Terrain Aided Localisation of Autonomous Vehicles in Unstructured Environments

by

Rajmohan Madhavan

ISBN: 1-58112-177-6

DISSERTATION.COM

Parkland, FL • USA • 2003

Terrain Aided Localisation of Autonomous Vehicles in Unstructured Environments

Dissertation.com
USA • 2003

ISBN: 1-58112-177-6
www.Dissertation.com/library/1121776a.htm

Terrain Aided Localisation of Autonomous Vehicles in Unstructured Environments

Rajmohan Madhavan

A Thesis submitted in fulfillment
of the requirements for the degree of
Doctor of Philosophy

Australian Centre for Field Robotics
Department of Mechanical and Mechatronic Engineering
The University of Sydney

January 2001

To my parents, my lovely wife Kalpana and our beloved daughter Priyanka.

Foreword

About the Author:

Rajmohan Madhavan was born in Madras, India, in 1973. He received the Bachelor of Engineering degree (Electrical and Electronics) in 1995 from the College of Engineering, Guindy, Anna University, India and the Master of Engineering (Research) degree (Control and Robotics) in 1997 from the Department of Engineering, The Australian National University, Australia. He received his Ph.D. degree (Field Robotics) in 2001 from the School of Aerospace, Mechanical and Mechatronic Engineering (Department of Mechanical and Mechatronic Engineering at the time of completion), The University of Sydney, Sydney, Australia. Since 2001, he has been a research associate with the Oak Ridge National Laboratory (ORNL) and is currently a guest researcher with the Intelligent Systems Division of the National Institute of Standards and Technology (NIST). His current research interests include autonomous vehicle navigation in large and complex environments, distributed heterogenous sensing and systems and control theory.

Author's Contact Email: `raj_madhavan@yahoo.com`

Declaration

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any other degree or diploma of the University or other Institute of higher learning, except where due acknowledgement has been made in the text.

Rajmohan Madhavan
January 08 2001.

Abstract

This thesis is concerned with the theoretical and practical development of reliable and robust localisation algorithms for autonomous land vehicles operating at high speeds in unstructured, expansive and harsh environments. Localisation is the ability of a vehicle to determine its position and orientation within an operating environment. The need for such a localisation system is motivated by the requirement of developing autonomous vehicles in applications such as mining, agriculture and cargo handling. The main drivers in these applications are for safety, efficiency and productivity. The approach taken to the localisation problem in this thesis guarantees that the safety and reliability requirements imposed by such applications are achieved. The approach also aims to minimise the engineering or modification of the environment, such as adding artificial landmarks or other infrastructure. This is a key driver in the practical implementation of a localisation algorithm.

In pursuit of these objectives, this thesis makes the following principal contributions:

1. The development of an Iterative Closest Point - Extended Kalman Filter (ICP-EKF) algorithm - a map-based iconic algorithm that utilises measurements from a scanning laser rangefinder to achieve localisation. The ICP-EKF algorithm entails the development of a map-building algorithm. The main attraction of the map-based localisation algorithm is that it works directly on sensed data and thus does not require extraction and matching of features. It also explicitly takes into account the uncertainty associated with measurements and has the ability to include measurements from a variety of different sensors.

2. The development and implementation of an entropy-based metric to evaluate the information content of measurements. This metric facilitates the augmentation of landmarks to the ICP-EKF algorithm thus guaranteeing reliable and robust localisation.

3. The development and adaptation of a view-invariant Curvature Scale Space (CSS) landmark extraction algorithm. The algorithm is sufficiently robust to sensor noise and is capable of reliably detecting and extracting landmarks that are naturally present in the environment from laser rangefinder scans.

4. The integration of the information metric and the CSS and ICP-EKF algorithms to arrive at a unified localisation framework that uses measurements from both artificial and natural landmarks, combined with dead-reckoning sensors, to deliver reliable vehicle position estimates.

The localisation framework developed is sufficiently generic to be used on a variety of other autonomous land vehicle systems. This is demonstrated by its implementation using field data collected from three different trials on three different vehicles. The first trial was carried out on a four-wheel drive vehicle in an underground mine tunnel. The second trial was conducted on a Load-Haul-Dump (LHD) truck in a test tunnel constructed to emulate an underground mine. The estimates of the proposed localisation algorithms are compared to the *ground truth* provided by an artificial landmark-based localisation algorithm that uses bearing measurements from a laser. To demonstrate the feasibility and reliability of both the natural landmark extraction and localisation algorithms, these are also implemented on a utility vehicle in an outdoor area within the University's campus. The results demonstrate the robustness of the proposed localisation algorithms in producing reliable and accurate position estimates for autonomous vehicles operating in a variety of unstructured domains.

Acknowledgements

"The human mind is not capable of grasping the Universe. We are like a little child entering a huge library. The walls are covered to the ceilings with books in many different tongues. The child knows that someone must have written these books. It does not know who or how. It does not understand the languages in which they are written. But the child notes a definite plan in the arrangement of the books - a mysterious order which it does not comprehend, but only dimly suspects"
- Albert Einstein

Let me start by thanking Dr. Gamini Dissanayake and Prof. Hugh Durrant-Whyte for their support and guidance throughout the Ph.D. program. Dissa taught me to think critically and to always keep the bigger picture in mind. I appreciate all his help and suggestions during my candidature. My sincere thanks are due to Hugh for giving me an opportunity to work with the ACFR and for his enthusiasm and direction. I found both of them to be very approachable and willing to assist in every step of the way. The years I have spent with the ACFR and the CMTE have been extremely rewarding and has been a time of personal and professional growth. Most of all, it's been a humbling experience.

The three field trials that are reported in this thesis would not have been possible but for the following people:

`Eric Nettleton for the 4WD trials` - Eric's help with data-logging and calibration trials is very much appreciated. Sorry about the freezing conditions in Forest Lodge in Brisbane!

`The CSIRO/CMTE/AMIRA team for the LHD trials` - I would like to thank the team and especially Jonathan Roberts. Jon was always willing to help every time. I am much obliged for his help. I would also like to thank the CSIRO team for their help in various issues during several visits to QCAT with Eric Nettleton and Ali Göktogan: Jock Cunningham, Graeme Winstanley, Stuart Wolfe, Leslie Overs, Reece McCasker, Elliot Duff, Peter Corke and Pavan Sikka.

`José "Señor" Guivant for the Utility vehicle trials` - Jose's help and time in gathering the field data is also much appreciated. He's been a good friend and long will I remember our conversations way past midnight on subjects ranging from the frivolous *maté* to the very existence of life. *Gracias, mi amigo.*

Thanks also goes to CMTE for its financial assistance in the form of a PhD studentship. CMTE's financial support to attend the SME Annual meeting to receive the outstanding student paper award is also much appreciated. I would like to thank CEO Mike Hood, Education Program Leader Dominic Howarth and Hal Gurgenci. ACFR's financial assistance for several conferences and visits is also gratefully acknowledged.

I would also like to thank:

- ACFR and the University of Sydney gang: Anna McCallan for helping me to sort out administrative issues especially in the last year among other things, Eduardo Nebot for advice on career plans, Som Majumder and Stefan Williams for being my cubicle-mates and for putting up with my self-cursing and incessant questions, Quang Nguyen and Miguel Santos for those intense ping-pong games, Mike Stevens for his system administrative help (who also does an almost perfect Indian accent), Ralph Koch for his help with my temperamental bike, Ben Grocholsky for discussions on entropy, Tim Bailey for discussions on Blues and in general on guitarists, Steve Scheding for his help during a visit to CMU, Salah Sukkarieh for all those wonderful Lebanese sweets, Trevor "Irishman" Fitzgibbons, Keith Willis, Tomonari Furukawa, Monica Louda, Jong-Hyuk Kim, Chris Misfud, Ikramy Garas, Ali Göktogan for a trip to Montville and the turkish delight, Gurce Isikyildiz for introducing me to the *Saz*, Maher "Sugar" Magrabi, Jeff "Hotpants" Leal for introducing me to the Dave Matthews Band, Julio Rosenblatt, Vinita Martin, Lyn Kennedy and Susan Gonzalez.

- Brisbane gang: Adrian and Sanda Bonchis, Bradley and Natalie Horton, Jasmine Banks, Yvonne Kolatschek, Shivakumar Karekal, Tricia Dolphin, Helene Marszalek and Lena Mete for their assistance and company during several visits to CMTE.

- *Visitors* to ACFR: Henrik "Ribs" Svedlund, Joachim "Sir Jake" Andersson, Sven Ronnbäck and Tobias Carlsson (from Sweden), Gerold Kloos and Stefan Baiker (from Germany), Rob Dawkins (from UK), Guillame Lebas (from France) for introducing me to the cultures and customs of their countries and teaching me a little bit of their languages (especially the words that I could have done without during visits to their countries!). These guys helped me to go and see places in Sydney and elsewhere that I, as a Sydneysider, for the best part of the last three years, never knew of their existence! Thanks fellas.

- My friends Rittwik Jana, Chandar Muthukrishnan, Rajesh Sambandam, Magesh Govindarajan, Venkatesh Babu, Rajkumar Kandasamy, Kangesh Gunaseelan, Devaraj Nagarajan and Sharon Laubach - These guys kept me sane during the later half of the Ph.D. (even though most of them are on the other side of the Pacific!), Simon Julier for discussions with the Unscented filter and life otherwise.

- Everyone else whom I have forgotten to mention.

I reserve the greatest thanks to my parents and to my brother. They have stood by me all these years and I am indebted for their love, support and prayers. At the end of it all, I can always go home.

Nomenclature

Only frequently used notations are included here.

x_v	x coordinate of the vehicle
y_v	y coordinate of the vehicle
ϕ_v	orientation of the vehicle
r_v	wheel radius of the vehicle
α_v	rear slip angle of the LHD
β_v	front slip angle of the LHD
ss_v	shaping state for the gyro drift
R	range of the scanning laser rangefinder
θ	bearing of the scanning laser rangefinder
θ_{gcs}	bearing of the bearing-only laser
$R^j_{n\ell}$	range of the jth natural landmark
$\theta^j_{n\ell}$	bearing of the jth natural landmark
(X_i, Y_i)	cartesian location of artificial landmark i
$(x^j_{n\ell}, y^j_{n\ell})$	cartesian location of natural landmark j
$\mathbf{x}_k$	vehicle state at discrete time instant k
$\mathbf{u}_k$	vehicle control input at discrete time instant k
$\mathbf{z}_k$	vehicle observations (measurements) available at discrete time instant k
$\mathbf{x}_{(k\|k-1)}$	predicted vehicle state at discrete time instant k based on the state estimate at time instant $(k-1)$
$\mathbf{P}_{(k\|k-1)}$	predicted vehicle covariance at discrete time instant k based on the covariance estimate at time instant $(k-1)$
$\mathbf{x}_{(k\|k)}$	updated vehicle state at discrete time instant k obtained by combining the predicted estimate at time instant k and the observation(s) at time instant k
$\mathbf{P}_{(k\|k)}$	updated vehicle covariance at discrete time instant k obtained by combining the predicted covariance estimate at time instant k and the observation(s) at time instant k
$\mathrm{E}[\cdot]$	mathematical expectation operator

Abbreviations

ACFR	Australian Centre for Field Robotics
AMIRA	Australian Mineral Industries Research Association
CSIRO	Commonwealth Scientific Industrial Research Organisation
CRC	Cooperative Research Centre
CMTE	CRC for Mining Technology and Equipment
CMST	CSIRO Manufacturing Science and Technology
E & M	CSIRO Exploration and Mining
EKF	Extended Kalman Filter
MMSE	Minimum Mean Squared Estimate
ICP	Iterative Closest Point
CSS	Curvature Scale Space
4WD	Four-Wheel Drive
LHD	Load-Haul-Dump
AGV	Automated Guided Vehicle
TOF	Time Of Flight
GPS	Global Positioning System
IMU	Inertial Measurement Unit
INS	Inertial Navigation System

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Aims and Objectives

THE aim of this thesis is to develop reliable and robust terrain aided localisation systems for autonomous land vehicles operating in unstructured, expansive and harsh environments. The resulting localisation system should be able to estimate the position of the vehicles to sub-metre accuracy and also minimise the requirement to add infrastructure to the operating environment.

Localisation is defined as the ability of a vehicle to determine its two-dimensional position (x_v,y_v) and orientation (ϕ_v) within the operating environment. For autonomous vehicle localisation and in the context of this thesis, the definitions of the salient terms in the above aim are:

- *Reliability* denotes the ability to deliver precise vehicle estimates over many trials in a variety of operating conditions. Thus reliability encompasses both *accuracy* and *repeatability*.
- *Robustness* indicates the competence to cope with external unmodelled disturbances.
- *Unstructured* implies a physical environment that does not have much regular structure or layout.
- An *expansive* environment is one which is much larger than the range of the sensors available on the vehicle.
- *Harshness* refers to the uneven undulatory terrain of the operating environment.

The objectives of this thesis are:

- To develop algorithms that will guarantee reliable and robust localisation with particular attention to challenges imposed by the operating environment.
- To investigate the feasibility of utilising landmarks that are naturally available in the operating environment with the intention of minimising modifications to the environment. Towards this, develop an algorithm that is capable of reliably detecting and extracting natural landmarks present in the environment.

- To develop an algorithmic framework that allows the utilisation of measurements (observations) from multiple sensors and has the potential to use measurements that provide maximum information in reducing localisation error. This requires an analysis and understanding of the limitations of employing landmark-based localisation systems in unstructured environments.
- To provide a demonstration of the algorithms developed using real data obtained from field trials conducted under varying operating conditions with attention to practical considerations.
- To examine and demonstrate the feasibility of these algorithms in other autonomous land vehicle applications.

1.2 Motivation and Background

There is a strong case for the use of autonomous systems in various domains, including planetary exploration, mining, nuclear waste remediation, military reconnaissance, agriculture and construction [143],[197],[238],[30],[193],[198] with potential advantages ranging from increased safety to productivity benefits to environmental reasons. The general problem of autonomous vehicle navigation can be summarised by the three questions: "Where am I?", "Where am I going?", "How do I get there?" [145]. The first question is termed *localisation*, the second as *goal recognition* and the third as *path planning*. Localisation is the process of determining the position and orientation (hereafter collectively referred to as the *pose*) of the vehicle within the operating environment. To realise the second and third tasks, given uncertainty in control and sensing, localisation is a fundamental ability required of an autonomous vehicle.

1.2.1 Sensor-based Localisation

The problem of autonomous vehicle localisation has received considerable attention from roboticists in recent years and many methods have been proposed. These methods vary significantly, depending on the environment in which the vehicle is to navigate and the type of sensors that are available. A survey of the vast body of literature [72],[48],[115],[26],[239],[20] on robot localisation indicates that the proposed approaches can be broadly classified as *relative* and *absolute* methods. Most implementations of the localisation algorithms combine the above two methods.

Dead-reckoning (derived from deduced reckoning) is the process of determining the position of a vehicle by incrementally integrating motion information over a period of time. Odometry and Inertial Navigation System (INS) are the widely used dead-reckoning methods for vehicle positioning. Odometric methods use encoders to measure wheel rotation and/or steering orientation. An inertial navigation system makes measurements of the accelerations and attitude rates of the vehicle using an inertial measurement unit (IMU) consisting of gyroscopes and accelerometers. Both odometry and inertial sensing provide high frequency, high bandwidth data.

The drawback of odometry is that the estimated vehicle position at a given instant depends on the previous estimate. This makes it difficult to eliminate errors associated with the previous cycle due to sensor inaccuracies, the assumption that the heading remains constant over the sampling interval, wheel slippage and quantisation effects. As a consequence, the

vehicle pose (position and especially the orientation) becomes less and less certain and the errors associated with the pose grow without bound. When the information from an IMU is appropriately integrated, the INS can be used to determine vehicle pose. However, integration of these noisy signals will result in a gradual growth or drift in pose estimate error.

The major sources of dead-reckoning errors can be grouped into two categories namely *systematic* and *non-systematic* errors [25]. Systematic errors accumulate constantly over time due to unequal wheel diameters, misalignment of wheels, encoder sampling rates, drifts associated with time etc. whereas non-systematic errors are as a result of wheel slippage, uneven terrain conditions and may occur unexpectedly. For a vehicle navigating on uneven terrain the non-systematic errors are more dominant than the systematic errors. Due to the accumulation of errors in relative sensor-based methods, they need to be periodically reset by applying external corrections from absolute positioning methods. Consequently, odometry and INS are almost always used in conjunction with some absolute position sensing system.

Examples of absolute positioning systems include beacons placed at known locations in the environment in combination with a sensor on board the vehicle, for example a sonar or a laser, or Global Positioning System (GPS) reported fixes. Such absolute systems can then be used in conjunction with odometry/INS to provide a non-divergent estimate of the vehicle position. The resetting may also be accomplished by utilising an *a priori* map and/or natural features[1] in place of artificial landmarks. A typical navigation loop consists of high frequency dead-reckoning sensors that predict vehicle pose and low frequency absolute sensors that provide external aiding information to bound positioning errors.

Localisation based on a combination of relative and absolute methods can be subdivided into three basic categories namely [99]:

- **Landmark-based Methods**:
  There are several definitions of landmarks reported in the literature [179],[149],[147]. In the context of this thesis, landmarks are taken to mean a localised, stationary, two-dimensional physical feature that the vehicle sensors can reliably and efficiently extract and recognise for determining vehicle position. The major drawback of landmark-based approaches is that the location of the landmarks must be known *a priori*. In particular, artificial landmark approaches require careful engineering of the environment whereas natural landmark approaches require reliable feature extraction algorithms.

- **Iconic Matching Methods**:
  Iconic methods do not require the extraction of geometric features. The matching works directly on sensed data rather than on a small set of features, thus eliminating the need to determine what constitutes a landmark and to extract features. Iconic methods attempt to globally minimise the discrepancies between sensed data and a model of the environment. Because the search is confined to small perturbations of the sensor scans, it is computationally efficient. The drawback of these methods is that they are sensitive to bad matches and incorrect error models.

- **Behaviour-based Methods**:
  The central theme of behaviour-based robotics is to directly couple a robot's sensors and actuators so as to obviate the need of maintaining an environment model. Behavioural approaches rely on the interaction of the robot with its environment in

[1] "Landmarks" and " Features" are used interchangeably except when an explicit distinction between natural and artificial landmarks is made.

which it is to navigate. Thus their ability to localise geometrically is limited since their navigation capability is implicitly dependent on their sensor/action history. There are some advantages in using a behaviour-based scheme. These schemes are more responsive to the environment and their performance degrades gradually and gracefully in the case of system failure rather than abruptly and catastrophically. Conversely, the interactions are not easily modelled, are less predictable and so are difficult to understand and modify.

Achieving localisation in structured environments is simpler than that in unstructured environments. In the context of this thesis, *structured* environments are taken to mean even-terrain feature-rich domains such as offices. An *unstructured* environment is taken to mean environments with uneven terrain susceptible to vehicle slipping and skidding and few distinguishable naturally occurring landmarks. From an autonomous vehicle navigation point of view, an unstructured environment presents the following unique challenges:

- Errors introduced due to distance travelled can be significant and unpredictable. This is a direct consequence of the undulatory nature of the terrain of travel and the uncertainties introduced into sensor data. Errors introduced by irregular terrain conditions can not be predicted or modelled. Higher vehicle speeds in such environments also exacerbate these errors.
- For large unstructured environments, it is extremely difficult to construct perfect *maps* for navigation as compared to small structured domains where construction of such maps is possible due to the underlying polygonal structure of the environment.
- Reliable detection and association of both artificial and natural landmark measurements from corrupted sensor data is difficult. The landmark detection and data association algorithms should be robust to false alarms so as to prevent catastrophic failure. Occlusions and outliers also complicate the process of determining a legitimate landmark.
- Sensor suite selection should take into consideration the challenges posed by the unstructured domain. For example, GPS can not be employed underground since the required line-of-sight can not be maintained. Furthermore, sensors that are commonly employed in a structured environment may not function properly in an unstructured environment due to various external environmental factors like humidity, dust and heat. Other detrimental effects include excessive vibration and pitching and rolling motions experienced by vehicles navigating in such domains. Sufficient care should be taken to ensure that the sensors have satisfactory range and attenuation to be useful in unstructured environments.

These challenges make it comparatively difficult to realise successful navigation in unstructured environments.

In mining, automation of underground operations is an important goal for reasons of personnel safety and lower operational costs. With automation of underground vehicles, productivity benefits come from the increase in operating hours, reduced wear and tear, elimination of travelling time to the machine and greatly reduced operator fatigue. This thesis is a wider part of a project aimed at achieving autonomy of Load-Haul-Dump (LHD) vehicles in underground mine tunnels. It was funded by the Australian Mineral Industries Research

Association (AMIRA) and was a collaborative venture between the Commonwealth Scientific Industrial Research Organisation (CSIRO) Manufacturing Science and Technology (MST) automation group, CSIRO Exploration and Mining (E & M) and the University of Sydney's Australian Centre for Field Robotics (ACFR) conducted under the auspices of the Cooperative Research Centre for Mining Technology and Equipment (CMTE).

Motivated by the above factors, this thesis develops localisation algorithms for autonomous navigation of vehicles operating in unstructured environments, like underground mines, by drawing upon the knowledge gained from the collaborative venture and work performed alongside underground miners and mining researchers. The localisation algorithms described in this thesis fall into the iconic and landmark-based category but are sufficiently generic to be used on a variety of other autonomous land vehicles. These methods were chosen because of their efficiency and robustness. The algorithms also seek to minimise the engineering or modification of the environment, such as adding artificial landmarks or other infrastructure. This is a key driver in the practical implementation of a localisation algorithm.

In this thesis, observations from natural and/or artificial landmarks need to be combined with dead-reckoning estimates to estimate vehicle pose. This requires a framework within which the sensor data can be fused synergically. The Kalman filter is a statistically efficient way of achieving this requirement. The Kalman filter has been widely employed to fuse low-level redundant data to provide an estimate of the states of a system [162] and is employed as the sensor fusion and estimation tool in this thesis. Sections A.1 and A.2 of the Appendix, respectively, describe the linear Kalman filter and the extended Kalman filter (EKF) from a navigation point of view to introduce notations that will be used throughout this thesis.

1.3 Contributions of the Thesis

This thesis makes the following theoretical and practical contributions:

- The development and demonstration of a robust framework that efficiently reduces infrastructure by combining measurements from artificial and natural landmarks to reliably estimate vehicle position.
- The development of an iconic map-based localisation algorithm for a four-wheel drive vehicle and a Load-Haul-Dump truck operating in mine tunnels. The localisation algorithm entails the development of a map-building algorithm.
- The analysis, development and adaptation of a view-invariant multiscale natural landmark extraction algorithm that is robust with respect to sensor noise.
- The development of a metric for the evaluation of information content of measurements towards the selection of a landmark from a given set of landmarks.
- The demonstration of a natural landmark localisation algorithm on a utility vehicle obviating the need for artificial landmarks.
- The implementation of an artificial landmark localisation algorithm that facilitates the comparison and analysis of the proposed algorithms, by providing the *ground truth*.

1.4 The Structure of the Thesis

Chapter 2 presents an artificial landmark-based localisation algorithm for two vehicles - a four-wheel drive vehicle and a Load-Haul-Dump truck. The algorithm is based on measurements obtained from a bearing-only laser that returns bearing values to artificial landmarks strategically placed in the operating environment of the vehicles.

The sensor suite, the calibration trials and the hardware and software aspects of the data logging system for the four-wheel drive vehicle are first described. This is followed by the description of the process and observation models and observation validation procedures that are required for the successful implementation of the localisation algorithms. The corresponding experimental results for two field trials for the four-wheel drive vehicle and the load-haul-dump truck are then presented and discussed.

A critical overview of various landmark-based navigation methods and previous attempts to automate underground mining vehicles is also included.

Chapter 3 develops a map-based iconic algorithm for localisation of the four-wheel drive vehicle and the load-haul-dump truck. The proposed algorithm requires the development of a map-building algorithm. Both the map-building and localisation algorithms use data from a time-of-flight range and bearing laser.

The shortcomings of the proposed iconic localisation algorithm are identified and a strategy to overcome these deficiencies is proposed. This leads to the development of a metric that embodies the information contained in a measurement. Based on this metric, the map-based iconic localisation algorithm is augmented with artificial landmarks. Experimental results of the proposed map-based localisation algorithm for both the vehicles are discussed.

An overview of map-based navigation methods is also provided.

Chapter 4 is concerned with the development of a multiscale natural landmark extraction algorithm. The Chapter starts out with a comprehensive review of various multiscale description techniques with particular attention to the concept of scale and its selection. The proposed algorithm has the advantage of examining laser range scans at varying levels of detail to arrive at a stable natural landmark. The need to decide upon what constitutes a landmark and the motivation behind the selection of a primitive for the extraction of a landmark in the operating environment are also elaborated. The proposed algorithmic procedures lead to the construction of a natural landmark map.

Chapter 5 develops a unified algorithmic framework for localisation of autonomous vehicles operating in unstructured environments by combining the ideas presented in Chapters 3 and 4. By using measurements from both artificial and natural landmarks, with the aid of the information metric developed in Chapter 3 and the natural landmark map and extraction algorithm detailed in Chapter 4, a minimal infrastructure localisation algorithm is demonstrated on two field trials for localisation of the load-haul-dump truck. The developed framework is shown to deliver reliable vehicle pose estimates proving the robustness of the proposed algorithms.

The Chapter then extends the multiscale natural landmark algorithm to demonstrate localisation of a utility vehicle without the use of any external infrastructure.

Chapter 6 provides a summary and lists the principal results of the thesis. Suggestions for further research and the extension of the theories developed in this thesis to other areas are also briefly discussed.