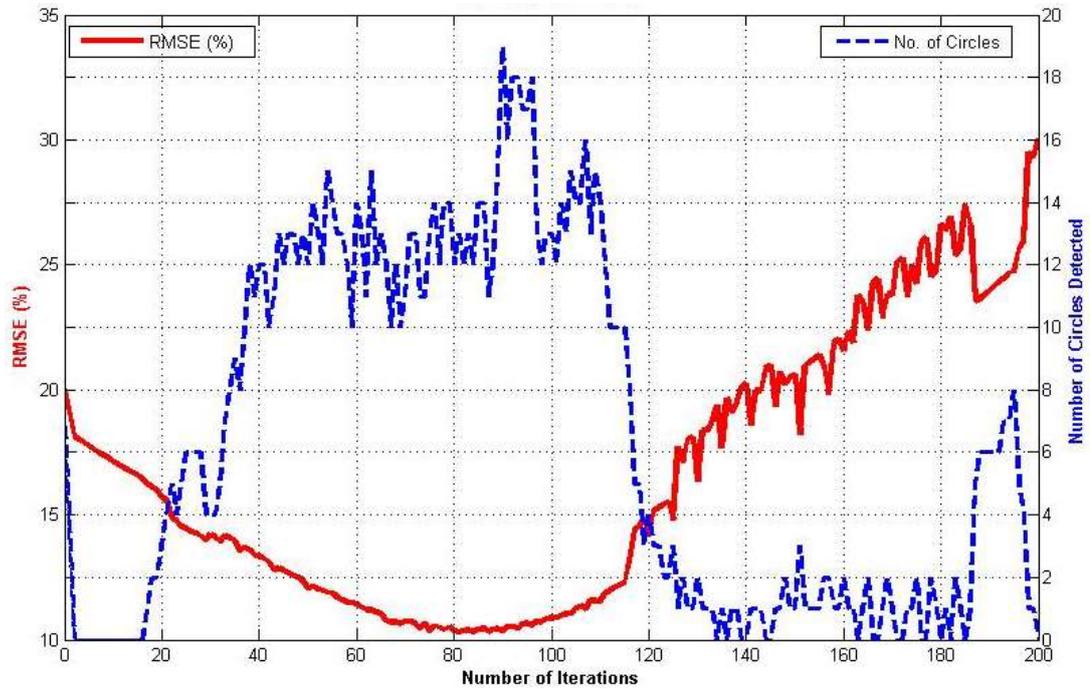


Figure 5.9. RMSE percentage & number of circles detected at each iteration

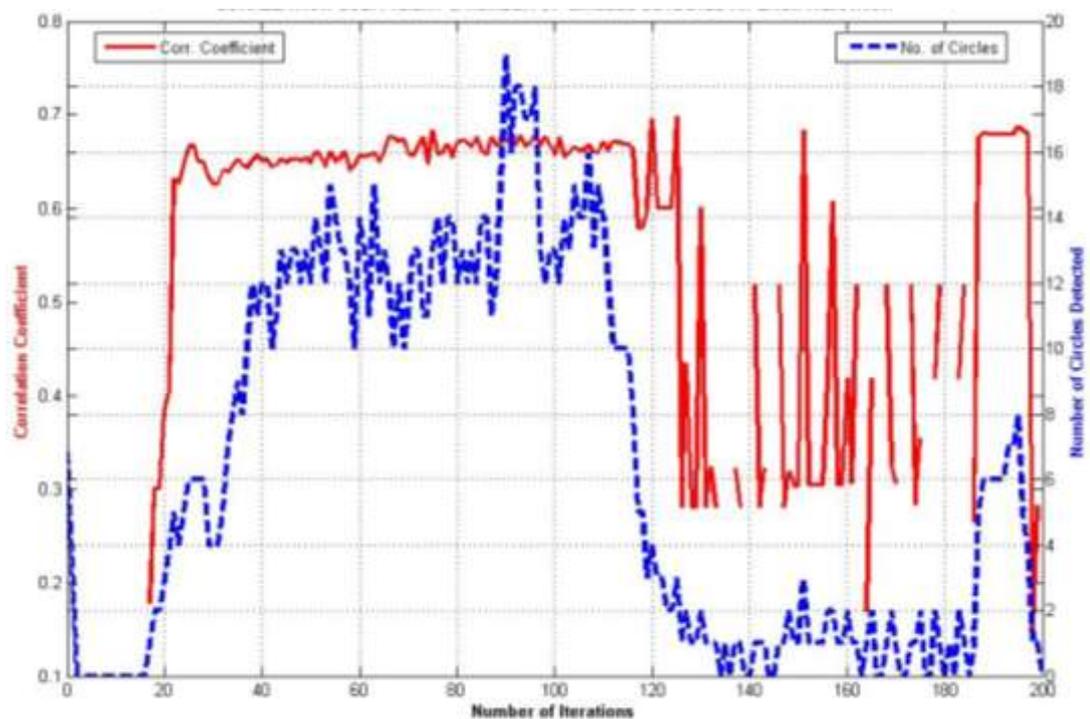


Figure 5.10. Correlation coefficient & no. circles detected at each iteration

Table 5.6. Additional statistical results from the 200 *HT+* iterations

ITERATION	# DETECTED CIRCLES	RMSE	RMSE (%)	MAE	CORRELATION COEFFICIENT
1	7	102.17	(20.03)	83.64	0.530
50	12	61.47	(12.05)	28.67	0.654
100	13	55.44	(10.87)	-15.66	0.666
150	1	105.18	(20.62)	-81.92	0.303
200	0	153.20	(30.04)	-136.68	Null

5.3.5 Discussion

The results of the object detection phase of the conceptual model applied to the Buncefield imagery demonstrated the advantage of an iterative modelling approach over the single use of an object detection algorithm. Application of the unmodified *HT-* in the first iteration generated several erroneous circles that were systematically dismissed as false positive results through modification of the accumulator space and iteration of the circle detection phase of the Hough Transform.

It was apparent that the model displayed the ability to accentuate accumulator values in areas where oil tanks existed in the image, whilst reducing the accumulator scores of falsely detected circles. This was achieved in an automated and unsupervised fashion, despite only having two input criteria, i) the need to detect circles, and ii) the approximate brightness values of the objects to detect, fulfilling much of the criteria described in objective number 4 (Table 1.1). The gradual manipulation of the accumulator space by creating and applying a weighted error image removed all false positive objects that were classified as a drum top that were observed to be shadow, or other background value. This was achieved without prior knowledge of the properties of such phenomena.

The application of the model to the chosen imagery identified several limitations that should be acknowledged. This implementation was constrained to objects with largely homogeneous spectral characteristics. The presence of multiple candidate circles detected adjacent to the three drums is due largely to the view angle at which the image was captured. Much of the south-eastern sides of the drums are visible and illuminated due to the Sun angle at the time of image capture, displaying large spectral brightness values (Figure 5.7). Between iterations 20 - 120

the reduction of the RMSE to below $\rho = 15\%$ (Figure 5.9), and the modest correlation (0.6; Figure 5.10) between the simulated and input images despite still detecting over 10 circles, is evidence of this view angle problem. This was due, in part, to false positive circles detected in areas that are overlapping the illuminated side of an oil drum and, consequently, the error levels remained low and correlation levels high despite over 10 circles being detected. Fundamentally, this was a limitation of the choice of imagery as there was no mechanism to automatically differentiate between the tops and sides of cylinders using brightness. As a result, the algorithm was prone to overestimation of footprints of the tops of the oil drums.

Examination of Figure 5.9 and Figure 5.10 shows a sharp decrease in detected circles from 10 to 3 at iteration 118. Simultaneously, the RMSE jumps from $R = 12$ to $\rho = 14\%$ and the correlation coefficient decreased from $c = 0.66$ to $c = 0.59$. The metric values continued to oscillate in irregular fashion until the end of the implementation as the numbers of detected circles fluctuates between zero and one. The total number of iterations was set to 200, in which time the algorithm reduced the total error to a minimum, before the described oscillations were observed. Multiple trials of the method showed that these oscillations began well before 200 iterations were reached, however this number of iterations was retained as the algorithm had not converged on an optimal solution of detected objects.

The method described in this section is based on uniform brightness properties of objects. Future modules will focus on geometric and statistical properties of objects which may mitigate such limitations. It is due to such issues that the RMSE values are greater than expected. An additional concern is due to the discrete nature of objects in the model. Iterations > 120 displayed few detected circles and the addition or subtraction of a whole circle at a time had a significant effect on the RMSE and residual error.

5.3.6 Summary of the HT+ Buncefield Exemplar

The application of the *HT+* phase has provided results showing the ability of the algorithmic extension to generate a scene model of objects from a remotely-sensed image. The iterative nature of the forward and inverse modelling technique reduced the number of false positive circle detections, reducing operator interaction with the model, and potentially reducing processing time of such algorithms.

The limitations of this initial application provided an opportunity to extend the model in two ways. The initial *HT+* algorithm was designed to detect two-dimensional objects in imagery, yet the choice of the Buncefield imagery with its oblique view angle added significant amounts of noise resulting in false positive detections. These were only “false positives” as the aim of the study was to solely detect tops of the oil drums.

To improve the effectiveness of the iterative processing of two-dimensional objects, the algorithm in its current form should be applied to imagery containing flat objects or on images with a sensor location close to nadir. The algorithm was further refined and tested on an image containing two-dimensional objects that did not display significant height characteristics. This application is described in the following section. In addition to this, the Buncefield imagery was revisited in order to extend the algorithm to detect three-dimensional cylindrical objects. This application is described in the following chapter.

5.4 APPLICATION TO SITE 3: CIRCULAR IRRIGATION FACILITIES, SAUDI ARABIA

5.4.1 Introduction

The previous section described the development of the extended *HT+* algorithm to detect tops of oil drums in remotely-sensed imagery. Limited in its ability to differentiate spectrally similar segments of the image, the application of the algorithm had limited success in the detection of the tops of cylindrical features in oblique-view imagery. Additional imagery was identified and used as a secondary test site for the two-dimensional object detection method, and is described in this section. This is an advancement of objective number **2**, as it further refined the method of object detection that is generalisable across multiple images and study areas.

Briefly, the HT algorithm was modified and extended to detect circular features in imagery of irrigation facilities in the desert of Saudi Arabia. The implementation of a single iteration of the chosen HT algorithm (*HT-*) is quantitatively compared to the use of an iterative accumulator *boosting* phase, the

HT+. The *HT+* algorithm used differs slightly from the Buncefield technique, with the differences outlined in section 5.4.3. The choice of input parameters is also outlined, with section 5.4.4 describing the results, and 5.4.5 describing a brief discussion. A short description of the data used in this implementation follows.

5.4.2 Study Site and Data

A more detailed description of the data used can be found in the description of site 2 data in chapter 4. The image used as an input can be seen in Figure 5.11 and shows several irrigation facilities in the Saudi Arabian desert. Its row \times column dimensions were 560×670 pixels, and, in accordance with the algorithm used, was firstly converted from RGB to greyscale brightness values before the algorithm was implemented.

5.4.3 Method

Following a review of the performance on the Buncefield data, several modifications were made to the algorithm. The initial modification was made to the edge detection phase. The HT algorithm used in sections 5.2 and 5.3 had used a MATLAB-defined *Gradient* function that used numerical gradients to gather edge information in the x and y directions across the input image. The previous edge detection method was replaced by a Canny edge detector for the Saudi implementation (described in detail



Figure 5.11. Subset of the Saudi Arabia image used as input to the model

in section 2.3.2) to advance the algorithm from the Buncefield implementation, to detect spectral gradients that might otherwise have been missed by the MATLAB gradient function used previously. The Canny method (Canny, 1986), was chosen over other edge detectors due to its ability to detect true weak edges using a dual-threshold approach. It is also the most powerful edge detector that is built into the MATLAB programming environment (Gonzalez *et al.*, 2004). The detector resulted in x and y direction edge strength information in the same format as the MATLAB Gradient function (Figure 5.12).

Input parameters

In general, the technique used in the Saudi Arabia application was similar to that used in the Buncefield example described in section 5.3. The input parameters were defined, with the range of possible circle radii constrained to be between 1 and 50 pixels. So that computation was not overly constrained, and a realistic range of radii values was maintained according to both the size of the objects of interest and the spatial resolution of the imagery used. The gradient threshold was set to 5% of the maximum brightness value in the image, with the local maxima filter retained at the default value of 8.

Table 5.7. Input parameters used for the Saudi image application

INPUT PARAMETERS
$1 < r < 50$ pixels
$M^t = 5\%$
$L^t = 8$
Iterations = 100

Algorithm application

The $HT+$ was applied for 100 iterations. During the first iteration, the Canny edge detector generated edge strength information (shown in Figure 5.12) from which the HT was able to generate voting information, stored in the accumulator space. The local maxima filter was applied to the accumulator space to identify the locations of the most probable candidate circles.

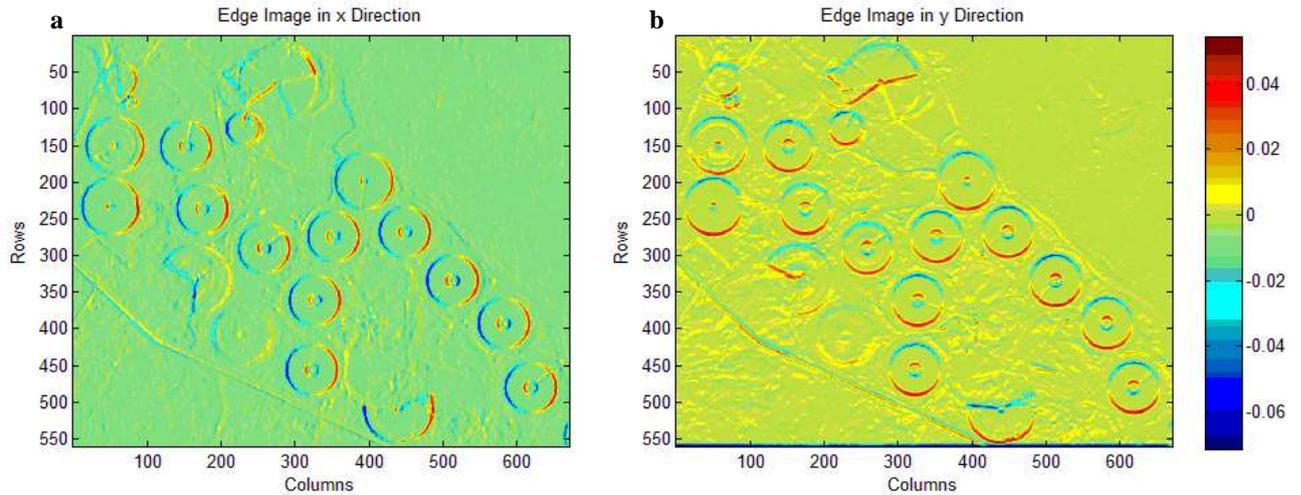


Figure 5.12. Results from the Canny edge detection in the x (a) and y (b) directions

Boosting Extension ($HT+$)

The simulated image brightness values were calculated by generating a vector of brightness values for the pixels within each detected circle, and taking the mode value of all entries in that vector. Conversely, the simulated brightness value for the background (every pixel not lying within a detected circle), was generated from the mean brightness value of every non-circle pixel. As described in section 3.3.2, the mode is chosen over the mean brightness value of the pixels to avoid using a skewed mean value arising from spectral anomalies, caused by mis-fitting of object boundaries, or false-positive candidate object detections.

For the first iteration, the brightness values were 5 for the intra-circle and 152.2 for the background values respectively. These values were used to create a simulated image (s) of equal dimensions to the input image (i), which was used to generate error information from the comparison of simulated and observed images (5-4). As with the Buncefield application, the control value was set at $\zeta = 0.2$ to restrict large changes in the accumulator space at each iteration. The local maxima filtering of the modified accumulator space \mathbf{A}_n (where n = the iteration number) was repeated iteratively, each time producing an updated accumulator score (\mathbf{A}_{n+1}), from which candidate circles were delineated.

5.4.4 Results

The performance over the first 100 iterations was assessed using RMSE, MAE and c , as before. Figure 5.12 shows the results of the Canny edge detection, with Figure 5.13 displaying the detected circles at various iterations. In this image, the circle centres and circle perimeters are overlaid on the input image with each circle allocated a number based on its location in the image to aid description in the discussion section that follows. Figure 5.14 and Figure 5.15 display the results of the RMSE and c for each iteration, with additional statistics described in Table 5.8.

Figure 5.13 describes the candidate circles for iterations 1, 25, 50, 75 and 100. It is apparent that the first iteration detected the darkest circles that were observed in the image, with the brighter circles (presumably disused irrigation facilities) not detected. Circle number 8 was the exception to this, with a detection of the inner, bright circle at the centre of a much larger irrigation facility characterised by a large dark circle. In the south of the image, there were two overlapping candidate circles (#13 and 15) on top of a segmented circle. After 25 iterations, this had been reduced to a single circle, closely matching the boundary of the incomplete circle. By this point in the iterations, the RMSE had increased from that seen in the first iteration, due in part to the false negative detection of the circle described as circle #3 in iteration 1.

By iteration 50, the RMSE had fallen well below the single-iteration level, mainly because of the correct detection of circle #3 and an expansion of the candidate circles overlapping the irrigation site centred at (258,290). The RMSE generally continued to fall as evidence was added to the centre of this circle, as is

Table 5.8. Additional statistical results from the 100 $HT+$ iterations

ITERATION	# DETECTED CIRCLES	RMSE	RMSE (%)	MAE	CORRELATION COEFFICIENT
1	18	39.0796	7.6627	-4.8291	0.7546
25	15	39.2489	7.6959	-4.4896	0.7490
50	17	38.1364	7.4777	-4.7064	0.7676
75	21	37.5728	7.3672	-4.7869	0.7764
100	26	38.3015	7.5101	-5.2659	0.7684

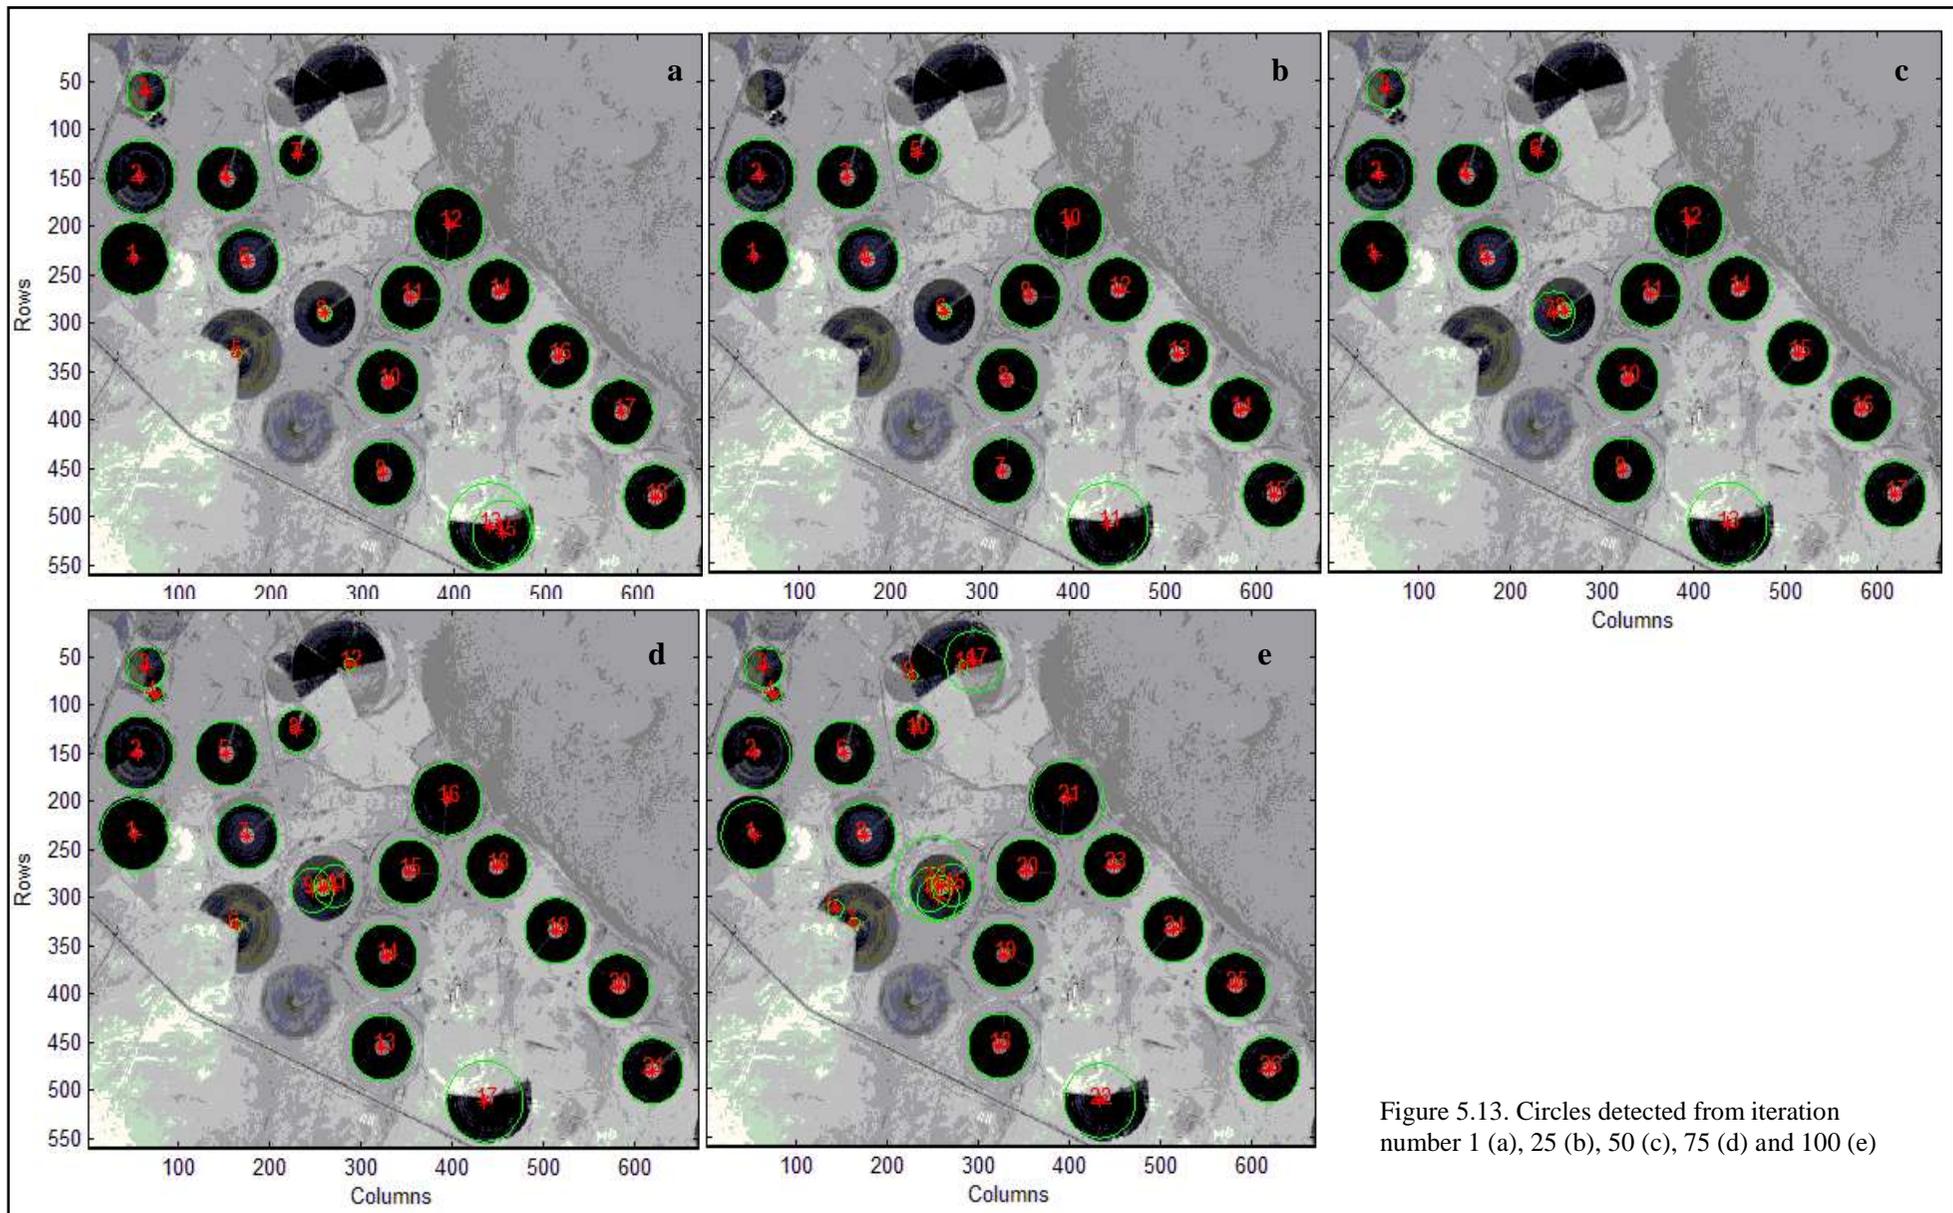


Figure 5.13. Circles detected from iteration number 1 (a), 25 (b), 50 (c), 75 (d) and 100 (e)

evident in Figure 5.13d. Combined with the additional detection (at 75 iterations) of circle #4, 6 and 12, the correlation between the simulated and observed images increased further by iteration 75 with this trend continuing before reaching a maximum value at iteration 97. Figure 5.13e shows the candidate circles after 100 iterations of the $HT+$, where there is evidence that the majority of the circles in the image have been detected to some degree. However, the precision of object

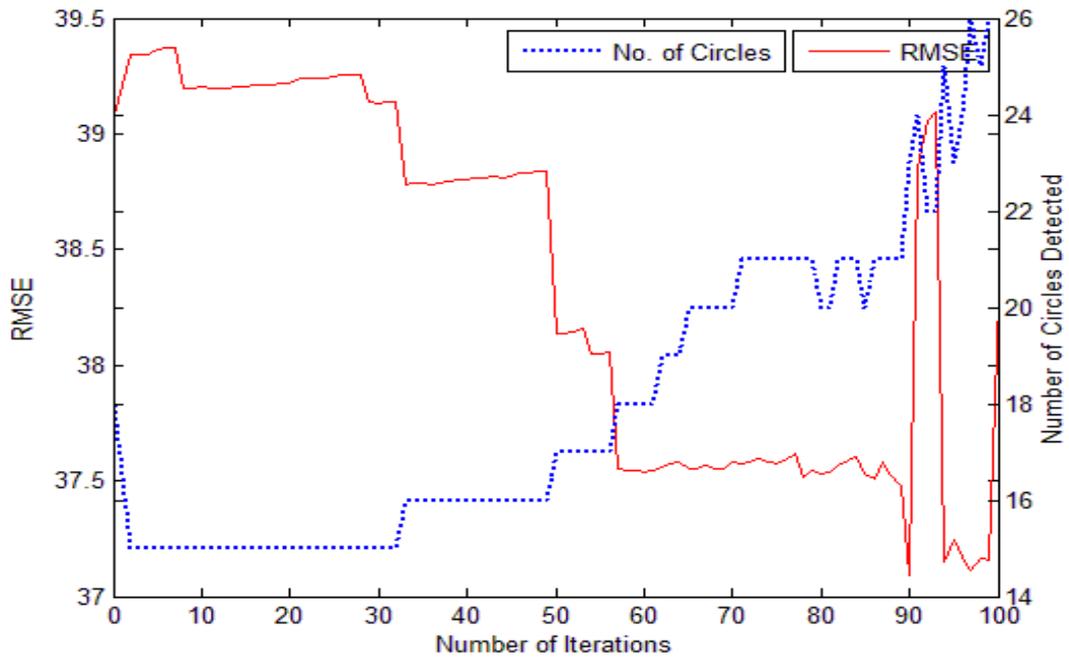


Figure 5.14. RMSE (grey-levels) & number of circles detected at each iteration

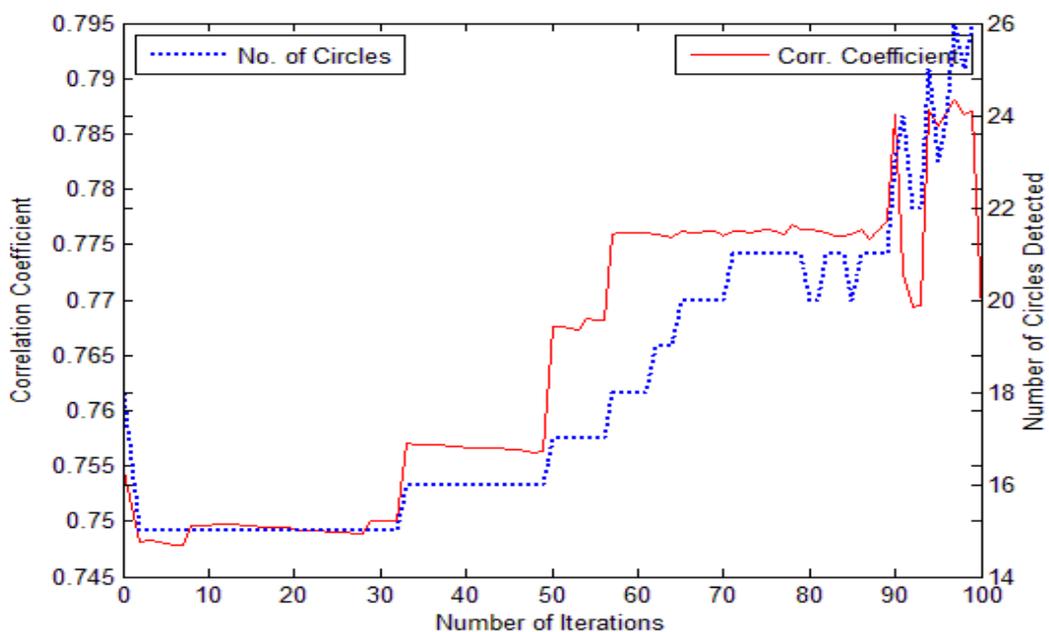


Figure 5.15. Correlation coefficient & number of circles detected at each iteration

delineation on many of the circles was reduced, as is apparent for circles 1 and 2. Combined with the undefined circle centred at (258, 290), the precision issue resulted in an increase of RMSE and a decrease in correlation for this iteration.

5.4.5 Discussion

100 iterations of the *HT+* algorithm resulted in an overall decrease in RMSE from the value obtained from a single iteration of the *HT-*. The process of adding additional evidence to the *HT-* resulted in detection of the majority of circles in the chosen imagery. The method was extended from the Buncefield example, allowing the *HT+* algorithm to automatically generate brightness values automatically from imagery, rather from a supervised, manual means. These values were used for the simulated image created from the circles detected in the previous iteration. This is thought to be one of the reasons for the lack of delineation of the brighter circles in the image, resulting in consistent delineation of the prominent dark circles.

There were several unique artefacts from the Saudi imagery that resulted in reduced performance of the *HT+*. Thirteen examples of the irrigation facilities were characterised by a small bright circle within a larger dark circle, sharing the same centre point. This concentricity gives these facilities characteristics more akin to the shape of a compact disc. In the case of circle #8 from Figure 5.13a, the object detection focused on the inner circle without detecting the more prominent larger feature. As a consequence, a large error was propagated through the simulated image into the accumulator space where votes were removed rather than boosted. Accumulator peaks were not formed in these instances, and therefore large circles were not delineated precisely. However, the boosting extension was seen to have been effective for the gradual modification of the accumulator space, allowing circles that were not detected from initial iterations to be identified in later iterations, as displayed by circles #4, 5, 7 and 9 in Figure 5.13e.

5.4.6 Summary of the *HT+* Saudi Arabia Exemplar

The imagery for Saudi Arabia was chosen in contrast to the Buncefield imagery to display the effectiveness of the *HT+* algorithm to detect two-dimensional image features. The irrigation facilities seen in the imagery were low to the ground and so the choice of data avoided the introduction of any 3-D elements into the scene, and as such, reduced the problems associated with oblique views of tall building structures, such as shadow effects. The algorithm detected fewer false positive circles after the first iteration, and was effective in detecting the centres of the false negatives through a process of iterative updating of the accumulator space.

The Saudi implementation showed that the *HT+* was capable of detecting dark image features in contrast to the spectrally bright features contained in the Buncefield imagery. The method also added a further element of automation (as outlined in objective number 4) with the unsupervised definition of the brightness values used in the forward modelling of the simulated image. This conceptual difference adds preliminary evidence to show the method could be made to be generalisable across image and feature types.

5.5 SUMMARY

The conceptual model described in Figure 3.1 has two central stages in its implementation – object detection and object comparison. This chapter has described the advancement of one possible method of object detection for the identification of circular features within aerial imagery. Using one of the key stages of the model, a widely used method of object detection from the field of computer science was identified, and open-source code for the algorithm was obtained. The chosen circular Hough Transform algorithm was tested in its raw form (*HT-*) on a series of three synthetic image-sets that were created solely for this purpose. These sets had a series of image distorting processes applied, with Section 5.2 describing the application of the HT, and section 5.3 describing an extension to *HT-* algorithm and a test on a more complex image acquired from an airborne sensor. The encouraging results from the Buncefield tests prompted additional testing on another

remotely-sensed image, this time acquired from Quickbird sensor satellite imagery from a publically available source.

The results from this chapter have highlighted the potential of the extended Hough Transform algorithm as a novel method of detection of circular objects that goes some way to fulfilling one of the key phases of the conceptual model. Evidence modification through an iterative forward and inverse modelling process showed potential for reducing false positive and false negative detection of circles. The application is suited for detection of two-dimensional circular image features, that is, objects that are close to the ground, or objects in imagery captured with the sensor directly overhead (nadir). The examples described have shown the potential of automatic evidence modification for detection of objects. However, the technique requires modification and expansion beyond circles before a practical use for real-world features can be recommended. These additional modifications are discussed in more detail in chapter 8.

The Buncefield application also identified a limitation in the algorithm's ability to identify features that have homogeneous spectral response with background features, and so could not effectively delineate the tops from the sides of the oil drums. This limitation was a key driver to further develop the method beyond detection of two-dimensional image features towards a technique that can identify three-dimensional objects from off-nadir imagery. With this aim, the *HT+* was further developed to iteratively identify 3-D objects in two-dimensional imagery. Its application on the Buncefield imagery is described in the following chapter.

CHAPTER 6
THREE-DIMENSIONAL OBJECT
EXTRACTION

6 THREE-DIMENSIONAL OBJECT EXTRACTION

6.1 INTRODUCTION

6.1.1 Background

The automatic creation of scene models from remotely-sensed imagery provides a possible basis for advancing current techniques of change detection for urban areas. As is evident in the empirical results of the previous chapter, it is possible to delineate two-dimensional shapes from images through an iterative forward and inverse modelling approach that utilises an existing low-level feature extraction technique, the Hough Transform. The HT was demonstrated to be effective for the delineation of objects that lacked the vertical height dimension. However, it was less effective when applied to imagery containing 3-D objects captured from oblique angles (i.e. when the sensor view angle is off-nadir).

This chapter describes an automatic method for generating scene models of cylinders in oblique aerial imagery. It builds on the Buncefield case study described in section 5.2, and further evolves the object detection phase of the conceptual model, from capturing two-dimensional object tops to modelling three-dimensional objects. The technique described focuses on the detection of 2-D primitives, and amalgamates these primitives to create a 3-D model.

6.1.2 Image Considerations

Imagery captured from an aerial sensor can be described as vertical or oblique, depending on the position of the sensor relative to the scene it captures. The angle between the sensor and the scene is known as the observational zenith angle (Z_s ; Figure 6.1). When $Z_s = 0$, the sensor is said to be nadir to the image, creating a vertical image. The Buncefield images described in Figure 4.3 and Figure 4.4 were captured from aerial sensors, and therefore are subject to several distortions resulting

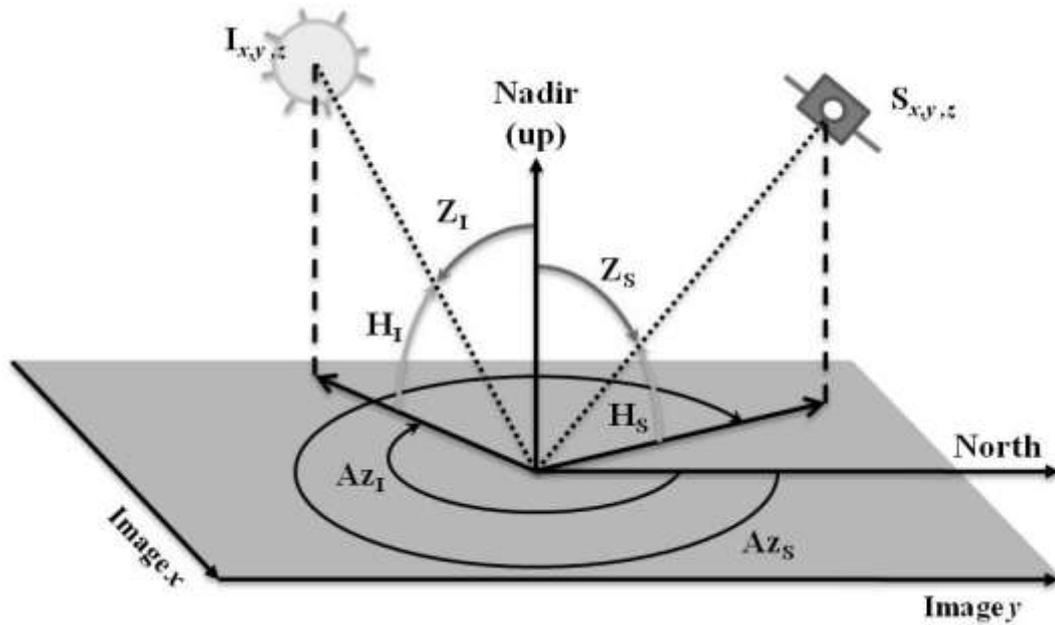


Figure 6.1. Image, sensor and illumination geometry.

Az = azimuth angle, H = elevation angle, Z = zenith angle for illumination (I) and sensor (S)

from the respective angle of the sensor to the image when captured. This section describes several of these considerations.

Raw oblique-view imagery has several associated distortions that should be acknowledged. The scale of features in this type of imagery differs throughout the image, with the features in the foreground appearing larger than objects towards the top of the image. The scale is only constant along lines that are parallel with the horizon, but scale will vary for all other lines in the image, with features appearing progressively larger the closer they are to the sensor (Moffitt, 1970).

The effect of terrain on the appearance of features in imagery should also be understood when using aerial imagery, especially if viewed from the perspective projection. Figure 6.2 displays the comparative geometry of a map with orthographic projection (a) and of an image captured from a perspective view (b) (Lillesand *et al.*, 2008). The map (a) shows the true positions of features, relative horizontally to each other in two dimensional space. In the photograph (b), the features that have a higher elevation are closer to the sensor, therefore appear larger than features on lower lying land. From this perspective view (b), the objects in the imagery will always appear to be tilting – the tops of the objects will be displaced horizontally in the image space away from the object bases.

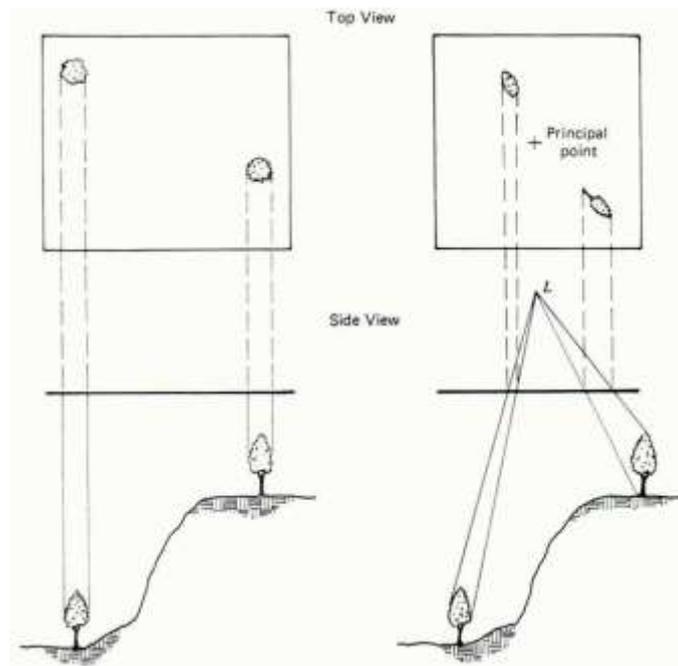


Figure 6.2. Comparative geometry of a map (orthographic projection: a), and vertical aerial photograph (perspective projection: b). Note the differences in size, shape and location of the objects in the image. (from Lillesand *et al.*, 2008)

The effect of relief displacement can be reduced through orthographically “correcting” the image to create orthophotographs. This digital process combines images with elevation information from a digital elevation model (DEM) to remove the distortions shown in Figure 6.2 – specifically scale and feature displacement. These processed photographs project the image from a perspective view into orthographic projection (Moffitt, 1970). Examples of how common image features differ when displayed using perspective and orthogonal projections can be seen in Figure 6.3 – features in orthophotographs would take the form of (b), with the perspective view provided in (a). In this diagram, there is a clear offset in object tops compared to their respective bases.

This chapter describes a method for extracting three-dimensional structures from a single remotely-sensed image. There are, however, several established methods for identifying and extracting 3-D objects from imagery, most notably in the field of photogrammetry. Photogrammetry is the study of measuring spatial data from photographs, such as distance, area, or deriving 3-D information from imagery, both in digital and softcopy form. For decades, scientists have used pairs of 2-dimensional stereoscopic images with overlapping areas for the derivation of 3-dimensional information. Originally an analogue process using equipment such as

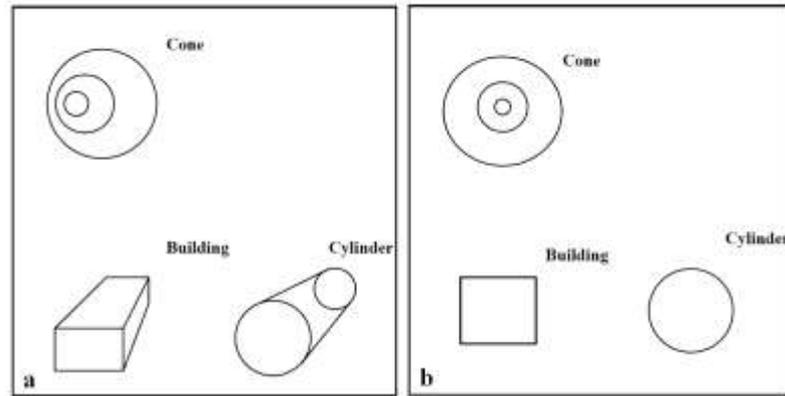


Figure 6.3. Examples of how common shapes differ in appearance when displayed in perspective projection (a) and orthogonal projection (b)

stereoscopes, stereocomparators and stereoplanigraphs, the process has largely migrated to digital format, with various systems affording stereo viewing using specialist terminals and goggles.

Photogrammetric techniques have been widely used for the extraction of land surface information, used in the creation of DEMs, and in the extraction of buildings and urban detail from imagery (Nevatia *et al.*, 1997; Shi *et al.*, 1997; Gruen, 1998; Zhang *et al.*, 2001). Height information can be derived from both stereo image pairs and single images through the geometric relationship between sensor location and the displacement of objects in oblique images, and also through the change in relative position of the same object viewed from 2 different angles (Moffitt, 1970; Lillesand *et al.*, 2008). Therefore, image processing can utilise geometric information captured from the sensor position, its location and the time the image was captured. These data can be used to constrain analysis of imagery sets such as single images, or stereo pairs.

The remainder of this chapter focuses on the detection of three-dimensional cylindrical objects from a single oblique-view aerial photograph. The described study uses imagery with a perspective projection, and so has not been orthorectified. The method therefore is designed to identify features based on their primitive components, before amalgamating them in an intelligent manner using filters and similarity to reduce the large number of possible configurations of primitives. Information relating to the sensor geometry and a series of similarity measures are used to constrain a large search space that is present when matching detected objects.

6.2 STUDY SITE AND DATA

6.2.1 Choice of Input Imagery

The Buncefield imagery was used for the development, testing and implementation of the 3-dimensional object detection algorithm. This imagery is described in detail in chapter 4, and Figure 5.7 shows the image subset used for these purposes.

6.2.2 Image Metadata Preparation

Before any analysis could take place, the geometric characteristics of the image were assessed. Description of the characteristics of the image, including its illumination and viewing geometry provided contextual information and combining these data with the extracted objects results in a more complete and descriptive scene model.

Sensor properties

The x , y , z location information of both the sensor and the image were provided within the imagery metadata. A series of simple trigonometric calculations were required to determine both the sensor azimuth angle and the observational zenith angle with respect to the image. Initially, the distance between the centre point of the image and the sensor was calculated as the hypotenuse of a right angled triangle between the x , y locations the image centre and point $S_{x,y}$ (the geographic x , y coordinates of the sensor), using Pythagoras' Theorem. The interior angles were subsequently calculated. These calculations were repeated to generate the observational zenith angle (Figure 6.1).

Illumination properties

The image metadata also included the date and time of the image capture. These data were fed into an online tool³ to calculate the solar zenith angle (Z_I) and the solar azimuth angle (Az_I), measured clockwise from the north of the image.

³ <http://www.largeformatphotography.info/Sunmooncalc>

6.3 METHOD

A multi-stage approach was designed and implemented for the automatic detection of three-dimensional objects from the Buncefield imagery. A workflow is shown in Figure 6.5 and the remainder of the methods section is divided into sub-sections accordingly.

6.3.1 Object Detection

Two main strategies were considered for cylindrical feature extraction: (i) a holistic approach, whereby the cylinder is detected in one step (Figure 6.4); or (ii) to reduce the cylinder shape into geometric primitives (one and two-dimensional shapes such as lines, circles and squares in Figure 6.4). The second of these two approaches was chosen for the following reasons. In accordance with the conceptual model ideals as described in chapter 3, the method should aim to be generalisable across images and have the ability to detect a variety of objects irrespective of their shape, hence, reducing reliance on *a priori* knowledge of object properties. Existing low-level feature extraction techniques can be used for efficient extraction of these geometric primitives from remotely-sensed imagery. Introducing a more complex shape (i.e. a bean-shape in this instance; Figure 6.4a), would make parameterisation increasingly complicated, for example, buildings with irregular shapes would require the supervised creation of template shapes to match the target objects. Such an approach would be *ad hoc* to the choice of image used, and, hence, does not generalise well. In contrast, having a pre-defined set of primitives to detect, followed by an amalgamation process to create whole objects, is more replicable across differing images and objects, and generalises well. This primitive amalgamation approach is also generalisable to images with sensor azimuth angles closer to nadir. In such instances, additional techniques (such as using shadow information, or photogrammetric techniques) would be required to calculate building height when moving from a 2-D to a 3-D scene model.

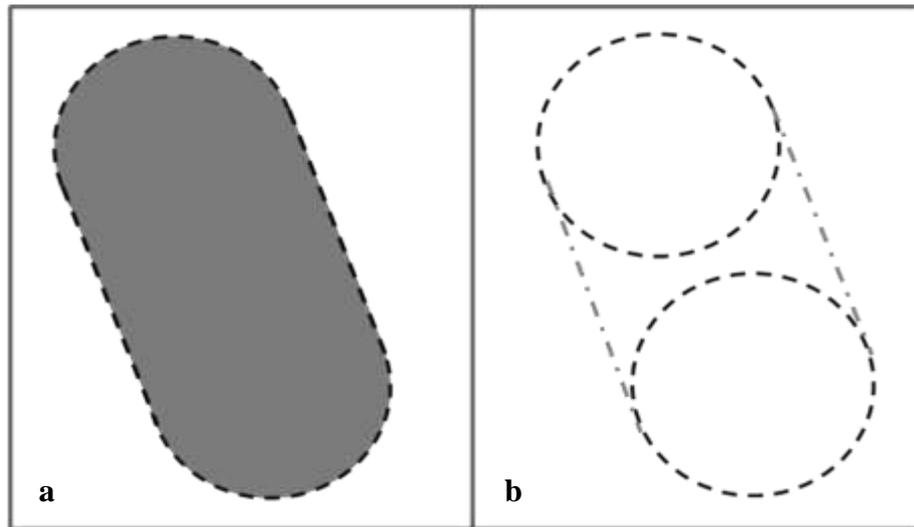


Figure 6.4. Two methods of automatic cylinder detection from perspective imagery.
a) detection of whole feature at once; b) detection of primitives – two circles combined with two lines

An additional processing stage was added to the conceptual model to generate a three-dimensional scene model of objects. The view of the buildings is affected by the characteristics of the image and sensor (Figure 6.1) and the orientation of objects in oblique aerial imagery is both a function of the sensor azimuth and observational zenith angles. The geometric properties of the input image can therefore be used to constrain the search space for possible objects to a manageable level. This information can also be used to determine the relative location of the tops and bases of objects: The tops of objects, when viewed in oblique imagery, will always appear to be further away from the sensor than their respective bases. The remainder of this chapter describes the application of an automatic method of detecting 3-D cylinder information from 2-D aerial imagery acquired with an oblique observational zenith angle.

6.3.2 The Hough Transform

The Hough Transform algorithm described in chapter 5 was used as primary method in the workflow. The improved *HT-* algorithm (containing the Canny edge detector) from section 5.4 was implemented for one iteration, without the forward modelling and accumulator space modification from the *HT+*. The *HT-* was initialised with the parameter values described in Table 6.1.

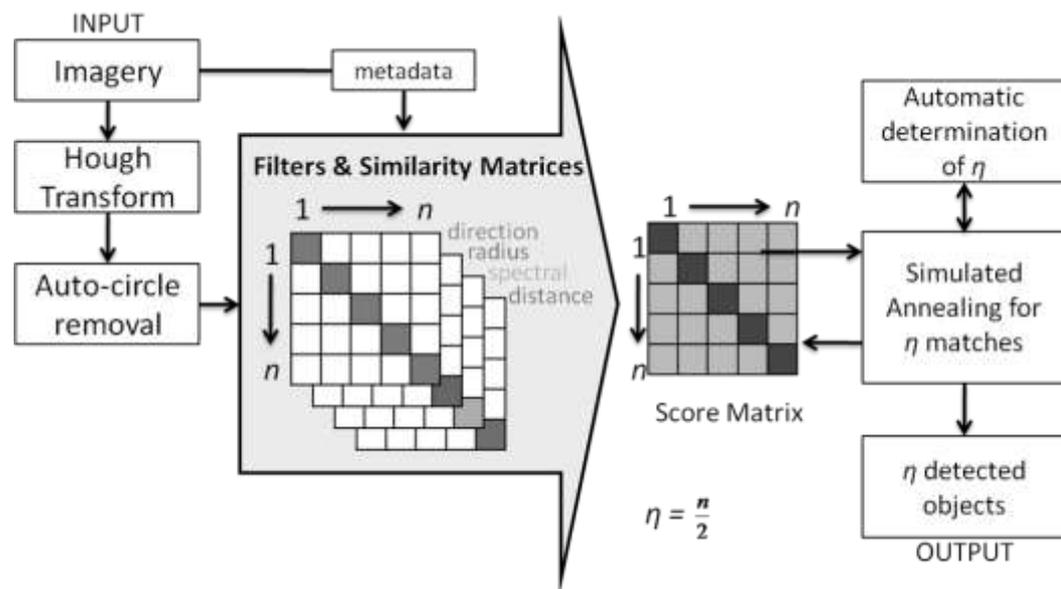


Figure 6.5. Workflow for automatic cylinder detection

Table 6.1. Parameters used in the *HT*- implementation

$1 < r < 50$ pixels
$M^t = 5\%$
$L^t = 10$
Iterations = 1

6.3.3 Automatic Circle Removal

The *HT*- was initialised with the parameters described in Table 6.1 to detect as many candidate circles as possible, including erroneous or false positive detections. Any erroneous detections were to be removed later in the application using filters, similarity matrices and simulated annealing optimisation processes. An additional assumption was made that the image would contain the full extent of each object to be detected. Therefore, a rule was created to exclude all circles detected that were not contained within the bounds of the image (an example of this rule can be seen in Figure 6.6). This was achieved by removing all circles that do not fit the following criteria:

$$\begin{aligned}
 &x_i + r_i < x_{\max} \text{ OR} \\
 &x_i - r_i > x_{\min} \text{ OR} \\
 &y_i + r_i < y_{\max} \text{ OR} \\
 &y_i - r_i > y_{\min}
 \end{aligned}
 \tag{6-1}$$

where (x,y) represents the pixel location of the circle centre, r represents the circle radius, i is the circle index, with $x_{\min}, y_{\min}; x_{\max}, y_{\max}$ representing the upper and lower bounds of the image respectively (Figure 6.6).

Following the removal process, each remaining candidate circle was compared with all other circles using a series of filters developed to compare the similarity of all circles to create a representation of error associated with each potential match. This was required because the algorithm was designed to create a series of cylindrical objects by automatically matching a top and base detected circle. This was achieved by the generation of several measures of circle-to-circle similarity stored in a range of matrices, described in the following section.

6.3.4 Similarity Matrices & Filters

Matrices were devised to store the result of similarity filters for each circle that was detected by the *HT*-. Each matrix had $n \times n$ dimensions where n was the number of circles detected by the *HT*-. Each circle was compared to every other circle and

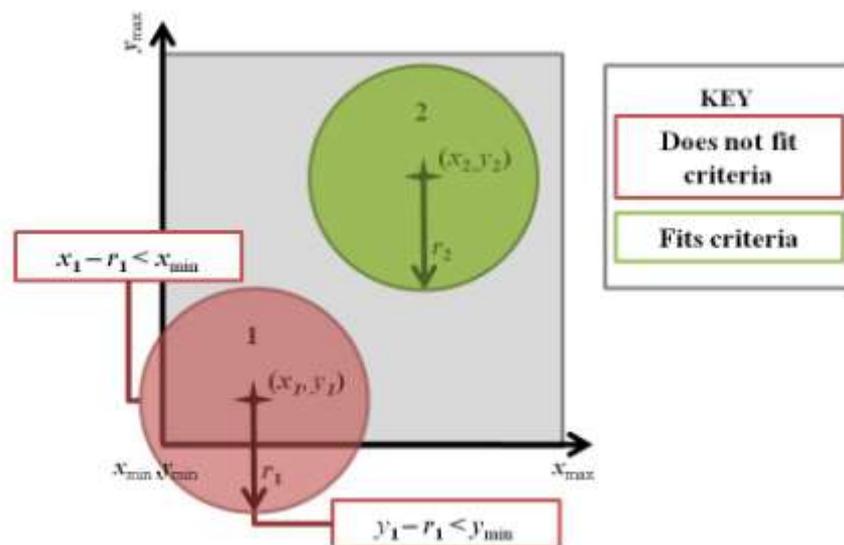


Figure 6.6. Example of the criteria for automatic removal of circles. In this instance circle 1 would be removed, with circle 2 retained

similarity measures were calculated for two circle characteristics, described below, with potential matches limited by the development of two filters to reduce the size of the search space for input into an optimisation algorithm. This *search space* is defined as the total set of potential circle-to-circle matches permitted by the algorithm. At the outset the search space and number of potential matches, γ , can be calculated by

$$\gamma = \frac{n!}{2!(n-2)!} \quad (6-2)$$

where 2 is the number of circles to be matched together to create a cylinder. The two circles to be matched are hereafter described as “circle₁” (α) and “circle₂” (β). The similarity measures used are:

Similarity measures

- i. Radius: the difference in radii of circle₁ and circle₂
- ii. Spectral: the difference between the mean brightness of all pixels within circle₁ and circle₂

Filters

- iii. Direction: absolute direction (in degrees) between the centre point locations of circle₁ and circle₂
- iv. Distance: the distance (in pixels) between the centre point locations of circle₁ and circle₂

Each of these similarity matrices and filters is now described in detail:

i. Radius

Using the assumption that the tops and bases of the cylinders to detect have equal dimensions (i.e. vertical sides of the cylinders), each circle radius (r) was compared with the radius of every other detected circle in turn and the ratio of r_α to r_β (expressed as a percentage) stored in the matrix \mathbf{R} , generated so that

$$\mathbf{R}_{\alpha,\beta} = \left(\frac{r_\alpha}{r_\beta} \right) \times 100 \quad (6-3)$$

where r_α and r_β are the radii of the two detected circles represented in row α and column β of matrix \mathbf{R} . The relationship between r_α and r_β was expressed as a percentage to allow the creation of a filter that retains high values (i.e. where similarity is high), and discount circle matches where radii similarity was low.

A filter was used to retain all the values within \mathbf{R} that fulfilled the criterion $\mathbf{R}_{\alpha,\beta}^2 > \sigma_{\mathbf{R}}^2$, where $\sigma_{\mathbf{R}}$ is the standard deviation of all values in \mathbf{R} . $\sigma_{\mathbf{R}}$ was chosen for this threshold to take into account the variance of the values within \mathbf{R} , and to only retain those circle matches with high similarity. Each value within \mathbf{R} was normalised by the optimal value (100), and stored in \mathbf{R}^1 , with unit variance, so that

$$\mathbf{R}_{\alpha,\beta}^1 = \left(\frac{\mathbf{R}_{\alpha,\beta} - 100}{\sigma_{\mathbf{R}}^2} \right)^2. \quad (6-4)$$

ii. Spectral

The spectral responses of each circle were compared with one another, using the assumption that all objects in the scene have a similar brightness response. There are several factors that would affect this response, resulting from physical phenomena such as occlusion (cloud, haze or shadow), or the calibration of the sensor. Because of variations in expected brightness values of detected objects, a threshold was generated to allow flexibility in the system by retaining spectrally similar objects in a comparative matrix and removing spectrally dissimilar objects.

A spectral similarity matrix, \mathbf{G} was created using a comparison of the mean brightness value (\bar{x}_g) of all pixels within circles α and β , expressed as a percentage in \mathbf{G} , where

$$\mathbf{G}_{\alpha,\beta} = \left(\frac{\bar{x}_{g\alpha}}{\bar{x}_{g\beta}} \right) \times 100 \quad (6-5)$$

Matrix \mathbf{G} was normalised in the same fashion as (6-4), replacing the radius values with percentage spectral values and using the standard deviation of all spectral values within \mathbf{G} , $\sigma_{\mathbf{G}}$, so that

$$\mathbf{G}_{\alpha,\beta}^1 = \left(\frac{\mathbf{G}_{\alpha,\beta} - 100}{\sigma_{\mathbf{G}}^2} \right)^2. \quad (6-6)$$

iii. Direction

The initial processing stage for the directional filter involved directional information being calculated from the sensor azimuth angle of every circle to the known location of the sensor, (derived from the image metadata). The mean value was calculated from the direction matrix and used as a synthetic sensor direction, measured in degrees (D_s).

Depending on the value of D_s , a directional range, δ , was set to provide a filter to constrain the choice of potential circle-to-circle matches. This range was based on the direction between the sensor and the extremities of the image with the bearing with the smallest angle for the Buncefield image equating to the image coordinates $x = 0, y = y_{\max}$. The direction with the largest angle was constrained to $x = x_{\max}, y = 0$. Figure 6.7 displays the creation of the range of acceptable image directions used for Buncefield. The observed direction between circle₁ and circle₂ (stored in matrix \mathbf{D}_1) was calculated by

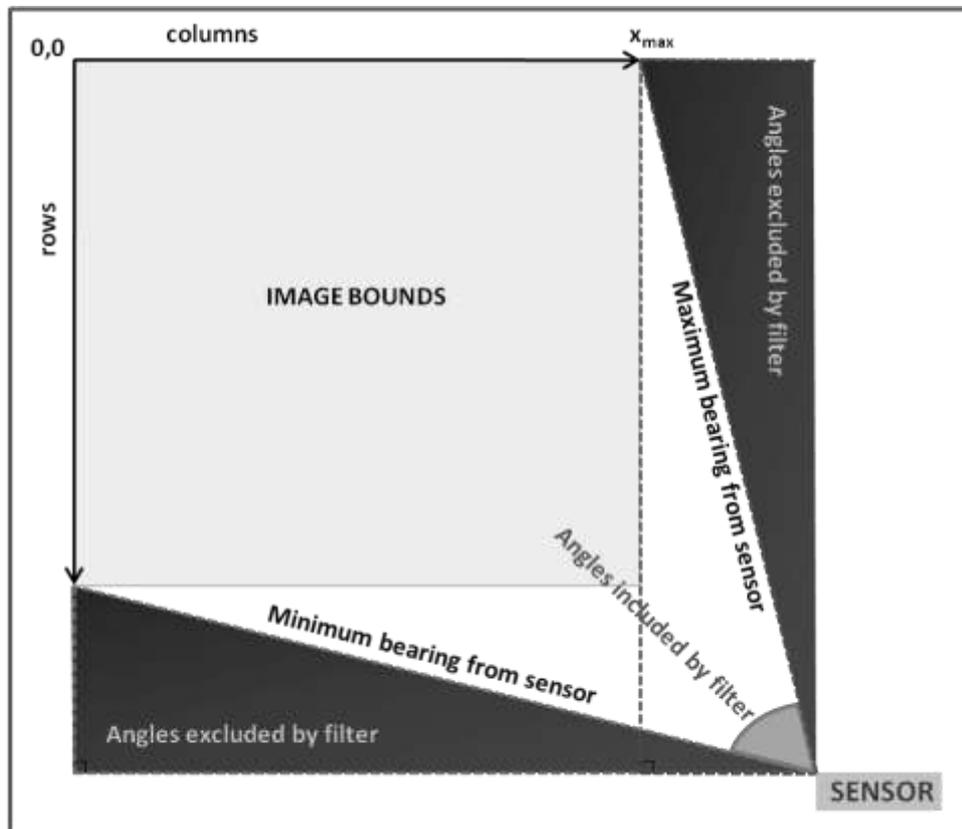


Figure 6.7. Creation of the direction filter for images where $90^\circ < D_s < 180^\circ$

$$\mathbf{D}_{1\alpha,\beta} = \tan^{-1}\left(\frac{\Delta x_{\alpha,\beta}}{\Delta y_{\alpha,\beta}}\right) \quad (6-7)$$

where Δx and Δy are the differences in x and y coordinates respectively for circle₁ and circle₂. These are filtered based on a comparison with the optimal direction between the base and the top of an object, D_s with those values within the direction matrix lying outside the set range were set to ‘not a number’, or “NaN”. This strength of directional fit was normalised in \mathbf{D}_1^1 as in (6-8), where $\sigma_{D_1}^2$ is the variance in all values stored in \mathbf{D}_1 .

$$\mathbf{D}_{1\alpha,\beta}^1 = \left(\frac{\mathbf{D}_{1\alpha,\beta} - D_s}{\sigma_{D_1}^2}\right)^2 \quad (6-8)$$

iv. Distance

The pixel distance from the centre points of every two potential circles from the complete set of potential matches was used to populate the $n \times n$ distance matrix, \mathbf{D}_2 (6-9)

$$\mathbf{D}_{2\alpha,\beta} = \sqrt{(x_\alpha - x_\beta)^2 + (y_\alpha - y_\beta)^2} \quad (6-9)$$

A threshold was used to filter those distances above an acceptable range of values expected for the target objects to be detected. In this case, the threshold was set to exclude circles matched at a distance greater than twice the mean diameter of all detected circles ($4\bar{r}$). This value was chosen using prior knowledge of the relationship between object width and height. In this case, the operator can filter the system based on the knowledge the oil tanks are not tall and narrow. All values greater than this threshold were set to *NaN* within matrix \mathbf{D}_2 . This supervised value was chosen due to the expected characteristics of the target objects.

6.3.5 The Score Matrix

Each of the aforementioned similarity matrices and filters were combined to create a single matrix describing the overall similarity of each possible circle-to-circle match. This matrix, known here as the ‘‘Score Matrix’’, Ψ , was created by firstly calculating the relative importance (or weight) of the four measures. A weighting factor, κ , was generated through a two-stage process for radius, \mathbf{R}^1 , spectral, \mathbf{G}^1 and direction \mathbf{D}_1^1 :

Stage 1:

Radius:

$$\kappa_{\mathbf{R}^1} = \left(\sigma_{\mathbf{R}^1} \right)^2 + \left(\kappa_{\mathbf{R}^1} \right)_{opt} \quad (6-10)$$

Spectral similarity:

$$\kappa_{\mathbf{G}^1} = \left(\sigma_{\mathbf{G}^1} \right)^2 + \left(\kappa_{\mathbf{G}^1} \right)_{opt} \quad (6-11)$$

(where $(\sigma_{\mathbf{R}^1})^2$ and $(\sigma_{\mathbf{G}^1})^2$ are the variances of \mathbf{R}^1 , and \mathbf{G}^1 respectively). These variances were added to $(\kappa_{\mathbf{R}^1})_{opt}$ and $(\kappa_{\mathbf{G}^1})_{opt}$ - the optimal values of $\kappa_{\mathbf{R}^1}$ and $\kappa_{\mathbf{G}^1}$.

Direction used the same process:

$$\kappa_{\mathbf{D}_1^1} = \left(\sigma_{\mathbf{D}_1^1} \right)^2 + D_s \quad (6-12)$$

where the variance is added to D_s - the mean direction from the sensor to the centre of each detected circle. $\kappa_{\mathbf{R}^1}$, $\kappa_{\mathbf{G}^1}$, and $\kappa_{\mathbf{D}_1^1}$ are all single values that are used to generate the relative importance of radius, spectral similarity and direction in stage 2, and thus, determine the weighting for each variable. The distance information was

not weighted, but simply used as a filter to exclude all unfeasible matches within the threshold set previously.

Stage 2:

The second stage of the weighting process involved the calculation of relative importance of each metric, W , so that

$$W_{\mathbf{R}^1} = \frac{\kappa_{\mathbf{R}^1}}{\kappa_{\max}}, \quad (6-13)$$

$$W_{\mathbf{G}^1} = \frac{\kappa_{\mathbf{G}^1}}{\kappa_{\max}}, \quad (6-14)$$

$$W_{\mathbf{D}_1^1} = \frac{\kappa_{\mathbf{D}_1^1}}{\kappa_{\max}}, \quad (6-15)$$

where κ_{\max} is the largest of the three weighting factor values: $\kappa_{\mathbf{R}^1}$, $\kappa_{\mathbf{G}^1}$, and $\kappa_{\mathbf{D}_1^1}$.

All matrices created in this section were combined using the weighting information from (6-13) to (6-15), to create the Score Matrix, Ψ :

$$\Psi = (W_{\mathbf{R}^1} \times \mathbf{R}^1) + (W_{\mathbf{G}^1} \times \mathbf{G}^1) + (W_{\mathbf{D}_1^1} \times \mathbf{D}_1^1) \quad (6-16)$$

The Score Matrix was then filtered using the distance information, and for every matrix entry where $\mathbf{D}_{2\alpha,\beta} > 4\bar{r}$, where \bar{r} is the mean radius of all detected circles.

$$\Psi_{\alpha,\beta} = \Psi_{\max} + 0.001. \quad (6-17)$$

where Ψ_{\max} is the maximum value stored in the Score Matrix.

6.3.6 Circle Matching

As stated in (6-2), there was a large combination of potential results if $\frac{n}{2}$ cylinders were to be created from n circles detected by the HT. For 20 detected circles, $\gamma = 190$, and for $n = 17$, $\gamma = 136$. These search spaces were reduced by the filtering process. However, a heuristic method was still needed to reduce the potential search space further to provide an optimal solution within a realistic timeframe. Each cell in the matrix provided a similarity score for each combination of circles, and accessing the data held within the Score Matrix provided an energy function, e , to be minimised by an optimisation method.

An initial combination of circle matches was created, providing a starting point for the optimisation process, generated by randomly matching the n detected circles to create $\frac{n}{2}$ matches. This was achieved using a random integer generator for values between 1 and n . For circle detections where n was an odd number, $\frac{n-1}{2}$ matches were created, also represented by η . The Score Matrix was accessed for each pair of circles matched and an initial energy value, e_{Ω} , was returned for the set of candidate cylinders where:

$$e_{\Omega} = \sum_{i=1}^{\eta} e_i \tag{6-18}$$

6.3.7 Optimisation by Simulated Annealing

The set of circle matches were retained as the best possible combination thus far, with the cumulative energy of the matched cylinders, e_{Ω} , used as the energy function to minimise. An iterative process of “swaps” then followed where two circles were randomly chosen from the n detected (even when n is an odd number). These circles were replaced in the candidate cylinder combination with error information gathered directly from the Score Matrix, and accumulated as described in Figure 6.8.

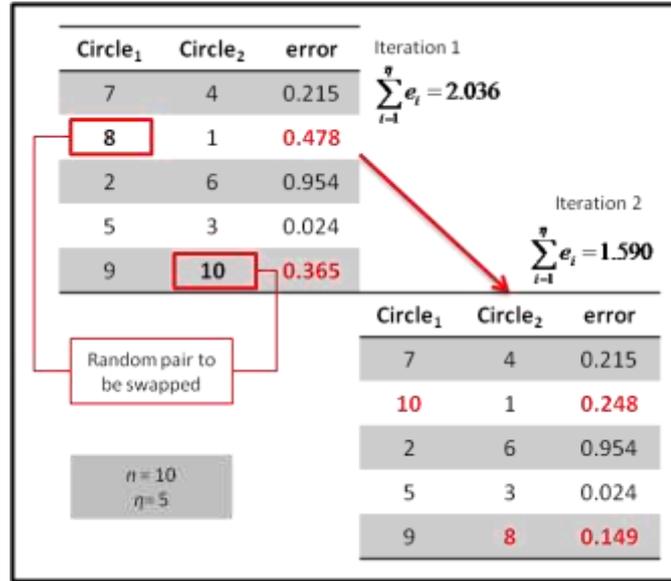


Figure 6.8. Example of the circle swapping process

After each swap, a new combination of e was created from Ψ , e_{new} , and varied according to the combination of circle-circle matches. An optimisation process was required to manage this sequence of swaps to reduce the energy function to an acceptable level. A simulated annealing (SA) algorithm was chosen for this purpose with the method described in the following sections.

Iterations

The SA algorithm was implemented in an iterative fashion. Initially only primed to detect η cylinders, the swapping process described earlier was programmed to continue for an indeterminate time (that differed according to the cooling schedule, described later), based on the results of the error information gathered from the swapping process. The new combination created by a swap was retained if the combined error value was reduced. A new pair of random circles was then swapped in the next iteration.

Probability function

If a bad swap was made (i.e. if $e_{\text{new}} > e_{\Omega}$), the SA algorithm accepted these changes only if a probability, p , was exceeded. In order for this to happen, a random value, v , between zero and one, was directly compared to p where:

$$p = \exp \frac{e_{\Omega} - e_{\text{new}}}{kT} \quad (6-19)$$

where $k = 0.8$: a constant included for the gradual reduction of the temperature, T . If $v > p$, the swap was rejected and the circle matches returned to the previous configuration. If $v \leq p$, the probability function was met, and the new cylinder configuration was retained despite $e_{\text{new}} > e_{\Omega}$. This process allowed the algorithm to escape local minima in order to achieve an alternative heuristically-derived minimum.

Cooling schedule

The SA process is consistent with its analogy with the process of cooling of metals. It included a cooling strategy, where a starting “temperature” (T) was progressively reduced to a minimum. An exponential cooling schedule was chosen for this implementation of the SA: starting at $T = 1$, the temperature was cooled at a rate of

$$T_{\text{new}} = T - \frac{T}{10000} \quad (6-20)$$

T was reduced only when a “good” swap was found (i.e. when the energy, e_{new} , of the new cylinder configuration was less than the old configuration). For swaps where there was no energy decrease, T remained unchanged until a good swap was achieved, irrespective of whether the probability function was met and the new cylinder configuration was retained. The SA process was repeated with T being gradually reduced, until it is unchanged for 5000 iterations, when T is set to zero. At this stage, the best configuration (i.e. the circle match combination with the lowest combined error) retained as the solution of circle match possibilities.

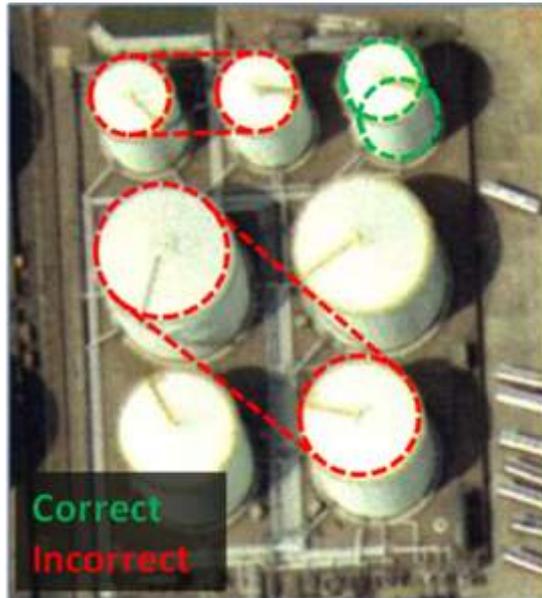


Figure 6.9. Potential erroneous matching of top circles. These potential errors can be reduced by reducing the distance threshold

6.3.8 Automatic Object Completion & Cylinder Orientation

The combined HT and SA approach focused on detection and matching of circles based on assessing the strength of fit between circles that potentially form the top and base components of candidate cylinders. It was also necessary to complete the cylindrical features in full and compare with the observed, input image. This was necessary as two potential circles that were detected on the tops and bases of image features may be matched by the optimisation algorithm and display a strong match based on the similarity metrics used. However, these matched circles have the potential to be components of two discrete objects and display characteristics outside the circles more akin to background image features. The whole feature was therefore simulated to identify potential circle mismatches (Figure 6.9).

The simulation of cylindrical features used the assumption that a cylinder with an oblique projection would appear as a two-dimensional image footprint similar to that displayed by two circles (the top and base of the cylinder) connected by two lines (l_1 and l_2 ; the two sides of the cylinder). Creation of the sides of the cylinders was achieved in an automated fashion using attributes of the detected circles gained during the development of the similarity matrices. The cylinder side information was generated through several key processes:

- i. Calculation of cylinder direction
- ii. Generation of top and base perimeter nodes
- iii. Generation of line information – cylinder sides
- iv. Creation of the simulated image

i. Cylinder orientation

Knowledge of the orientation of the features in the image was gained from the image metadata used when initialising the algorithm. The *direction* matrix created in section 6.3.3 held information on the bearing of all circles in relation to all other circles detected by the HT. This information was used by the algorithm to automatically determine the base and top circles of cylinders. Figure 6.10 displays the relation of the top and base of a circle using the image location with respect to the location of the sensor. This is true for the Buncefield image used, which is displayed with perspective projection, rather than orthorectified.

ii. Generation of top and base nodes

Two points, or nodes, were identified on the circumference of both the top and base circles and were joined by vectors to create the cylinder sides. Points T_1 and T_2

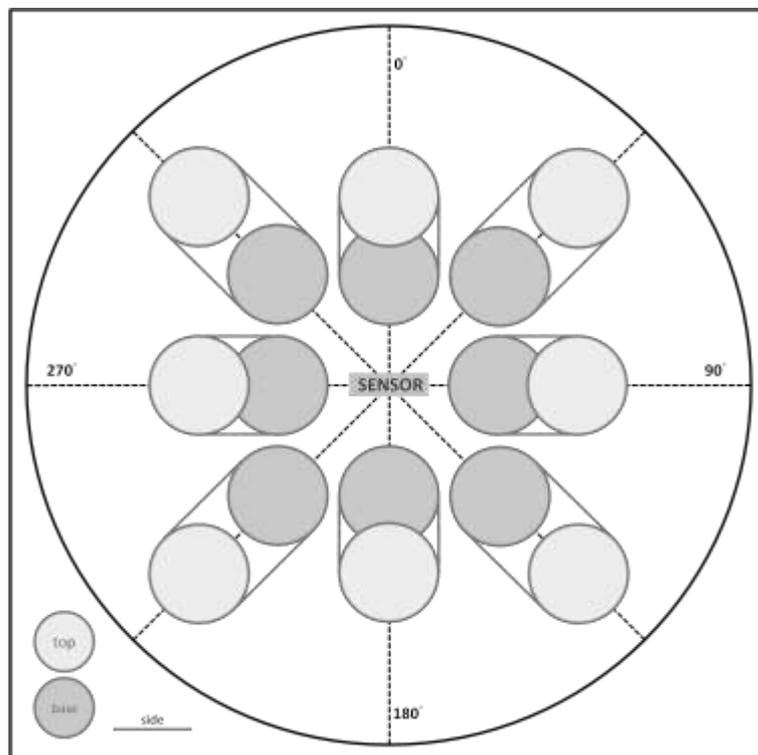


Figure 6.10. Orientation of cylinders with respect to sensor location
(perspective projection)

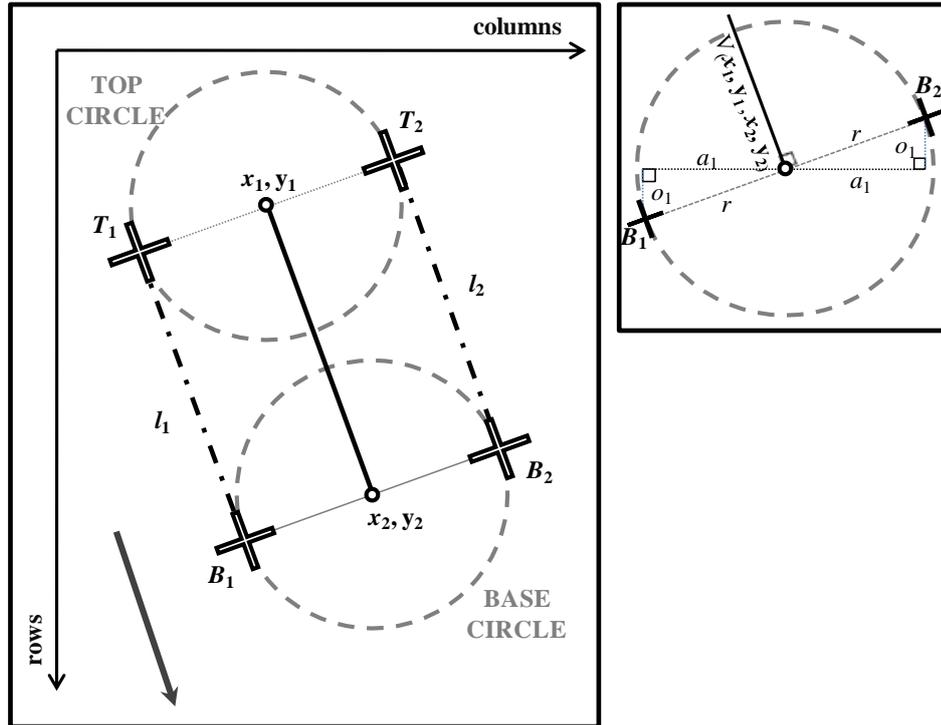


Figure 6.11. Generation of top and base node information

differentiated the respective left and right points on the top and B_1 and B_2 were the equivalent base points (Figure 6.11). The locations of $T_1, T_2, B_1,$ and B_2 were calculated using the following process:

Firstly, the matrix of gradient directions (\mathbf{G}) used to develop the *direction* similarity matrix was accessed, returning the angle between the centres of the top circle, (x_1, y_1) , and the base circle, (x_2, y_2) . Vectors perpendicular to the vector connecting both top and base centre points, $V_{(x_1, y_1, x_2, y_2)}$, were created with lengths twice the radius of top and base circle respectively, with the midpoint of each vector at their respective circle centre (Figure 6.11). The centre point location of each circle was used as a location for a triangle apex with a line perpendicular to the $V_{(x_1, y_1, x_2, y_2)}$ vector forming the hypotenuse for the triangle, with length equal to the circle radius, r . The triangle was completed using trigonometric functions to determine the interior angles of the triangle and subsequently the lengths of both the opposite (o_1) and adjacent (a_1) sides. The apex where the opposite meets the hypotenuse is used as the location for one of the perimeter nodes. The process was replicated for all four perimeter locations necessary for the creation of a cylinder (T_1, T_2, B_1, B_2).

iii. Generation of line information

The gradients of each of the cylinder sides were calculated from the equation

$$m_{(x_1, y_1, x_2, y_2)} = \frac{\Delta y}{\Delta x} \quad (6-21)$$

where Δy and Δx represent the difference in y and x locations, calculated from points T_1, B_1 and T_2, B_2 for lines l_1 and l_2 . The equations for these lines were generated from equation (6-22, where λ is the y -intercept of the line.

$$\lambda = y - mx \quad (6-22)$$

iv. Creation of the Simulated Image

The information generated in steps ii and iii provided the ability to gauge whether any pixel within the image was located within or outside the internal rectangle (bounded by T_1, B_1, B_2 and T_2) of the candidate cylinder. With the ability to simulate all cylinder primitives (the top, base, and both sides), it was possible to create a mask to identify all pixels lying within the boundary of a potential cylinder. At this stage of the analysis, discrete circles that were detected by the Hough Transform were matched and classified as *candidate* cylinders, using cylinder matching through the described method of automated quantification of similarity and generation of cylinder features.

A forward model assigned a brightness value of 235 to all pixels inside the image footprints of the features, with a background value of one assigned to pixels outside. These values were generated manually to represent approximate spectral values of cylinder objects and background pixels in the input image. This simulated image (**s**) was compared to the observed image (**i**) and error information (**e**) for the combined cylinder scene calculated as

$$\mathbf{e} = \mathbf{i} - \mathbf{s} \quad (6-23)$$

The RMSE (3-2) for the simulated image was calculated giving a single value to represent the combination of all detected objects (ρ_{all}), rather than the per-pixel error information represented by \mathbf{e} .

6.3.9 Automatic Determination of ‘ η ’

An algorithm that automatically matches sets of two circles to create cylinders requires a method for automatically determining the number of cylinders in an image, to prevent multiple false positive cylinder creations, or to avoid forcing of the SA algorithm to detect η cylinders. This section describes the initial method devised for this purpose.

Relative importance of features and determination of η

The ρ_{all} value described previously is a metric to determine the error associated with the combination of cylinders when all detected objects are simulated together. Each cylinder was also ranked in terms of relative importance against all others to gauge per-object error information. η simulated images were forward modelled with ($\eta, \eta-1, \eta-2 \dots 1$) cylinders from the all-object set and \mathbf{e} and RMSE were generated for each combination.

The per-object error data were used to generate F simulated scenes showing the top F –ranked cylinders plotted together where F is a value that ranges from $1:\eta$, and the first scene would show the single cylinder with the smallest error, the second scene would combine the two cylinders with the smallest error value etc.

6.3.10 Repeated SA Using New Value of η

Comparing the errors generated by combinations of the cylinders with the lowest energies effectively shows the optimal number of objects to detect in a scene. If the combination of the best F objects does not contain erroneous candidate objects, the RMSE value would be smaller than when a false positive is contained in the combination. Using this basis as a rule, the smallest RMSE value of the combinations was used as an input for one further iteration of the SA process. If the

RMSE for the three best cylinder matches is consistently lower than all other combinations of the F strongest matches, the SA is repeated to pair only three sets of circles.

6.3.11 Monte Carlo Simulations

To gauge the effectiveness of the method described, the whole process thus far was repeated 50 times. The results of the 50 iterations were used to estimate the optimal value for η according to the method described in section 6.3.9. All parameters remained unchanged from those in Table 6.1, and no supervised input was required for the 50 Monte Carlo simulations of the *HT*- and SA method.

6.4 RESULTS

Results for the described method are given in Figure 6.12 to Figure 6.17.

Table 6.2 shows the calculated image, sensor and illumination geometry for the Buncefield image subset used, with the sensor attribute values used to create the direction filter matrix. These were manually calculated using trigonometric functions.

Results from the initial circle detection using the standard Hough Transform technique (*HT*-) are shown in Figure 6.12a, with 20 circles detected in total. The number of circles was automatically reduced to 17 using a filter that excludes those circles not wholly inside the image Figure 6.12b). An example of η circles matched by one implementation of the simulated annealing algorithm is shown in Figure 6.13, alongside the creation of cylinder sides and a simulated image showing the

Table 6.2. Sensor and illumination geometry for the Buncefield image subset

Sensor Azimuth Angle (A_z)	124.8°
Sensor Zenith Angle (Z_s)	10.4°
Illumination Azimuth Angle (A_i)	236°
Illumination Zenith Angle (Z_i)	41°

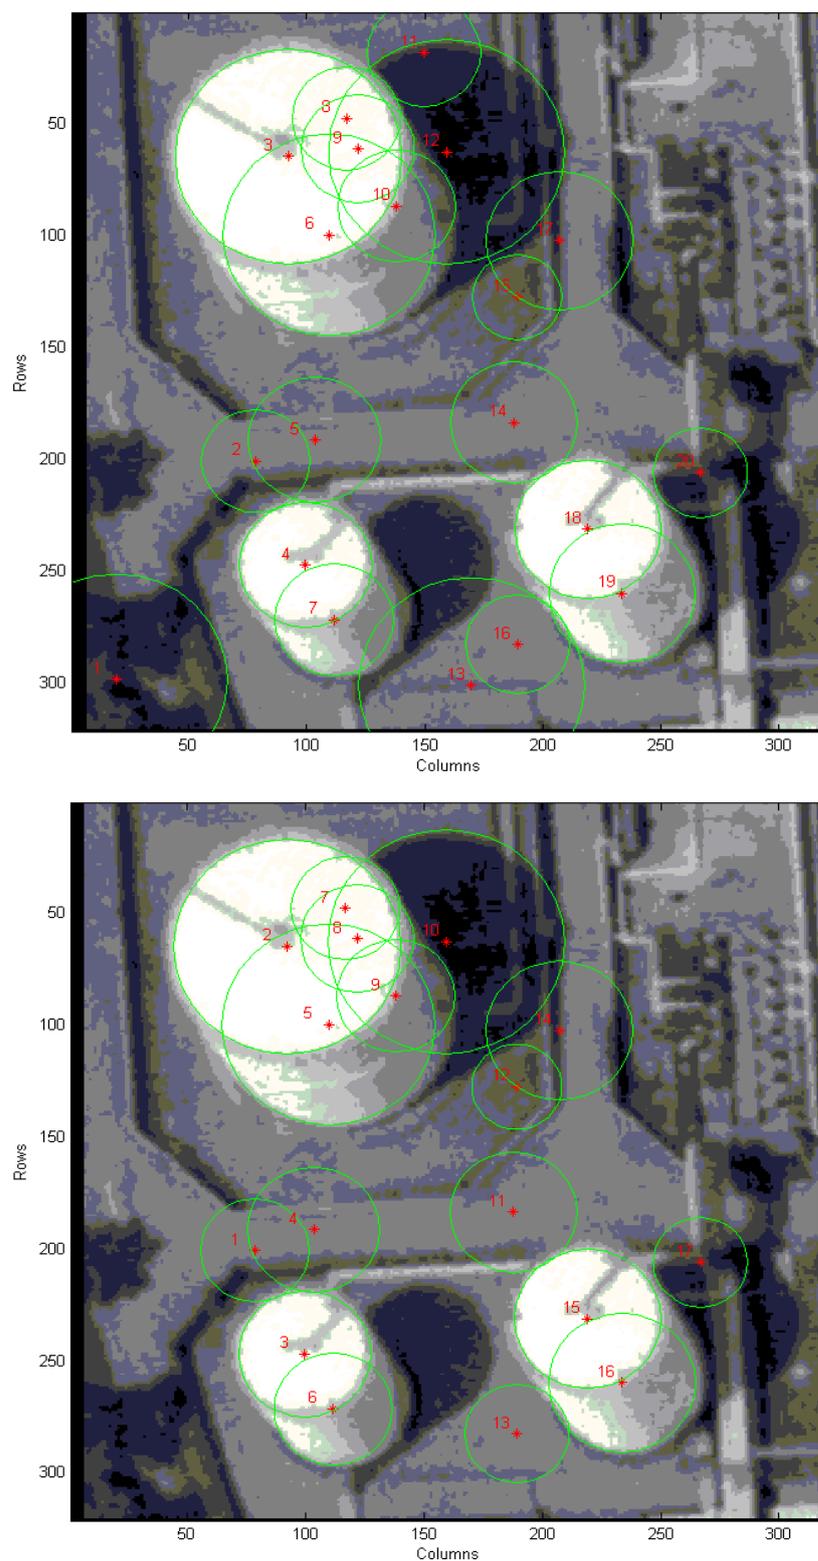


Figure 6.12. Circles detected by the Hough Transform algorithm
a) Detected by *HT*-; b) Circle numbers reduced by image filtering algorithm

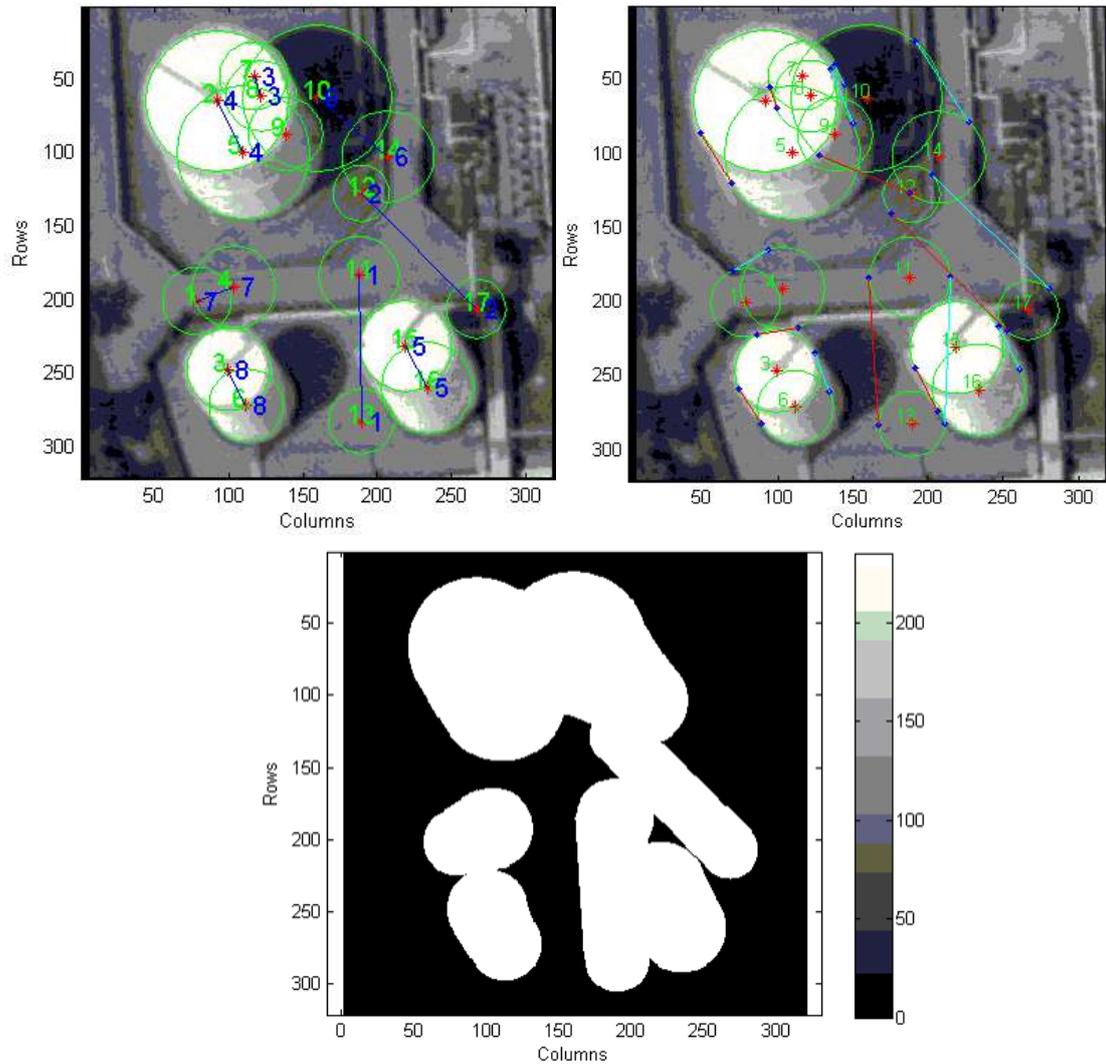


Figure 6.13. Circle matching and simulating cylinders
 a) top & base circle matches; b) creation of cylinder sides;
 c) simulated image for n_2 cylinders (units: grey-levels)

masked representation of these cylinders depicted in the same figure. The SA process for the described iteration took 204,234 iterative swaps before the temperature value (T) converged. The energy fluctuations generated at each step of the SA process are displayed in relation to the lowest energy value for the equivalent iteration (Figure 6.14), with the relationship between T and the lowest produced energy value displayed in Figure 6.15.

Figure 6.18 displays a scatter-plot of the RMSE against the grouped F circle matches with the lowest energy, and shows the grouped RMSE at its smallest when three cylinders were detected. When $\eta = 3$ was returned into the system (i.e. it was automatically determined that three matches should be made), the SA optimisation process was run for 50 Monte Carlo simulations. Figure 6.16 shows the circles that

were matched and simulated to generate RMSE information (Figure 6.17) and Table 6.3 displaying the frequency of circle matches for the 50 SA-generated solutions. The most frequent match was for circles [8,7] (pertaining to [base, top] circles), with only seven combined matches from a potential 136 possibilities. The RMSE for each MC iteration of the SA process is plotted alongside a supervised analysis of ‘correct’ and ‘incorrect’ detections (Figure 6.17). The most common combination of 3-cylinder candidates was when [5,2], [8,7] and [16,15] were combined (a 66% success rate for detection of correct cylinders). The correct match of [6,3], [5,2] and [16,15] was only converged on once during the 50 repetitions of the process (Figure 6.17). Figure 6.18 displays the plotted combined RMSE results from the relative importance calculations to automatically compute a value for η . The mode result of the object detection is displayed alongside the correct combination of detected features in Figure 6.19 and Figure 6.20, respectively.

Table 6.3. Most frequent circle matches and Score for 50 MC implementations of SA where $\eta=3$

Matched Circles	Frequency of Match		<i>e</i>
	(count)	(%)	
[8, 7]	48	96	0.0017
[5, 2]	39	78	0.0096
[16, 15]	23	46	0.0101
[13, 11]	17	34	0.0141
[6, 3]	14	28	0.0128
[11, 4]	8	16	0.0208
[11,9]	1	2	0.0232

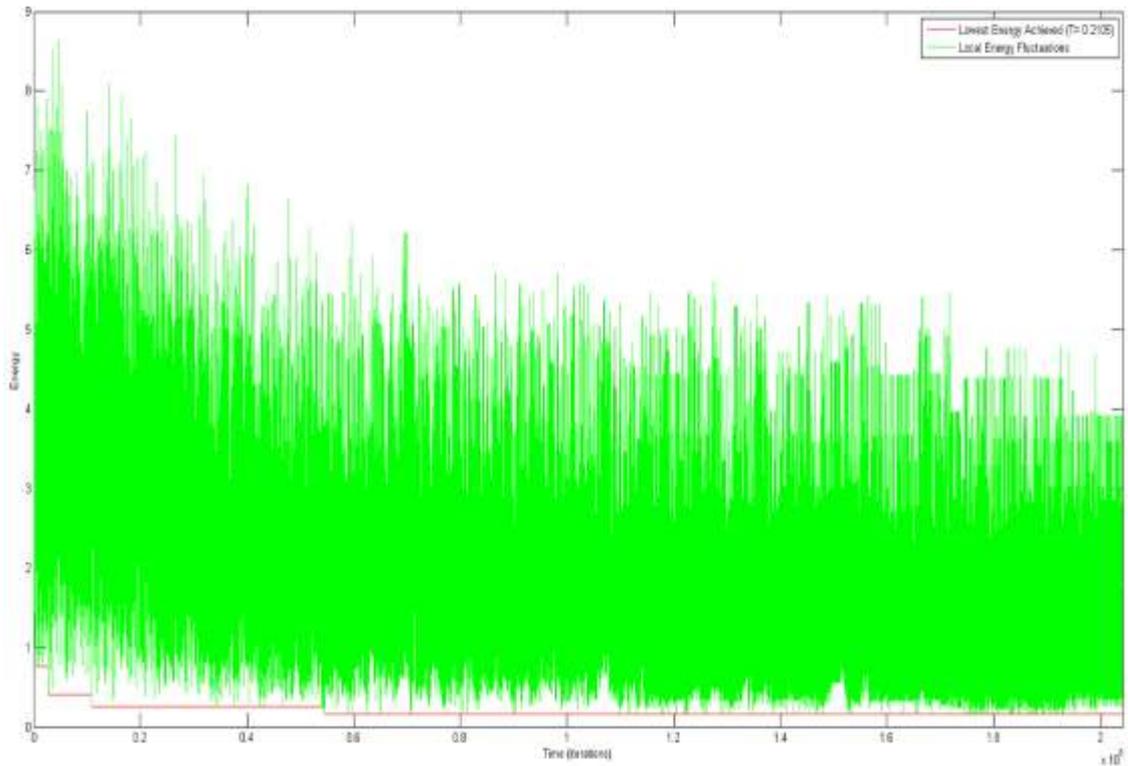


Figure 6.14. Local energy fluctuations for each iteration of SA

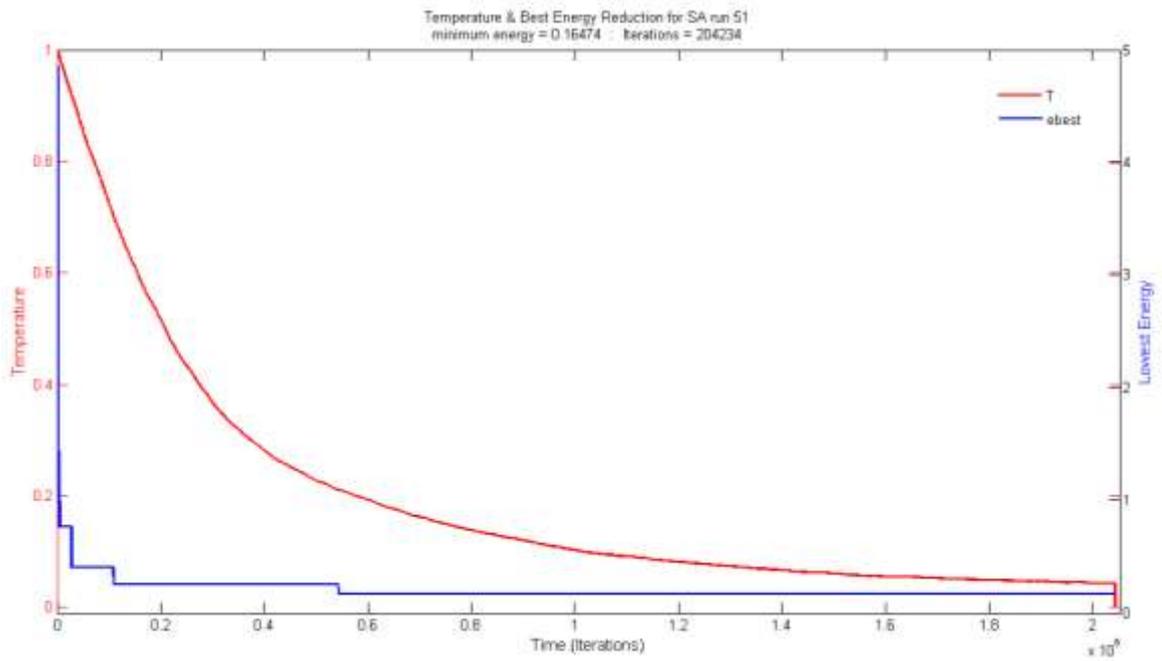


Figure 6.15. Temperature and lowest energy reduction for 1 run of SA

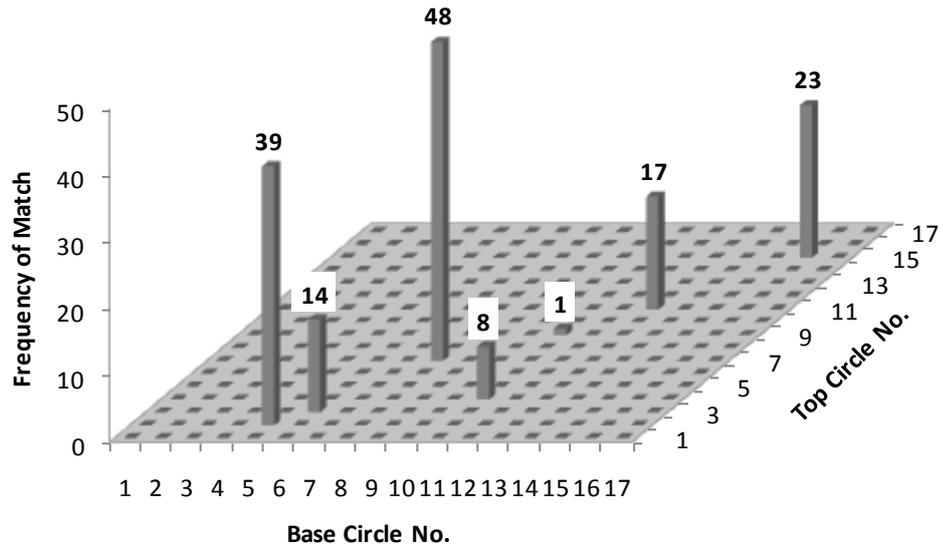


Figure 6.16. Top & base circle matches for 50 iterations of three-cylinder creation

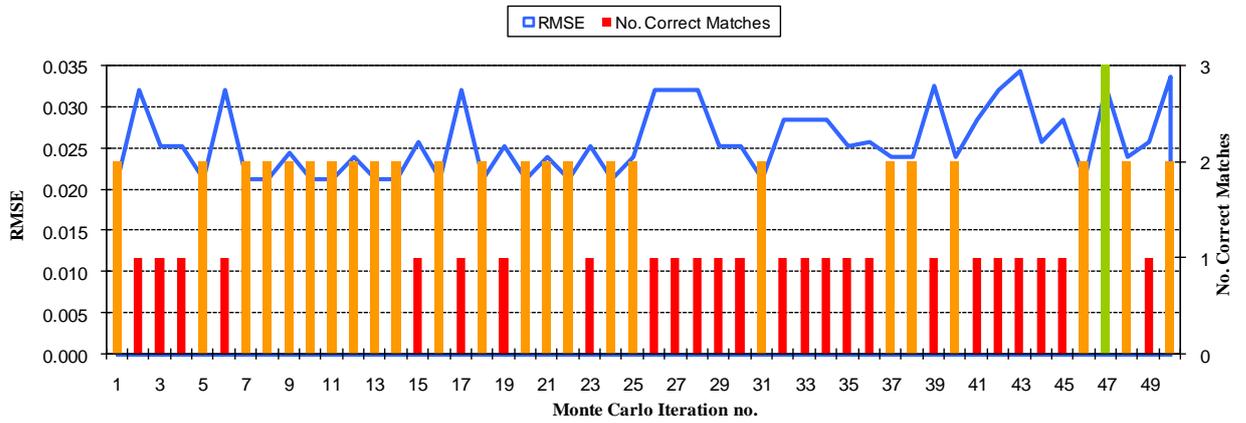


Figure 6.17. Total RMSE and number of correct cylinder detections for 50 iterations of three-cylinder creation. Red, orange and green columns show one, two and three correct matches respectively

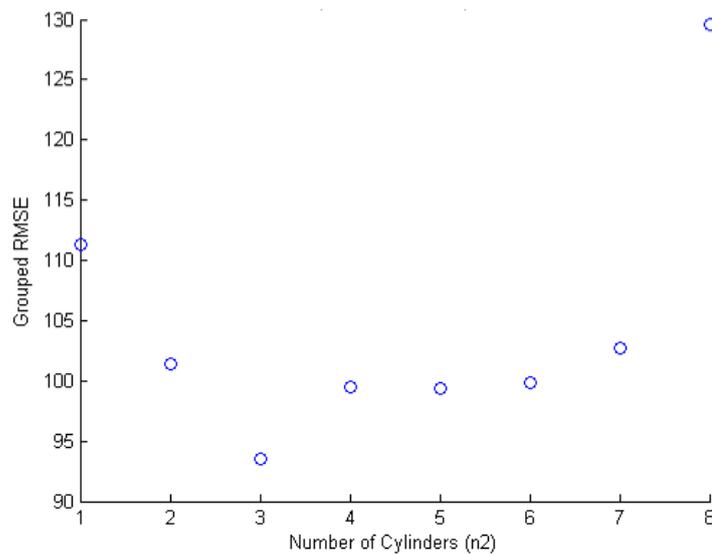


Figure 6.18. Grouped RMSE of best F matched cylinders

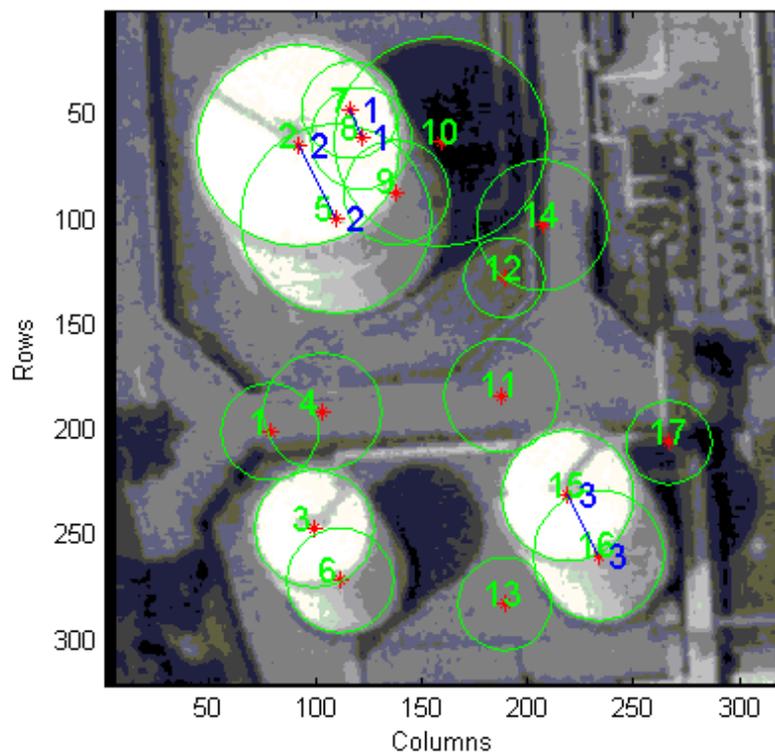


Figure 6.19. The modal combination of cylinders for $\eta = 3$

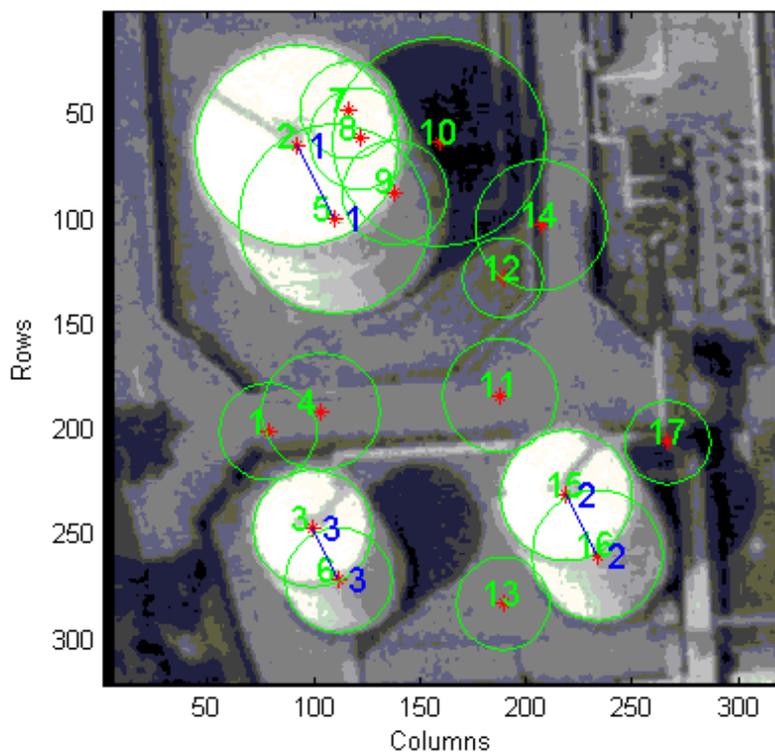


Figure 6.20. Correct cylinder detection from Monte Carlo iteration 47

6.5 DISCUSSION

The results described and displayed in the previous section showed 17 circles remained after the HT and circle filtering. The initial filter relied on the prior knowledge that the target objects had their projected footprint completely within the image. As a result, three falsely detected circles were removed, reducing the potential search space for circle matches from 190 (20 circles) to 136 (17 circles). This was reduced further through the development of metrics that scored each potential match and provided an energy value to minimize through an optimization process.

The similarity matrices and filters were aided by geometric information from the image metadata providing the image illumination and location of the sensor, relative to the target objects. This provided contextual information describing how a three-dimensional feature should look within the image space (Figure 6.10). External input was required for only one similarity metric described in section 6.3.3: the potential matches were constrained using logical knowledge of the characteristics of an oil storage tank. In this instance, a directional filter of $h \leq 4\bar{x}r$ (i.e. no extremely tall cylinders in relation to their base radius should be found in the image) was used. This is a valid assumption to make, as the user should know the basic projected shape of the target features. Also included was the premise that the tops and bases of objects should be spectrally similar. This was true for the objects in the Buncefield image, however it is noted that this is often not the case for other buildings, where differences in spectral response can be much greater. In these cases, the spectral variable can be excluded from the similarity matrices, with correspondence between objects achieved through comparison of the remaining variables 0 size of object base vs. height, distance and direction from base-top of the feature. These rules add an additional element of user input to the object detection process, however the remaining filters and weightings used were generated automatically, adhering to one of the key objectives (#4) set out in chapter 1 of this thesis.

The simulated annealing process showed potential to be an effective heuristic to find a low energy combination of randomly swapped circles, reaching a

combination of eight cylinders $\left(\frac{17-1}{2}\right)$ in 46.9 minutes, and three cylinders in 29.5 minutes, with approximately 200,000 and 140,000 iterations respectively. Figure 6.14 shows the fluctuations of energy at each iteration of SA, with Figure 6.15 showing the relationship between lowest energy and the cooling temperature. It is evident from these that due to the cooling schedule chosen, that the system cooled rapidly at first, and hence, lowered the energy rapidly.

An exponential cooling schedule was chosen for cooling the temperature (T) from $T = 1$ to $T \approx 0$, allowing a progressive reduction of T , reducing the probability of a bad swap being accepted as T decays. The area beneath the red line (T) in Figure 6.15 essentially shows the reduction in probability, p , as T is cooled. The energy decreased rapidly during early iterations, with the rate of reduction slowing as the potential for accepting a bad swap was reduced with T . An exponential schedule was chosen over a linear cooling schedule to allow a gradual ‘cooling’ of the system. It has been seen in previous applications of the SA algorithm that linear schedules are also effective reduction agents and either a linear or exponential schedule are preferable over logarithmic schedules (Nourani and Andresen, 1998). Choosing an exponential schedule meant the temperature would never reach zero, and stopping criteria was used to stop the system if 5000 iterations occurred without finding a combination that lowered the energy in the system. This was a fair assumption to make, as the MC results consistently showed convergence on a solution with the lowest energy. The simulated annealing algorithm was designed to avoid the premature cooling of T (which would result in settling into a local minimum solution), and so a constant, k , was included with a value of 0.8 as used in previous implementations of the algorithm (Kirkpatrick *et al.* 1983). The presence of k reduced the effect of each reduction of T , allowing the optimisation to be extended, and providing convergence on a more probable solution.

It was assumed that a 3-D cylindrical feature in a 2-D image will display a certain geometric profile and can be degraded into four geometric primitives: two circles representing the top and base of a cylinder and two lines representing the sides of the cylinder (the chosen method described in Figure 6.4b). The SA algorithm showed its effectiveness for generating results for $\frac{n}{2}$ cylinders made up of two circles only. Creation of a simulated image of candidate cylinders required

simulation of the sides of a cylinder combined with two matching circles. This method models the cylinder as a contiguous object made up of combined primitives rather than accepting two circles to represent the top and base of a cylinder when they could easily be unrelated in the original image (Figure 6.9). It is possible that two unrelated circles detected from the image have a potentially high similarity score, and hence, may be matched. However, when modelled in a simulated image, the error associated with that candidate cylinder will be large due to spectrally dissimilar properties in the middle of the object (in the area between the two matched circles) between the simulated and observed images. It is therefore necessary to simulate the ‘body’ of the cylinder and, hence, derive error information for all pixels within, and between the matched circles.

The use of sensor geometry for constraining the search space has been used previously in photogrammetric applications. However, the similarity matrices and filters-based method for generating score information is valid for several reasons: Firstly, the method is reliant on a model-based method of feature extraction. In the exemplars, circles were detected by the HT, but there is potential that the method may not precisely detect the outlines of circular features. By allowing a range of directions to be acceptable in the search space, slight offsets in the direction between circle α and circle β (6-7) can be tolerated, and objects can be modelled.

The nodes on the perimeter of the top and base circles were efficiently identified and used as start and end-points for the cylinder sides. It should be noted that the sides of the cylinders were not assumed to be parallel, and the method did not take into account the co-relationship of the two sides. If an object had parallel vertical sides, the top and base circle radii would be identical. Therefore, the method was able to simulate every possible cylinder, irrespective of top/base radial differences, allowing the automatic detection and removal of poor candidate matches. This method has been shown to be effective for the identification and modelling of cylinders within imagery captured at an oblique viewing angle. However, it should be noted that the technique may not be applicable to imagery captured with a low observational zenith angle (i.e. when the sensor is nadir to the objects). In these instances, the bases and sides of the objects would mostly be obscured, reducing the confidence of the automatically generated height information. To extract three-dimensional object information would therefore

require the use of additional information, such as shadow, or *a priori* knowledge of the objects (e.g. contained in a building inventory database).

Figure 6.18 displays the RMSE generated when the best F cylinders are simulated together. This method was used to determine the appropriate number of objects to detect automatically and was implemented effectively in this instance. A more elegant solution to determine η may be available, however, it should be acknowledged that no human intervention was required and the described method remained true to objective number 4 of the thesis. In each instance of the 50 MC iterations, the combination of the three best matches displayed the smallest RMSE value.

One of the key observations of the data displayed in Figure 6.17 and Table 6.3 is that the most frequent circle-circle match was not one of the target cylinders from the image. Rather, the SA matched two circles that were detected on the top of one of the oil storage tanks (combination [8,7]). These were not identified automatically as false positive detections because they exhibited similar spectral values to the target. The two circles fitted the required distance, direction and radius filter specifications set out in the similarity matrices and filters. This pairing displayed the lowest observed value in the Score Matrix, and was detected in 96% of the 50 three-cylinder MC simulations. It can therefore be concluded that the automated SA method effectively fulfilled its required function of repeatedly returning this combination.

From the 50 MC runs of the SA with a potential 136 matches, only seven matches were identified (Table 6.3). This highlights the potential of the similarity and optimisation approach described in this chapter for constraining the search space using basic object characterisation and geometric image metadata. Whilst the correct combination of cylinders was returned only once in the 50 MC runs, on 23 occasions two of the correct cylinders were detected alongside the modal detection [8,7]. There is potential that if a further constraint was added to avoid objects that have overlapping footprints in the imagery, the correct combination of cylinders would be detected more frequently. This would require simulating each detected object individually, and gathering RMSE information at the object level. Subsequent comparison of the overlapping objects is likely to display larger RMSE values for candidate objects that only overlap a portion of the correct object footprint. In the described implementation, the SA optimisation constrains the problem well;

however, to give the method generalisability to other oblique image sets, additional constraints may need to be included.

6.6 MODELLING THE THREE-DIMENSIONAL SCENE

This chapter has described the methodology associated with automatic detection and extraction of objects from remotely-sensed imagery, using cylinders from an oil depot as a test case. However, what happens when these circle matches are made and objects successfully identified? As has been seen in recent years with the expansion of online mapping services such as Microsoft Bing Maps, Google Maps and Google Earth, geographic data users have the ability to simply and efficiently generate maps for dissemination. Digital two-dimensional maps from point, line or polygon data can be overlain on top of publically-available fine spatial resolution satellite sensor images. This capability is extended when height information is included, allowing the generation of three-dimensional datasets, that can be rendered using oblique angle images. Preliminary investigations have been carried out into this approach using the information generated in the object detection phase of the conceptual model. These are described in the following section.

The objectives for modelling objects are not simply for trivial viewing of the detected objects, but rather, to be able to simulate the geometry of the input image, its objects and the external influences on the image, such as illumination and sensor locations; physically modelling the pre-change image in this way allows the further modification of the scene model to inherit the geometric, illumination and sensor characteristics of the post-event image, and hence provide a platform from which to gather per-object change information. This process is described in more detail in chapter 7.

6.6.1 POV-Ray 3-D Models

A primitive scene model can take the form of the data produced using the described methodology in section 6.3 (e.g. object location, height, spectral characteristics).

Knowledge has been gained for the image coordinates of the base of the cylinders, with the heights of the objects relating to a geometric transform of the distance between the base and top circle centre points. This 3-D modelling investigation relied on the prior knowledge of a flat-topped object, and added the assumption of vertical sides of the drums. The radius of the top of each cylinder was therefore used for both the top and base radii of the 3-D drums. The centre point of the top was used for the x , y coordinates of the cylinder, with the base location adjusted to be directly below the top.

The first investigation into presenting the results used a freeware ray tracing software package, the Persistence of Vision Raytracer (POV-Ray, 2008), to render the cylinders from the successful object detection implementation (Figure 6.20). An extension was written in the MATLAB code to automatically generate the following image information:

- geographic coordinates of the centres of the cylinder base circles (x , y coordinates)
- radius (r) and height (h) information for the detected cylinders
- observational zenith angle
- illumination zenith and azimuth angles

The geometric coordinates were automatically generated from a MATLAB function *pix2map*, which calculated an image to geographic coordinate transformation based on the world file metadata from the input image. Information on r and h was generated for each detected cylinder through the following equations

$$r_g = r_i \mu \tag{6-24}$$

$$h_g = h_i \mu \tag{6-25}$$

where r_g and h_g refer to the cylinder radius and height in metres for each cylinder, and r_i and h_i are the image dimensions in pixels for radius and height of the cylinders, and μ is a scaling factor equal to the spatial resolution of the image, measured in metres (see section 1.1.1, p4, for a note on spatial resolution).

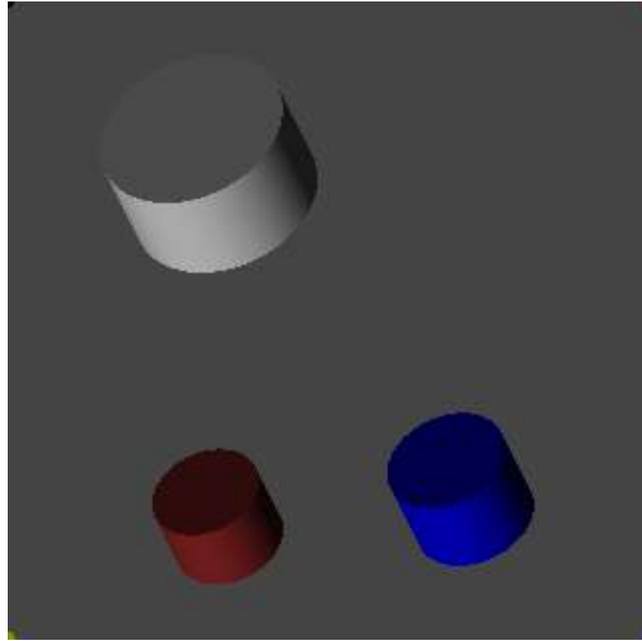


Figure 6.21. POV-Ray rendering of the three correct cylinders

In order to test the methodology, a manual approach selected the three correctly detected cylinders displayed in Figure 6.20. These were rendered in POV-Ray, producing the basic scene shown in Figure 6.21, with the cylinders rendered in different colours to differentiate their positions within the scene. Manipulation of the scene information within POV-Ray proved to be troublesome and would require an in-depth knowledge of this coding environment. However, it has shown early potential as a method of synthetically creating an image using the detected objects.

6.6.2 Google Earth Models

Another method of displaying and manipulating the information stored in the scene model used the publically available Google Earth software. The geographic information generated for the three cylinders (the geographic location and radii of cylinders) was used to create vector GIS shapefiles and Google Earth KML files for the detected top and base circles. This also allowed validation of the strength of fit for the detected circles, compared to geo-registered remotely-sensed imagery. Figure 6.22 shows the tops and base circles that were generated automatically from the statistics stored within the scene model. It shows a strong fitting of the cylinder tops with a slight offset to the southwest of the base circles.

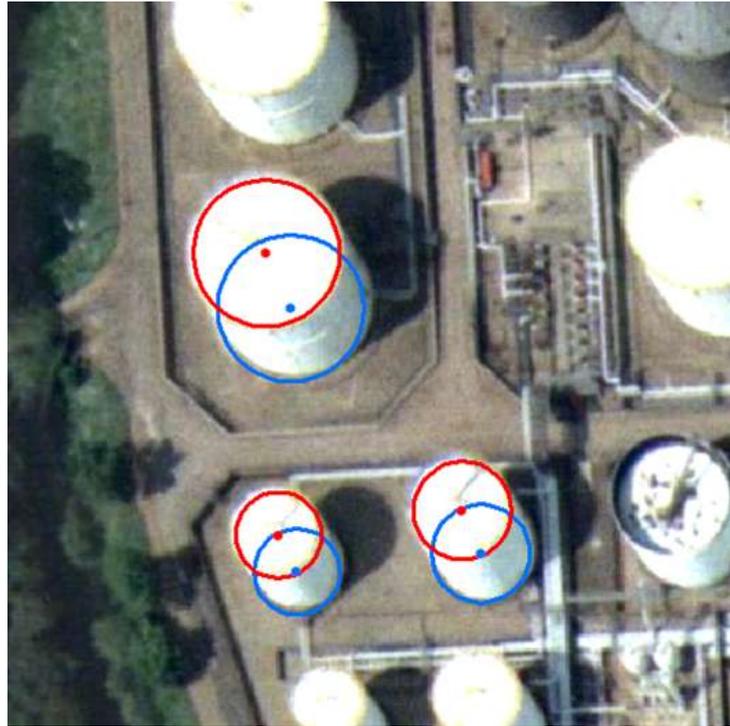


Figure 6.22. GIS-produced map of top (red) & base (blue) circles in geographic coordinates. Information needed to create map was automatically generated from the described HT & optimisation approach.

Knowledge of the base and top footprints allowed the generation of three-dimensional models of the cylinders. The base circles were imported into Google Earth and projected at a height equivalent to the h information from the scene model. Figure 6.23 shows an oblique-angle view of the study site with the 3-D cylinder objects displayed in yellow.

Using Google Earth or an equivalent software package is an effective method for dissemination of results to a worldwide audience as currently the use of publically available online mapping resources is generally increasing. Rendering the objects in Google Earth allows the user to zoom around the scene and simulate almost an unlimited variety of view angles and illumination effects, and includes functionality allowing display of images from more than one date. Figure 6.23 shows Google Earth archive imagery from 2000, while the most recent imagery available (after the explosion and fire) is shown with the detected cylinders at an angle close to nadir in Figure 6.24. Both these figures are displayed with sensor zenith and azimuth angles comparable to the Buncefield imagery used in the implementations previously described. The yellow cylinders are created with vertical sides, and displayed to mimic similar view angles to both the pre- and post-

event Buncefield imagery. It should be noted that the Google pre-event imagery was acquired with a sensor azimuth angle closer to nadir than the pre-event data described in chapter 4.



Figure 6.23. Oblique-angle view of Google Earth-generated 3-D objects. Viewing position was selected using information on sensor location at time of capturing the 1999 Buncefield image (Figure 4.3). Imagery shown ©2009 Google Earth



Figure 6.24. Nadir view of post-event Buncefield image with detected cylinders overlaid ©2009 Google Earth

6.7 SUMMARY

This chapter has advanced the method described in chapter 5 for two-dimensional object detection to automatically detect cylinders from remotely-sensed imagery with an oblique observational zenith angle. The conceptual model described in chapter 4 has been enhanced to provide a semi-automated, model-based approach for the extraction of three-dimensional objects from an oblique two-dimensional image, fulfilling several of the key aims of this thesis (Section 1.5.1); namely to iteratively identify objects of interest from imagery, using a model-based platform to guide the approach. The number of detected circles was automatically reduced using a rule stipulating that whole footprints of objects should be within the bounds of the image. Target features were identified by amalgamating primitive geometric shapes using a combination of similarity matrices and filters optimised using a simulated annealing heuristic.

Preliminary investigations were carried out to model the results of the final set of objects stored in the scene model (Section 6.6). If a realistic scene can be created, taking into account sensor and illumination geometry, the results of the object detection phase of the conceptual model can easily be disseminated to expert and non-expert users. The resulting data can also be input into a change detection process. This is investigated further in the following chapter using the results of the correct detection of the three objects seen in Figure 6.20.

CHAPTER 7

DETECTION OF OBJECT CHANGE

7 DETECTION OF OBJECT CHANGE

7.1 INTRODUCTION

Previous chapters in this thesis have described the development of the first phase of the conceptual model, *object detection*. This phase has the requirement to extract target objects from fine spatial resolution remotely-sensed imagery. Inputs to the model are images of the same geographic area captured before and after a change event, providing observations of normal changes in urban morphology, or a more sudden change event, such as a natural or anthropogenic disaster. The pre-event image was used to extract a scene model using the method outlined in chapters 5 and 6.

This chapter describes the initial development and testing of the object comparison phase of the conceptual model. The approach uses the scene model information created section 6.3 (SM_1) as an input, which is then modified using the image geometry from the *after* image to represent what the objects in SM_1 would look like in unchanged form in the post-event image. This process results in a second, post-event scene model (SM_2) that can be compared with SM_1 to generate change information. In contrast to conventional approaches to change detection, the method produces change information that can be described in both quantitative and qualitative terms, providing a holistic description of object-oriented change between the two time periods. The following section describes the methodology in full, with the results presented and discussed in sections 7.3 and 7.4 respectively.

7.1.1 The Input Scene Model

Chapter 6 described the method adopted for the semi-automated extraction of three-dimensional object information. This process resulted in the generation of SM_1 which was used to describe the location and spectral attributes of the detected

objects. The scene model describing the cylindrical objects detected from the Buncefield image, contains the following data:

1. The locations of the centres of the top and base circles. These are provided in both image and geometric coordinates.
2. The radius of each detected circle.
3. The horizontal displacement of the top and base circles in the image space.
4. The object heights in metres calculated from 3, above.
5. The spectral properties of every pixel within the projected object footprint.

The methodology in this chapter was developed and tested using the correctly identified cylinders detected in the object detection phase of the conceptual model (Figure 6.20), and uses the premise that further refinements to the object detection phase of the model (described in section 6.5) may improve the frequency of correct identification of these objects. Alternatively, the SM_1 used as input to this component of the conceptual model can be generated from vector data describing the objects in their pre-changed state. Section 6.6 described how the information can be accumulated to form 3-D models that can be used in several tangible real-world applications, such as building inventory creation. The pre-event scene model was also used to predict what unchanged objects would look in an image with comparable viewing geometry to the post-event image.

7.2 METHOD

Figure 7.1 describes the workflow adopted for the generation of change information from comparison of the pre- and post-event scene models. This section describes the workflow and each processing stage in detail.

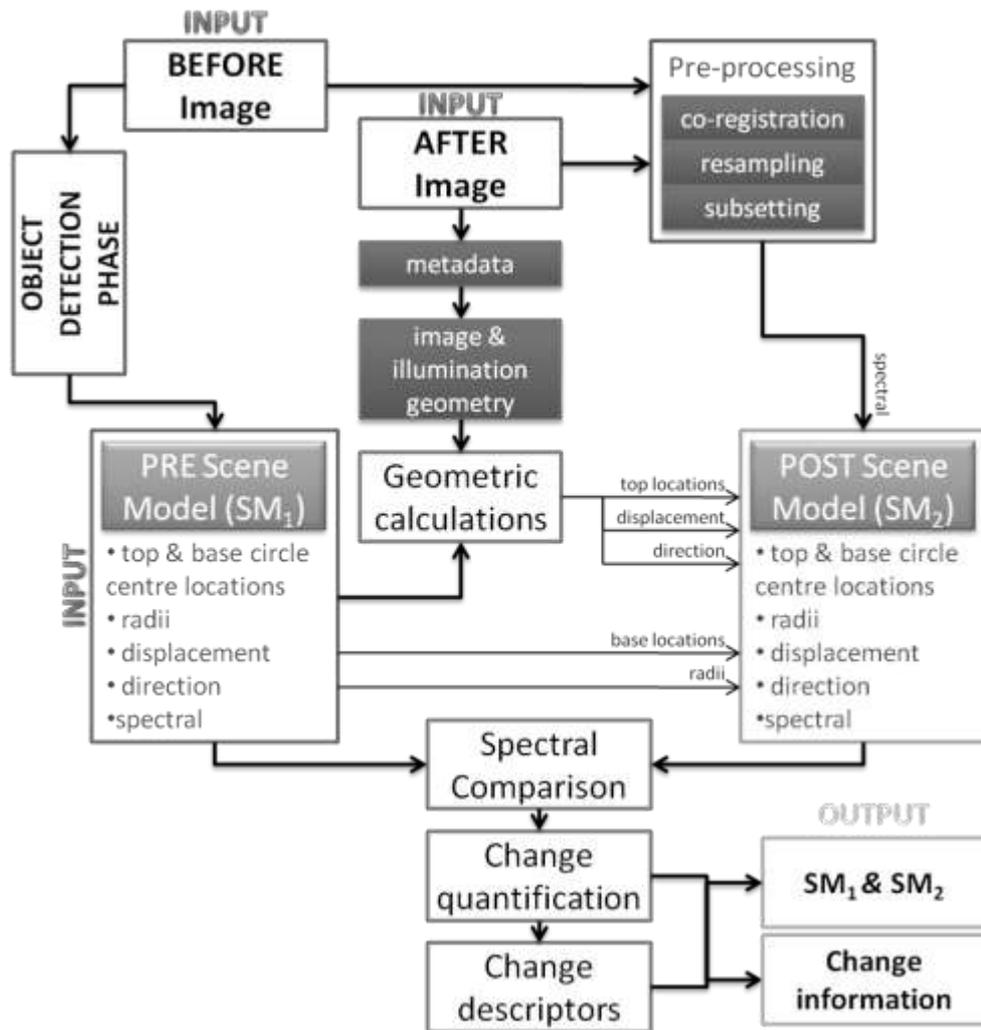


Figure 7.1. Workflow of change detection methodology

7.2.1 Image Preparation and Pre-processing

Generation of change information is reliant on comparable geographic extents of the datasets for both the pre-and post-event images. In this case, both images were delivered in a geo-registered state in British National Grid Projection, as described in chapter 4. The images were checked and found to have residual error in the geo-registration of some image features. The images were subsequently co-registered: the 2006 pre-event image was warped to match the post-event image to a root mean square error of 1.08 m. Visual assessment of the registered image showed adequate registration of the area of interest. The pre-event image subset used in chapters 5 and 6 was used to test the change detection methodology in Figure 7.1, alongside the corresponding post-event image subset (Figure 7.2). The post-event image was



Figure 7.2. Subset of the resampled post-explosion image of the Buncefield oil depot (2006)

resampled to 25 cm spatial resolution to match the pre-event image, facilitating direct comparison between the two time periods.

7.2.2 Grounding the Object Bases

The locations of the bases of the objects from the pre-event scene model (SM_1) were used to project the footprints of the object bases in the post-event scene model (SM_2). Base circles were located in SM_2 using the image coordinates from SM_1 to position the circle centres with the radius from the top circle in SM_1 used to project circle perimeters.

7.2.3 Projection of Top Circle

The base information of each cylinder was used to locate the relative position of the top of the cylinder in SM_2 . This projection required a dual stage approach that generated information on how far away from the base circle it should be projected, and also, in which direction. The base information from SM_1 was required as input into this process.

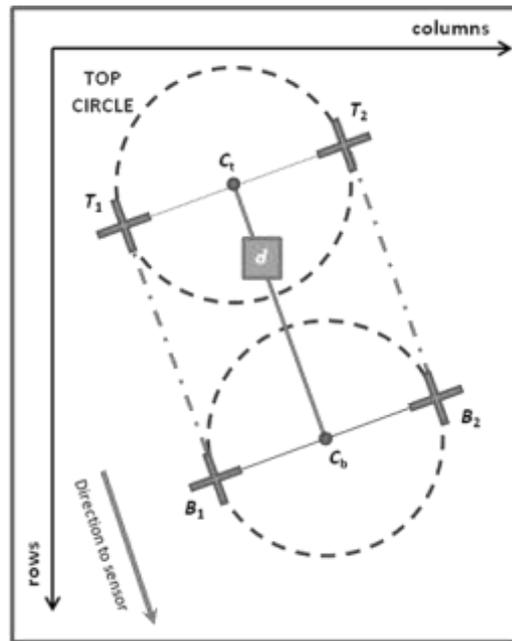


Figure 7.3. Cylinder geometry showing displacement (d) of top and base circles

Displacement

Firstly, it was necessary to calculate the horizontal displacement, d , between the top and base circles (C_t and C_b respectively) as observed in the image (Figure 7.3). The magnitude of d differs between images depending on several image properties, most significantly the observation zenith angle. As the angle increases (i.e. the image's 'obliqueness' becomes greater), the displacement between projected positions of C_t and C_b increases. Figure 7.3 shows the geometric calculations needed to determine d using the height information generated from SM_1 . The height, h , is equivalent to the hypotenuse of the right-angled triangle with angle θ equal to the observational zenith angle. The side d , lies on the image plane with its length calculated by:

$$d = h \cos \theta \quad (7-1)$$

Displacement information was generated for each cylinder in turn with the resulting data used in tandem with a direction variable to calculate the position of the top circles in the new SM_2 .

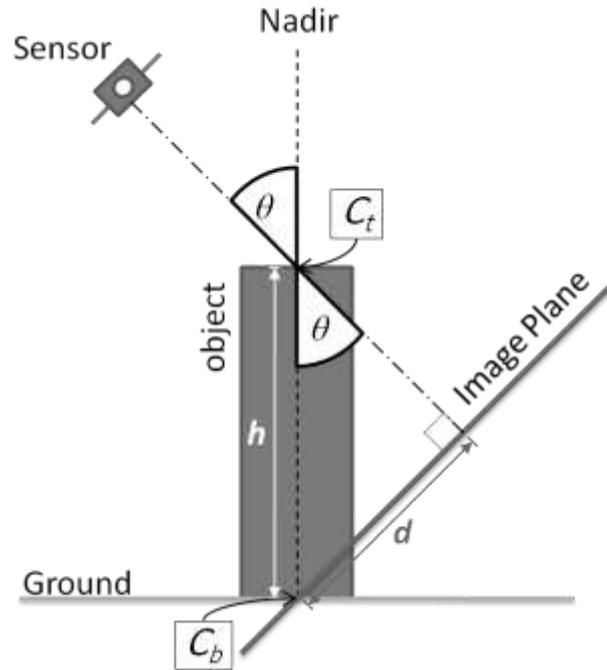


Figure 7.4. Calculation of the magnitude of projected displacement (d) between top, C_t and base, C_b of an object. Also shown is the object height, h

Direction

Secondly, directional information to describe the horizontal displacement of the top circle with respect to the base was calculated using metadata from the post-event image. The ground position of the sensor (x, y) was converted to image coordinates using a MATLAB function *map2pix*, enabling the generation of object-specific sensor azimuth angles (i.e. the direction from the sensor to each object). Figure 7.5 outlines the direction geometry, with each object sensor azimuth angle (D), calculated by:

$$D = 180 - \theta \quad (7-2)$$

where

$$\theta = \tan^{-1} \left(\frac{\Delta x}{\Delta y} \right) \quad (7-3)$$

when Δx and Δy are the differences in x and y positions respectively for the geographic locations of the sensor and object.

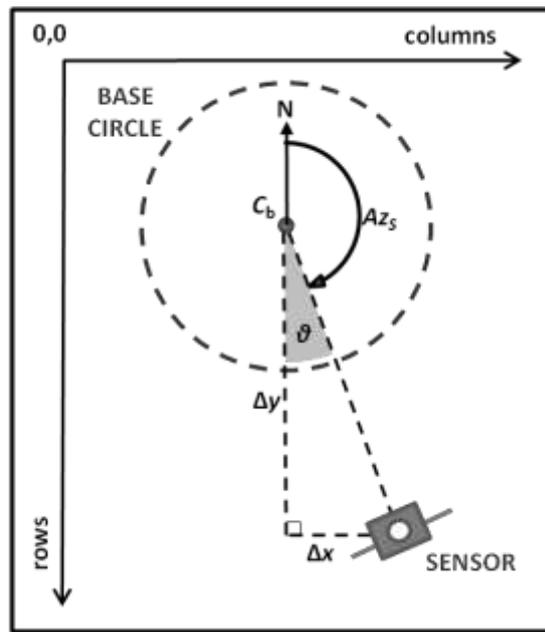


Figure 7.5. Calculation of the direction of displacement

Positioning the cylinder tops

Using the calculated distance and direction information, the top circles for each object were projected into their new locations in SM_2 , based on the post-event image metadata.

7.2.4 Forward Modelling

The positions of the top and base circles of each object were used to generate the sides of cylinders using the methodology described in section 6.3.7. The sides of each cylinder were equal in length to the object displacement, d , calculated in 7.2.3 and were projected in the same direction as the sensor azimuth angle, Az_s , to each object. The forward-model simulation of the post-event scene was designed to project unaltered objects into a simulated image with equivalent image geometry to the post-event image.

A group of primitive features (sides, top and base circles) for each cylinder were modelled as one object with a mask created for individual objects, as well as the collective group of all objects. These object masks were used to create

simulated images by assigning spectral values to each pixel: A forward model assigned a value of zero to those pixels outside each object, with the internal pixels assigned the spectral values equal to the post-event image. Simulating each cylinder in SM₂ allowed intra-object statistics to be generated to describe all pixels within the object image footprint. These statistics facilitated the generation of per-object change information between SM₁ and SM₂.

7.2.5 Quantification of Change

Summary statistics

Statistical change information was generated from a comparison of SM₁ and SM₂ as described in Figure 7.1. The simulated images created in the previous section were used to generate the following statistics for each detected cylinder:

- i. Minimum and maximum spectral values
- ii. Mean and modal spectral value within each object footprint
- iii. Standard deviation, σ_o , of intra-object pixel brightness values,

$$\sigma_o = \left(\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \right)^{\frac{1}{2}} \quad (7-4)$$

where x_i is the brightness value for pixel i , and \bar{x} is the mean brightness value of all n pixels within the object. These statistics were also calculated for the pre-event simulated image with the results of both described in section 7.3. The differences in the statistics gained for SM₁ and SM₂ generated change statistics for all objects correctly detected by the automated 3-D object detection described in chapter 6.

Texture metrics

The objects stored in SM₁ and SM₂ were used to calculate intra-object texture describing the variability in neighbouring pixels within each projected footprint. Firstly, a 3×3 window was passed across the whole of the *before* and *after* input images calculating the range and standard deviation of the 8 pixels adjacent to the

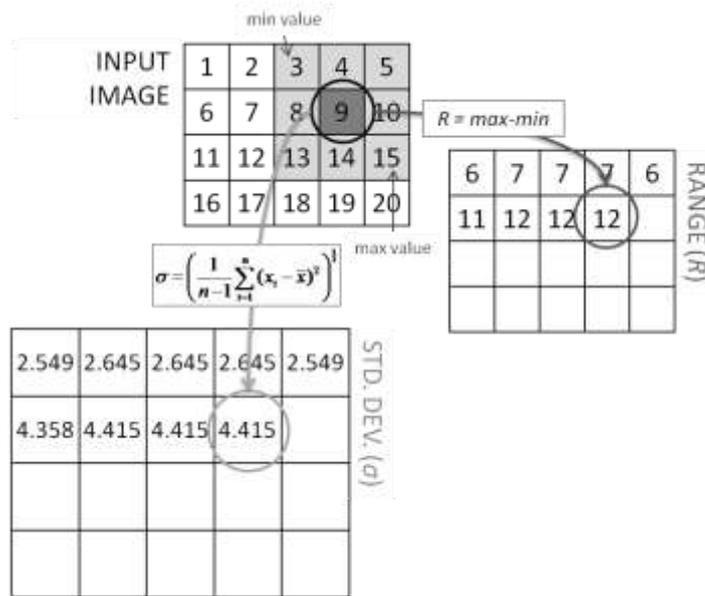


Figure 7.6. Calculation of per-pixel texture information

centre pixel (Figure 7.6). This resulted in per-pixel values for the range, R , of $[i_{\min}, i_{\max}]$, and standard deviation, σ , describing the homogeneity of pixel brightness values in each of the objects. These parameter spaces were independently masked for each object and the intra-object pixel texture values accumulated as a measure of per-object pixel brightness variability.

7.2.6 Describing Change

Whilst quantitative statistics of brightness change were generated automatically by the coded MATLAB model, a manual approach was used to describe the object change over time in qualitative terms. The aim of this process was to provide a thorough interpretation of the change in objects between the 1999 and 2006 Buncefield images. The detected change was classified according to several existing damage scales: the European Macroseismic Scale (EMS-98) and the combined Wind and Flood Scale (WFS). Both scales have been used in the literature to assess building damage in the aftermath of natural disasters, using remotely-sensed imagery as a data source (Miura *et al.*, 2005; Womble *et al.*, 2006). These scales were chosen to investigate how amenable the resulting data from the object

comparison phase is to established damage scales, albeit for natural disasters, rather than the anthropogenic event seen at Buncefield. This resulted in assigned damage classes for each object in the scene model.

7.3 RESULTS

The change statistics generated using the methods described in the preceding section show

- i. Object changes due to differences in image geometry
- ii. Object differences described in quantitative and qualitative terms

Firstly, Figure 7.7 shows the position of the sensors in relation to the study area. Figure 7.8 shows the projections of the object footprints on both the (a) pre-event and (b) post-event images. Base coordinates and radii were used to locate the bases of the objects in SM_2 and as such, share common image coordinates as SM_1 . Spectral values from the respective pre- or post-event images were applied to the object pixels in Figure 7.8. It is apparent that the differences in observational zenith angles between the images (shown in Figure 7.7) had a dramatic effect on how the objects were projected in both images.



Figure 7.7. Locations (nadir points) of the *before* (green) and *after* (red) image sensors in relation to the study area (blue)

Table 7.1. Object location data from the pre- (a) and post-event (b) Scene Models

a	Object number	Base x	Base y	Top x	Top y	Radius	Direction	Displacement
		(image coordinates)				(pixels)	(degrees)	(pixels)
PRE-EVENT SCENE MODEL	1	109.62	99.81	92.26	64.86	48	153.59	39.02
	2	111.64	271.61	99.53	247.07	28	153.73	27.36
	3	233.72	259.85	219.16	231.30	31	152.97	32.05

b	Object number	Base x	Base y	Top x	Top y	Radius	Direction	Displacement
		(image coordinates)				(pixels)	(degrees)	(pixels)
POST-EVENT SCENE MODEL	1	109.62	99.81	107.79	95.96	48	154.61	4.26
	2	111.64	271.61	109.40	269.63	28	131.42	2.99
	3	233.72	259.85	233.33	256.37	31	173.65	3.50

Table 7.1 describes object properties for both SM_1 and SM_2 , with object number 1, 2 and 3 relating to the top left, bottom left and bottom right cylinders respectively (in Figure 7.8). The results of the statistical analysis described in section 7.2.5 are shown in Table 7.2 and Figure 7.9, and characterise objects in their post-event state. This figure is presented with the post-event objects normalised against their pre-event states. A textural comparison of objects in SM_1 and SM_2 is

Table 7.2. Per-object spectral data from SM_1 and SM_2 . Units are greyscale brightness values.

SM	Object number	Minimum	Maximum	Mode	Mean	Range	Standard Deviation
SM_1	1	67.00	252.00	251.00	213.28	185.00	47.1706
SM_1	2	62.00	253.00	250.00	201.41	191.00	48.3515
SM_1	3	97.00	253.00	250.00	210.36	156.00	41.8736
SM_2	1	22.00	234.00	54.00	81.84	212.00	39.7417
SM_2	2	26.00	249.00	47.00	72.95	223.00	39.4096
SM_2	3	19.00	215.00	42.00	69.74	196.00	41.8338

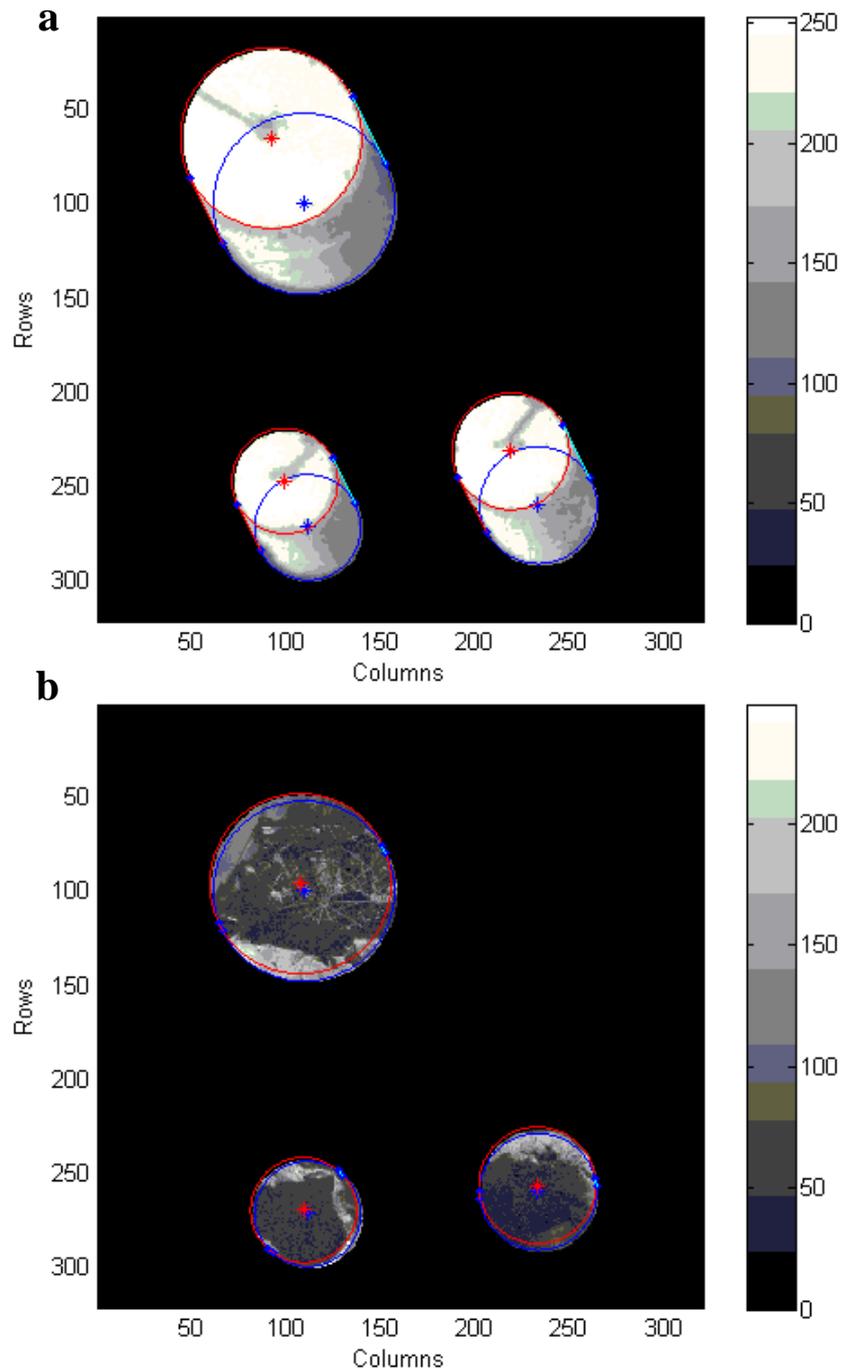


Figure 7.8. Pre-event (a) and post-event (b) Scene Model objects projected in greyscale with 1999 and 2006 image geometry respectively, and including intra-object spectral properties. Object bases (blue) and tops (red) are overlaid with their sides.

described in Figure 7.10, in terms of the standard deviation of both the (a) range and (b) standard deviation parameter spaces created by the method described in Figure 7.6. Qualitative descriptors generated from existing natural disaster damage scales are shown in Table 7.3.

Table 7.3. Qualitative descriptors of change from a manual interpretation using the EMS-98 and the WFS Damage Scales

Scene Model	Object number	EMS-98 Damage Level	Description	WF Scale Damage Level	Description
SM ₂	1	Grade 5	Destruction (very heavy structural damage):	WF-4	Destruction:
SM ₂	2	Grade 5	Collapse of ground floor or parts (e.g. wings) of buildings.	WF-4	Complete roof failure and/or failure of wall frame. Loss of more than 50% of roof sheathing
SM ₂	3	Grade 5		WF-4	

A visual inspection of the imagery captured both before and after the Buncefield explosion shows significant differences in both the spectral and geometric attributes of the target objects. The three oil drums that existed in the pre-event image had a bright spectral response, with a mean brightness value ranging from 201 to 213, and a modal brightness value for all objects of over 250. The explosion and fire resulted in a significant amount to damage to each drum, resulting

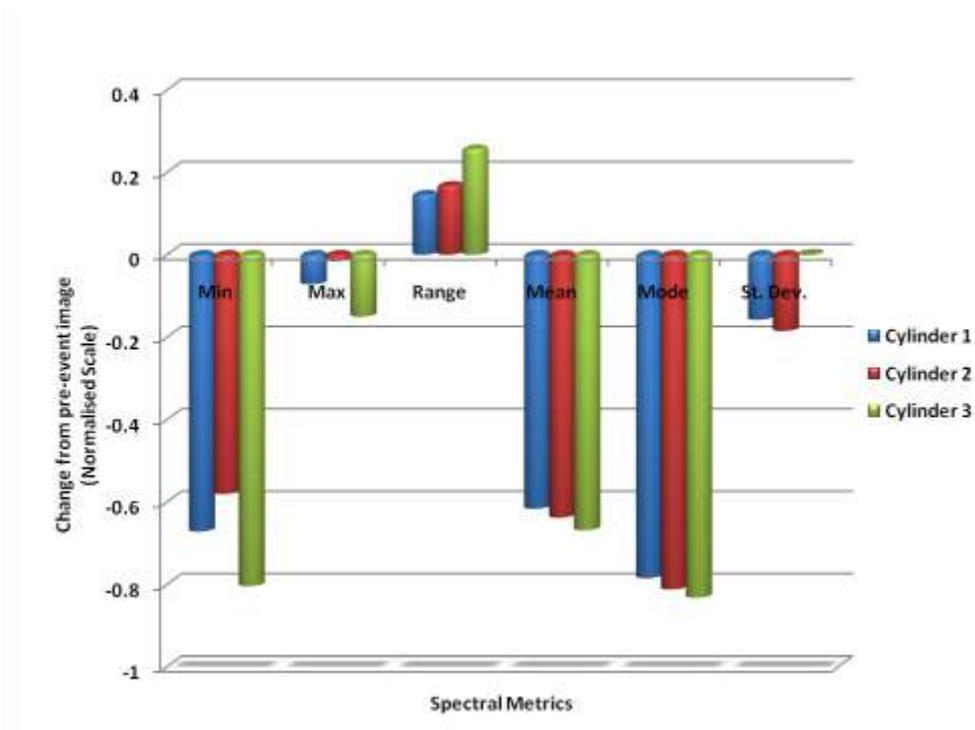


Figure 7.9. Display of the six change measures used for description of intra-object pixel brightness changes

in much darker objects with more fragmented object boundaries.

The following section provides a discussion of the change detection results, formatted to follow format described in section 7.3.

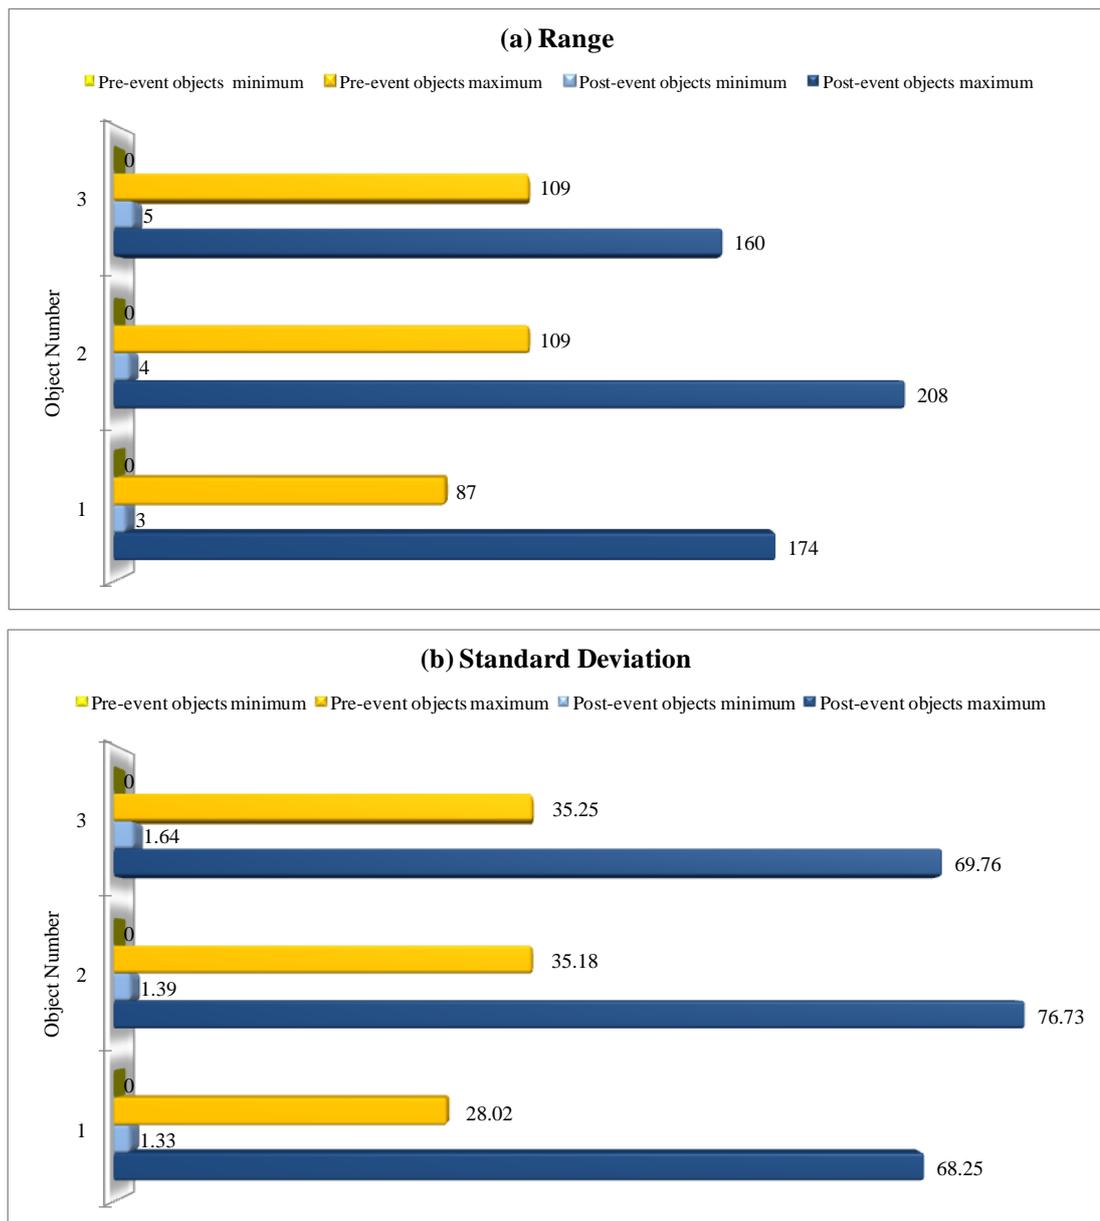


Figure 7.10. Intra-object textural comparison.

Pre-event objects are shown in yellows, with post-event objects in blue shades. The minimum and maximum values are shown for: (a) Range: Variation in the range of per-pixel values within each cylinder. (b) Standard deviation of pixels within each of the three cylinders.

7.4 DISCUSSION

7.4.1 Object Changes due to Differences in Image Geometry

Changes in the image viewing geometry affected the displacement between top and base circles; decreasing as the observational zenith angle decreased from 10.4° to 1.1° from the 1999 to 2006 image. The range of sensor azimuth angles in relation to the objects is shown in Table 7.1 and shows a variation of 0.66° between all pre-event objects and the sensor location of the before image. A considerably greater range of 42.4° was observed from the post-event image. Figure 7.7 shows the relative locations of both sensors to the study area, and this discrepancy in ground distances accounts for both the more oblique observational zenith angle in the pre-event image and, hence, the greater range of observational azimuth angles of the post-event objects. Despite these geometric differences between the pre- and post-event images, it is apparent the model-based method was able to compare the same objects in the imagery, irrespective of the size or orientation of the projected object footprints. This is a potentially very useful capability as it allows the use of imagery captured at oblique angles for damage assessment, or inventory update. Following natural disasters, there are often occlusions to the affected area (e.g. smoke), that are consequential of the event type. The ability to capture imagery using optical satellite sensors is reduced during the occluded period, and initial damage assessments use aerial imagery captured from helicopters or aeroplanes. A modelled methodology insensitive to sensor zenith and azimuth angles has a greater ability to utilise these valuable datasets at a time when data is time-critical. Further testing of the described method is necessary before its generalisation to a variety of zenith and azimuth angles can be fully established.

7.4.2 Object Differences Described in Quantitative and Qualitative Terms

Quantitative results

Change statistics are shown in Table 7.1 and Table 7.2, with the comparison of objects between pre- and post-event image dates shown in Figure 7.9. This shows the comparison of objects in their pre- and post-event states normalised with the unchanged pre-event cylinder state plotted as zero. The range in values is the only

metric to have increased from its pre-event value, with an increase of between 15% and 26% in spectral ranges for cylinder 1 and cylinder 3 respectively. Large decreases in spectral values within object footprints are apparent from greater than 50% decreases for minimum, mean and mode spectral values for all cylinders, while maximum spectral values are observed to be 7-16% less than the pre-event object states. An increased range of pixel values between the before and after image dates suggests greater variability in pixels within objects. This trend is also observed in the results of the texture metrics. Figure 7.10 shows the variability in range and standard deviation texture descriptors. Local variation between neighbouring pixels in the SM_1 objects was consistently $\geq 50\%$ less for the changed objects. The range in the nine neighbouring pixels (as calculated from Figure 7.6), is also consistently greater for post-event objects.

Figure 7.10 shows a greater than 200% increase in average standard deviation texture values from pre- to post-event states due to the change in the oil drums from bright, spectrally homogeneous cylinders, to a tangle of scorched metal and assorted debris seen in the image location where cylinders should have been found. The spectral gradient between the cylinders in the *after* image has been reduced as the white paint on the cylinders was burnt off during the fire. A discrepancy in image dates should also be acknowledged as the influence of solar illumination is less in the January 2006 image than is seen in the May 1999 image. Much of the centre of the damaged cylinders is obscured by shadow and is partly responsible for the reduction of minimum, mean and mode brightness metrics. The standard deviation and range metrics used to describe intra-object texture were implemented automatically on the objects stored in SM_1 and SM_2 . The model-based approach facilitated this object-oriented analysis of two primary measures of spectral variability within each object. These measures are generalisable across image types, and could be used for any object-based comparison using RS imagery.

The complexity of features within the object footprints in the post-event image resulted in increased textural values for both metrics described. Again, the damage from the explosion and fire increased the geometric complexity of the objects, and therefore the homogeneity of neighbouring pixels was lower than for the pre-event image.

Qualitative results

Qualitative damage classes are shown in Table 7.3. As discussed at the outset of this section, the three objects changed dramatically in both spectral and geometric terms, and were therefore assigned to the “Destruction” damage classes of both the EMS-98 and WFS damage scales. In both cases, this equates to complete collapse of buildings, as observed in the aftermath of the Buncefield Oil Depot fire. No damage scales have been found that explicitly describe *building* damage in the aftermath of industrial accidents. However, a European Scale of Industrial Accidents (ESIA) does exist, and ranks multiple effects of industrial accidents in terms of the following criteria:

- i. Dangerous materials released
- ii. Human and social consequences
- iii. Environmental consequences
- iv. Economic consequences

In terms of these four factors, the Buncefield accident was described by the following classifications on the ESIA:

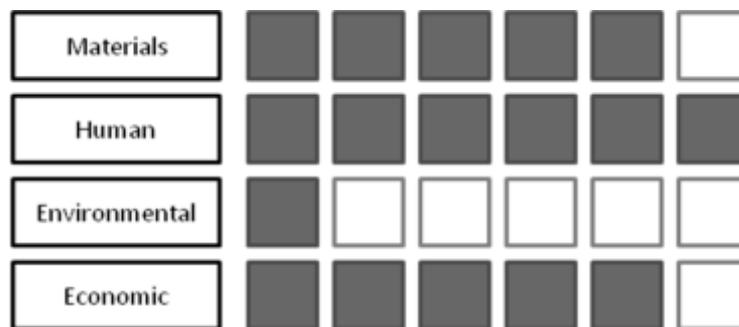


Figure 7.11. Damage to the Buncefield Oil Depot as classified by the European Scale of Industrial Accidents (ARIA, 2009)

The *Human Consequences* and *Economic Consequences* are the classes that most closely relate to structural damage. Both cases were assigned the most severe damage levels (6 and 5 for Human and Economic, respectively) and so provide a form of validation for the EMS and WFS change classifications. It should be noted however, that the Buncefield accident appears to have been over-estimated in this classification, especially in terms of human and economic consequences, as there were no fatalities, and only 43 injuries from the event (ARIA, 2009).

Two scales were used to describe the extent of change in the post-event objects. In reality, the scale of damage seen in Buncefield was catastrophic, and using proxy scales such as EMS and WFS was appropriate in this instance. However, when describing subtle changes in objects, it is uncertain whether the proxy scales would be suitable, despite their effective usage in remote sensing applications described in the literature. The qualitative results resonate with the recommendations of Huyck *et al.* (2004a), who stated that distinguishing total collapse can be achieved from remote sensing image analysis alone. However, it is expected that more moderate levels of damage would be more difficult to distinguish solely from remote sensing. The lack of a damage scale developed for anthropogenic disasters is notable, however this is understandable because of the wide variety of catastrophic events that could befall a city.

7.5 SUMMARY

The objects detected through the methods described in chapters 5 and 6 have been used as the input into a forward modelling method described in this chapter. Objects stored in the pre-event scene model were represented in unchanged form, using the metadata associated with the post-event Buncefield image. This image was very different to the pre-event image in terms of the geometric and spectral responses of the cylinders in both images. A method was detailed to quantify the change and also generate descriptors through the use of pre-defined damage scales widely used in natural disaster damage assessment, and so fulfil rationale number 6 (Table 1.1), which described the need for providing automated methods for describing change, both by quantitative means and through the use of existing descriptors.

The analysis of change presented in this chapter focused on a high-level quantification and description of spectral and textural differences both at an object level and also in terms of intra-object pixel variability. As a proof-of-concept for the novel idea of modelling the pre-event objects as they would appear were they unchanged in the post-event image, the results displayed show great potential for future advancement of the forward and inverse modelling approach detailed in the conceptual model. There is scope for furthering this analysis to include the

description of geometric changes in objects, allowing the update of the post-event scene model to display a geometric representation of the objects in their post-event state. This is discussed in greater detail in the following chapter.

CHAPTER 8

DISCUSSION

8 DISCUSSION

8.1 BACKGROUND

The driving forces for the research detailed in previous chapters were described at the outset of this thesis, with two major application areas identified: monitoring the change and development of urban areas, and identifying damage following natural or anthropogenic disasters. These applications are extremely relevant in today's world, as global populations are increasingly clustered in urban centres, which can be at risk of planned or unexpected change events. The thesis investigated whether a single model could be used to detect change for both applications using exemplars to test the potential of the conceptual model as a viable method of object detection and change assessment.

The preceding chapters have introduced and described the evolution of a conceptual model for change detection with two central phases. The model was designed to detect objects from optical aerial or satellite sensor images (from the earliest image used), using methods grounded in the fields of machine vision and image processing. These objects can represent buildings or other phenomena of interest and are extracted to and stored within a pre-event scene model (SM_1). A forward model is applied to this scene model, simulating their appearance if unchanged, and viewed using the sensor geometry of the post-event image (stored in SM_2). The comparison of SM_1 and SM_2 generates a variety of change measures, which can be described using qualitative terms, and existing hierarchies of change descriptors, such as damage scales.

A discussion follows of the merits, assumptions and limitations of the proposed conceptual model and the exemplars used to test the model phases. Initial discussion focuses on a detailed analysis of the conceptual model described in chapter 3, followed by an analysis of each aim and objective described in chapter 1, to determine the extent to which the aims were met. This chapter concludes with an

examination of the requirements necessary to advance the model towards full and robust implementation.

8.2 THE CONCEPTUAL MODEL

The conceptual model described in chapter 3 was designed with two major components in mind: object detection and object comparison. The need for change assessment information may arise in locations where no geographic data exist, save for archived fine spatial resolution images captured from the numerous satellite sensors currently in orbit. The object detection phase of the conceptual model was designed for such circumstances to facilitate the generation of pre-event object information. Subsequently, a second, post-event image could be captured following a change event (e.g. after a sudden disaster, or simply at a later date than the *before* image) for use in the object comparison phase of the model to describe the changes of the target objects. This section examines the evolution of the model through the aforementioned studies in the earlier chapters of the thesis.

8.2.1 Object Detection

Chapter 5 described the development of a feature extraction technique for use in phase 1 of the conceptual model. The model was designed to utilise an existing algorithm from the field of machine vision, and extend its functionality to allow a greater degree of automation than existed previously. In the described exemplars, a Hough Transform was adopted to detect circular features in imagery, and tested on a series of datasets. The HT algorithm was chosen due to its widespread implementation in a variety of previous applications, as detailed in chapter 2, ranging from biscuit manufacture (Davies, 1984) to the detection of lunar craters (Sawabe *et al.*, 2006). A robust and well developed method of feature extraction, it relies on using edge information to delineate object boundaries. The raw algorithm (*HT*-) would be capable of delineating objects in an urgent situation (such as a

natural disaster) due to its model-based format. However, it requires prior definition of several parameters that affect its performance, and *a priori* understanding of these parameters would help optimise the method, yet time would be lost through learning and refining parameters. The modified *HT+* method provides a development for the method as initial parameters are not required to be optimal due to the iterative updating process of the evidence-gathering accumulator space. This “boosted” the evidence through a forward and inverse modelling process designed to reduce the total evidence for false positive detections and synthetically increase the evidence for false negatives.

The initial implementation of the *HT+* extension on the Buncefield image was limited by the input imagery characteristics of the pre-event image and the definition of the objective goals: The primary objective was the detection of the tops of the oil tanks and was limited in its effectiveness by the oblique viewing angle (observational zenith angle of 10.4°), and the spectrally homogeneous properties of the objects. The forward modelling simulation of the candidate circles used spectral brightness values from the 8-bit greyscale input images. The modal brightness value was used in the simulated image for comparison with the observed, pre-event, image. This segmentation relied on the objects having steep spectral gradients across the edges of the tops of the objects; without which, the HT was less likely to correctly identify the cylinders.

In the Saudi Arabia example, the extension of the *HT+* was able to detect 8 circles that had not been identified by one iteration of the *HT-*. The *boosting* algorithm simulated enough additional evidence for the algorithm to detect partial circles, despite the objects not having totally homogeneous spectral characteristics. This result is encouraging, as it shows how the *HT+* algorithm can be used to detect false negative objects in the previous iteration through an iterative forward and inverse modelling method.

The choice of the Buncefield images exposed the limitations of attempting to detect solely the tops of objects when the observational zenith angle is such that the sides of the objects are also visible. This issue was noted by Lin and Nevatia (1995), who observed that the contrast between roof and wall may be lower than the object-to-background contrast in oblique images. The Buncefield data were, however, valuable in extending the functionality of the conceptual model from the detection of

two-dimensional to three-dimensional objects. Again, the HT was used for this process, with the iterative boosting extension that had previously updated the accumulator space, replaced with a scoring and optimisation process. The score was based on candidate matching of random pairs of circles and assessing the relationship between top and base circles in terms of distance, direction, radii and spectral properties. This approach has merit as it used the geometric image properties (incorporated in most image metadata) to constrain the number of potential matches that could be made. This reduced the search space for solutions, and therefore, contributed to speed up the processing required for the heuristic to converge on a minimum value. This stage can also be tailored, depending on the objects to be detected in the imagery, where variables can be removed or added without disruption to the system.

Thus far, the implementations relied on several fundamental assumptions that were adopted to progress the development of the conceptual model. These were apparent in the detection of the tops of the Buncefield oil tanks, where the object detection method relied on the target objects displaying homogeneous surface brightness characteristics. Using the HT as a feature extraction method was dependent on the method's ability to distinguish edge boundaries of objects. The latter limitation was reduced with the *boosting* extension, as the minimal evidence gathered for false negative features was adjusted artificially, allowing the objects to be detected in later iterations of the *HT+* algorithm.

The three-dimensional object identification described in chapter 6 assumed that the target cylinders had both vertical sides and flat tops; this was a slight simplification of the reality. However, the Score Matrix developed to match top and base circles in the cylinder creation process can be adjusted to allow variation in the top:base radius ratio, to detect objects that have non-vertical sides. The 3-D object detection described for the Buncefield example was effective due to the off-nadir observational zenith angle of the 1999 image. The method in its current form is designed for use on this type of oblique imagery. However, it has potential to be extended using additional image features such as shadow information for gathering height evidence from the image (Sarabandi *et al.*, 2008b).

8.2.2 Object Comparison

The research described in chapter 7, for the object comparison phase of the conceptual model, was limited in scope. However, it was included in this thesis to highlight the potential of the scene model approach for comparing detected objects. One of the strengths of the conceptual model is its ability to compare features captured within two images that are viewed from different observational zenith and azimuth angles. Chapter 7 showed that despite irregularities in the sensor observational angles, the objects were modelled effectively using the object characteristics stored in the scene model. This object information was attributed with the sensor geometry of the *after* image, and in doing so, create a ‘perfect’ scene model using the assumption that the objects are unchanged during the time period between the images. This provided baseline information that was compared to the post-event objects in the *after* image, enabling the generation of basic statistics describing the changes between the objects as they appear in the pre-and post-event imagery. To humans, understanding and describing the scenes offered by the Buncefield image pair comes intuitively. However, the conceptual model allows the automated extraction of these basic statistics, which, when extended to a larger geographic study area, has potential to be used to describe major changes in buildings between two time periods.

8.3 AIMS AND OBJECTIVES

Several aims and objectives were detailed in the rationale points set out in chapter 1, and have formed a reference for all methodological discussions to this point. These objectives were created from a review of the current state-of-the-art described in the literature (chapter 2). This section focuses on each point in turn and investigates the extent to which the research described through exemplars in chapters 5-7 has met these objectives.

8.3.1 Objective 1: Provide a framework or guidelines for processing of RS images

The first objective of this thesis was to develop a conceptual framework for the processing of remotely-sensed images. In time-pressing scenarios where results are required under often stressful conditions, a set of guidelines delivered through a methodological framework can streamline processing workflows. The conceptual model detailed in chapter 3 demonstrates a dual-phase workflow designed to be adaptable to meet individual project requirements. The results demonstrated in chapters 5-7 show the candidature of such a framework for change detection analysis of multi-temporal remotely-sensed images, and it has been tested for fundamental geometric shapes amalgamated to create more complicated objects within a scene model. The choice of imagery allowed the development of the method for extracting two and three-dimensional object information, however the image subset used in the Buncefield example only contained three objects to detect.

The conceptual model is designed to be flexible in many dimensions (e.g. the choice of object detection algorithm, the type of input imagery used, and the geometric properties of the input images). The exemplars described in this thesis have demonstrated how circular and cylindrical features can be automatically detected and modelled to capture rudimentary change information. These implementations were shown in chapters 5-6, with automated measures of per-object, and intra-object, changes described in chapter 7. Few prior assumptions were necessary for these implementations, with a basic knowledge of the shapes of the objects of interest (i.e. circles and cylinders) guiding the choice of algorithm used (i.e. circular HT). A full implementation of the framework could be envisaged to include a toolbox of algorithms for the detection of a range of object complexities, dependent on the event-specific application of the model. Similarly, the object comparison module has the potential for including more comprehensive change indicators and classes, as described further in section 8.4.2.

8.3.2 Objective 2: Provide methods that identify objects of interest and provide pre-event information as a basis for change detection

This rationale point provided motivation to develop and test methods that can be used to identify and extract features from imagery. Pre-event vector information for building footprints is not always available, and alternative sources of baseline information are therefore required. Remotely-sensed imagery can provide a viable source, if suitable extraction techniques exist. The fields of machine vision and computer science offered several extraction techniques, some with previous evidence of application to Earth observation images: The HT has been used in the literature to detect circular geologic structures (Cross, 1988), hydrocarbon seepages (van der Werff, 2006) and man-made structures (Wang *et al.*, 2004; Liu *et al.*, 2005). Methods that utilise deformable models, such as active contours were used to detect buildings in previous applications (Trinder, 1995; Mason and Baltasvias, 1997; Neuenschwander *et al.*, 1997), while basic edge detection approaches are widely used in the pre-processing stages of feature extraction from remotely-sensed imagery (Frère *et al.*, 1995; Lin and Nevatia, 1995; Ali and Clausi, 2001).

The examples of object extraction in this thesis have used a Hough Transform-based method. This circular feature extraction algorithm was chosen specifically for the Buncefield and Saudi Arabia imagery. Techniques have been demonstrated through exemplars showing the potential for object detection techniques to be modified and extended to increase the detection rates of objects through automated and iterative updating techniques, with the *HT+* providing an example of this for circle detection. Evidence has also been shown that there is much potential for existing feature extraction techniques to be used within the confines of the model framework, although additional testing is required with other feature extraction algorithms before the robustness of the model framework can be truly understood. Disaggregating objects into primitive features allows the use of simple geometric identification algorithms such as edge-based methods, and avoids the complex parameterisation of shape matching methods that would be limited in their ability to identify multiple object shapes. The scope for testing additional object extraction techniques is considerable, and would be dependent on the target features of interest within the image. However, it is clear that the development of

suitable feature extraction methods will go some way to unlocking the latent and ever-increasing archive of worldwide remotely-sensed imagery.

8.3.3 Objective 3: Provide a model-based platform for per-building analysis

This objective aimed at focusing the methodological development of the feature extraction phase on the identification and amalgamation of primitive geometric features. The utility of these methods is well established in the literature and was seen in the detection of buildings through the identification of lines and edges (Frère *et al.*, 1995; Brenner *et al.*, 2001). Objects could be detected using multi-dimensional methods that detect complex features as a whole, however, a primitive amalgamation approach was chosen to allow generalisability across object and image types, with the application to circular and cylindrical features demonstrated. Developing a system generalisable for use in multiple images required a method that could reduce each feature of interest into multiple parts that can be described by one or two parametric dimensions only. This objective was met in the implementations discussed in chapter 6, through the detection of circular features and complementary linear components which together created a representation of cylinders. This primitive amalgamation approach allows the adjustment of the relationship between top and base circle radii for the detection of objects whose sides do not have to be vertical, and as such, makes this method more flexible than methods that detect the whole object at once using shape matching methods (as described in Figure 6.1). This is apparent, as the final spatial combination of primitive features does not need to be known *a priori*, and can be converged on using logical measures such as similarity matrices and optimisation.

The complexity of the images used in chapters 5 and 6 was reduced through the modelling of binary image features and the masking of the image background. Combined with the simulation through the forward modelling of object features, this allowed the generation of per-object statistics used for the quantification of change observed between the two sampled time periods. This reduction in image complexity allows the increase in efficiency of image processing techniques as the proportion of the input image used in the analysis is reduced compared to analysing the complete

image extent. Describing a combination of features and objects within a scene model also has the benefit of utilising knowledge of the shapes of each feature. Khoshelham (2007), describes these forms of models as being more robust to the effects of partial occlusions. This was seen in the Buncefield imagery with the sides of the objects having similar reflectance to the tops of the drums, and in the Saudi Arabia imagery that contained several disused irrigation fields that appeared faded in the image. As such, partial occlusions were inherent in the imagery used, and through an iterative model-based approach, were identified despite the low initial evidence of each object.

The strength of the scene model lies in its simplicity; it is essentially a set of metrics describing the location, spectral and geometric properties of each object. It allows data to be used from different sources irrespective of the viewing geometry of the image. Creating scene models in this thesis used an iterative approach, and used more processing time than single iterations, however results showed that the results from the iterations were consistently better than from the single iteration *HT*.

8.3.4 Objective 4: Provide semi-automated and iterative methods

The provision of automated and iterative methods was the remit generated by the objective number **4** to reduce the pre-application emphasis on determining the optimal parameters and to provide a method that is parametrically generalisable. This aim sought to replace this ‘perfect parameterisation’ with methods that are self-regulating through means of iterative update. By doing so, the method avoids the need to exhaustively search through multiple input parameters when a new feature extraction method was used, allowing the implementation of methods previously excluded due to complex parameterisation requirements. Iterative updating of the Hough Transform was the exemplar used to provide evidence that objective **4** was a sound conceptual goal. The process described in chapter 5 for the detection of two-dimensional objects was automated, with the exception of retaining the default input parameters of the algorithm. The method almost converged on the correct solution before image complications restricted this convergence from occurring. In this instance, the hypothesis was set to detect the circular tops of the oil drums. Due to

the oblique viewing angle, the tops and sides of the drums (with similar brightness values to the drum tops) were detected as a whole. The accumulator modification process penalised these detections, and removed the circle in the next iteration. The following iteration re-added this circle, and these fluctuations continued; as apparent from iterations >120 in Figure 5.9 and Figure 5.10. This was resolved for the Buncefield case by modifying the methodology to detect three-dimensional objects using oblique imagery.

The overarching objective was to enable the use of alternative feature extraction methods that will be generalisable across image and object types. Reducing the input required from the operator during the feature detection process will facilitate running batch processing operations on powerful computer clusters, thus, reducing the efficiency issues that may have constrained the method described in this thesis. Brenner *et al.* (2001) acknowledged the difficulty in splitting buildings into geometric primitives (objective 3), but the use of an iterative modelling approach, as described in 4, facilitates the iterative generation of the object features, and the gradual fitting of the pre-event scene model to the image features observed in their post-event state.

8.3.5 Objective 5: Provide a change detection method for capturing urban morphological change

Rationale point R5 focused on the creation and provision of a change detection method for use in understanding the morphological changes associated with ‘natural’ urban changes as well as consequential changes arising from natural or human-induced disasters. Urban change detection has been the focus of numerous remote sensing studies in the past, and range from automated land cover change techniques (Seto *et al.*, 2002), to object-orientated quantification of per-building change (Agouris *et al.*, 2000).

The model framework described in this thesis facilitates the comparison of objects detected from the pre-event image through modelling their unchanged appearance in the after image. The simulation of the objects using the post-event image geometry is a novel approach that facilitates the direct comparison of each

building in turn, irrespective of the observational zenith angle or azimuth angle of the post-event image. This method has great potential for application to oblique-angle imagery, such as the freely available Pictometry data included in web applications such as the Bird's Eye imagery in Microsoft Bing Maps⁴.

The quantitative change descriptors outlined in the previous chapter provide a fundamental explanation of the change associated with the Buncefield event, and early promise has been shown for linking qualitative descriptors to quantitative change statistics. These can come in several forms, ranging from damage statistics to proprietary change classes adopted by mapping agencies such as Ordnance Survey. Proxy damage scales that have been effectively used to assess damage from remotely-sensed imagery were used in chapter 7. In future applications, it is envisioned that change scales could be used to describe both planned and unexpected change, and ontologies may be created to link the quantitative statistics in the model to these established sets of qualitative descriptors.

8.3.6 Objective 6: Provide results in a form that facilitates the generation of qualitative and quantitative change descriptors

The overarching aim of the object comparison module of the conceptual model was to provide the user with information on object-based change that has occurred between two time periods sampled by remote sensing images. Information generated at the object level is of use in many remote sensing applications (Agouris *et al.*, 2000; Gusella *et al.*, 2005a; van der Werff, 2006). Many metrics could be used to describe the four major types of object changes (i) objects can change class category; (ii) objects can expand, shrink, or change shape; (iii) objects can change position; and (iv) objects can fragment or unite (Coppin *et al.*, 2004). The method described requires prior identification of the measures most suited for change description.

Change information was generated in two forms for the Buncefield example: Using statistics of per-object change, and also the provision of qualitative descriptors using pre-existing damage scales. Once the metrics required for the

⁴ <http://www.bing.com/maps>

analysis were chosen, statistics of per-object change were attained in an automated process involving no user iteration. Objects from the three-dimensional object detection were described in terms of their intra-object variability using first-order descriptors such as minimum, maximum, mean, mode, range and standard deviation of brightness values. Object texture was measured using range and standard deviation of a 3×3 floating window, (**Figure 7.10**). These measures were useful descriptors to test the automated generation of object change descriptors, and the choice of measures could be further advanced for future applications to include metrics describing the complexity of the object shapes (e.g. perimeter:area ratio), complexity of the object scene (e.g. using fractal dimension measures), or density of object features (e.g. nearest neighbourhood calculations).

The qualitative descriptors involved human interaction with the image to assign objects to damage classes. The key to the generation of these statistics was providing an extraction method that produced results that were conducive to detecting change. The method does not rely on finding matching features in the post-event image, but instead models the pre-event objects as viewed in the post-event image. This technique uses the power of a forward model to describe the location and spectral characteristics of the image features as if they were unchanged. Change information can subsequently be determined through the comparison of the extracted image objects stored in SM_1 and SM_2 .

The modelling approach adds to the current literature on the subject whose primary focus is on the detection of corresponding image features in multi-temporal images. If the objects are not found in the post-event image using the method described in this thesis, the algorithm results in statistics showing significant changes, allowing the user to apply a supervised method to exclude the feature from the analysis, and to update the vector records of the scene for the post-event time period. One example of the everyday usage of such techniques could be the prioritisation of areas of interest for change assessment, through the exclusion of unchanged areas; with specific applications in both urban morphology for map updates, and post-disaster damage assessment. This application was noted in literature describing Ordnance Survey applications of remote sensing change detection (Holland and Tompkinson, 2003).

8.4 POTENTIAL MODEL ENHANCEMENTS

The research described in this thesis has detailed the development of a potentially useful framework for measuring and describing change between images of the same geographic area, captured at two disparate points in time. The conceptual model in its current form has been tested for the detection and modelling of circular and cylindrical features. The following section describes several refinements to the model that have been identified that could improve its overall functionality and extend its capacity to be generalisable across a variety of image types and event scenarios.

The imagery used in the exemplars described in this thesis provided a suitable scenario for the development and testing of the conceptual model. However, the Buncefield disaster was a relatively unique event, with the dominant physical impacts focussed on the oil drums. Natural disasters do not discriminate impacts according to building shape, and so it is recognised that the model requires additional development and testing on a variety of images before it could be used operationally. For the exemplar applications, the model was tailored to detect circles and cylinders automatically to create a scene model from which to detect per-object and scene-wide changes. Development of the model would focus on both the Object Detection and Object Comparison modules.

For scenarios where pre-event GIS data exists for buildings (e.g. using OS MasterMap), or where manual approaches are used to delineate the pre-event scene model (e.g. the GEO-CAN post-Haiti earthquake remote sensing damage assessment: Bevington *et al.*, 2010), the Object Comparison module could be implemented directly using the building footprint information and height attributes for each structure. The model would provide rudimentary change information, and identification of those areas that have changed more than others. Where this is not the case, the Object Detection module would require additional development to expand the number of building shapes that could be detected.

The most practical way to extend the functionality of the Object Detection phase of the model is to incorporate algorithms to detect rectilinear and rectangular features as primitives which can be amalgamated to form building outlines. These are the most frequently observed shapes of buildings' ground footprints. The

conceptual model relies on the availability of a voting space that gathers evidence on the probability of a pixel being part of a shape. Techniques do exist for linear features, most notably the Hough Transform for lines and rectangles. Modification of the model to include alternative algorithms would be time consuming, however would lead to a more robust toolkit of primitive types that could be tailored for each implementation.

The research team at Ordnance Survey is actively investigating automated and semi-automated methods of building detection for use in updating geospatial products. The methods most commonly researched involve the use of imagery to extract objects, which are then compared to existing vector GIS data. These are often used to generate simplified 3-D models of buildings (i.e. without chimneys, dormers etc.), that are attributed with building information. Image-to-image comparison techniques are not actively used by the agency currently, due in part to the lack of comparable images captured in-house. The current digital image archive is expanding though and, with the ability to combine the extracted data with common image centres and object footprints, the interest in image-to-image comparison techniques may likely increase.

Previously investigated methods of object detection have been useful for identifying simplified building shapes, but in practise, very few buildings in GB are accurately matched by these shapes. Therefore a method that could extract and amalgamate multiple primitive features in an intelligent fashion may go some way to reducing this problem. However, it would be beneficial to include quality assurance metrics such as an easily-interpretable measure of the confidence in the identified change. The conceptual model, if shown to generalise to other building shapes, such as rectilinear buildings, could be useful to Ordnance Survey commercially. However, the transition of research products into production is often a significant obstacle in the agency (Tompkinson *et al.*, 2003).

8.4.1 Enhancement of the Object Detection Module

- One of the primary enhancements of the object detection phase of the model would be the inclusion of a rule that forbids objects from having overlapping footprints in the image space. Addition of this rule would potentially reduce the
-

number of false positive detections that were seen in both the two and three-dimensional object detection examples. The three-dimensional Buncefield implementation resulted in several false positive detections, characterised by a lower score value than the correct combination of top and base circles. Inclusion of a non-overlap rule in this instance that directly compares conflicting feature pairings could potentially increase the likelihood of convergence on the correct solution during the Monte Carlo simulations.

- The object detection module was developed to allow the substitution of feature extraction algorithms, depending on the application. Additional emphasis could be placed on the adoption of alternative methods (e.g. linear or generalisable HT methods) that can benefit from the *boosting* extension used in the *HT+* to converge iteratively on the objects in the pre-event image. Complex image scenes may also require the use of multiple primitive detectors within the framework to capture the target objects within the scene model. Information could be used beyond the traditional edge-based extraction methods; for example, texture could be used to add additional object target characteristics. Using texture also fits the ideals of the model, as it would allow the forward modelling simulation of features for comparison with the RS image.

8.4.2 Enhancement of the Object Comparison Module

Measuring Geometric Change

It is envisaged that the three-dimensional scene created from the Buncefield image could be extended to include modelling of the geometric changes in objects. Early investigations have been made into this process during this project. One method which has great potential in this area is the active contour (*snakes*) algorithm to capture and model the geometric changes associated with the Buncefield disaster, and other applications of the model. The enhancements include:

- The use of a pre-event scene model created in the first half of the conceptual model described in chapter 3 will provide *a priori* location knowledge for active contour initialisation. Using a dual snake approach could reduce the frequency of the model settling on local minima, which are often characterised by areas
-

with spectral properties similar to the target features. The algorithm will not converge until a global minimum is achieved and both contours have converged on the boundary of the target feature.

- The use of a snakes approach has been envisaged to result in the iterative geometric morphing of SM_1 . Initialised on the location of objects within the pre-event scene model, snakes could make small modifications to the geometric contours of each object using dynamic programming techniques to fit the snake to the footprints of the objects in the post-event image. This change could be described by a set of change vectors around the perimeter of each object footprint describing in detail the direction and magnitude of the geometric change of an object's boundary. This results in the provision of metrics describing the magnitude of geometric change that would complement the brightness and textural change characteristics described in chapter 7.

Enhancement of change descriptors

- Iterative geometric morphing of the scene model objects (described above) will generate change statistics at each iteration. These can be amalgamated to quantify the change in objects in greater detail than is already provided in the literature.
 - An additional advancement in object-based change descriptors could see the introduction of ontologies that could provide an additional level of change understanding. Ontologies can be developed to describe changes at multiple levels, ranging from a per-object morphological description, to the automatic generation of change ontologies that describe the whole image scene. This would require the development of training data from multiple events that link the statistics generated automatically from the conceptual model (change detection) with a set of descriptive words and statements. The ultimate goal of such a study would be the automatic generation of descriptors for image pairs, allowing the efficient semantic retrieval of matching image sets using keywords.
-

Enhancement of hardware and software

The model created during this thesis was designed, written and coded using the MATLAB development environment, and the model implementation was achieved using a dual-core Intel Pentium 4, 3.60 GHz processor with 2GB of RAM memory. There are several areas identified for improvement, increasing the viability of the conceptual model for future larger-scale implementations.

- Update the specifications of the PC used for the running of the model, using a grid or cluster computer setup that allows code to be allocated to multiple processors simultaneously, greatly reducing the time taken for the processing tasks essential to the model. Such a setup would facilitate batch processing of multiple images allowing the model to be run on much larger geographic regions.
- Translate the programs from MATLAB into a compiled programming language such as C++ which would result in more efficient iteration of machine code, hence reducing the processing time, and expanding the capacity of model for larger datasets.

8.5 SUMMARY

At the beginning of this thesis, six objective points were set out to govern the scope of this thesis. This chapter has re-evaluated these benchmarks based on the development of a dual-phase conceptual model that has shown ability for use as a feature extraction and comparison method and which, therefore, adds to the current state-of-the-art in the field of change detection using remotely-sensed imagery. The exemplars described in this research are based on a small image subset, but were designed to investigate the potential of the thesis objectives stated earlier. Several assumptions that were in place during this work have also been acknowledged and justified. The combination of components within the conceptual model provides a range of potential enhancements, making an encouraging basis for the future development of a toolbox of methods that can be customised within the model framework, based on the needs of the specific study.

CHAPTER 9

CONCLUSIONS

9 CONCLUSIONS

Imagery acquired at two or more time periods from remote sensors onboard satellites or aircraft, and analysed over a period of time, can provide information on the magnitude of change in phenomena that has occurred in a specific region of the world. The collection of multi-temporal imagery from two or more points in time for the same geographic area can yield important insights into phenomena as varied as environmental change or the evolution of urban areas. National mapping agencies, such as Ordnance Survey in Great Britain, use applied techniques in the latter field to facilitate rapid and accurate identification of areas requiring updating in their flagship vector dataset – OS MasterMap[®]. Conversely, similar change detection methods can be used for post-disaster damage assessment allowing rapid investigation of affected areas worldwide in the days following a major catastrophe. The increase in data availability, due largely to the revolution in commercially available satellite sensors, the development of micro-satellite technology, and the boom in PC performance and usage have led to a reduction in price of fine spatial resolution optical remotely-sensed imagery. This has subsequently increased the number of potential applications for the data captured and requires a necessary parallel initiative through the development of analytical techniques to detect and monitor objects within the imagery.

This thesis has described the development of a novel framework for the processing of fine spatial resolution satellite and aerial imagery. This framework is designed for change detection for monitoring urban morphology and post-disaster damage assessment. A conceptual model was designed, firstly to facilitate the extraction of buildings or other structures as objects from the pre-event image, enabling the identification of archive imagery to be accessed following a significant natural or anthropogenic disaster. The second key phase focused on modelling the pre-event objects as they would look in unchanged form in the post-event image. A subsequent comparison of this *simulated image* with the observed post-event image revealed information on the extent of change associated with each object.

Preliminary testing of the conceptual model focused on several exemplars. Most notably, an image-set for the Buncefield oil depot in Hemel Hempstead, UK was used as the main test data for model development and implementation. A pre-event aerial image from 1999 was paired with an aerial image captured one month after the explosion and fire that destroyed much of the depot on 11 December 2005. Initial testing had particular emphasis on the object detection phase of the conceptual model: a widely used algorithm from the field of machine vision, the Hough Transform was adopted to detect the tops of the circular oil tanks in the 1999 image. A novel boosting extension (dubbed the *HT+*) was developed to extend the functionality of the *HT-*, with mixed success. The *HT+* initially over-predicted the number of circles in the image and the boosting extension reduced the number of false positive results. However, it failed to converge on the correct number of objects, due to the homogeneous spectral nature of the object properties in the oblique-view image.

The two-dimensional object detection technique was further tested on imagery of circular irrigation features in the Saudi Arabian desert. The boosting mechanism of the *HT+* reduced the number of false positive responses and also reduced the number of false negative circle detections through the modification of the evidence-gathering accumulator space. The accumulator space was updated iteratively using a simulation of the detected objects through a forward model and subsequently compared to the input image, producing per-pixel error information that informed the following iteration. The results showed the ability of the model to reduce the number of false positive detections seen by an unaltered *HT* run, whilst boosting the missing information for false negative detections.

The Buncefield imagery was also used to test the detection of three-dimensional objects, using the *HT-* alongside an iterative optimisation approach which matched the detected circles as top and base pairs for each candidate cylinder. Several matrices were generated to compare the strength of similarity between the circles with a matrix of score information produced that included each potential top and base match. The total score for the combination of circle matches was reduced and the lowest combination was delivered consistently using a simulated annealing algorithm. A further iteration automatically determined the number of objects in the image, thus, allowing the convergence on a combination of the correct number of

objects with the lowest overall score. The final phase of the thesis demonstrated the potential of a forward modelling technique to compare objects at two image dates. The creation of a “perfect” scene model allowed the generation of quantitative and qualitative descriptors of change of the objects from their pre-event state.

The exemplars described have shown the potential of the conceptual model for providing an iterative, model-based framework for change detection. Although limited in its application, and choice of target building features, the model showed potential for object delineation and assessment of change in two images with unequal sensor geometry. Further developments of the model are required to take the model beyond detection of circular and cylindrical features, towards rectangular shapes, for which there are more tangible operational applications for monitoring both urban and post-disaster change.

This thesis has provided intellectual merit through the preliminary investigation of a novel method of change detection. The model builds on current literature by expanding and amalgamating a variety of existing algorithms implemented within a model-based framework for the purpose of change detection from remotely-sensed imagery. One of the products of this thesis is a conceptual model that utilises an iterative forward and inverse modelling loop to automatically detect circular objects in imagery and provides a potential method for the comparison of these objects from their pre-event to post-event states. This research provides the starting point for a new method that could have an impact on the current state-of-the-art: the methods described show candidature for further development into a model that has considerable potential to be robust and generalisable across numerous image formats, and allowing the detection and modelling of increasingly complex geometric shapes through the intelligent amalgamation of geometric primitives.

The rapid proliferation of online mapping systems in recent years has widened the gap between data availability and image processing algorithms, and, as such, there is an abundance of underutilised data in the public realm. Datasets such as the Bird’s Eye view supported by Microsoft Bing offers huge potential for the use of a model-based method that is insensitive to variations in observational zenith angles. In addition, the increase in computer processing power has limited previous algorithmic explorations. However, technologies such as grid and cloud computing

are becoming more realistic and accessible, and the availability of automated change detection methods that facilitate batch processing of numerous images, thus, covering a larger geographic area, is becoming more evident.

The ultimate aim of the model is to have a series of algorithms that are replaceable components within the conceptual model. Such a suite of tools would offer a valuable and flexible tool for use in updating vector datasets by mapping agencies. These tools and techniques could be expanded over time for routine usage to identify areas within images that have changed in map products. The modelling approach targets the objects of interest, reducing the complexity seen in the image space, to a more fundamental abstraction of reality in the scene model.

In the event of an unexpected catastrophe, an appropriate algorithm would be used within the model framework to extract and model the features of interest, resulting in the identification of the areas that have experienced significant change. It is hoped that having investigated these potential approaches and developing the conceptual model, this thesis will make a contribution to the increase in the efficiency of post-disaster data collection and urban monitoring.

REFERENCES

- ADAMS, B. J., HUYCK, C. K., MANSOURI, B., EGUCHI, R. T., and SHINOZUKA, M., 2004, *Application of High-Resolution Optical Satellite Imagery for Post-Earthquake Damage Assessment: The 2003 Boumerdes (Algeria) and Bam (Iran) Earthquakes*. Research Progress and Accomplishments 2003-2004. State University of New York, Buffalo, NY, USA.
- AGOURIS, P., BEARD, K., MOUNTRAKIS, G., and STEFANIDIS, A., 2000, Capturing and Modeling Geographic Object Change: A spatiotemporal gazetteer framework. *Photogrammetric Engineering and Remote Sensing*, **66**(10): 1241-1250.
- ALI, M., and CLAUSI, D., 2001, Using the Canny Edge Detector for Feature Extraction and Enhancement of Remote Sensing Images: *Geoscience and Remote Sensing Symposium, 2001. IGARSS '01. IEEE 2001 International*. 2298-2300.
- AOKI, H., MATSUOKA, M., and YAMAZAKI, F., 1998, Characteristics of Satellite SAR Images in the Damaged Areas Due to the Hyogoken-Nanbu Earthquake *Proceedings of the 1998 Asian Conference on Remote Sensing*, Manila, Philippines.
- APLIN, P., 2003, Comparison of Simulated IKONOS and SPOT HRV Imagery for Classifying Urban Areas. In *Remotely Sensed Cities*, edited by V. Mesev (London: Taylor & Francis).
- ATIQUZZAMAN, M., 2005, Coarse-to-Fine Search Technique to Detect Circles in Images. *International Journal of Advanced Manufacturing Technology*, **15**(2): 96-102.
- BAILLOEUL, T., DUAN, J., PRINET, V., and SERRA, B., 2003, Urban Digital Map Updating From Satellite High Resolution Images Using GIS Data as A-Priori Knowledge: *Proceedings of the 2nd GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas (URBAN 2003)*, Berlin. IEEE.
- BALLY, P., BEQUIGNON, J., ARINO, O., and BRIGGS, S., 2005, Remote Sensing and Humanitarian Aid - A life-saving combination. *ESA Bulletin-European Space Agency*(122): 36-41.
- BAMFORD, P., and LOVELL, B., 1998, Unsupervised Cell Nucleus Segmentation with Active Contours: An Efficient Algorithm Based on Immersion Simulations. *Signal Processing*, **71**: 203-213.
- BARNESLEY, M. J., BARR, S. L., HAMID, A., MULLER, J.-P. A. L., SADLER, G. J., and SHEPHERD, J. W., 1993, Analytical Tools to Monitor Urban Areas. In *Geographical Information Handling: Research and applications*, edited by P. M. Mather (Chichester: John Wiley & Sons).
- BAUDOT, Y., 2001, Geographical Analysis of the Population of Fast-growing Cities in the Third World. In *Remote Sensing and Urban Analysis*, edited by J.-P. Donnay, M. J. Barnesley, and P. Longley (London: Taylor & Francis), 225-242.
-

-
- BBC, 2007, *Can Shanghai Turn Green and Grow?*, British Broadcasting Corporation, <http://news.bbc.co.uk/1/hi/business/6683103.stm>, Accessed 31 May 2007
- BELWARD, A. S., STIBIG, H. J., EVA, H., REMBOLD, F., BUCHA, T., HARTLEY, A., BEUCHLE, R., KHUDHAIRY, D., MICHIELON, M., and MOLLICONE, D., 2007, Mapping Severe Damage to Land Cover Following the 2004 Indian Ocean Tsunami using Moderate Spatial Resolution Satellite Imagery. *International Journal of Remote Sensing*, **28**(13): 2977-2994.
- BEVINGTON, J. S., ADAMS, B. J., and EGUCHI, R. T., 2010, *GEO-CAN Debuts to Map Haiti Damage*, In *Imaging Notes*, Spring 2010. (Denver, CO.: Blueline Publishing).
- BIEVER, C., 2006, *See it, snap it, and let your phone find it on the web*, In *New Scientist*, 2533, 7 January 2006.
- BITELLI, G., CAMASSI, R., GUSELLA, L., and MOGNOL, A., 2004, Image Change Detection On Urban Area: The Earthquake Case: *Proceedings of ISPRS XXth Congress, Commission 7*, Istanbul, Turkey.
- BMIIB, 2006, The Buncefield Investigation: Third Progress report. *Buncefield Major Incident Investigation Board*(21 February 2006).
- BMIIB, 2008, The Buncefield Incident 11 December 2005: The final report of the Major Incident Investigation Board. *Buncefield Major Incident Investigation Board*, **1**.
- BORST, W. N., 1997, Construction of Engineering Ontologies for Knowledge Sharing and Reuse University of Twente, Enschede, The Netherlands.
- BOSCOLO, R., BROWN, M., and MCNITT-GRAY, M., 2002, Medical Image Segmentation with Knowledge-guided Robust Active Contours. *Radiographics*, **22**(2): 437-448.
- BRENNER, C., HAALA, N., and FRITSCH, D., 2001, Towards Fully Automated 3D City Model Generation. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 47-57.
- BUCKHAM, B. J., and LAMBERT, C., 1999, *Simulated Annealing Applications*. University of Victoria, http://www.me.uvic.ca/~zdong/courses/mech620/SA_App.PDF, Accessed 13 May 2008
- CAMPBELL, J. B., 2002, *Introduction to Remote Sensing*, 3rd Edition (London: Taylor & Francis).
- CARR, R., 2008, *Simulated Annealing*. MathWorld - A Wolfram Web Resource, created by Eric W. Weisstein, <http://mathworld.wolfram.com/SimulatedAnnealing.html>, Accessed 3 August 2008
- CCRS, 2008, *Tutorial: Fundamentals of Remote Sensing*. Natural Resources Canada: Canada Centre for Remote Sensing, http://www.ccrs.nrcan.gc.ca/resource/tutor/fundam/chapter2/03_e.php, Accessed 23 May 2010
- CHAN, R., and SIU, W. C., 1991, New Parallel Hough Transform for Circles. *IEEE Proceedings Computers and Digital Techniques*, **138**(5): 335-344.
- CHAN, T. F., and VESE, L. A., 2001, Active Contours Without Edges. *IEEE Transactions on Image Processing*, **10**(2): 266-277.
-

-
- CHIROIU, L., and ANDRE, G., 2001, *Damage Assessment using High Resolution Satellite Imagery: Application to 2001 Bhuj, India, earthquake*. RiskWorld., <http://www.riskworld.com/Nreports/2001/Bhuj,India,earthquake2001.PDF>
- CIBELLI, M., NAPPI, M., and TUCCI, M., 2004, ABI: Analogy-based Indexing for Content Image Retrieval. *Image and Vision Computing*, **22**(1): 23-34.
- CONRAD, J., 2009, *The Sun/Moon Calculator*. largeformatphotography.info, <http://www.largeformatphotography.info/sunmooncalc/>, Accessed 31/03/09
- COPPIN, P., JONCKHEERE, I., NACKAERTS, K., MUYS, B., and LAMBIN, E., 2004, Digital Change Detection Methods in Ecosystem Monitoring: A Review. *International Journal of Remote Sensing*, **25**(9): 1565-1596.
- CROSS, A. M., 1988, Detection of Circular Geological Features using the Hough Transform. *International Journal of Remote Sensing*, **9**(9): 1519-1528.
- DAI, X. L., and KHORRAM, S., 1998, The Effects of Image Misregistration on the Accuracy of Remotely Sensed Change Detection. *IEEE Transactions on Geoscience and Remote Sensing*, **36**(5): 1566-1577.
- DAVIES, E. R., 1984, Design of Cost-effective Systems for the Inspection of Certain Food Products during Manufacture: *Proceedings of the Fourth International Conference on Robot Vision and Sensory Control*, London. 437-446.
- DE VICENTE, J., LANCHARES, J., and HERMIDA, R., 2003, Placement by Thermodynamic Simulated Annealing. *Physics Letters A*, **317**(5-6): 415-423.
- DEER, P., 1999, *Digital Change Detection Techniques: Civilian and Military Application* (London: Taylor & Francis).
- DILLEY, M., CHEN, R., DEICHMANN, U., LERNER-LAM, A. L., and ARNOLD, M., 2005, *Natural Disaster Hotspots: A global risk analysis*. Disaster Risk Management Series. The World Bank,
- DMCII, 2009, *Satellite Image Products*. DMC International Imaging Ltd., <http://www.dmcii.com/products.htm>, Accessed 6 March 2009
- DONNAY, J.-P., BARNESLEY, M. J., and LONGLEY, P. A., 2001, *Remote Sensing and Urban Analysis* (London: Taylor & Francis).
- DRIGGERS, R. G., 2003, *Encyclopedia of Optical Engineering* (New York: Marcel Dekker).
- DUDA, R. O., and HART, P. E., 1972, Use of the Hough Transform to Detect Lines and Curves in Pictures. *Communications of the ACM*, **15**(1): 11-15.
- DURBHA, S. S., and KING, R. L., 2004, Knowledge Mining in Earth Observation Data Archives: A domain ontology perspective: *Proceedings of IEEE International Geoscience and Remote Sensing Symposium, 2004*, Anchorage, AK, USA. IEEE. 172-173.
- EGUCHI, R. T., HUYCK, C. K., and ADAMS, B. J., 2005, An Urban Damage Scale Based on Satellite and Airborne Imagery: *Proceedings of the Third International Workshop on Remote Sensing for Post-Disaster Response*, Chiba, Japan. EERI / MCEER.
- EGUCHI, R. T., HUYCK, C. K., ADAMS, B. J., MANSOURI, B., HOUSHMAND, B., and SHINOZUKA, M., 2003, *Resilient Disaster Response: Using Remote Sensing Technologies for Post-Earthquake Damage Detection*. Research Progress and Accomplishments 2001 - 2003. State University of New York, Buffalo, NY, USA.
-

-
- EMS, 1998, *European Macroseismic Scale*. European Seismological Commission, [http://www.gfz-potsdam.de/portal/-?%\\$part=binary-content&id=1883158&status=300](http://www.gfz-potsdam.de/portal/-?%$part=binary-content&id=1883158&status=300), Accessed 28 March 2009
- ESA, 2007, *GMES: The Living Planet Programme*. European Space Agency, http://www.esa.int/esaLP/SEM3B60DU8E_LPgmes_0.html, Accessed 15 May 2007
- FOODY, G. M., 2003, Remote Sensing of Tropical Forest Environments: Towards the Monitoring of Environmental Resources for Sustainable Development. *International Journal of Remote Sensing*, **24**(20): 4035-4046.
- FRÈRE, D., HENDRICKX, J., VANDEKERCKHOVE, T., MOONS, L., and VAN GOOL, L., 1995, On the Reconstruction of Urban House Roofs from Aerial Images. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (I)*, edited by A. Gruen, O. Kuebler, and P. Agouris (Berlin: Birkhäuser Verlag), 87-96.
- FRIEDLAND, C. J., ADAMS, B. J., and LEVITAN, M. L., 2006, Remote Sensing and Advanced Technology for Estimating Post-Hurricane Structural Storm Surge Damage: *Proceedings of the Fourth International Workshop on Remote Sensing for Disaster Response*, Cambridge, UK. EERI / MCEER.
- FROST, V. S., STILES, J. A., SHANMUGAN, K. S., and HOLTZMAN, J. C., 1982, A Model for Radar Images and Its Application to Adaptive Digital Filtering of Multiplicative Noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-4**(2): 157-166.
- GEO, 2009, *Group on Earth Observations: About GEO*. Group on Earth Observations, http://earthobservations.org/about_geo.shtml, Accessed 22 May 2010
- GERIG, G., 1987, Limiting Image-space and Accumulator-space: A new approach for object recognition: *Proceedings of the First International Conference on Computer Vision*, London. 112-117.
- GERIG, G., and KLEIN, F., 1986, Fast Contour Identification through Efficient Hough Transform and Simplified Interpretation Strategy: *Proceedings of the Eighth International Conference on Pattern Recognition*, Paris, France. 498-500.
- GERKE, M., HEIPKE, C., and STRAUB, B.-M., 2001, Building Extraction from Aerial Imagery using a Generic Scene Model and Invariant Geometric Moments: *Proceedings of the IEEE/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas*, Rome, Italy.
- GOLDBERG, D., 1988, *Genetic Algorithms in Search, Optimisation and Machine Learning*: Addison-Wesley).
- GONZALEZ, R. C., WOODS, R. E., and EDDINS, S. L., 2004, *Digital Image Processing using MATLAB*, Second Edition (New Jersey, USA: Pearson).
- GOPAL, S., and WOODCOCK, C., 1996, Remote Sensing of Forest Change using Artificial Neural Networks. *IEEE Transactions on Geoscience and Remote Sensing*, **34**(2): 398-404.
- GRUEN, A., 1998, TOBAGO - A Semi-automated Approach for the Generation of 3-D Building Models. *ISPRS Journal of Photogrammetry and Remote Sensing*, **53**: 108-118.
- GRUEN, A., and WANG, X., 2001, News from CyberCity-Modeler. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited
-

- by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 93-101.
- GUDIVADA, V. N., and RAGHAVAN, V. V., 1995, Content-Based Image Retrieval Systems. *IEEE Computer*, **28**(9): 18-22.
- GÜLCH, E., and MÜLLER, H., 2001, New Applications of Semi-automatic Building Acquisition. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 103-114.
- GUIL, N., and ZAPATA, E. L., 1997, Lower Order Circle and Ellipse Hough Transform. *Pattern Recognition*, **30**(10): 1729-1744.
- GUNN, S. R., 1996, Dual Active Contour Models for Image Feature Extraction University of Southampton, Southampton.
- GUNN, S. R., and NIXON, M. S., 1997, A Robust Snake Implementation; A Dual Active Contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**(1): 63-68.
- GUNN, S. R., and NIXON, M. S., 1998, Global and Local Active Contours for Head Boundary Extraction. *International Journal of Computer Vision*, **30**: 43-54.
- GUOPU, Z., QINGSHUANG, Z., and CHANGHONG, W., 2006, Dual Geometric Active Contour for Image Segmentation. *Optical Engineering Letters*, **45**(8): 080505.
- GUSELLA, L., ADAMS, B. J., BITELLI, G., HUYCK, C. K., and MOGNOL, A., 2005a, Object-Oriented Image Understanding and Post-Earthquake Damage Assessment for the 2003 Bam, Iran, Earthquake. *Earthquake Spectra*, **21**(S1): S225-S238.
- GUSELLA, L., HUYCK, C. K., ADAMS, B. J., CHO, S., and CHUNG, H., 2005b, Damage Assessment with Very-High Resolution Optical Imagery following the December 26 2003 Bam Iran Earthquake: *Proceedings of the Third International Workshop on Remote Sensing for Post-Disaster Response*, Chiba, Japan. EERI / MCEER.
- HALL, F. G., BOTKIN, D. B., STREBEL, D. E., WOODS, K. D., and GOETZ, S. J., 1991, Large-Scale Patterns of Forest Succession as Determined by Remote Sensing. *Ecology*, **72**(2): 628-640.
- HAMARNEH, G., CHODOROWSKI, A., and GUSTAVSSON, T., 2000, Active Contour Models: Application to Oral Lesion Detection in Color Images: *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, Nashville, TN, USA. 2458-2463.
- HART, G., and GREENWOOD, J., 2003, A Component Based Approach to Geo-ontologies and Geodata Modelling to Enable Data Sharing: *Proceedings of the Sixth AGILE*, Lyon, France.
- HAYASHI, H., HASHITERA, S., KOHIYAMA, M., MATSUOKA, M., MAKI, N., FUJITA, H., and ELVIDGE, C. D., 2000, International Collaboration for the Early Damaged Area Estimation System Using DMSP/OLS Nighttime Images: *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Honolulu, HI, USA. IEEE.
- HELLER, A. J., LECLERC, Y. G., and LUONG, Q.-T., 2001, A Framework for Robust 3-D Change Detection: *Proceedings of the International Symposium on Remote Sensing*, Toulouse, France.
-

-
- HENDERSON, F. M., and XIA, Z.-F., 1997, SAR Applications in Human Settlement Detection, Population Estimation and Urban Land Use Pattern Analysis: A Status Report. *IEEE Transactions on Geoscience and Remote Sensing*, **35**: 79-85.
- HEROLD, M., SCEPAN, J., and CLARKE, K. C., 2002, The use of Remote Sensing and Landscape Metrics to Describe Structures and Changes in Urban Land Uses. *Environment and Planning A*, **34**: 1443-1458.
- HESE, S., and SCHMULLIUS, C., 2003, Forest Cover Change Detection in Siberia: *Proceedings of ISPRS & EARSeL Joint Workshop " High Resolution Mapping from Space 2003"*, Hannover, Germany.
- HINDSON, D., and MCCARTHY, J., 1994, *Defining and Gauging the Problem*. In *Here to Stay: Informal Settlements in KwaZulu Natal*, edited by D. Hindson, and J. McCarthy (Durban: Indicator Press).
- HOLLAND, D., and TOMPKINSON, W., 2003, Improving the update of geospatial information databases from imagery using semi-automated user-guidance techniques: *Proceedings of Geocomputation 2003*, Southampton.
- HOLLAND, J. H., 1975, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (Ann Arbor, MI, USA: University of Michigan Press).
- HOUGH, P. V. C., 1962, *A Method and Means for Recognising Complex Patterns* (US Patent No. 3, 069, 654, 1962).
- HUYCK, C. K., and ADAMS, B. J., 2002, *Emergency Response in the Wake of the World Trade Center Attack: The Remote Sensing Perspective*. Engineering and Organizational Issues Related to the World Trade Center Terrorist Attack. State University of New York, Buffalo, NY, USA.
- HUYCK, C. K., ADAMS, B. J., CHO, S., EGUCHI, R. T., MANSOURI, B., and HOUSHMAND, B., 2004a, *Methodology for Post-Earthquake Building Damage Detection using SAR and Optical Remote Sensing: Application to the August 17, 1999 Marmara, Turkey Earthquake*. Technical Report MCEER-04-004. MCEER,
- HUYCK, C. K., ADAMS, B. J., and GUSELLA, L., 2004b, *Damage Detection Using Neighborhood Edge Dissimilarity of Very High-Resolution Optical Data: Proceedings of the Second International Workshop on Remote Sensing for Post-Disaster Response*, Newport Beach, CA, USA. EERI / MCEER.
- INGRAM, K. I., KNAPP, E. M., and ROBINSON, J. W., 1981, *Change Detection Technique Development for Improved Urban Area Delineation. Final Report*. Computer Sciences Corporation, Silver Spring, MD, U.S.A.
- ITO, A., 2005, Issues in the Implementation of the International Charter on Space and Major Disasters. *Space Policy*, **21**(2): 141-149.
- JIANG, J., and PERRIE, W., 2007, The Impacts of Climate Change on Autumn North Atlantic Midlatitude Cyclones. *Journal of Climate*, **20**(7): 1174-1187.
- JOHNSON, R. D., and KASISCHKE, E. S., 1998, Change Vector Analysis: A Technique for the Multispectral Monitoring of Land Cover and Condition. *International Journal of Remote Sensing*, **19**(3): 411-426.
- KÄLVIÄINEN, H., HIRVONEN, P., LEI, X., and ERKKI, O., 1995, Probabilistic and Non-probabilistic Hough Transforms: Overview and comparisons. *Image and Vision Computing*, **13**: 239-252.
- KÄSER, C., 2001, Requirements of Users from a Public Mapping Organisation. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images*
-

- (III), edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 77-82.
- KASS, M., WITKIN, A., and TERZOPOULOS, D., 1988, Snakes: Active Contour Models. *International Journal of Computer Vision*, **1**: 321-331.
- KESIDIS, A. L., and PAPAMARKOS, N., 1999, On the Inverse Hough Transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**(12): 1329-1343.
- KHOSHELHAM, K., 2007, Extending Generalized Hough Transform to Detect 3D Objects in Laser Range Data: *Proceedings of the ISPRS Workshop on Laser Scanning 2007*, Espoo, Finland.
- KIRKPATRICK, S., 1984, Optimization by Simulated Annealing - Quantitative Studies. *Journal of Statistical Physics*, **34**(5-6): 975-986.
- KIRKPATRICK, S., GELATT, C. D., JR., and VECCHI, M. P., 1983, Optimization by Simulated Annealing. *Science*, **220**(4598): 671-680.
- LABEN, C., 2002, Integration of Remote Sensing Data and Geographic Information System Technology for Emergency Managers and their Applications at the Pacific Disaster Center. *Optical Engineering*, **41**(9): 2129-2136.
- LAMBIN, E. F., and STRAHLER, A. H., 1994, Change-Vector Analysis in Multitemporal Space - a Tool to Detect and Categorize Land-Cover Change Processes Using High Temporal-Resolution Satellite Data. *Remote Sensing of Environment*, **48**(2): 231-244.
- LEE, J.-S., 1980, Digital Image Enhancement and Noise Filtering by Use of Local Statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-2**(2): 165-168.
- LILLESAND, T. M., KIEFER, R. W., and CHIPMAN, J. W., 2008, *Remote Sensing and Image Interpretation*, Sixth edition (Hoboken, NJ, USA: Wiley).
- LIN, C., and NEVATIA, R., 1995, 3-D Descriptions of Buildings from an Oblique View Aerial Image: *Proceedings of the International Symposium on Computer Vision*, Coral Gables, FL, USA. IEEE. 377-382.
- LIU, J. H., SHAN, X. J., and YIN, J. Y., 2004, Automatic Recognition of Damaged Town Buildings Caused by Earthquake using Remote Sensing Information: Taking the 2001 Bhuj, India, earthquake and the 1976 Tangshan, China, earthquake as examples. *Acta Seismologica Sinica*, **17**(6): 686-696.
- LIU, Z. J., WANG, W., and LIU, W. P., 2005, Building Extraction from High Resolution Imagery based on Multi-scale Object Oriented Classification and Probabilistic Hough Transform: *Proceedings of Geoscience and Remote Sensing Symposium, 2005*, Seoul, Korea. 2250-2253.
- LO, C. P., and SHIPMAN, R. L., 1990, A GIS Approach to Land-use Change Dynamics Detection. *Photogrammetric Engineering and Remote Sensing*, **56**: 1483-1491.
- LOPES, A., NEZRY, E., TOUZI, R., and LAUR, H., 1990, Maximum A Posteriori Speckle Filtering And First Order Texture Models In Sar Images: *Geoscience and Remote Sensing Symposium, 1990. IGARSS '90. 'Remote Sensing Science for the Nineties'*, 10th Annual International. 2409-2412.
- LU, D., MAUSEL, P., BRONDÍZIO, E., and MORAN, E., 2004, Change Detection Techniques. *International Journal of Remote Sensing*, **25**(12): 2365-2401.
- LYON, J. G., DING, Y., LUNETTA, R. S., and ELVIDGE, C. D., 1998, A Change Detection Experiment using Vegetation Indices. *Photogrammetric Engineering and Remote Sensing*, **64**(2): 143-150.
-

-
- MASON, S., and BALTSAVIAS, E., 1997, Image-Based Reconstruction of Informal Settlements. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, edited by A. Gruen, E. Baltsavias, and O. Henricsson (Berlin: Birkhäuser Verlag), 97-108.
- MATHWORKS, 2009, *MATLAB Central: An open exchange for the MATLAB and Simulink user community*.
<http://www.mathworks.com/matlabcentral/fileexchange>, Accessed 15 November 2009
- MAYUNGA, S. D., COLEMAN, D. J., and ZHANG, Y., 2007, A Semi-automated Approach for Extracting Buildings from QuickBird Imagery Applied to Informal Settlement Mapping. *International Journal of Remote Sensing*, **28**(10): 2343-2357.
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H., and TELLER, E., 1953, Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, **21**(6): 1087-1092.
- MITCHELL, M., 1998, *An Introduction to Genetic Algorithms*: MIT Press).
- MIURA, H., WIJEYEWICKREMA, A. C., and INOUE, S., 2005, Evaluation of Tsunami Damage in the Eastern Part of Sri Lanka Due to the 2004 Sumatra Earthquake Using High-Resolution Satellite Images: *Proceedings of the Third International Workshop on Remote Sensing for Post-Disaster Response*, Chiba, Japan. EERI / MCEER.
- MIURA, H., YAMAZAKI, F., and MATSUOKA, M., 2006, Building Damage Mapping of the 2006 Central Java, Indonesia Earthquake Using High-Resolution Satellite Images: *Proceedings of the Fourth International Workshop on Remote Sensing for Post-Disaster Response*, Cambridge, UK.
- MOFFITT, F. H., 1970, *Photogrammetry*, Second edition (Scranton, PA, USA: International Textbook Company).
- MONTOYA, L., 2003, Geo-Data Acquisition Through Mobile GIS and Digital Video: An urban disaster management perspective. *Environmental Modelling & Software*, **18**(10): 869-876.
- MUCHONEY, D. M., and HAACK, B. N., 1994, Change Detection for Monitoring Forest Defoliation. *Photogrammetric Engineering and Remote Sensing*, **60**: 1243-1251.
- NEUENSCHWANDER, W., FUA, P., IVERSON, L., SZÉKELY, G., and KÜBLER, O., 1997, Ziplock Snakes. *International Journal of Computer Vision*, **3**: 191-201.
- NEVATIA, R., LIN, C., and HUERTAS, A., 1997, A System for Building Detection from Aerial Images. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, edited by A. Gruen, E. Baltsavias, and O. Henricsson (Berlin: Birkhäuser Verlag), 77-86.
- NIEDERÖST, M., 2001, Automated Update of Building Information in Maps using Medium-scale Imagery. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 161-170.
- NIXON, M., and AGUADO, A., 2002, *Feature Extraction and Image Processing*, 1st Edition (Oxford: Elsevier).
- NOURANI, Y., and ANDRESEN, B., 1998, A Comparison of Simulated Annealing Cooling Strategies. *Journal of Physics A: Mathematical and General*, **31**: 8373-8385.
-

-
- OED, 2004, *Ontology*. In Oxford English Dictionary.
- OLSEN, B. P., KNUDSEN, T., and FREDERIKSEN, P., 2002, Hybrid Raster/Vector Change Detection for Map Database Update: *Proceedings of the 11th Danish Conference on Pattern Recognition and Image Analysis*, Copenhagen.
- ORDNANCESURVEY, 2007, *Who are we? What do we do?*, Ordnance Survey, <http://www.ordnancesurvey.co.uk/aboutus/>, Accessed 6 August 2008
- ORDNANCESURVEY, 2009a, *Research: Data capture*. Ordnance Survey, <http://www.ordnancesurvey.co.uk/oswebsite/partnerships/research/research/remote.html>, Accessed 6 March 2009
- ORDNANCESURVEY, 2009b, *What is OS MasterMap®?*, Ordnance Survey, <http://www.ordnancesurvey.co.uk/oswebsite/products/osmastermap/osmmexplained.html>, Accessed 6 March 2009
- PARKER, C., and STILEMAN, M., 2005, Disaster Management: The Challenges for a National Geographic Information Provider. In *Geo-Information from Disaster Management*, edited by P. Zlatanova, E. M. Fendel, and P. Van Oosterom (Berlin: Springer), 191-214.
- PARKER, D. C., EVANS, T. P., and MERETSKY, V., 2001, Measuring Emergent Properties of Agent-Based Land-cover / Land-use Models Using Spatial Metrics: *Proceedings of the Seventh Annual Conference of the Society for Computational Economics*, Yale University.
- PETROU, M., and KITTLER, J., 1988, On the Optimal Edge Detector: *Proceedings of the Alvey Vision Conference*, Manchester. 191-196.
- PITIÉ, F., KOKARAM, A. C., and DAHYOT, R., 2004, *Oriented Particle Spray: A New Probabilistic Contour Tracing with Directional Information*. In Proceedings of the Irish Machine Vision and Image Processing (IMVIP '04) (Dublin, Ireland).
- POV-RAY, 2008, *The Persistence of Vision Raytracer*. <http://www.povray.org/>, Accessed 23 November 2009
- PRINCEN, J., YUEN, H. K., ILLINGWORTH, J., and KITTLER, J., 1989, A Comparison of Hough Transform Methods: *Proceedings of the Third International Conference on Image Processing and its Applications*. 73-77.
- QINETIQ, 2007, *TopSat*. QinetiQ, <http://www.qinetiq.com/home/commercial/space/topsat.html>, Accessed 2 May 2007
- RABBANI, T., and VAN DEN HEUVEL, F., 2005, Efficient Hough Transform for Automatic Detection of Cylinders in Point Clouds: *Proceedings of the 11th Annual Conference of the Advanced School for Computing and Imaging (ASCI '05)*, Het Heijderbos, Heijen, The Netherlands.
- RAO, C. R. N., STOWE, L. L., and MCCLAIN, E. P., 1989, Remote Sensing of Aerosols over the Oceans using AVHRR data: Theory, Practice and Applications. *International Journal of Remote Sensing*, **10**(4): 743-749.
- REJAIE, A., and SHINOZUKA, M., 2004, Reconnaissance of Golcuk 1999 Earthquake Damage using Satellite Images. *Journal of Aerospace Engineering*, **17**(1): 20-25.
- RESPOND, 2004, *RESPOND Mission Statement*. RESPOND, http://www.respond-int.org/Respond/public/html/what_is_respond/missionStatement.html, Accessed 15 May 2007
-

-
- ROBERTS, L. G., 1965, Machine Perception of Three-dimensional Solids. In *Optical and Electro-Optical Information Processing*, edited by J. T. Tippett (Cambridge, MA, USA.: MIT Press), 159-197.
- ROGERSON, P. A., 2002, Change Detection Thresholds for Remotely Sensed Images. *Journal of Geographical Systems*, **4**(1): 85-97.
- ROWLAND, T., and WEISSTEIN, E. W., 2009, *Genetic Algorithm*. MathWorld - A Wolfram Web Resource, created by Eric W. Weisstein, <http://mathworld.wolfram.com/GeneticAlgorithm.html>, Accessed 21 March 2009
- RUTENBAR, R. A., 1989, Simulated Annealing Algorithms: An overview. *Circuits and Devices Magazine, IEEE*, **5**(1): 19-26.
- RUTHER, H., MARTINE, H. M., and MTALO, E. G., 2002, Application of Snakes and Dynamic Programming Optimisation Technique in Modeling of Buildings in Informal Settlement Areas. *ISPRS Journal of Photogrammetry and Remote Sensing*, **56**(4): 269-282.
- SAITO, K., BROWN, D., BEVINGTON, J., ADAMS, B., PLATT, S., CHENVIDYAKARN, T., SPENCE, R., CHUENPAGDEE, R., KHAN, A., and SO, E., 2009, Measuring, Monitoring and Evaluating Recovery – Towards standardised indicators for post-disaster recovery: *Proceedings of the Seventh International Workshop on Remote Sensing for Disaster Management*, Austin, TX. USA.
- SAITO, K., and SPENCE, R., 2004, Rapid Damage Mapping using Post-Earthquake Satellite Images: *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Anchorage, AK, USA. IEEE. 2272-2275.
- SALINGER, M., 2005, Climate Variability and Change: Past, Present and Future - An Overview. *Climatic Change*, **70**(1-2): 9-29.
- SARABANDI, P., KIREMIDJIAN, A. S., EGUCHI, R. T., and ADAMS, B. J., 2008a, *Building inventory Compilation for Disaster Management: Application of Remote Sensing and Statistical Modeling*. Buffalo.
- SARABANDI, P., KIREMIDJIAN, A. S., EGUCHI, R. T., and ADAMS, B. J., 2008b, *Building Inventory Compilation for Disaster Management: Application of Remote Sensing and Statistical Modeling*. Technical Report MCEER-08-0025. Buffalo.
- SAWABE, Y., MATSUNAGA, T., and ROKUGAWA, S., 2006, Automated Detection and Classification of Lunar Craters using Multiple Approaches. *Advances in Space Research*, **37**(1): 21-27.
- SCHMITT, S. R., 2004, *About Simulated Annealing*. Abecedrical Systems, http://home.att.net/~srschmitt/sa_demo/SA-demo_1.html, Accessed 3 August 2008
- SETO, K. C., WOODCOCK, C. E., SONG, C., HUANG, X., LU, J., and KAUFMANN, R. K., 2002, Monitoring Land-use Change in the Pearl River Delta using Landsat TM. *International Journal of Remote Sensing*, **23**(10): 1985-2004.
- SHI, Z., SHIBASAKI, R., and MURAI, S., 1997, Automated Building Extraction from Digital Stereo Imagery. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, edited by A. Gruen, E. Baltsavias, and O. Henricsson (Berlin: Birkhäuser Verlag), 119-128.
- SHINOZUKA, M., and REJAIE, S. A., 2000, Analysis of Remotely Sensed Pre- and Post-Disaster Images for Damage Detection: *Proceedings of SPIE's Seventh*
-

-
- Annual International Symposium on Smart Structures and Materials*, Newport Beach, CA, USA.
- SINGH, A., 1989, Review Article: Digital Change Detection Techniques Using Remotely-Sensed Data. *International Journal of Remote Sensing*, **10**(6): 989-1003.
- SOHL, T. L., 1999, Change Analysis in the United Arab Emirates - An investigation of techniques *Photogrammetric Engineering and Remote Sensing*, **65**(4): 475-484.
- SOHN, G., and DOWMAN, I. J., 2001, Extraction of Buildings from High-resolution Satellite Data. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 345-355.
- SONG, C., WOODCOCK, C. E., SETO, K. C., LENNEY, M. P., and MACOMBER, S. A., 2001, Classification and Change Detection using Landsat TM Data: When and how to correct atmospheric effects? *Remote Sensing of Environment*, **75**(2): 230-244.
- SONKA, M., HLAVAC, V., and BOYLE, R., 1999, *Image Processing, Analysis and Machine Vision*, 2nd Edition (Pacific Grove, CA: PWS Publishing).
- SOUZA, I. M., PEREIRA, M. N., and KURKDJIAN, M. L. N. O., 2002, Evaluation of High Resolution Satellite Images for Urban Population Estimation: *Proceedings of the Third Symposium for Remote Sensing of Urban Areas* INPE, Brazil.
- SRINIVAS, M., and PATNAIK, L. M., 1994, Genetic Algorithms: A Survey. *Computer*, **27**(6): 17-26.
- SUI, H., ZHOU, Q., GONG, J., and MA, G., 2008, Processing of Multitemporal Data and Change Detection. In *Advances in Photogrammetry, Remote Sensing and Spatial Information Sciences: 2008 ISPRS Congress Book*, edited by Z. Li, Chen, J. and Baltsavias, E. (London: CRC Press), 227-247.
- SUN, H., LI, S. X., LI, W. J., MING, Z., and CAI, S. B., 2005, Semantic-based Retrieval of Remote Sensing Images in a Grid Environment. *IEEE Geoscience and Remote Sensing Letters*, **2**(4): 440-442.
- SURVEY, O., 2009, *Welcome to Ordnance Survey Research* Ordnance Survey, <http://www.ordnancesurvey.co.uk/oswebsite/partnerships/research/> Accessed 6 March 2009
- SWEETING, M. N., and UNDERWOOD, C. I., 2003, Small-Satellite Engineering and Applications. In *Spacecraft Systems Engineering*, edited by P. Fortescue, J. Stark, and G. Swinerd (Chichester, UK: Wiley).
- TAUD, H., and PARROT, J.-F., 1992, Detection of Circular Structures on Satellite Images. *International Journal of Remote Sensing*, **13**(2): 319-335.
- TOMPKINSON, W., SEAMAN, E., HOLLAND, D., and GREENWOOD, J., 2003, An Automated Change Detection System: Making Change Detection Applicable to a Production Environment: *Proceedings of the Second International Workshop on the Analysis of Multi-temporal Remote Sensing Images* Joint Research Centre, Ispra, Italy.
- TRALLI, D. M., BLOM, R. G., ZLOTNICKI, V., DONNELLAN, A., and EVANS, D. L., 2005, Satellite Remote Sensing of Earthquake, Volcano, Flood, Landslide and Coastal Inundation Hazards. *ISPRS Journal of Photogrammetry and Remote Sensing*, **59**(4): 185-198.
-

-
- TRALLI, M., 2001, *Assessment of Advanced Technologies for Loss Estimation* (State University of New York at Buffalo: Multidisciplinary Center for Earthquake Engineering Research).
- TRINDER, J. C., 1995, Semi-Automatic Feature Extraction by Snakes. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (I)*, edited by A. Gruen, O. Kuebler, and P. Agouris (Berlin: Birkhäuser Verlag), 95-104.
- UN, 2007, *World Urbanization Prospects: The 2007 Revision and World Urbanization Prospects*. Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat, <http://esa.un.org/unup>, Accessed 22 October 2009
- UNOSAT, 2007, *UNOSAT: Who We Are*. UNOSAT, <http://unosat.web.cern.ch/unosat/>, Accessed 15 May 2007
- VAN DER WERFF, H., 2006, Knowledge-Based Remote Sensing of Complex Objects: Recognition of spectral and spatial patterns resulting from natural hydrocarbon seepages University of Utrecht, Enschede, The Netherlands.
- VOSELMAN, G., GORTE, B. G. H., SITHOLE, G., and RABBANI, T., 2004, Recognising Structure in Laser Scanner Point Clouds. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, **46**(8/W2): 33-38.
- VU, T. T., MATSUOKA, M., and YAMAZAKI, F., 2005, Preliminary Results in Development of an Object-based Image Analysis Method for Earthquake Damage Assessment: *Proceedings of the Third International Workshop on Remote Sensing for Post-Disaster Response*, Chiba, Japan. EERI / MCEER.
- VU, T. T., MATSUOKA, M., and YAMAZAKI, F., 2007, Dual-scale Approach for Detection of Tsunami-affected Areas Using Optical Satellite Images. *International Journal of Remote Sensing*, **28**(13): 2995-3011.
- WANG, Y., TANG, M., TAN, T., and TAI, X., 2004, Detection of Circular Oil Tanks Based on the Fusion of SAR and Optical Images: *Proceedings of the Third International Conference on Image and Graphics*, Hong Kong, China. ICIG.
- WIKANTIKA, K., SINAGA, A., HADI, F., and DARMAWAN, S., 2007, Quick Assessment on Identification of Damaged Building and Land-use Changes in the Post-tsunami Disaster with a Quick-look Image of IKONOS and Quickbird (a Case Study in Meulaboh City, Aceh). *International Journal of Remote Sensing*, **28**(13): 3037-3044.
- WIKIPEDIA, 2009, *2005 Hertfordshire Oil Storage Terminal fire* Wikipedia, The Free Encyclopedia, http://en.wikipedia.org/w/index.php?title=2005_Hertfordshire_Oil_Storage_Terminal_fire&oldid=281484824, Accessed 4 April 2009
- WOMBLE, J. A., 2005, Remote-sensing Applications to Windstorm Damage Assessment Texas Tech University, Lubbock, TX, USA.
- WOMBLE, J. A., GHOSH, S., ADAMS, B. J., and FRIEDLAND, C. J., 2006, *Advanced Damage Detection for Hurricane Katrina: Integrating Remote Sensing and VIEWS Field Reconnaissance*. MCEER Special Report Series: Engineering and Organisational Issues Before, During and After Hurricane Katrina. MCEER, Buffalo.
-

-
- XIAO, J., LI, J., and MOODY, A., 2003, A detail-preserving and flexible adaptive filter for speckle suppression in SAR imagery. *International Journal of Remote Sensing*, **24**: 2451-2465.
- YAMAZAKI, F., YANO, Y., and MATSUOKA, M., 2005, Visual Damage Interpretation of Buildings in Bam City using QuickBird Images Following the 2003 Bam, Iran, Earthquake. *Earthquake Spectra*, **21**(Special Issue 1): S329-S336.
- YANG, M. D., SU, T. C., HSU, C. H., CHANG, K. C., and WU, A. M., 2007, Mapping of the 26 December 2004 tsunami disaster by using FORMOSAT-2 images. *International Journal of Remote Sensing*, **28**(13): 3071-3091.
- ZHAN, X., DEFRIES, R., TOWNSHEND, J. R. G., DIMICELI, C., HANSEN, M., HUANG, C., and SOHLBERG, R., 2000, The 250 m Global Land Cover Change Product from the Moderate Resolution Imaging Spectroradiometer of NASA's Earth Observing System. *International Journal of Remote Sensing*, **21**: 1433-1460.
- ZHANG, J.-F., XIE, L., and TAO, X., 2003, Change Detection of Earthquake-damaged Buildings on Remote Sensing Image and its Application in Seismic Disaster Assessment: *Proceedings of 2003 IEEE International Geoscience and Remote Sensing Symposium (IGARSS '03)*. 2436-2438.
- ZHANG, Y., 2001, Detection of Urban Housing Development by Fusing Multisensor Satellite Data and Performing Spatial Feature Post-classification. *International Journal of Remote Sensing*, **22**(17): 3339-3355.
- ZHANG, Z., ZHANG, J., and HU, X., 2001, Semi-automatic Building Extraction from Stereo Image Pairs. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*, edited by E. Baltsavias, A. Gruen, and L. Van Gool (Abingdon: A.A. Balkema), 115-122.
-