

A Comparison of Data Association Techniques for Simultaneous Localization and Mapping

by

Aron J. Cooper

B.A.E.M Aerospace Engineering and Mechanics, University of Minnesota, 2003

B.S. Astrophysics, University of Minnesota, 2003

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN AERONAUTICS AND ASTRONAUTICS

AT THE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2005

© Aron J. Cooper, 2005. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author: _____

Department of Aeronautics and Astronautics
May 20, 2005

Certified by: _____

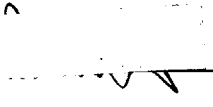


Nicholas Roy, Ph.D.

Assistant Professor of Aeronautics and Astronautics

Thesis Supervisor

Certified by: _____

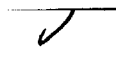


Don Gustafson, Ph.D.

Distinguished Member of the Technical Staff, C. S. Draper Laboratory

Thesis Supervisor

Certified by: _____

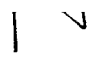


Marc McConley, Ph.D.

Principal Member of the Technical Staff, C. S. Draper Laboratory

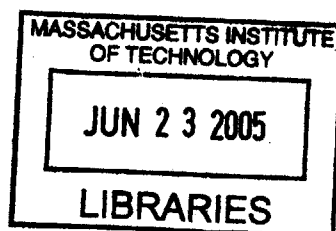
Thesis Supervisor

Accepted by: _____



James Peraire, Ph.D.

Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students



AERO

Handwritten scribbles and marks.

[This page intentionally left blank.]

Handwritten scribbles and marks.

SEARCHED
SERIALIZED
INDEXED
FILED

Handwritten scribbles and marks.

A Comparison of Data Association Techniques for Simultaneous Localization and Mapping

by

Aron J. Cooper

Submitted to the Department of Aeronautics and Astronautics
on May 20, 2005, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

The problem of Simultaneous Localization and Mapping (SLAM) has received a great deal of attention within the robotics literature, and the importance of the solutions to this problem has been well documented for successful operation of autonomous agents in a number of environments. Of the numerous solutions that have been developed for solving the SLAM problem many of the most successful approaches continue to either rely on, or stem from, the Extended Kalman Filter method (EKF). However, the new algorithm FastSLAM has attracted attention for many properties not found in EKF based methods. One such property is the ability to deal with unknown data association and its robustness to data association errors.

The problem of data association has also received a great deal of attention in the robotics literature in recent years, and various solutions have been proposed. In an effort to both compare the performance of the EKF and FastSLAM under ambiguous data association situations, as well as compare the performance of three different data association methods a comprehensive study of various SLAM filter-data association combinations is performed. This study will consist of pairing the EKF and FastSLAM filtering approaches with the Joint Compatibility, Sequential Compatibility Nearest Neighbor, and Joint Maximum Likelihood data association methods. The comparison will be based on both contrived simulations as well as application to the publicly available *Car Park* data set.

The simulated results will demonstrate a heavy dependence on geometry, particularly landmark separation, for the performance of both filter performance and the data association algorithms used. The real world data set results will demonstrate that the performance of some data association algorithms, when paired with an EKF, can give identical results. At the same time a distinction in mapping performance between those pairings and the EKF paired with Joint Compatibility data association will be shown. These EKF based pairings will be contrasted to the performance obtained for the FastSLAM-Sequential Nearest Neighbor marriage. Finally, the difficulties in applying the Joint Compatibility and Joint Maximum Likelihood data association methods using FastSLAM 1.0 for this data set will be discussed.

Thesis Supervisor: Nicholas Roy, Ph.D.

Title: Assistant Professor of Aeronautics and Astronautics

Thesis Supervisor: Don Gustafson, Ph.D.

Title: Distinguished Member of the Technical Staff, C. S. Draper Laboratory

Thesis Supervisor: Marc McConley, Ph.D.

Title: Principal Member of the Technical Staff, C. S. Draper Laboratory

Acknowledgments

First, I would like to thank the world class institutions of Draper Laboratories, MIT, and the Department of Aeronautics and Astronautics for allowing me the opportunity to live, study, and perform research in such an intellectually stimulating environment. The talented people at these fine institutions create an environment with direction, purpose, and credentials.

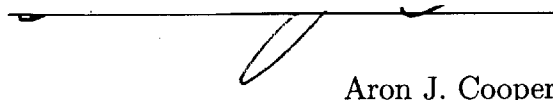
I would like to thank a number of people that have made my experience at MIT and Draper Labs unforgettable. The person with whom I have had the lengthiest interaction, and whom I have gotten to know both on a personal and professional level, is Don Gustafson. Don is not only a very talented and brilliant engineer, but also a terrific advisor. Above all else Don is an inspiring teacher. It has been an honor to work with Don for the time that I have been at Draper; given the opportunity I would not hesitate to do so again. In addition to Don, I have had the privilege of working with Marc McConley as a second Draper advisor. Marc has been a terrific program manager for the IR&D project that I have worked on during my time at Draper. In addition, he has been a great advisor and more recently a very talented proofreader. He has offered me invaluable feedback and assistance on my thesis. I would also like to offer thanks to Tom Thorvaldsen, who has acted as a de facto problem solver for my computer problems. He is someone who was always willing to offer me assistance.

During the course of my thesis research and writing process I was privileged to have Professor Nick Roy as my MIT thesis advisor. Without question the quality of my thesis benefited greatly from Professor Roy's dedication to the highest standards of research and an unwavering determination for me to hold similar standards during this process. I am also very thankful to Professor Roy for his willingness to perform a number of iterations on making corrections and edits to my thesis, the finished product will certainly reflect his painstaking assistance. I would also like to thank Professor John Deyst for his willingness to include me in the Draper-MIT research group meetings, which helped to influence and motivate my research. Additionally, I would like to thank Professor Eric Feron who has influenced my thesis work by allowing me the pleasure of participating in his classes. It

is my belief that Professor Feron is a fantastic teacher whose lectures at times are so powerful that they transcend the subject matter in such a way that one cannot help but to be inspired in their own work.

In an effort, such as the research for and writing of a Master's thesis, which involves so much time and energy, it is not possible to overstate the importance of my personal relationships. These relationships have served to balance out my life and provide me with a great deal of motivation to constantly push forward. I would like to thank all of the Draper Fellows and friends I have made while at MIT, whom have made the experience of being here that much more enjoyable. The Friday treat and social gatherings were a good deal of fun and will be remembered for some time. In particular, I would like to thank my good friend Peter Lommel, with whom I share such similar life experiences it is almost scary. Peter is one of the most intelligent people that I know, and it has been my good fortune to have a friend such as Peter who is always willing to lend an ear to my latest hair-brained idea, or whatever it is that I have on my mind.

Finally, with great love that I would like to thank my beautiful wife Heather and my amazing son Gideon for their contributions to my effort, which has come in the form of incredible patience, forgiveness, tolerance, understanding, and love. Additionally, I would like to thank my parents, Sandy and Gary Cooper, my brother and sister, Seth and Sara, as well as my close friends, Jeff Dotson and Chad Geving, whose unwavering support in my academic efforts has given me significant motivation to keep going during those long days and nights of work.



Aron J. Cooper

Assignment

Draper Laboratory Report Number CSDL-T-1524

In consideration for the research opportunity and permission to prepare my thesis by and at The Charles Stark Draper Laboratory, Inc., I hereby assign my copyright of the thesis to The Charles Stark Draper Laboratory, Inc., Cambridge, Massachusetts.

Aron J. Cooper

Date

[This page intentionally left blank.]

Contents

1	Introduction	19
1.1	Why SLAM?	20
1.2	SLAM Applications	21
1.3	SLAM Methods	22
1.4	Need for Data Association Algorithms	23
1.5	EKF vs. Particle Filter Data Association	27
1.6	Data Association Methodologies	27
1.7	Thesis Statement	28
1.8	Thesis Overview	29
2	Simultaneous Localization and Mapping	31
2.1	Extended Kalman Filter SLAM	32
2.1.1	The EKF Formulation	33
2.1.2	Major Issues on the Use of EKFs for SLAM	37
2.2	Particle Filter SLAM	38
2.2.1	Particle Filter Advantages	39
2.2.2	Particle Filter Formulation	39
2.2.3	Rao-Blackwellized Particle Filter	41
3	Data Association	45
3.1	Individual Measurement Data Association	46
3.1.1	Maximum Likelihood	46
3.1.2	Individual Compatibility	48

3.1.3	Combined Individual Compatibility and Maximum Likelihood	48
3.2	Batch Data Association	49
3.2.1	Sequential Compatibility Nearest Neighbor	49
3.2.2	Joint Maximum Likelihood	52
3.2.3	Joint Compatibility	54
3.2.4	Multiple Hypothesis Data Association	61
3.2.5	Delayed Assignment Data Association Algorithms	63
4	Simulated Results for Filter-Data Association Marriages	65
4.1	Assumptions and Simulation Set-up	66
4.1.1	Numerical Values Used for Simulations	67
4.2	Landmark Pairs Outside of Trajectory Simulated Experiment	68
4.2.1	Results for Perfect Data Association	69
4.2.2	Description of Data Format Used for Filter-Data Association Algo- rithm Marriages	71
4.2.3	Results for Sequential Nearest Neighbor Data Association	72
4.2.4	Results for Joint Compatibility Data Association	77
4.2.5	Results for Maximum Likelihood Data Association	80
4.2.6	A Quick Performance Comparison of the Various Filter-Data Asso- ciation Marriages	84
4.3	Landmarks Inside of Trajectory	87
4.3.1	Results for Perfect Data Association	88
4.3.2	Description of Data Format Used for Filter-Data Association Algo- rithm Marriages	89
4.3.3	Results for Sequential Nearest Neighbor Data Association	90
4.3.4	Results for Joint Compatibility Data Association	93
4.3.5	Results for Maximum Likelihood Data Association	96
4.3.6	A Quick Performance Comparison of the Various Filter-Data Asso- ciation Marriages	98
4.4	Summary of Results	100

4.4.1	Interpretation of Data Association Results	101
4.4.2	Interpretation of Measurement Rejection Results	103
4.4.3	Filter Dependent Effects	104
5	Results from Application of Filter-Data Association Marriages to Ex-	
	perimental Data	105
5.1	Experimental Set-up	106
5.2	Limitations of FastSLAM 1.0 for the Car Park Data Set	106
5.3	Landmark Localization	108
5.4	Agent Localization	111
6	Conclusions and Future Work	115
6.1	Future Work	117

[This page intentionally left blank.]

List of Figures

1-1	Data Ambiguity as a Result of Measurement Uncertainty	24
1-2	Data Ambiguity as a Result of Pose Uncertainty	25
1-3	Data Ambiguity as a Result of Landmark Uncertainty	26
3-1	Joint Compatibility Branch and Bound Search for a two landmark - two measurement scenario.	57
3-2	Joint Compatibility as a Constraint Satisfaction Problem	59
4-1	Trajectory and Landmark Pairings Used in Outside of Trajectory Case for Simulated Results.	69
4-2	Performance of EKF and FastSLAM for Perfect Data Association Case with Landmarks Outside of Trajectory	70
4-3	Sequential Nearest Neighbor Data Association Performance Variation Versus Time and Landmark Separation for Landmarks Outside of Trajectory.	73
4-4	Sequential Nearest Neighbor Data Association Measurement Rejection Characteristics for EKF and FastSLAM where the Landmarks are Outside of the Trajectory.	74
4-5	Average Measurement Rejection and Data Association Errors for Sequential Compatibility Nearest Neighbor using EKF and FastSLAM where the Landmarks are Outside of the Trajectory	75
4-6	EKF and FastSLAM performance using Sequential Nearest Neighbor Data Association for the Case where Landmarks are Outside of the Trajectory.	76
4-7	Joint Compatibility Data Association Performance Variation Versus Time and Landmark Separation for Landmarks Exterior to Trajectory.	78

4-8	Joint Compatibility Data Association Measurement Rejection Characteristics for EKF and FastSLAM where the Landmarks are Exterior to the Trajectory.	79
4-9	Average Measurement Rejection and Data Association Errors for Joint Compatibility using EKF and FastSLAM where the Landmarks are Outside of the Trajectory	80
4-10	EKF and FastSLAM performance using Joint Compatibility Data Association where the Landmarks are Outside of the Trajectory.	81
4-11	Joint Maximum Likelihood Data Association Performance Variation Versus Time and Landmark Separation for Landmarks Exterior to the Trajectory.	82
4-12	Joint Maximum Likelihood Data Association Measurement Rejection Characteristics for EKF and FastSLAM where the Landmarks are Located Outside of the Trajectory.	84
4-13	Average Measurement Rejection and Data Association Errors for Joint ML where the Landmarks are Exterior to the Trajectory	85
4-14	EKF and FastSLAM performance using Joint Maximum Likelihood Data Association where the Landmarks are Exterior to the Trajectory.	86
4-15	Trajectory and Landmark Pairings Used in Inside of Trajectory Case for Simulated Results.	90
4-16	Performance of EKF and FastSLAM for Perfect Data Association Case with Landmarks Inside of Trajectory.	91
4-17	Average Measurement Rejection and Data Association Errors for Sequential Compatibility Nearest Neighbor using EKF and FastSLAM with Landmarks Inside of Trajectory	92
4-18	EKF and FastSLAM performance using Sequential Nearest Neighbor Data Association with Landmarks Inside of Trajectory.	93
4-19	Average Measurement Rejection and Data Association Errors for Joint Compatibility using EKF and FastSLAM with Landmarks Inside of Trajectory	94

4-20	EKF and FastSLAM performance using Joint Compatibility Data Association.	95
4-21	Average Measurement Rejection and Data Association Errors for Joint ML with Landmarks Inside of Trajectory	97
4-22	EKF and FastSLAM performance using Joint Maximum Likelihood Data Association with Landmarks Inside of Trajectory.	98
5-1	Path and Map Estimate Superimposed on GPS-Truth for Various SLAM Filter-Data Association Marriages	109
5-2	Agent Localization Errors with Respect to GPS Path for Various SLAM Filter-Data Association Marriages	112
5-3	Root-Squared Agent Position Error with Respect to GPS Path for Various SLAM Filter-Data Association Marriages	113

[This page intentionally left blank.]

List of Tables

4.1	Pose, Landmark, and Measurement Uncertainties Used in Simulations . . .	67
4.2	Quantities Used in the Motion Model for the Simulations	67
4.3	Performance Comparison of Filter-Data Association Marriages for Agent Position Error for the Case where the Landmarks are Outside of the Tra- jectory	87
4.4	Performance Comparison of Filter-Data Association Marriages for Land- mark Position Error for the Case where the Landmarks are Outside of the Trajectory	88
4.5	Performance Comparison of Filter-Data Association Marriages for Aver- age Data Association Errors Made for the Case where the Landmarks are Outside of the Trajectory	89
4.6	Performance Comparison of Filter-Data Association Marriages for Agent Position Error for the Case where the Landmarks are Inside of the Trajectory	99
4.7	Performance Comparison of Filter-Data Association Marriages for Land- mark Position Error for the Case where the Landmarks are Inside of the Trajectory	100
4.8	Performance Comparison of Filter-Data Association Marriages for Average Data Association Errors Made for the Case where the Landmarks are Inside of the Trajectory	101
5.1	Map Building Characteristics of the Various SLAM Filter-Data Association Combinations for the Car Park Data Set	110

[This page intentionally left blank.]

Chapter 1

Introduction

The problem of robotic mapping, building a globally consistent of a robot's spatial surroundings, becomes easier to solve if the true path of a robot is known [40, ?]. At the same time, if a perfectly accurate map of the robotic agent's surroundings is available then the problem of navigation, or localization, is easier to solve [8]. However, in most realistic applications involving mobile agents neither of these situations, the true path or map information, is available. Additionally, the environmental and motion measurements used by robots are noisy and this noise creates uncertainty in the system. The operation of these sensors correlates the uncertainty, and this correlation is what makes it possible to simultaneously solve for both robot position and map of its surroundings [40]. This approach has come to be known as SLAM or CML, Simultaneous Localization and Mapping [9, 13] or Concurrent Mapping and Localization [25, 41], respectively.

One of the key issues in the SLAM problem is the data association problem [40]: deciding which noisy measurement corresponds to which feature of the map. Noise and partial observability can make the relationship between measurements and the model highly ambiguous. This problem is further complicated by considering the possible existence of previously unknown features in the map and the possibility of spurious measurements. The problem of data association in SLAM has recently received a great deal of attention within the Artificial Intelligence and robotics literature [3, 4, 31, 32, 21, 27, 28, 34]. Prior to this, a substantial amount of work had been done in the tracking domain, as in radar or target tracking [6, 5].

The data association problem is considered by some to be one of the hardest, if not the hardest problem, in robotic mapping [40] and a problem for which there exists "no sound solution" [39]. In particular, of the numerous data association algorithms that have been proposed in the literature it is not clear how these methods work with the various SLAM algorithms. Much of the recent literature on SLAM has emphasized the study of two types of filtering techniques: Extended Kalman Filters and Particle Filters, and the bulk of the research that has come out on data association emphasizes the use of one or the other of these SLAM algorithms. Different data association algorithms may have different effects on these filters, and it is the aim of this thesis to study these effects.

1.1 Why SLAM?

In situations where a robot, or even a human with a suite of instruments, is moving around an environment with the goal of both localizing itself and building up location information about its surroundings, e.g. mapping or target localization, the need for solving the SLAM problem becomes obvious.

However, utilizing a solution to the full SLAM problem may not be obvious in situations in which an agent is concerned with only self-localization or target localization and map building: the question of whether or not solving the full SLAM problem is necessary when only a portion of the result is needed is of concern. In certain instances, such as those in which a reliable set of GPS signals are available, which are found in air-based systems or those that operate in open-field environments on the ground, solving the full SLAM problem would not be necessary to maintain tightly bounded localization since an external sensor eliminates the need for inference over robot pose. Unfortunately, there are numerous environments where good GPS signals are not available. For example, GPS signals are not available or dependable in urban canyon environments, which are found in most highly populated areas, as well as for most indoor environments.

Assume a situation in which an agent is operating in GPS-denied environments and its only concern is with self-localization. If it is not going to attempt to solve the full SLAM problem it is left with only motion measurements or an Inertial Navigation Sys-

tem (INS) such as accelerometers, gyroscopes, wheel encoders, doppler, or image based pseudo-inertial measurements. Even with the most accurate combinations of these devices, localization using on-line algorithms becomes difficult, as very small motion errors propagate themselves forward in time causing error growth. This problem along with the lack of state observability using only an INS will lead to a navigation filter that will eventually diverge; in other words the accuracy of the state estimate will quickly worsen [10, 40]. In certain situations many of these problems can be alleviated, or at least delayed, by enforcing certain constraints on portions of the navigation filter. This topic is also discussed in reference [10].

If instead the goal of an agent is to only build a map or localize a target using relative measurements, such as range and bearing, a map building solution may not even be possible without knowing the robot's location. In situations where good *a priori* information is known, in regards to the robot's location, a solution to the map building problem can be attained. However, as soon as the robot begins to move around, the uncertainty in its location will begin to grow without bound for the previously mentioned reasons. Solving the SLAM problem is the logical resolution to these issues.

Further motivation for solving the SLAM problem is that it lends itself to use of instruments that are lightweight and cheap, as opposed to attempting to solve the mapping and localization problems separately which may in fact require more costly, complex instrumentation (e.g. GPS). If implemented correctly a solution to the SLAM problem should be with a minimal amount of instrumentation. Additionally, solving the SLAM problem can be done in an autonomous fashion using online algorithms; therefore regardless of the user, be it robotic or human, it is anticipated that little or no user input will be required for use of an instrument suite that makes use of the correct SLAM algorithms.

1.2 SLAM Applications

There are numerous applications for which successful implementations of SLAM algorithms can be used to produce substantial benefits for both human and robotic agents. First, it is a prerequisite for successful operation of mobile robots in almost every realistic

situation [40], as most applications of mobile robots involve some form of a SLAM implementation. Some of the more interesting applications are found in undersea autonomous vehicles [33] and robotic exploration of mines [42], with additional possibilities of future application to extraterrestrial planet exploration [30]. For many of these applications the SLAM solutions are only used for navigation purposes, however some do make use of the created map information for future use [42].

In addition to robotics applications, SLAM solutions may be used to enhance the capabilities of humans in localization and mapping, which is potentially a very powerful concept. In particular, it is hypothesized that the correct implementations of SLAM solutions combined with the proper instrument suite should enable one to produce systems that could be used by humans, yet operate autonomously without human intervention. Such systems could be used in the place of those currently in operation that involve complex procedures to be useful and also require very heavy, expensive instrumentation. One example of this is found in the expensive and bulky equipment currently used by some U.S. military forces to localize targets at a distance using ground personnel.

These types of SLAM based systems could have numerous applications in the defense realm from forward operations involving the mapping of areas where future operation may occur, to tracking of future troop movements, and real-time mapping, as well as the localization of specific targets. The need filled by SLAM based systems would be for environments in which GPS signals either do not exist or are not dependable, such as the urban canyon or indoor environments.

Non-military related uses for these type of technologies may also exist in the areas of human extra-terrestrial planet exploration, as well as in terrestrial applications such as for use by sportsman or search and recovery teams where GPS signals are not useful or available.

1.3 SLAM Methods

There are many SLAM algorithms that exist in the literature at this point in time. Of these algorithms there are many different approaches for the various aspects of the prob-

lem. There are those that use a grid-based approach for representing the world such as DP-SLAM [14, 15], as well as methods that attempt to describe the geometry of objects in the map with various methodologies [44, 2]. Unfortunately, many of these methodologies have serious limitations. For example, the DP-SLAM algorithm is formulated for use with a specific type of measurement device. In the case of the methods that attempt to describe the complex geometries of the environment, there is a tendency to have difficulty computing a solution in real-time especially for higher-dimensional or geometrically complex situations, i.e. numerous landmarks of various shape.

In the context of this thesis the methods considered will be those that maintain point representations of the world in continuous space, can be implemented as an online algorithm, have incarnations that scale well to high-dimensional problems, and can be flexible for use with any known type of environmental location measurements. From these constraints the two SLAM solutions that will be considered are the Extended Kalman Filter and Particle filter, specifically Rao-Blackwellized Particle filters much like FastSLAM [30].

1.4 Need for Data Association Algorithms

Until recently much of the work on solving the SLAM problem neglected the data association issue by assuming that it was always known *a priori* for all measurements [40]. In any situation where the map of the world may contain more than one feature there will be some amount of data association ambiguity. In fact, even in the case where it is known that only one object exists in the world there will still be data association ambiguity if the possibility of spurious measurements is allowed. In most applications where SLAM solutions are used the number of dimensions required to describe the world is quite high, as the map must contain enough prominent features to sufficiently describe an agent's surroundings. At the very least a minimal two-dimensional representation of an agent's environment would have on the order of magnitude $O(10^1)$ objects for a very simple problem, as in the *Car Park* data set where the environment is made up of 15 unique features or landmarks, along with the additional object of the agent itself [1].

Beyond the ambiguity created by multiple assignment possibilities for individual mea-

measurements, or assignment sets for multiple measurements, is the ambiguity created by uncertainties in measurements, pose, and landmark locations. To effectively describe how each type of uncertainty leads to greater ambiguity in data associations, a number of figures will be examined. In these plots the following variables will be used: F_i denotes the current estimate of the location of feature i , Z is used to describe the measurement (or the location a measurement is indicating a feature to be), and S_t denotes the current estimate of the agent's pose.

Measurement uncertainty creates greater data association ambiguity since greater measurement uncertainty means a larger number of assignments must be considered. This occurs because more data associations will have non-negligible likelihoods, which is a metric used by many data association algorithms, see figure 1-1.

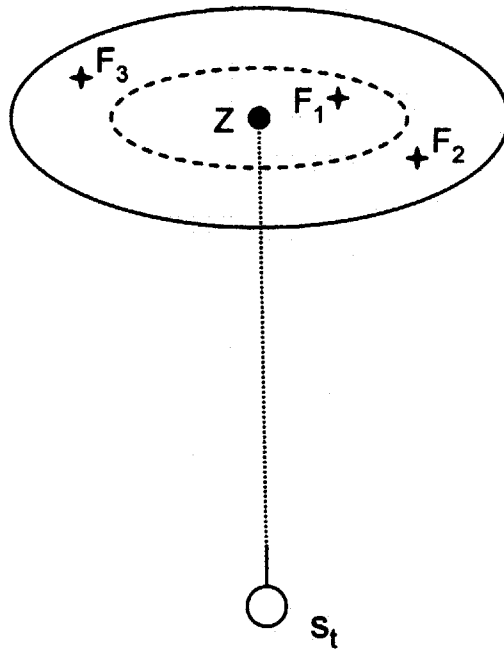


Figure 1-1: Data ambiguity as a result of measurement uncertainty. Here the inner dotted ellipse denotes a "small" measurement uncertainty while the outer solid ellipse denotes a "large" measurement uncertainty. Notice the increase number of features included within the measurement uncertainty with the larger value.

Pose uncertainty has a similar effect but in this case it is unclear to the agent which

object(s) is being observed as it is not entirely certain where it is or where it is facing, as illustrated in figure 1-2.

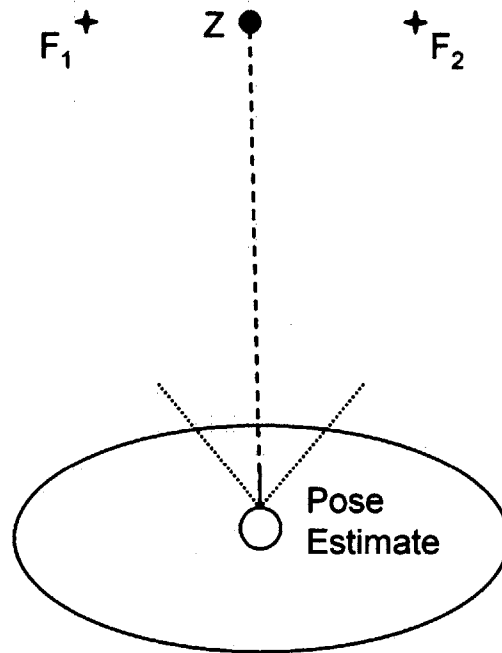


Figure 1-2: Data ambiguity as a result of pose uncertainty. Here the ellipse surrounding the "pose estimate" indicates the uncertainty in robot position, while the two dotted lines symmetric about the line pointing to the measurement indicate the heading uncertainty. This indicates that the pose and heading may be anywhere within this uncertainty, therefore allowing for confusion of which feature, F_1 or F_2 , the measurement should be associated with

Finally, landmark uncertainty leads to assignment ambiguity by making it possible, in a probabilistic sense, for multiple measurements to be associated with a single measurement. This can be seen by the situation depicted in figure 1-3.

In practice all of these errors are usually affecting the agent all at once, in a superposition of errors. Additionally, it is often the case that each measurement must also be considered to be one of a never-before-seen object or of a spurious nature that must be removed. This combination of uncertainties makes data association very difficult.

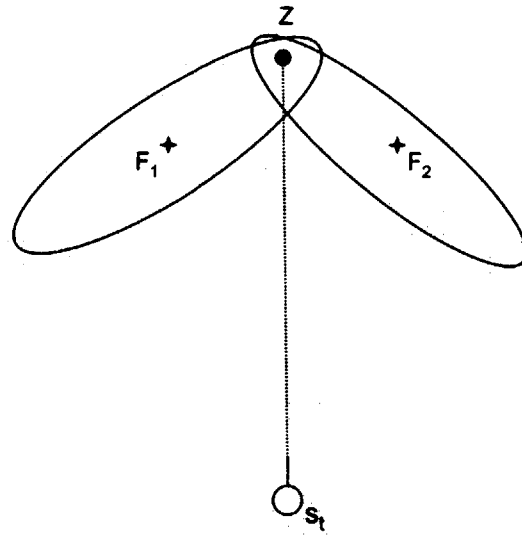


Figure 1-3: Data ambiguity as a result of landmark uncertainty. Here the uncertainty in the location of features F_1 and F_2 makes it equally likely that the measurement Z should be associated with either one of them.

1.5 EKF vs. Particle Filter Data Association

In this thesis, we will look at two particular inference algorithms, each with an inherently different treatment of the data association problem. In the EKF a single data association is used for each measurement in a very stiff manner, such that the assignment hypothesis made at the current time cannot be altered in the future, if future information indicates it was incorrect. This is a result of the EKF committing to an *a posteriori* distribution over robot poses and landmarks that is maintained as a single unimodal Gaussian.

In contrast, the Rao-Blackwellized particle filter represents the posterior over poses and landmarks non-parametrically, sampling particles from the posterior. The FastSLAM 1.0 particle filter allows for per-particle data association, thereby allowing for consideration of M independent data association hypotheses where M is the number of particles used in the filter. Since the uncertainty in the states is carried along in the dispersion of particles and their map knowledge as opposed to a covariance matrix, the effect of data association assignments differs substantially from the EKF. Finally, the operation of this particular method is such that with a non-zero probability the filter will eliminate those particles, or hypotheses, that have made incorrect data correspondences. This fact has the potential to have the greatest impact on the way ambiguous data associations will affect it.

1.6 Data Association Methodologies

A number of data association methodologies have been proposed and studied in the SLAM literature. It has generally been the case that these methods have only been applied to either EKF or FastSLAM without direct comparisons being made between the two. One such method that was recently proposed is known as *Joint Compatibility* (JC) and was applied to an EKF SLAM solution [31, 4]. In the original paper [31] that proposed this method it was compared to another data association algorithm *Sequential Compatibility Nearest Neighbor* (SCNN); however it too was only applied to an EKF implementation. The majority of the first papers to explore the FastSLAM algorithm used a maximum likelihood approach for data association [30, 28, 45], while one paper also explored a

particle splitting method of data association for FastSLAM [34]. However, none of these papers mentioned above explore the performance of both EKF and FastSLAM using the same data association methods.

Another data association method for SLAM that has been proposed in the literature, is *Combined Constraint Data Association (CCDA)*, which is formulated and explored in [4]. However, according to [30] this methodology should perform similarly to JC with the additional property that it can be used to determine pose information if none is available *a priori*.

One additional method that is mentioned in [30] is the concept of randomly sampling data associations on a per-particle basis. As mentioned in this reference the sampling would be done from a Probability Mass Function, or PMF, as defined by the normalized likelihoods of measurement associations. It would of course be possible to sample from some other PMF if so desired. While this method makes some intuitive and even mathematical sense for the FastSLAM formulation, it is unclear if this method would be a useful approach for the EKF.

There are other data associations that have recently been mentioned in the literature that involve the use of the Hough Transform and the RANSAC algorithm. These methods, as discussed in the available literature, perform data association in a delayed manner. In doing this there are some possible gains in data association performance, however the end result is that measurements will not be processed in real-time.

1.7 Thesis Statement

It is the intention of this document to advance the following thesis:

”The performance of a given data association algorithm for SLAM, in terms of the propensity of data association errors made and the subsequent effect on localization and mapping accuracy, depends on whether it is used in conjunction with an Extended Kalman Filter or a Rao-Blackwellized Particle Filter formulation as the solution to the SLAM problem.”

1.8 Thesis Overview

The research done in this thesis can be described by the following items:

- Three batch data association algorithms will be examined and described in detail: Sequential Nearest Neighbor, Joint Compatibility, and Joint Maximum Likelihood
- The operation of the three data association algorithms will be studied with regards to their operation when married to both EKF and FastSLAM filter implementations.
- The performance of the SLAM filter-data association combinations will be examined using both simulated and real world data.

[This page intentionally left blank.]

Chapter 2

Simultaneous Localization and Mapping

The Simultaneous Localization and Mapping, or SLAM, problem first arose when it was recognized that problems of robotic mapping and navigation are very difficult, and in some situations impossible, to solve as separate problems. The methods that have been the most successful in solving the SLAM problem have been ones that make use of statistical methods for taking into account the uncertainty in the system states as well as measurement uncertainty. Such a method was first proposed and popularized in a series of seminal papers by Smith, Self, and Cheeseman [37, 38] in which the solution took the form of the Extended Kalman Filter.

Solutions to the SLAM problem can be very effective in situations when global measurements such as GPS are not available, for instance in the case when a robot or agent that operates indoors or in the urban canyon. In these situations, the uncertainty in both the location of objects in the environment as well as the robots own pose become tightly linked. This is due to the fact that in these situations the agent must rely on relative measurements relating its current pose with objects in its surroundings, such as range and bearing measurements. This is what leads to the concept of solving the Simultaneous Localization and Mapping or SLAM problem, which has also been called Concurrent Mapping and Localization, or CML, in the literature [32].

2.1 Extended Kalman Filter SLAM

Framing the SLAM problem as a partially observable Markov chain allows for a compact, probabilistic notation of the SLAM problem which is known as the Bayes Filter Equation and is derived from the SLAM posterior [30]. For the Bayes Filter equation and the SLAM posterior to be true statements, the Markov assumption must hold. This assumption states that if the current state is known then the previous and future data are conditionally independent [11]. By making some very restrictive assumptions, the Bayes Filter equation will be equivalent to a Kalman Filter [39, 30].

In order to describe this problem a number of variables must be defined. First, the states of the system are represented by s_t the robot pose at time, t , and Θ the location of the objects in the environment that are known to the robot, which is also referred to as the map. The robot pose consists of N_{ag} the agent North position, E_{ag} the agent East position, and ϕ_{ag} the agent heading. The map information, Θ , as it exists within the state vector consists of the North and East positions of all landmarks in the environment, N_{l_N} and E_{l_N} respectively for the N -th landmark. The environmental measurements as considered across time are given by $z^t = \{z_0, z_1, \dots, z_t\}$ and the motion measurements or controls as taken over time are given by $w^t = \{w_0, w_1, \dots, w_t\}$. The quantity η is a normalization constant that is a byproduct of applying the Bayes rule.

What we wish to maintain at each time, t , is a distribution over the current robot pose and landmarks. We call this the SLAM Posterior:

$$p(s_t, \Theta | z^t, w^t) \tag{2.1}$$

Using Bayes Rule this can be re-written as:

$$p(s_t, \Theta | z^t, w^t) = \eta \cdot p(z_t | s_t, \Theta, z^{t-1}, w^t) \cdot p(s_t, \Theta | z^{t-1}, w^t) \tag{2.2}$$

After some manipulation this can be converted into the Bayes Filter Equation:

$$p(s_t | z^t, w^t) = \eta \cdot p(z_t | s_t, \Theta) \cdot \int p(s_t, \Theta | s_{t-1}, z^{t-1}, w^{t-1}) \cdot p(s_{t-1} | z^{t-1}, w^t) ds_{t-1} \tag{2.3}$$

This is a tractable formulation because it allows us to compute $p(s_t, \Theta | z^t, w^t)$ recursively, from the known sensor model $p(z_t | s_t, \theta)$, the known motion model $p(s_t | s_{t-1}, w_t)$, and the SLAM posterior of the previous time step $p(s_{t-1}, \theta | z^{t-1}, w^{t-1})$. In the filter formulations discussed here, the landmarks are thought of as point objects with single locations in two-dimensional space, such as the corner of a building or the center of a small tree. While this is an approximation, it is sufficient for the purposes of the algorithm development discussed here. Additionally, the error in this approximation can, for the most part, be absorbed in measurement error and uncertainty in feature locations.

2.1.1 The EKF Formulation

The SLAM problem can easily be framed as a state space system and therefore fits into the Kalman filter framework. In real-world applications there are often non-linearities involved in the physics of the situation, which is why the EKF is used. This formulation is essentially the same as the traditional Kalman filter architecture with linearization performed to allow for the real world physics of the system dynamics and the measurement models.

The solution offered by the EKF is attained by assuming a multi-variate Gaussian distribution used to describe $p(\mathbf{x})$. In addition the noise in the IMU and environmental measurements must be assumed to be white. Finally, it also assumes that substantial errors are not incurred by making use of linear approximations of the dynamics of the state propagation in time, as well as in the physics of the state relationships for environmental measurements.

The building blocks of the EKF are the state vector and its associated covariance matrix.

$$\mathbf{x} = \begin{bmatrix} s_t \\ \Theta \end{bmatrix} \quad (2.4)$$

Here s_t and Θ are defined as:

$$\mathbf{s}_t = \begin{bmatrix} N_{ag} \\ E_{ag} \\ \phi_{ag} \end{bmatrix} \quad (2.5)$$

$$(2.6)$$

$$\Theta = \begin{bmatrix} N_{l_1} \\ E_{l_1} \\ \vdots \\ N_{l_N} \\ E_{l_N} \end{bmatrix} \quad (2.7)$$

Then the covariance matrix is defined as:

$$\mathbf{P} = \begin{bmatrix} \sigma_{N_{ag}}^2 & \sigma_{N_{ag}E_{ag}} & \cdots & \cdots & \cdots & \cdots & \sigma_{N_{ag}E_{l_N}} \\ \sigma_{N_{ag}E_{ag}} & \sigma_{E_{ag}}^2 & \ddots & \cdots & \cdots & \cdots & \vdots \\ \vdots & \ddots & \sigma_{\phi_{ag}}^2 & \ddots & \cdots & \cdots & \vdots \\ \vdots & \cdots & \ddots & \sigma_{N_{l_1}}^2 & \ddots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \ddots & \sigma_{E_{l_1}}^2 & \ddots & \vdots \\ \vdots & \cdots & \cdots & \cdots & \ddots & \ddots & \vdots \\ \sigma_{N_{ag}E_{l_N}} & \cdots & \cdots & \cdots & \cdots & \cdots & \sigma_{E_{l_N}}^2 \end{bmatrix} \quad (2.8)$$

In general this matrix is initialized as a diagonal matrix, where the variances which lie on the diagonal represent the *a priori* knowledge with regard to the uncertainty in the individual states, which consist of features and the robot pose. The off-diagonal terms are built up over time by the filter and represent the correlations between these filter states.

Finally, knowledge of the measurements being considered by the filter must be understood. This includes both a model of the physical measurements, as well as the noise characteristics of these measurements. To be consistent with what is used throughout the analysis done here, range and bearing measurement pairs will be considered as a simultaneous measurement of a single landmark. This is a reasonable for a number of instruments that are used in SLAM implementations, such as the SICK laser.

The physical measurements for the range and bearing can be represented as:

$$\mathbf{h}_k(\mathbf{x}) = \begin{bmatrix} \sqrt{(N_{l_j}(k) - N_{ag}(k))^2 + (E_{l_j}(k) - E_{ag}(k))^2} \\ \arctan\left(\frac{E_{l_j}(k) - E_{ag}(k)}{N_{l_j}(k) - N_{ag}(k)}\right) - \phi_{ag} \end{bmatrix} \quad (2.9)$$

However, the actual measurement is noisy and is modeled in the following way:

$$z_k(x) = h_k(x) + u_k \quad (2.10)$$

Where the noise is represented by the vector of random variables u_k , which is zero-mean gaussian white noise.

$$u_k \sim N(0, R_k) \quad (2.11)$$

The knowledge of the measurement uncertainty is represented by the measurement covariance matrix, \mathbf{R}_k .

$$\mathbf{R} = \begin{bmatrix} \sigma_r^2 & \sigma_{rb} \\ \sigma_{rb} & \sigma_b^2 \end{bmatrix} \quad (2.12)$$

The motion model used is very simple: it is assumed that the robot issues controls of translational and rotational velocity in two-dimensions.

$$v_{k-1} = v_{true} + w_{v,k-1} \quad (2.13)$$

$$\Omega_{k-1} = \Omega_{true} + w_{\Omega,k-1} \quad (2.14)$$

Here the additive noise term w_{k-1} is defined as:

$$\mathbf{w}_{k-1} = \begin{bmatrix} w_{v,k-1} \\ w_{\Omega,k-1} \end{bmatrix} \sim N(0, Q) \quad (2.15)$$

$$Q = \begin{bmatrix} \sigma_v^2 & 0 \\ 0 & \sigma_{\Omega}^2 \end{bmatrix} \quad (2.16)$$

Now the operation of the EKF formulation can be laid forth, where z_k represents the actual measurement received and \hat{z}_k represent the estimated measurement as determined by the current values of the state estimate \hat{x}_k [18, 22].

Propagation Step:

1. Propagate the state estimate in time:

$$\hat{x}_k(-) = g(\hat{x}_{k-1}(+), v_{k-1}, \Omega_{k-1}) \quad (2.17)$$

Where $k - 1$ indicates the most current information available for that variable.

2. Propagate the state covariance matrix: ¹

$$P_k(-) = \Phi_{k-1} P_{k-1}(+) \Phi_{k-1}^T + G_{k-1} Q G_{k-1}^T \quad (2.18)$$

$$G_{k-1} = \begin{bmatrix} \cos(\hat{\phi}_{ag,k-1}) & 0 \\ \sin(\hat{\phi}_{ag,k-1}) & 0 \\ 0 & 1 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \end{bmatrix} \cdot dt \quad (2.19)$$

Here, the quantity dt is the time step over which state propagation is occurring.

The state transition matrix, Φ_{k-1} , is defined for this case as

¹This equation is the result of the following derivation, where the $(-)$ terms have been dropped for simplicity. Here, x_k is the truth state and \hat{x}_k is the filter's estimate of the state: $\tilde{x}_k = x_k - \hat{x}_k =$ error in state estimate, $P_k = E[\tilde{x}_k \tilde{x}_k^T]$, and $\tilde{x}_k \tilde{x}_k^T = (\Phi_{k-1} \tilde{x}_{k-1} - G_{k-1} w_{k-1})(\Phi_{k-1} \tilde{x}_{k-1} - G_{k-1} w_{k-1})^T$. After eliminating terms that go to zero, this can be rewritten as: $\tilde{x}_k \tilde{x}_k^T = \Phi_{k-1} \tilde{x}_{k-1} \tilde{x}_{k-1}^T \Phi_{k-1}^T + G_{k-1} w_{k-1} w_{k-1}^T G_{k-1}^T$, the expectation of this result can then be taken to obtain the propagation equation for P_k presented here.

$$\Phi_{k-1} = \Phi_k = \Phi = I_{(2N+3) \times (2N+3)} \quad (2.20)$$

an identity matrix of size $(2N + 3) \times (2N + 3)$ since there are no states carried along that explicitly describe the dynamics of the world, i.e. velocity or acceleration states.

Measurement Update Step:

1. Covariance matrix update:

$$P_k(+)= [I - K_k H_k] P_k(-) \quad (2.21)$$

$$K_k = P_k(-) H_k^T [H_k P_k(-) H_k^T + R_k]^{-1} \quad (2.22)$$

The matrix H_k , known as the measurement matrix, is the Jacobian of the measurement model $h_k(x)$.

2. State estimate update:

$$\hat{x}_k(+)= \hat{x}_k(-) + K_k [z_k - \hat{z}_k] \quad (2.23)$$

$$H_k = \nabla h_k(x) = \left. \frac{\partial h_k(x)}{\partial x} \right|_{\hat{x}_k} \quad (2.24)$$

2.1.2 Major Issues on the Use of EKF's for SLAM

One of the major road blocks that has faced the EKF method in its application to sophisticated problems has been the growth in complexity of the filter with the number of objects in its map. This is primarily due to the fact that this method, in its original formulation, relies on a covariance matrix of size $O(N^2)$, where N is the number of objects in the map.

There has been significant effort within the AI community to address this problem during recent years. One method which has been shown to resolve the complexity issues surrounding the use of EKF formulations and is considered the current state-of-the-art by many, is known as *Atlas* [7]. This method promises constant time, or at least bounded

time, operation and is therefore not dependent on the size of the map. However, it should be noted that the use of the *Atlas* methodology is not limited to EKF solutions of the SLAM problem, because in its general form it is a framework for using filtering methods that aims to reduce complexity, rather a filtering method itself. There are other EKF related methods which also promise either constant time or improved time operation and have been shown to produce good results, such as Sparse Extended Information Filters (SEIFs) [43, 27] and compressed EKFs [19, 20].

While these approaches do address computational limitations of the EKF formulation they do not, in their basic formulations, fully address the data association problem. The implementation of these methods and their robustness to data association errors will not be discussed in this thesis since there is no evidence in the literature, as well as no logical reason to believe otherwise, that these methods would be more robust to data association errors than the traditional EKF-based SLAM approach.

2.2 Particle Filter SLAM

Particle filters consist of a large class of Monte Carlo estimation methods that are applicable to problems that can be posed as partially observable Markov chains [39]. The formulation of the particle filter has been around since the mid-to-late 1990's, originally proposed in papers by Kitagawa as well as Liu and Chen [24, 26]. Indeed the applications of particle filters to robotics problems are very widespread, but some of the greatest achievements made by their use comes in the area of localization and mapping. The particle filter is credited with having solved the global localization and the kidnapped robot problem, both of which were previously unsolved and considered to be important for robust mobile robot operation [39].

One of the greatest drawbacks of the particle filter, as it was originally formulated, is that it does not scale well to high dimension state spaces. This is a result of the exponential time behavior of this implementation of the particle filter, which is not acceptable for problems such as SLAM which must be solved in real-time while maintaining numerous states. A development in the literature that has offered some resolution to this problem

is the Rao-Blackwellized particle filter [12]. This method has not only showed promise to greatly increase computational efficiency, but also improve estimation accuracy. The concept of a Rao-Blackwellized particle filter for application to the SLAM problem has been greatly developed in a number of papers by Montemerlo and Thrun [28, 29, 45, 30].

2.2.1 Particle Filter Advantages

Particle filters have a number of useful properties when compared to other methodologies for solving the SLAM problem, such as the EKF. First, the particle filter can approximate arbitrarily complex probability distributions, where as the EKF is restricted to Gaussian descriptions at all levels of uncertainty. Additionally, particle filters are not adversely affected by significant non-linearities in the motion and measurement models. This is because linearization is not required in the propagation of the state uncertainty, as this information is carried along in the distribution of the particles.

2.2.2 Particle Filter Formulation

For application of particle filters to the SLAM problem it is possible to begin at the same place as was done for the EKF formulation, that is the SLAM posterior.

$$p(s_t, \Theta | z^t, w^t) \tag{2.25}$$

As in the Kalman Filter formulation, this probabilistic relationship is valid if the Markov assumption holds. The Markov assumption states that if the current state is known then the previous and future data are conditionally independent [11].

As was done previously for the EKF the SLAM posterior can be converted into the Bayes Filter equation by use of the Bayes Rule.

$$p(s_t | z^t, w^t) = \eta \cdot p(z_t | s_t, \Theta) \cdot \int p(s_t, \Theta | s_{t-1}, z^{t-1}, w^{t-1}) \cdot p(s_{t-1} | z^{t-1}, w^t) ds_{t-1} \tag{2.26}$$

This equation can be solved in an approximate manner using a particle filter by ap-

proximating the continuous probability densities defined in equation (2.26) with discrete samples. The SLAM posterior may be thought of as a belief state, where a single *belief*(i) is defined as a hypothesis of agent pose, the map of the environment, and an associated weighting that defines the probability of the given belief being correct. In the particle filter this is represented by M samples of a continuous probability distribution along with the associated weighting for each sample.

$$belief(i) = p(s_t(i), \Theta(i) | z^t(i), w^t(i)) = \{s_t(i), \Theta(i), p(i)\}_{i=1, \dots, m} \quad (2.27)$$

At initialization this belief state, *belief*(i), is defined by whatever probability distribution is known to define the uncertainty in s_t and Θ [11, 39].

$$p(s_0, \Theta | z^0, w^0) = p(s_0, \Theta) \quad (2.28)$$

First, if *a priori* information describing the uncertainty in the states is available it is used to define the initial distributions. These distributions are then sampled from, to create the particle representation of the state space $x_0(i), p_0(i)$, where $p_0(i)$ is the particle weighting. The particle weighting is generally initialized to $\frac{1}{m}$ [11].

The estimation of the posterior can then be done in the following recursive fashion:

1. Propagation Step:

Obtain the new pose $s_t(i)$ using $p(s_t | s_{t-1}, w_{t-1})$ where this is equivalent to propagating each s_{t-1} using independent samples of w_{t-1} . This is defined by a motion model.

$$s_t(i) = g(s_{t-1}(i), w_{t-1}(i)) \quad (2.29)$$

Here, the variable w_{t-1} represents the noisy motion information of the agent as in equations 2.13 and 2.14, however in this case the noise can be represented by any probability distribution that can be sampled from. This process approximates the following predictive density:

$$p(s_t|s_{t-1}, w_{t-1})p(s_{t-1}, \Theta|z^{t-1}, w^{t-1}) \quad (2.30)$$

2. Measurement Step:

Now alter the belief state by weighting the particles using the likelihood of the particle occurrence given the measurement, z_t .

$$p(z_t|s_t(i), \Theta_i) \quad (2.31)$$

3. Re-sample Step:

Re-sample the M particles with replacement based on the normalized weighting, where

$$\sum_{i=1}^M p(i) = 1 \quad (2.32)$$

Loop back to step 1.

This procedure approximates the creation of the following posterior probability distribution [11].

$$\frac{p(z_t|s_t, \Theta)p(s_t|w_{t-1}, s_{t-1})p(s_{t-1}, \Theta|z^{t-1}, w^{t-1})}{p(z_t|z^{t-1}, w_{t-1})} \quad (2.33)$$

Note, that in actuality the distribution that would be ideal to sample from is the desired posterior, $p(s^t, \Theta|z^t, u^t)$, however this target function is unavailable. Using this ratio to determine the particle weighting along with the sampling procedure described in the section above, allows for the creation of an approximation to the desired posterior. This procedure is an example of a sampling importance re-sampling (SIR) algorithm [35].

2.2.3 Rao-Blackwellized Particle Filter

The Rao-Blackwellized particle filter is a variation of the traditional particle filter that scales well to problems of higher dimension. The Rao-Blackwellized particle filter is very general in its formulation and can be applied to problems other than SLAM [17, 23]. However, the formulation discussed here will only explore this type of filter as it applies

to the SLAM problem, in particular the incarnation known as FastSLAM 1.0 will be discussed [45]. In the FastSLAM formulation, a slightly different form of the SLAM posterior is used then what was presented in sections 2.2.2 and 2.1.1:

$$p(s^t, \Theta | z^t, w^t, n^t) \quad (2.34)$$

The differences exist in the variable $s^t = \{s_0, s_1, \dots, s_t\}$ which represents the entire robot path (not just the current pose) and the variable $n^t = \{n_0, n_1, \dots, n_t\}$ which represent the correct data associations over all time. The reason for maintaining the entire path, s^t , in the equation will be explained below. However, the presence of the data association assignment vector in the formulation does not seem to be dealt with in a satisfying manner in this formulation.

The formulation invokes the Rao-Blackwellization concept of marginalizing out variables from the posterior equation by implementing the proper conditioning. This was first developed in [12] to factor the SLAM posterior into the following:

$$p(s^t, \Theta | z^t, w^t, n^t) = p(s^t | z^t, w^t, n^t) \prod_{i=1}^N p(\Theta_i | s^t, z^t, u^t, n^t) \quad (2.35)$$

In words this means that it is possible to factor the problem into $N+1$ estimators. The first estimator represented by $p(s^t | z^t, w^t, n^t)$ aims to determine the posterior of path and the other estimators $p(\Theta_i | s^t, z^t, u^t, n^t)$ are used to determine the location of the N landmarks.

The path posterior $p(s^t | z^t, w^t, n^t)$ is estimated using a particle filter. The landmark estimators are obtained using EKFs, where each particle of the FastSLAM filter maintains N independent Kalman filters for estimating the N landmarks. The independence of the EKFs are a result of conditioning each landmark estimator on the robot path. This has the benefit of only having to maintain $N \ 2 \times 2$ covariance matrices for each particle as opposed to a full $(2N + 3) \times (2N + 3)$ covariance matrix.

The result is a set of particles where the i -th particle is defined as:

$$S_t(i) = [s^t(i), \mu_{1,t}(i), P_{1,t}(i), \dots, \mu_{N,t}(i), P_{N,t}(i)] \quad (2.36)$$

Here each particle contains a path along with the mean, $\mu_{m,t}(i)$, and covariance matrix, $P_{m,t}(i)$, for each of the N landmarks the particle is trying to estimate the location of, using an EKF.

For initialization the pose portion of the particles is found in a manner similar to that of the traditional particle filter.

$$p(s^0|z^0, w^0, n^0) = p(s^0) \quad (2.37)$$

Therefore the initial particles are sampled from the pose probability distribution defined by the available *a priori* information.

Using this framework the operation of the FastSLAM filter is now very much the same as that laid out in the section on particle filters 2.2.2. Here the aforementioned weight is calculated as follows [45].

$$p_t(i) = \frac{\text{targetdistribution}}{\text{proposaldistribution}} = \frac{p(s^t(i)|z^t, w^t, n^t)}{p(s^t(i)|z^{t-1}, w^t, n^{t-1})} \quad (2.38)$$

In actuality the distribution that would be ideal to sample from is the desired posterior, $p(s^t|z^t, w^t, n^t)$, however this target function is unavailable. Using this ratio to determine the particle weighting along with the sampling procedure described in the section on particle filters, allows for the creation of an approximation to the desired posterior. This procedure is an example of a sampling importance re-sampling (SIR) algorithm [35].

Using the Bayes Rule and a Markov assumption the weighting function can be written as [30].

$$p_t(i) = p(z_t|s^t(i), z^{t-1}, w^t, n^t) \quad (2.39)$$

This form allows for advantage to be taken of the EKF form of the landmark estimator. Now the weight per-particle can be written as the likelihood function for the measurement.

$$p_t(i) = \frac{1}{(2\pi)^{n/2} \sqrt{|Z_{n^t,t}|}} \exp\left(-\frac{1}{2}(z_t - \hat{z}_{n^t,t})^T Z_{n^t,t}^{-1} (z_t - \hat{z}_{n^t,t})\right) \quad (2.40)$$

[This page intentionally left blank.]

Chapter 3

Data Association

The problem of data association refers to the concept of relating the states of nature from which a measurement or set of measurements originate. There are many applications when data associations are known *a priori* or the problem of choosing the correct ones becomes trivial. However, for the problem of autonomous agents attempting to solve the SLAM problem it is a key component. In fact it is one of the most difficult aspects of developing a full solution to the SLAM problem [40, 4].

The study of data association has historically received the most attention within the literature on target tracking [5, 6]. However, with increasing attention in recent years being given to obtaining full solutions to the SLAM problem for robotic agents there has begun to be increasing attention to the data association problem within the AI literature [31, 3, 32, 21, 27, 4].

Much of the work that has come out of the AI/SLAM literature has sought to apply the ideas previously laid out in the tracking literature to the SLAM problem. At the same time some approaches to data association that have come out more recently appear to be unique to the SLAM domain such as *maximum common subgraph* (MCS), *combined constraint data association* (CCDA), and to a lesser extent *joint compatibility branch and bound* (JCBB) [4, 3, 31]. None-the-less many of these approaches have been demonstrated to perform successfully under specific conditions and implemented with particular data association-SLAM filter pairings.

One of the pieces that appears to be missing from the literature is a comprehensive

comparison of the performance of the most prevalent of these algorithms. In particular, how these algorithms perform when paired with the most popular filters for solving the SLAM problem and if the various data association picking algorithms will perform the same independently of the filter they are paired with.

3.1 Individual Measurement Data Association

The simplest form of the data association is that of having to associate a single measurement with the appropriate feature in the environment, where this could also include the possibility of either a spurious measurement or a previously unknown feature. This form of the problem, as opposed to a set of batch measurements that can be considered as being received simultaneously, appears to be very common in the tracking literature, from which many of the data association ideas originate.

3.1.1 Maximum Likelihood

One of the most basic and simplest methods for performing data association is to consider the measurement likelihoods. This is done by calculating the likelihood that each landmark known to the agent, i.e., that exists in the SLAM filter, is associated with the individual measurement being considered. This method is also known as Nearest-Neighbor data association in the literature [5].

In general the likelihood or Nearest-Neighbor calculation can be made for any probability distribution, as long as the probability density can be calculated for any possible measurement. For every case considered in this thesis the following assumptions will be made, thus allowing for a straightforward analytic representation of the likelihood.

- The true measurements at the present time are normally distributed, i.e. the measurement noise is gaussian.
- *A priori* knowledge of the characteristics of the measurement noise is available in terms of standard deviation $\sigma_{measure}$ and the mean.

Now the data association is chosen by determining the maximum likelihood.

$$e_l = \arg \max_j (f_{ij}) \quad (3.1)$$

$$f_{ij} = \frac{1}{(2\pi)^{n/2} \sqrt{|S_j|}} \exp\left(-\frac{1}{2} \nu_{ij}^T S_j^{-1} \nu_{ij}\right) \quad (3.2)$$

The variable $e_l = \{i, j\}$ is defined as the l -th data association for the measurement, i , and the feature j . The variable ν_{ij} is the innovation for the $\{i, j\}$ pair, S_j is the innovation covariance matrix, and n is the dimension of the innovation vector which is defined as the difference between the true measurement and the estimated measurement give the data association $e_l = \{i, j\}$.

This can also be chosen based on the maximum log-likelihood or alternatively the minimum normalized log-likelihood N_k , also known as the normalized distance [4].

$$\ln(f_{ij}) = -\frac{1}{2} \nu_{ij}^T S_j^{-1} \nu_{ij} - \ln\left(\frac{1}{(2\pi)^{n/2} \sqrt{|S_j|}}\right) \quad (3.3)$$

$$e_l = \arg \max_j (\ln(f_{ij})) \quad (3.4)$$

$$N_{ij} = \nu_{ij}^T S_j^{-1} \nu_{ij} + \ln|S_j| \quad (3.5)$$

$$e_l = \arg \min_j (N_{ij}) \quad (3.6)$$

Maximum likelihood can easily be made to include the ability to either reject spurious measurements or allow for the possibility of taking measurements of a previously unknown landmark. If only one or the other of these is going to be allowed as a possibility a single threshold can be used as either the likelihood of a spurious measurement or the measurement of a previously unknown landmark.

This can be done somewhat ad hoc by looking at the statistics of the measurement errors and using this information to calculate reasonable thresholds of what would constitute a measurement that is not consistent with the current knowledge of the world. Alternatively, a more rigorous method can be used whereby one uses the fact that the exponent of the likelihood calculation is χ^2 distributed. More about this will be

discussed in the next section involving individual compatibility.

3.1.2 Individual Compatibility

The method of individual compatibility by itself is not a method by which data associations can be strictly chosen. Instead those data associations that are statistically unlikely, or incompatible, can be eliminated from consideration. In [4] this type of method is referred to as "ambiguity reduction" as opposed to "ambiguity management" which refers to those methods such as *maximum likelihood* which explicitly make data association choices.

This can be accomplished by noting that the exponent of the likelihood function describing the distribution of innovations is in fact χ^2 distributed.

$$M_{ij} = \nu_{ij}^T S_j^{-1} \nu_{ij} \quad (3.7)$$

Where M_{ij} is χ^2 distributed, and is also referred to as *Mahalanobis distance* and *normalised innovation square* (NIS) in the literature [5, 4].

The individual compatibility test is then given by the following inequality:

$$M_{ij} < \gamma_n \quad (3.8)$$

The quantity γ_n is determined from two parameters, n , the dimension of the innovation vector and a free parameter which defines the expected percentage of correct associations that will be expected. In this study that parameter is always taken to be 95% or 0.95, which is commonly used in the literature [5, 4].

3.1.3 Combined Individual Compatibility and Maximum Likelihood

One possible approach to data association for the single measurement case is to use both individual compatibility and maximum likelihood together. This can be done by first processing the measurement via individual compatibility and therefore eliminating those

associations which fail individual compatibility. Then, within the set of remaining data associations, choose the one which has the maximum likelihood. The only difference in results between this method and using maximum likelihood by itself is that it allows for elimination of spurious measurements. In other words there will be cases where this method will choose to not make any data associations, thereby ignoring the measurement and treating it as faulty.

3.2 Batch Data Association

The problem of batch data association is very similar to that of the single measurement case and many of the algorithms developed for that case can be applied to this problem. Batch data association refers to the situation where the measurement devices being used, or the dynamics of the situation, are such that it is possible to treat a set of measurements as a batch process that can be processed simultaneously. This situation can arise when the device being used takes measurements of a portion of the environment in snapshot form, such as a digital camera or some laser range finders. An equivalent case can occur when the scan frequency of a measurement device is much faster than the dynamics of the vehicle, so the motion between the first measurement of a set and the last is negligible.

Posing the data association problem as a batch process has some nice properties which, if used correctly, can help to better define the proper data associations. The first is greedy mutual exclusion, which states that within a given batch of measurements no two measurements can be associated with a single landmark. If this is a good assumption, it alone will allow for the pruning of a large subset of all measurement-landmark pairing possibilities. The second property that this concept allows for is the possibility for comparing the all pairings within the entire data association set.

3.2.1 Sequential Compatibility Nearest Neighbor

One of the most basic algorithms that can be implemented to resolve data association ambiguities in the case where a batch of measurements can be considered as

simultaneous or near simultaneous is Sequential Compatibility Nearest Neighbor (SCNN) [31]. This method is nearly equivalent to the methods described in section 3.1.2, however the greedy mutual exclusion property for the batch data association requires a change in implementation.

In this case the algorithm must be varied to exclude the possibility of a particular landmark from being associated with more than one measurement within a single batch. This can be implemented in a number of ways and here are two possible implementations of this.

Implementation 1: Process measurements as a batch.

1. Remove as possibilities all associations which fail the individual compatibility test.
2. Calculate the likelihoods for all remaining data associations $e_l = \{i, j\}$.
3. Choose the data association that maximizes the likelihood function, as in equation 3.1 except consider all measurements i .
4. Eliminate from future consideration within the current measurement batch all hypotheses e_l that include either the measurement, i , or the feature, j .
5. Repeat process starting with step 2 until all of the measurements have been assigned data associations, or in other words until we have achieved a maximal hypothesis set, $E_m = \{e_1, e_2, \dots, e_m\}$.

Implementation 2: Process measurements sequentially but independently while preserving mutual exclusion rule.

1. Randomly pick one of the measurements, i_{fixed} , that has not been processed from the batch.
2. Eliminate from consideration those features, j , that do not satisfy the individual compatibility requirement with the chosen measurement i_{fixed} .
3. Calculate the likelihoods over all remaining associations, $e_l = \{i_{fixed}, j\}$, for the fixed i_{fixed} .

4. Choose the data association, $e_l = \{i_{fixed}, j\}$ from this set with the maximum likelihood.
5. Eliminate from future consideration within the current measurement batch all hypotheses e_l that include either the measurement, i_{fixed} , or the feature, j .
6. Repeat process starting with step 1 until all of the measurements have been assigned data associations, or in other words until we have achieved a maximal hypothesis set, $E_m = \{e_1, e_2, \dots, e_m\}$.

A benefit of using either of these implementations is that measurement updates can be made after each data association choice is made for individual measurements.

The first implementation may be beneficial because it will sequentially pick the global maximum that is still available via mutual exclusion. This means it has the property of at least considering likelihoods across multiple measurements.

The second implementation is advantageous because of its simplicity and therefore ease of implementation. Additionally, this second method has a random component to the way it chooses which measurements to process first, second, etc. This is particularly beneficial if there is an ordering effect, where a certain association tends to always be chosen by the first measurement processed. If there is a random component then there is more of a chance that the association chosen will be the correct one. This random ordering property may be more beneficial to the SLAM algorithm than it is to the EKF. Sequential compatibility does have two obvious drawbacks. First, it does not allow for the possibility that the data association which appears to be correct, when viewed through the possible assignments of a single measurement of the batch, may not be correct if all measurements and possible associations are considered jointly.

Second, a notion of optimality over an entire batch of data associations does not exist in the traditional SCNN implementation, whereby measurements are processed one at a time and therefore data association choices are based only on the compatibility of a given measurement and the greedy mutual exclusion requirement. The second SCNN implementation as introduced here does maintain a notion of optimality over the entire batch, where optimality in this situation refers to finding those associations that have

globally maximum likelihoods. However, the notion used in this method is relatively weak as it seeks to find the Maximum Likelihood of individual measurements over all data associations while enforcing individual compatibility and mutual exclusion.

3.2.2 Joint Maximum Likelihood

The method of joint maximum likelihood aims to find the maximal set of data associations, E_k which maximizes the product of the likelihoods [4].

$$E_k = \arg \max_{E_l} \prod_{\{e_m \in E_l\}} f_{e_m} \quad (3.9)$$

where $E_l = \{e_1, e_2, \dots, e_m\}$ is a possible data association set and $e_m = \{i, j\}$ is a data association of measurement i with feature j . The equation for the likelihood $f_{e_m} = f_{ij}$ is given in 3.2.

This problem can also be formulated as a maximization of the sum of log-likelihoods, which can be beneficial due to increased numerical stability.

$$E_k = \arg \max_{E_l} \sum_{\{e_m \in E_l\}} \ln(f_{e_m}) \quad (3.10)$$

Finally, this problem can equivalently be stated as a minimization of the sum of the normalized *Nearest Neighbor* distance [4]. Here the variable N_{e_m} is equivalent to N_{ij} as defined in equation 3.5.

$$E_k = \arg \min_{E_l} \sum_{\{e_m \in E_l\}} N_{e_m} \quad (3.11)$$

A known algorithm for finding the maximum log-likelihood set is the maximum-weight bipartite graph [4]. Additionally, some authors [31] suggest that any data association that does not satisfy individual compatibility should not be considered in joint

maximum likelihood data association. However, there is nothing in the formulation of this method that requires this. Under certain circumstances making this addition will lead to elimination of data associations that are in fact correct, so this concept will not be put into effect for the implementations of this method as discussed here.

One way to think about this method which can help to give one intuition into how it operates is to consider the likelihoods as probabilities. This can be accomplished by normalizing the likelihoods over all possible feature assignments for each measurement such that:

$$\sum_{j=1, i=fixed}^N P_{like}(measurement = i, feature = j) = 1 \quad (3.12)$$

Continuing this logic, the probability that is desired is the joint probability over data associations of measurement i to feature j , where each i and j can only occur once.

$$P(E_k) = P(e_1, e_2, \dots, e_m) \quad (3.13)$$

Where $E_k = \{e_1, \dots, e_m\}$ is a maximal data association set and a $e_l = \{i, j\}$ represents the assignment of measurement z_i with landmark j . Now if an independence assumption between data associations is made this probability can be re-written as:

$$P(E_k) = P(e_1)P(e_2)\dots P(e_m) = \prod_{l=1}^m P(e_l) \quad (3.14)$$

From here it seems that the most desirable E_k would be the one that corresponds to the highest probability, which would correspond to the maximum likelihood. Therefore the chosen data association set is:

$$E = \arg \max_k P(E_k) = \arg \max_k \prod_{l=1}^m P(e_l) \quad (3.15)$$

Herein lies an intuitive view of maximum likelihood data association, which could in fact be implemented in this way if one so desired.

3.2.3 Joint Compatibility

The concept behind Joint Compatibility (JC) data association is to take the individual measurement, chi-squared, gating test into the realm of batch data association. While the Sequential Compatibility Nearest Neighbor method maintains a notion of compatibility amongst an entire batch set of measurements by enforcing the individual compatibility requirement on a sequential basis, it does not consider simultaneous compatibility of the data associations for the entire batch of measurements. Sequential compatibility also fails to maintain a concept of joint optimality across the entire data association set. In a sense, SCNN is a naive application of individual measurement association to the batch data set situation, and has serious failures to take advantage of characteristics of a batch measurement set.

Joint compatibility resolves both of the major pitfalls of SCNN. At the core of the JC concept is the idea that it is possible to expand the individual compatibility concept to a set of many measurements and their landmark assignments. In this method the idea is to find the maximum likelihoods data association sets E_k that are simultaneously jointly compatible. Joint compatibility also maintains a concept of optimality for the entire data association set E , as the solution that will be chosen is the maximal, or nearest maximal, data association set that has the maximum likelihood. The Joint Compatibility method works in the following way:

1. Of all possible data association pairs, $e_l = \{i, j\}$, those that do not satisfy individual compatibility are eliminated.
2. Using the remaining e_l an appropriate search method is used to find the maximal data association set, or nearest maximal set, that satisfies the Joint Compatibility requirement. In the case that multiple, maximal data association sets exist the one with the maximum joint likelihood should be chosen.

Consider, ν_{E_k} which is the innovation sequence for a batch of measurements using the chosen data association set E_k .

$$\nu_{\mathbf{E}_k} = \begin{bmatrix} \nu_{e_1} \\ \nu_{e_2} \\ \vdots \\ \nu_{e_m} \end{bmatrix} \quad (3.16)$$

$$\nu_{e_l} = z_{e_l} - \hat{z}_{e_l} = z_i - \hat{z}_j \quad (3.17)$$

Here, z_i is measurement i and \hat{z}_j is the estimated measurement considering the landmark j is the one being associated with measurement z_i .

In the case of range and bearing measurement pairs being considered, as is the case here, \hat{z}_j would be:

$$\hat{z}_j = \mathbf{h}_j(\hat{\mathbf{x}}) = \begin{bmatrix} \sqrt{(\hat{N}_j - \hat{N}_{ag})^2 + (\hat{E}_j - \hat{E}_{ag})^2} \\ \arctan\left(\frac{\hat{E}_j - \hat{E}_{ag}}{\hat{N}_j - \hat{N}_{ag}}\right) - \hat{\phi}_{ag} \end{bmatrix} \quad (3.18)$$

To determine the joint innovation covariance the Jacobian $\nabla h_x = \frac{\partial h_{\mathbf{E}_k}}{\partial x} \big|_{\hat{x}}$ must be calculated. Here:

$$\hat{z}_{\mathbf{E}_k} = \mathbf{h}_{\mathbf{E}_k}(\hat{\mathbf{x}}) = \begin{bmatrix} h_{e_1}(\hat{x}) \\ \vdots \\ h_{e_l}(\hat{x}) \end{bmatrix} \quad (3.19)$$

This Jacobian is equivalent to building up the following matrix.

$$\mathbf{H}_{\mathbf{E}_k} = \begin{bmatrix} H_{e_1} \\ \vdots \\ H_{e_l} \end{bmatrix} \quad (3.20)$$

The terms H_{e_l} are what is traditionally referred to as the measurement matrix for $e_l = \{i, j\}$ in the Kalman Filter literature. To complete the formulation the matrix R_{E_k} must also be constructed, which is a block diagonal matrix of copies of the measurement covariance matrix.

$$\mathbf{R}_{E_k} = \begin{bmatrix} R_{e_1} & 0 & \dots & 0 \\ 0 & R_{e_2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & R_{e_l} \end{bmatrix} \quad (3.21)$$

$$\mathbf{R}_{e_1} = \mathbf{R}_{e_2} = \mathbf{R}_{e_l} = \begin{bmatrix} \sigma_r^2 & \sigma_{rb} \\ \sigma_{rb} & \sigma_b^2 \end{bmatrix} \quad (3.22)$$

In the cases considered here $\sigma_{rb} = 0$, where σ_{rb} is the range-bearing measurement cross term, σ_r^2 is the variance of the noise in the range measurements, and σ_b^2 is the variance of the noise in the bearing measurements.

These quantities are then used to compute S_{E_k} , the covariance of the innovations.

$$S_{E_k} = H_{E_k} P H_{E_k}^T + R_{E_k} \quad (3.23)$$

The chosen data association set, E_k , must satisfy Joint Compatibility in the following way:

$$\nu_{E_k} = z_{E_k} - \hat{z}_{E_k} \quad (3.24)$$

Then the joint innovation is subject to the following test, or Joint Compatibility requirement:

$$M_{E_k} = \nu_{E_k}^T S_{E_k}^{-1} \nu_{E_k} < \gamma_n \quad (3.25)$$

Here n of γ_n is the dimension of the innovation sequence and γ_n is also a function of the free parameter which determines the probability of not rejecting a good measurement. This is generally taken as 95% or 99%. The requirement of satisfying the $M_{E_k} < \gamma_n$ must be met at all levels in joint compatibility from the individual measurement case to the maximal data association set.

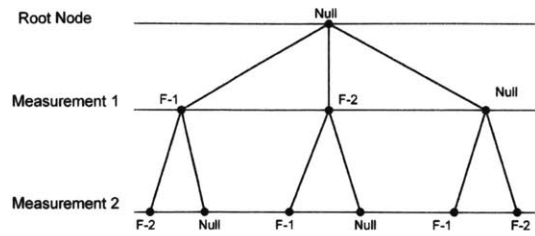


Figure 3-1: Joint Compatibility Branch and Bound Search for a two landmark - two measurement scenario.

Branch and Bound

In order to search all allowable data associations, so that the best Jointly Compatible set E_k can be found, specific search algorithms must be implemented or developed. It is the goal of the search algorithms to search all viable solutions while doing so in as little time as possible or with the least computational complexity. The method that has been presented in the literature to perform this function is called *Branch and Bound* [31, 4].

This search method is uses a depth first search tree where each level of the tree represents a measurement within the batch being considered and each leaf of the tree represents pairing of landmark to the measurement at that level. These leaves would include the null assignment, signifying the possibility of a spurious measurement, see figure 3-1.

In the *Branch and Bound* search methodology the decision of which node is chosen to be searched next is determined by which association given the largest likelihood.

At each level of the tree the current E_l is checked to verify that it passes joint compatibility. If the current association set does not satisfy the compatibility requirement the search down that path is halted and pruned from consideration. The search then moves back up the tree to expand and consider the next available addition to the association set. In this way the methodology "bounds" its search to only compatible sets of data associations.

Joint Compatibility as a Constraint Satisfaction Problem

Another possible method for posing the Joint Compatibility problem is to view it as a *Constraint Satisfaction Problem* (CSP). In framing the problem in this way a more intuitive representation of the problem can be attained, and additionally search methods that are used for solving CSPs become easily applied to the problem at hand [46, 36].

- It will be demonstrated how JC can be framed as a CSP.
- A search method will be presented for solving the CSP formulation, namely *Backtrack Search with Forward Checking* (BT-FC) using *Dynamic Variable Ordering* (DVO).
- The search method will be at worst, as efficient as *Branch and Bound*. In most cases it is expected to be more efficient.

The framing of JC as a CSP is straightforward. Using the framework laid out in [46] a CSP is defined by the triple $\langle V, D, C \rangle$, where V represents the variables, D represents domains, and C represents constraints. In this problem these are defined as:

- Domains == Measurements
- Variables == Landmarks, Spurious Measurement
- Constraints:
 - $V_i \neq V_j$ (Mutual Exclusion Constraint)
 - $M_{[V_1, V_2, \dots, V_l]} < \gamma_n$ (Joint Compatibility Constraint)

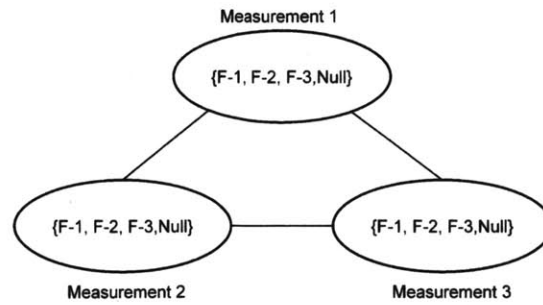


Figure 3-2: Joint Compatibility as a Constraint Satisfaction Problem. Where $\{F-1, F-2, F-3, \text{Null}\}$ represent the possible variable assignments and the F-i's correspond to features in the filter and Null represents the assignment of measurement rejection or new landmark hypothesis depending on the given situation.

Pictorially, this can be represented as in figure 3-2 which shows a three measurement, three landmark case.

Backtracking Search with Forward checking using Dynamic Variable Ordering

Once the Joint Compatibility method of data association is framed as a CSP pre-existing search methods for solving CSPs can be implemented to efficiently find the optimal Jointly Compatible data association set, E_k . One method is *Backtracking Search with Forward Checking*, (BT-FC), using *Dynamic Variable Ordering*, (DVO) [46, 36]. Use of this method is advantageous as it makes it possible to take advantage of the structure of the problem, in particular the sparsity of those data associations that are allowed after accounting for the often large number that are eliminated by use of

individual data association, or measurement, compatibility.

The method of *Dynamic Variable Ordering* is described in [46], and is summarized here for completeness. The premise of this method is that the order of the domains to expand next within the BT-FC search tree and which of the available variables in that domain to assign should be done dynamically as the search proceeds. This is in contrast to the traditional BT-FC search process where the ordering of domains to expand and variables to assign is fixed.

Dynamic Variable Ordering methodology:

- Most constrained domain: Meaning the domain with the fewest allowable variable assignments, in this case that would mean the measurement with the fewest individually compatible landmarks.
- Least constraining variable assignment: Choose the variable assignment that rules out the fewest variables in other domains, which in the JC case this means choose the landmark assignment that is present in the fewest other domains.

The BT-FC method using DVO works in the following way:

- Prune domains using constraint propagation, where the constraints used are defined in 3.2.3.
- Push onto a Queue: DVO chosen children of the domain which is also chosen using DVO.
- Loop: while Queue is not empty
 - Pop queue, thus making a variable assignment to a particular domain.
 - Prune domains based on the above assignment, if a domain has become empty then backtrack and go to 1. This is the Forward Checking step.
 - Check the Joint Compatibility of the current data association set (partial or maximal), E_l , based on variable assignments to the domains as depicted in figure 3-2. If the set fails the compatibility test then backtrack and go to 1.

- If the data association set, E_l , is maximal and has a larger Joint Likelihood than the previous candidate set, E_k , then save as current candidate. Now backtrack and go back to 1.
- Push onto the Queue: DVO chosen children of the domain which is also chosen using DVO, and loop back to 1.

The benefit of this method is that it speeds up the search process by reducing the branching factor at the highest levels of the search tree to the minimal allowed while still searching the entire tree. Considering that the Branch and Bound method can be thought of as comparable to either pure backtracking search or at best BT-FC, the literature allows for a simple comparison of the Branch and Bound method to BT-FC using DVO. In comparing these methods for determining a solution to the N-queens problem, using comparable computational resources, the capability of these methods is about 15-30 queens for the non-DVO based methods and about 1,000 queens for the DVO based method [46]. While this is a rough comparison it does give an order of magnitude sense of how the performance of these two methods would stand up.

3.2.4 Multiple Hypothesis Data Association

The basic concept behind this method is that at the time of receiving an observation and prior to processing it, the way to deal with ambiguity in the data association is to consider all reasonable data association possibilities simultaneously. The success of this method relies on the idea that information about which data association is correct will come in the form of future information, therefore allowing the filter to decide which past data associations are correct.

In an EKF framework this method requires the maintenance and update of an entire bank of EKFs where each has committed to a different set of data associations at each measurement step. Unfortunately, such a bank of filters can grow in size in an exponential fashion since every filter in the bank may have to split into multiple filters after each measurement is received. Even in the case where pruning methods may be available to limit the number of EKFs needed to explore the data association ambiguity,

with a large number of landmarks this method will quickly become very computationally expensive. Additionally, there is the problem of deciphering how to use the information available in the numerous filters being maintained for a useful purpose, as the different filters will contain conflicting information. In essence, this methodology for EKF's is very difficult to implement as a real-time algorithm.

For FastSLAM the multiple hypothesis concept is already ingrained in the formulation as data association is done on a per-particle basis. The particles are not only a representation of the uncertainty of the pose, but are also individual hypotheses of the pose states. Therefore, it may be advantageous to allow for each particle to try each data association hypothesis that is plausible, as well as the measurement rejection and new landmark hypothesis. This can be accomplished by splitting each particle and forcing it into enough copies to consider each of these hypotheses. Such a method has been explored in the literature by Thrun, Nieto, Guivant, and Nebot in [45].

This method is interesting because it allows for a simple method by which both the possibility of measurement rejection and new landmark initializations can be considered simultaneously. Additionally, this method appears to have good SLAM results [45]. There are some major problems with this method however. First, it is not clear how to interpret the information in the filter when these multiple hypothesis of particles that have been split are floating around. Additionally, in order to allow this method to continue to work in real-time some ad hoc limits must be put in place when the number of particles are reduced to the original number and these limits will be heavily dependent on changing variables such as the number of landmarks in the filter and the number of measurements being received.

Because of the way that the FastSLAM filter works in term of sampling particles based on their weighting this method would intuitively seem to have the same results as simply making data associations by sampling the likelihoods of the data associations. According to one of the original authors of the FastSLAM method this is indeed the result of applying this type of data association method to this filter [30].

3.2.5 Delayed Assignment Data Association Algorithms

More recently a number of new data association algorithms have appeared in the literature that take a different approach to the ones previously discussed. The basic premise of these methods are to put more importance on making the correct data association algorithms and less importance on processing measurements as they arrive. The benefit of doing this is that it allows of future data to resolve data association ambiguities that may otherwise be unresolvable. Some methods that incorporate this method make use of the RANSAC algorithm [16] to match data from the multiple measurement sets, while others make use of the Hough Transform to accomplish a similar goal of improving data association assignments by delaying the assignments made [32]. More recently there has been a paper [21] that discussed a maximum likelihood data association method that also sought to improve data association assignments by being "lazy" about making assignments. This method has the additional property of being able to go back in time and repair previous data association assignments. These methods are certainly useful and quite powerful in their ability to discern the correct data association in ambiguous situations. However, they will not be studied in-depth here as they would require a change in paradigm with respect to the real-time filtering applications that are being considered.

[This page intentionally left blank.]

Chapter 4

Simulated Results for Filter-Data Association Marriages

Through the use of simulations, a comparison will be made of the operation of the EKF and FastSLAM filters when married to three different batch data association methods: Sequential Nearest Neighbor, Joint Compatibility, and Joint Maximum Likelihood. To make these comparisons, relatively simple scenarios will be used. However, the problems posed by such scenarios will be shown to be challenging in certain cases and the same problems will persist in more complex situations.

The basic simulation involves an agent moving around a square trajectory for eight minutes of simulation time. During this time the agent takes measurements of two landmarks every five seconds. As a part of this analysis the same set of simulations were run for eight different landmark pair separations. These eight landmark pair separations were used both for a scenario where the landmark pairs are outside of the agents trajectory, as examined in section 4.2 and for a case where the pairs are inside, or partially inside, the trajectory as discussed in section 4.3. The purpose of this is to offer problems of varying difficulty to the data association algorithms.

4.1 Assumptions and Simulation Set-up

In order to interpret the results presented here it is important to understand the major assumptions made in the problem set-up. First, it is assumed that only two landmarks are present in the world at a given time, and that the filter begins with an initial estimate of the location of the landmarks and a notion of uncertainty in the landmarks position. It is assumed that the agent has an *a priori* estimate of its own pose as well as a reasonably correct notion of the uncertainty in the pose, in other words the initial error between the distribution mean and the truth state is bounded.

Additionally, it is assumed that all measurements are received as range-bearing pairs and both landmarks are seen at every measurement time step. The possibility of a spurious measurement is still allowed within the data association algorithm. These assumptions mean that the data association problem for these simulations can be reduced to the determination of which of the two landmarks each measurement should be associated with. At the same time the possibility that one or both of the measurements received should in fact be rejected as spurious is also allowed.

The greatest simplifying assumption that is made here is that the algorithms do not need to deal with adding new landmarks that may not have been known about *a priori*. If both the possibility of measurement rejection and a new landmark exist simultaneously the data association problem becomes much more difficult. Most solutions to this problem that would make use of the algorithms discussed here would involve some *ad hoc* procedures to be functional. This is because the criteria that is used here to reject measurements is the same one that would be used to make the decision to initialize a new landmark, if that was allowed as a possibility.

Finally, for each of the landmark separations that is studied for each filter-data association pair, including the perfect data association case, a set of 20 Monte Carlo runs were performed. It should be noted that while this is expected to be a sufficient number of runs to determine trends in the behavior of filter-data association marriages, it may not be enough to determine true average behavior for a given situation.

4.1.1 Numerical Values Used for Simulations

This section will briefly describe the relevant numerical values used in the simulation in terms of uncertainties, the motion model, and the data association thresholds. The initial pose and landmark uncertainties, as well as the measurement uncertainties, used in the simulation are given in table 4.1. The implication here is that the initial distribution of both landmark and pose states is assumed to be Gaussian. While this is knowingly an incorrect assumption for angular states (heading), it was made to allow for direct comparisons between EKF and FastSLAM performance.

$\sigma_{\{N,E\}_{ag}}$	$\sigma_{\phi_{ag}}$	$\sigma_{\{N,E\}_{F_{1,2}}}$	σ_{range}	$\sigma_{bearing}$
10 m	0.393 rad	20 m	5 m	8×10^{-3} rad

Table 4.1: Initial pose and landmark uncertainties used in simulations, as well as the range and bearing measurement uncertainties used.

The motion model used for the agent is relatively simple: it involves a slip scale factor and a skid error both for the translational velocity and rotational velocity. The translational velocity was held constant at 3.0 meters-per-second, while the rotational velocity was nominally zero except when the agent turned the corner on the trajectory. The numerical values used for these quantities are given in table 4.2.

v_t	v_r	$slip_{v_t}$	$skid_{v_t}$	$slip_{v_r}$	$skid_{v_r}$
3.0 m/s	0 rad/sec	0.1	0.01 m/s	0	0.01 rad/sec

Table 4.2: The slip scale factor and skid error used for calculating the translational and rotational velocity errors, as well as the translational and rotational velocities used.

These values are used to calculate the translational and rotational velocity errors in the following way:

$$\sigma_{v_t} = slip_{v_t} \cdot |v_t| + skid_{v_t} \quad (4.1)$$

$$\sigma_{v_r} = slip_{v_r} \cdot |v_r| + skid_{v_r} \quad (4.2)$$

The errors in both components of velocity are assumed to be normally distributed, such that:

$$v_{t,measure} \sim N(v_{t,true}, \sigma_{v_t}) \quad (4.3)$$

$$v_{r,measure} \sim N(v_{r,true}, \sigma_{v_r}) \quad (4.4)$$

For the Joint Compatibility and Sequential Nearest Neighbor methods of data association there is a free parameter which must be chosen so that a threshold for measurement rejection can be established, as previously stated in 3.1.2, this will always be taken to be 95%. This is a result of using the following rule:

$$M_{ij} = v_{ij}^T S_j^{-1} v_{ij} < \gamma_n \quad (4.5)$$

For a single measurement of range and bearing, n is equal to 2 and therefore γ_n is taken to be 5.9915. However, for Joint Compatibility the value of n will very depending on the portion of the data association set that is being tested.

4.2 Landmark Pairs Outside of Trajectory

Simulated Experiment

The trajectory used along with the landmark pairing separations considered for the *outside of trajectory* case, including the numbers assigned to these pairings, are depicted in figure 4-1.

While this case is seemingly very simplistic in nature the results obtained are important and can generalize to many more complex situations. When it comes to picking the correct data associations the challenge of overcoming ambiguity will often boil down to resolving the data association ambiguity for just two landmarks.

This is best understood by considering the case where two landmarks that are relatively close together are within view of a measurement instrument and because of the relative distance and the accuracy of the instrument some ambiguity exists. If then a third landmark comes into view but is further away than the separation of the previous two

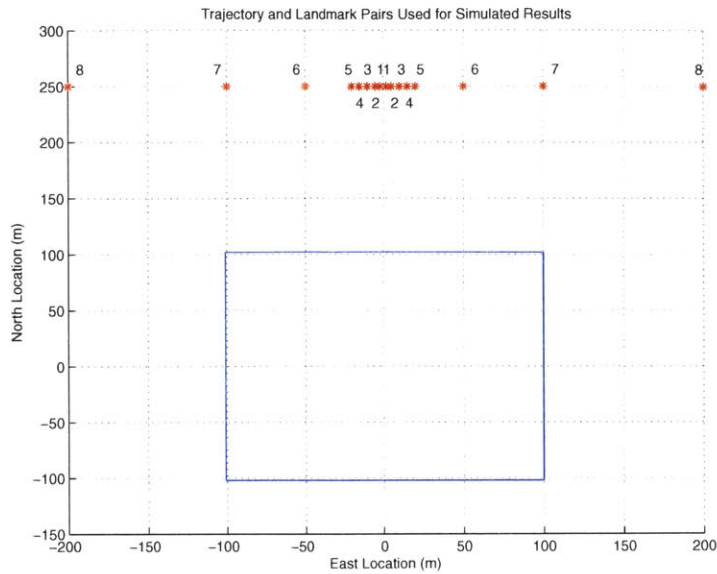


Figure 4-1: The trajectory and landmark pairings used in computing the simulated results for the outside of trajectory case. Landmark Pair Numbers correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

landmarks which are still in view, the data association problem can in-effect be decoupled so that the problem is still one of resolving the ambiguity between the closer two (of the three) landmarks. This case will not always hold, but in many cases it does and the results from this experiment are therefore informative.

4.2.1 Results for Perfect Data Association

In order to set a performance benchmark against which to compare the performance of the data association-filter marriages, a set of twenty simulations for each landmark separation were run where data association assignments are known *a priori*. The results from these simulations are expressed using average RMS of the state errors for both an EKF and FastSLAM approach. The RMS errors are plotted as a function of landmark separation in figure 4-2.

Intuitively, the results shown here make a good deal of sense, particularly for the EKF. Namely, for the EKF the average RMS errors for the pose states (agent position and heading) decrease with increased landmark separation. This is because more information can be gained about the agent's pose when the landmarks are further away

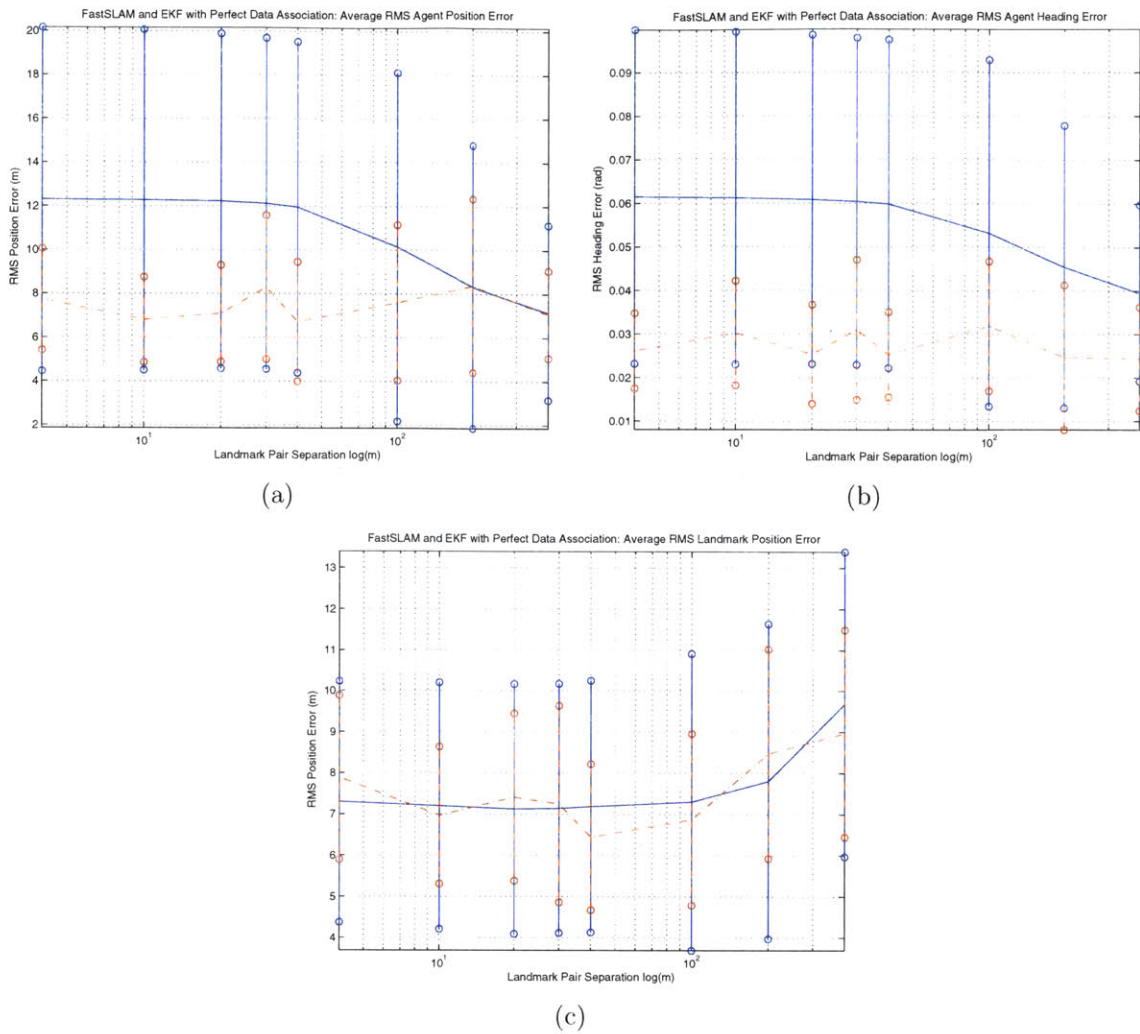


Figure 4-2: Performance of the EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, filters for the perfect data association case where the landmarks are located outside of the trajectory. Subfigures (a) and (b) show the agent's RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

from the observer and one another. Additionally, the average RMS error in the landmark states increases for large separations which is also intuitively sensible. This behavior is due mainly to the fact that for situations where the landmarks are far away, small bearing measurement errors have a large effect on where the landmarks are interpreted to be via measurements. Conversely, this effect is not seen in the agent states because the two landmarks, whose positions are being measured, are located in opposite directions which helps to cancel this effect out.

Similar behavior to that described for the EKF is seen for the FastSLAM filter as well. However, it is obvious from figure 4-2 that the FastSLAM filter has much tighter average performance than that of the EKF, as the EKF has much larger variability in its performance which is indicated by the error bars. The average performance of the FastSLAM filter for localizing the pose states is shown to be significantly better than the EKF for all but one landmark separation, where the biggest difference appears for landmarks that are closer together. Overall the error in the pose states for FastSLAM is relatively flat with changing landmark separation, where there is about a 1.5 meter difference between the best and worst cases for position and about 0.07 radians for heading.

The characteristics of the average RMS error in landmark position for FastSLAM is similar to that of the EKF with slightly more variability. One thing that may be significant to note is the decreased performance of FastSLAM and EKF for the 200 meter landmark separation case, which is the one where the landmarks are lined up at the end of the trajectory, pair seven in figure 4-1.

This same behavior will come up again in later cases of FastSLAM-data association algorithm marriages. It is likely an artifact of particle depletion issues with the FastSLAM 1.0 filter used here, which may come into play in this situation since the "wrong" particles may be sampled away when the agent is moving directly towards (or away from) one of the landmarks. This would be caused by a decrease in observability that occurs under this situation, since lateral motion becomes more difficult to determine.

4.2.2 Description of Data Format Used for Filter-Data Association Algorithm Marriages

For each of the data association methods considered there are four different sets of plots that are shown. The first is a three-dimensional plot which shows the average data association error made for the particular data association algorithm when married to a particular SLAM filter. This is plotted as both a function of time and landmark

separation. For the EKF the average is of an ensemble nature, as it is taken over twenty Monte Carlo runs. For FastSLAM two results are shown: the first is the ensemble average of the number of data association errors averaged over all particles and the second is the ensemble average for just the maximum likelihood particle at each measurement step.

The second set of plots is calculated in exactly the same manner as for the data association errors. However, the statistic considered by these plots is the average number of measurements rejected.

The third set of plots consolidates the information shown in the first two sets of plots by averaging this data over time. These plots then show the average number of measurement rejections and data association errors only as a function of landmark separation.

The final set of plots attempt to depict the performance of the SLAM filters. All three of the plots in this set show average RMS errors for different quantities: agent position, agent heading, and landmark position. This is the same format that was used to show the performance of the filters for the perfect data association case in 4-2.

Additionally, a set of three tables will be used to compare the results for the three different data association methods applied to the two filtering methods considered.

These tables will summarize RMS agent position error, RMS landmark position error, and average data association errors. This will be done in a separate section where all methods will be compared simultaneously.

4.2.3 Results for Sequential Nearest Neighbor Data Association

Due to the characteristics of Sequential Nearest Neighbor data association, as discussed earlier, particularly its lack of a notion of optimality and its inability to consider associations to multiple measurements in a joint manner, it would be expected that this method would perform worse than the Joint Compatibility and Joint Maximum Likelihood methods. The results shown here demonstrate that, in terms of the number

data association errors made, this is the case for larger landmark separations when compared to the other methods studied. At the same time this method makes fewer data association errors at the smallest landmark separations when compared to the other two data association algorithms considered.

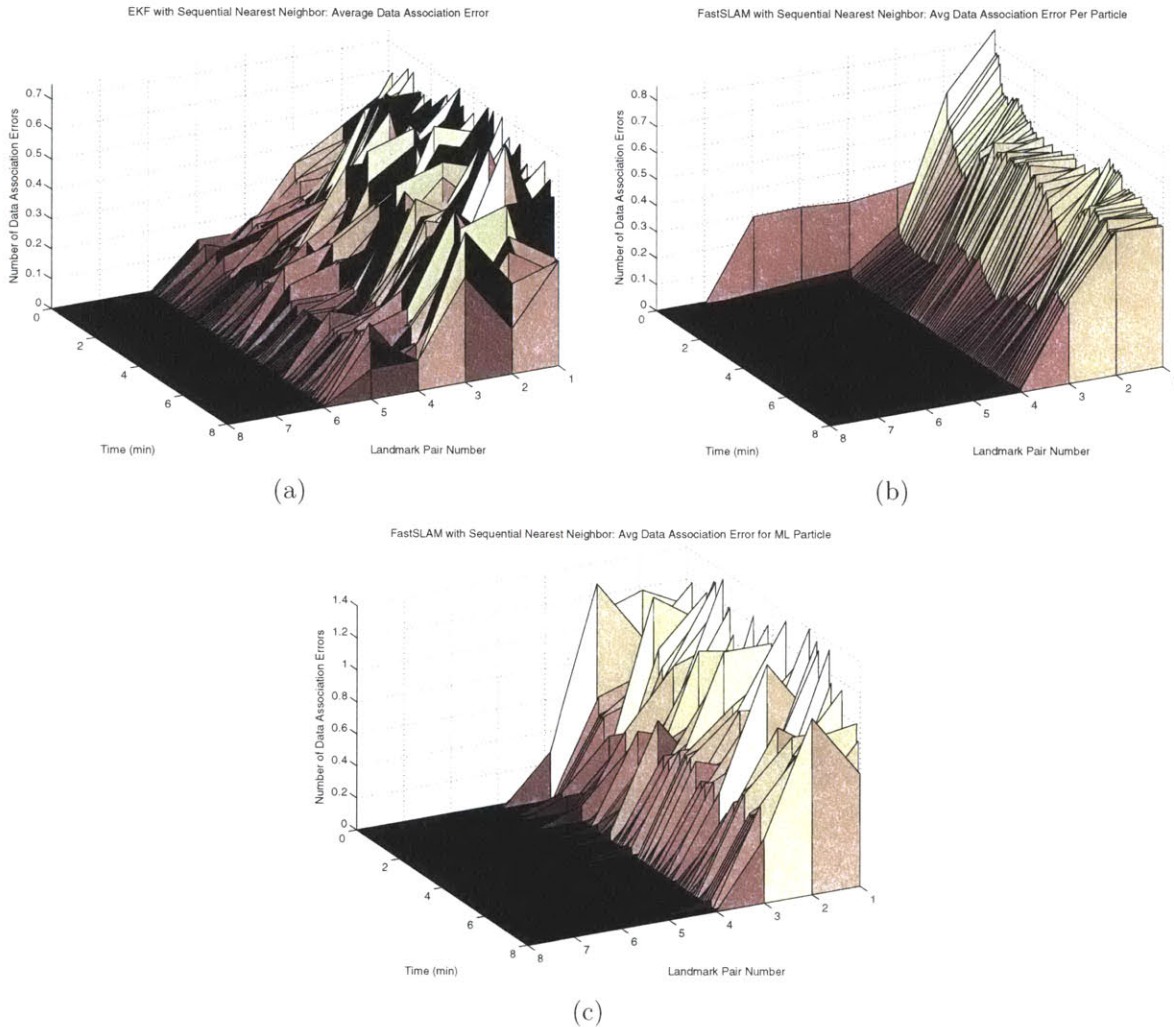


Figure 4-3: Sequential Nearest Neighbor data association performance variation over time and landmark separation where the landmarks are located outside of the trajectory. The average number of data association errors versus time and landmark separation for EKF (a) and FastSLAM, where (b) shows the average per-particle errors and (c) shows just the errors for the maximum likelihood particle. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

In figure 4-3 it can be seen that for the EKF-SCNN marriage, data association errors are being made regularly, even at landmark separations of 40 meters. This is much worse than the FastSLAM-SCNN marriage, which makes very few data association

errors for landmark separations greater than 20 meters. It will also be shown that other data association algorithms perform better, when married to the EKF, than the SCNN algorithm does for landmark separations greater than 30 meters.

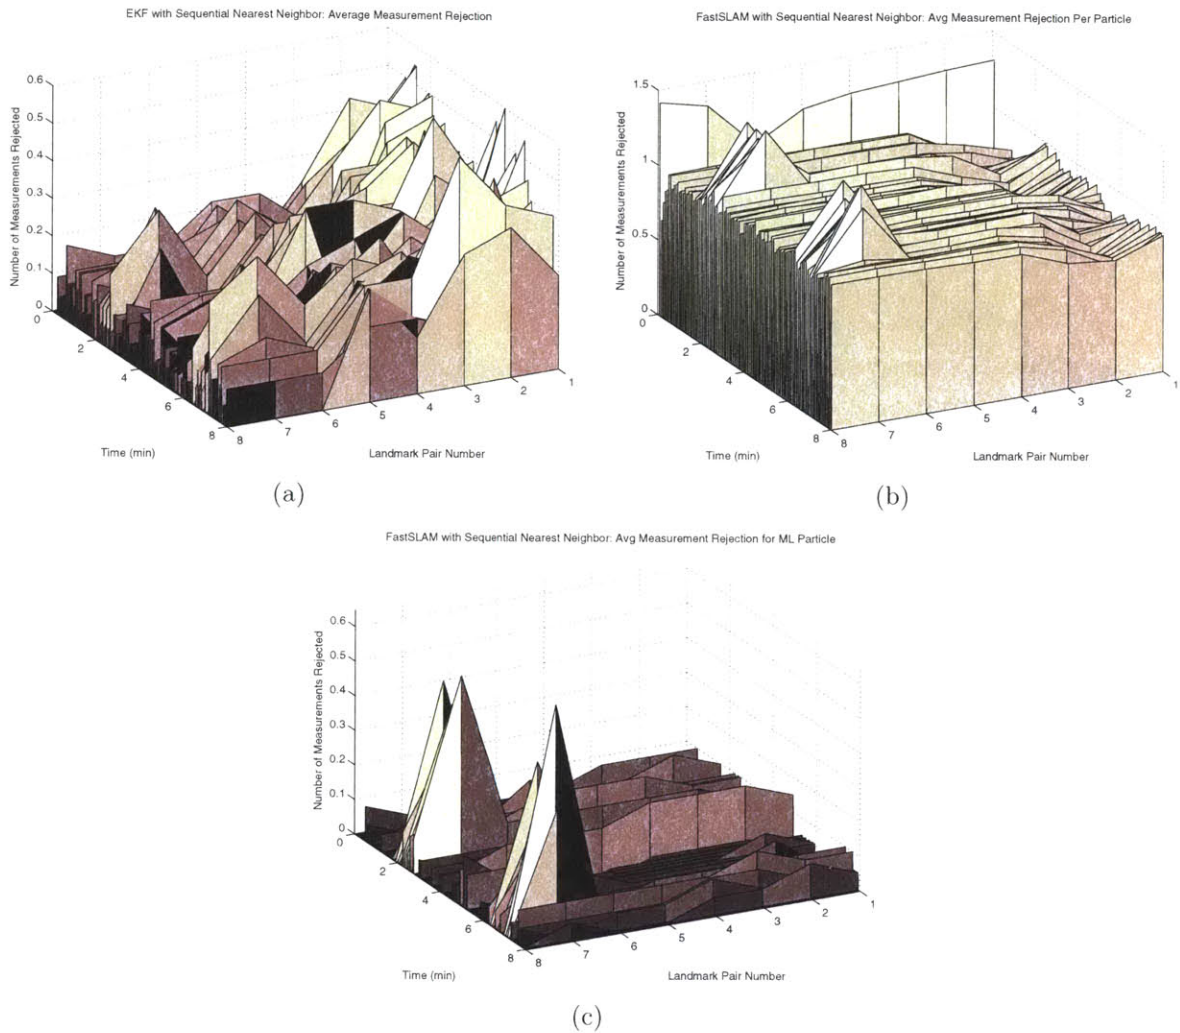


Figure 4-4: Sequential Nearest Neighbor data association measurement rejection characteristics where the landmarks are located outside of the trajectory. The average number of measurements rejected versus time and landmarks separation for EKF (a) and FastSLAM, where (b) shows the average per-particle rejection and (c) shows results for just the maximum likelihood particle. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

In the Figure 4-4, the most notable aspect is the large number of measurements that get rejected using this method. It is particularly interesting to see that the average number of measurements rejected per-particle for FastSLAM is close to one. This would be of greater concern; however, the lower plot also shows that the maximum likelihood

particle makes far fewer rejections, which due to re-sampling, means that the measurement information is still being absorbed into the filter. At the same time this could become a significant issue for real world applications which will be discussed in chapter 5. The most notable feature of these plots are the characteristics seen in subfigure 4-4(c), which shows that for FastSLAM the maximum likelihood particle only rejects measurements for the 200 meter landmark separation case and only for certain spots in the trajectory. This is notable because it appears to correspond to a noticeable improvement in filter performance over the perfect data association case.

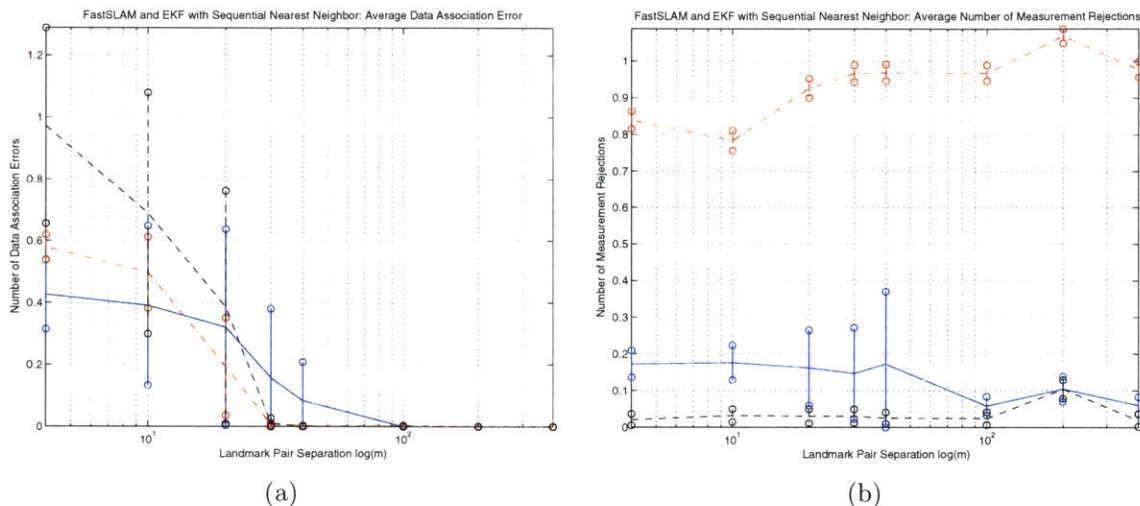


Figure 4-5: Average measurement rejection, (a), and data association errors, (b), for Sequential Nearest Neighbor data association for both EKF and FastSLAM for the case where the landmarks are outside of the trajectory. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

The plots in figure 4-5 shows some interesting characteristics in comparing EKF and FastSLAM for SCNN. The average data association error plot, 4-5(a), shows that at small landmark separations less than 20 meters, the EKF is actually performing better than FastSLAM. It is also interesting to note that the average number of data association errors made by the maximum likelihood particle at the smallest separations 4 and 10 meters performs far worse than SCNN with the EKF. The measurement rejection plot, 4-5(b), once again demonstrates that the average particle in FastSLAM rejects many more measurements than the EKF implementation or the FastSLAM

maximum likelihood particle.

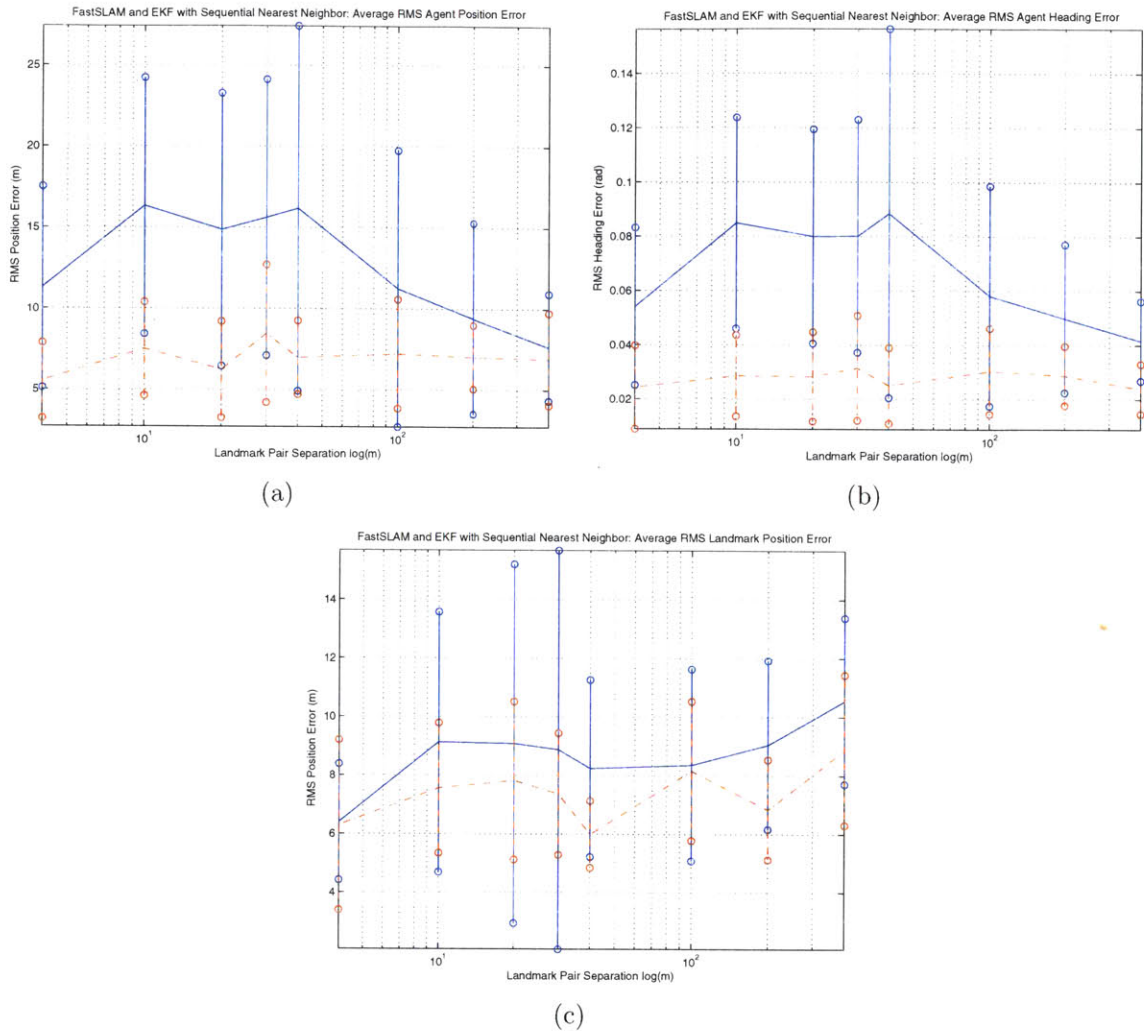


Figure 4-6: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Sequential Nearest Neighbor data association where the landmarks are located outside of the trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

In comparing figures 4-2 and 4-6 for landmark separations from 10 meters to 30 meters it can be seen that the EKF has experienced a significant decrease in performance, especially in localization. This is the range in which SCNN makes the most data association errors and measurement rejections. It is interesting to see that for the 4 meter separation case there is not a large change in performance even though this is also the case for which data association errors are most prevalent. This is because the

landmarks are quite close in proportion to the range measurement error and therefore data association errors have less of an effect.

Additionally, it is interesting to observe that the performance of the FastSLAM filter, while slightly degraded in some cases, is not significantly affected by the ambiguous data association situation. The most notable difference here is that the spike in position error for the 200 meter landmark separation case, that was seen in the perfect data association case, figure 4-2, is not present here in figure 4-3. This behavior is the result of two factors, the first is that the SCNN algorithm appears to have the ability to act as a pre-filter, only allowing the best measurements to be processed. The second factor is that particles which reject measurements receive the lower weights in this scheme, which is based on the gating threshold, and are therefore more likely to be sampled away.

4.2.4 Results for Joint Compatibility Data Association

The results for data association errors made using Joint Compatibility can be seen in figure 4-7. For the EKF these results are similar to those using SCNN, however it appears that at smaller landmark separations Joint Compatibility performs worse. In comparing the FastSLAM results it appears that for the average particle the number of data association errors made is slightly larger with JC than it was with SCNN, however with JC there is an increased dependence on time, or location within the trajectory. Also the average particle appears to make a single data association error at the first measurement. The most significant difference, in data association errors, between SCNN and JC for FastSLAM can be seen in the maximum likelihood particle performance. Here, the maximum likelihood particle makes very few data association errors for landmark separations greater than 10 meters.

The results for measurement rejections using Joint Compatibility are shown in figure 4-8. It appears that the characteristics are not that much different than was found with SCNN; however, the average number of rejections being made has been reduced for both the EKF and FastSLAM with both the average per-particle and maximum likelihood particle.

The average measurement rejection and data association error performance is

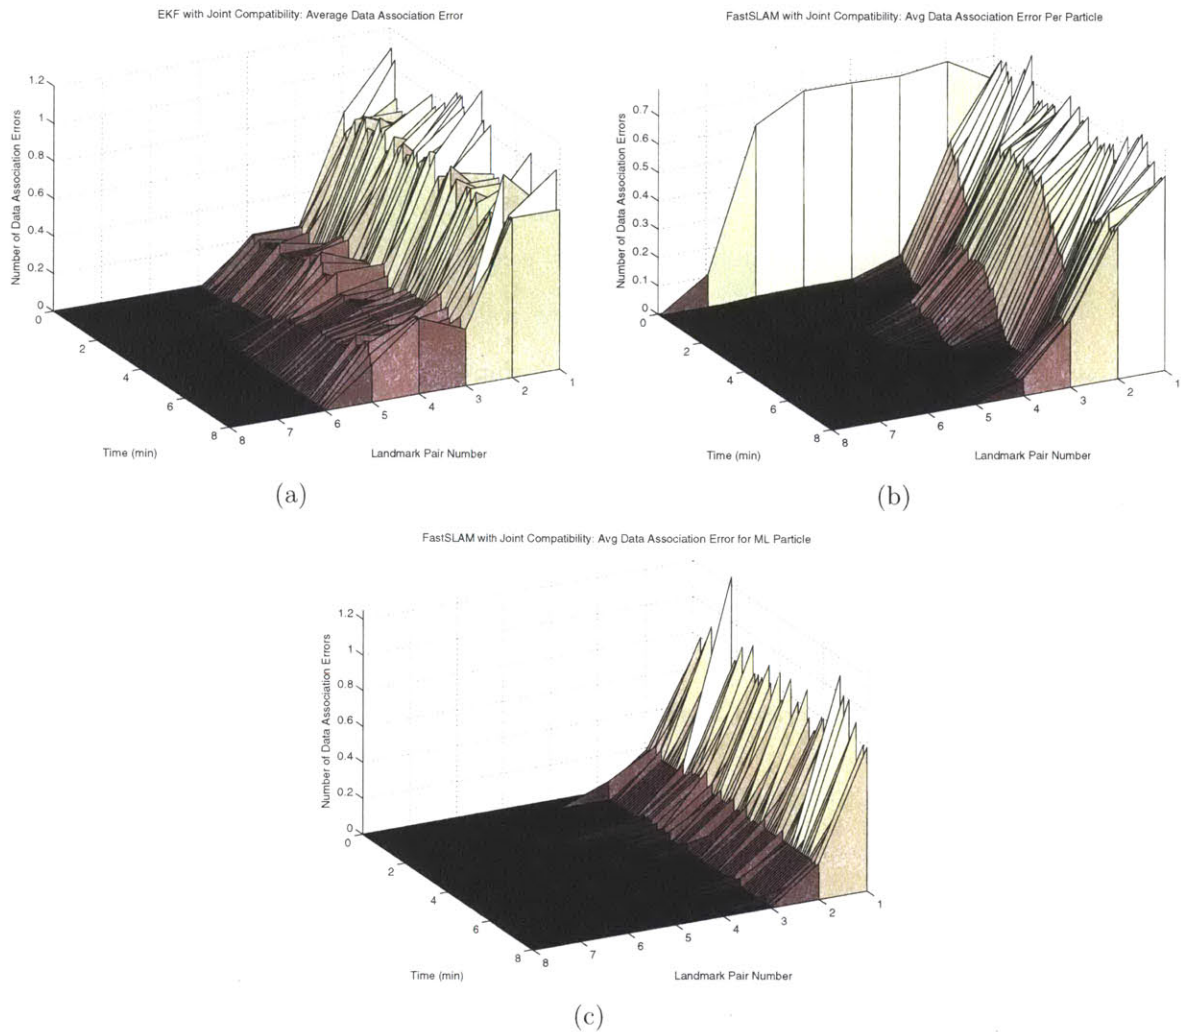


Figure 4-7: Joint Compatibility data association performance variation over time and landmark separation where the landmarks are located outside of the trajectory. The average number of data association errors versus time and landmark separation for EKF (a) and FastSLAM, where (b) shows the average per-particle error and (c) shows just errors made by the maximum likelihood particle. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

summarized in figure 4-9. The most significant result that can be seen here for the EKF is that at the landmark separations of 4 and 10 meters this method performs far worse than SCNN. For the FastSLAM filter the most significant item to note about these plots is that the maximum likelihood particle makes fewer data association errors and the average particle makes fewer measurement rejections than was found using the SCNN. The performance of the EKF and FastSLAM filters when paired with Joint Compatibility is shown in figure 4-10. In examining the results for the EKF, it appears

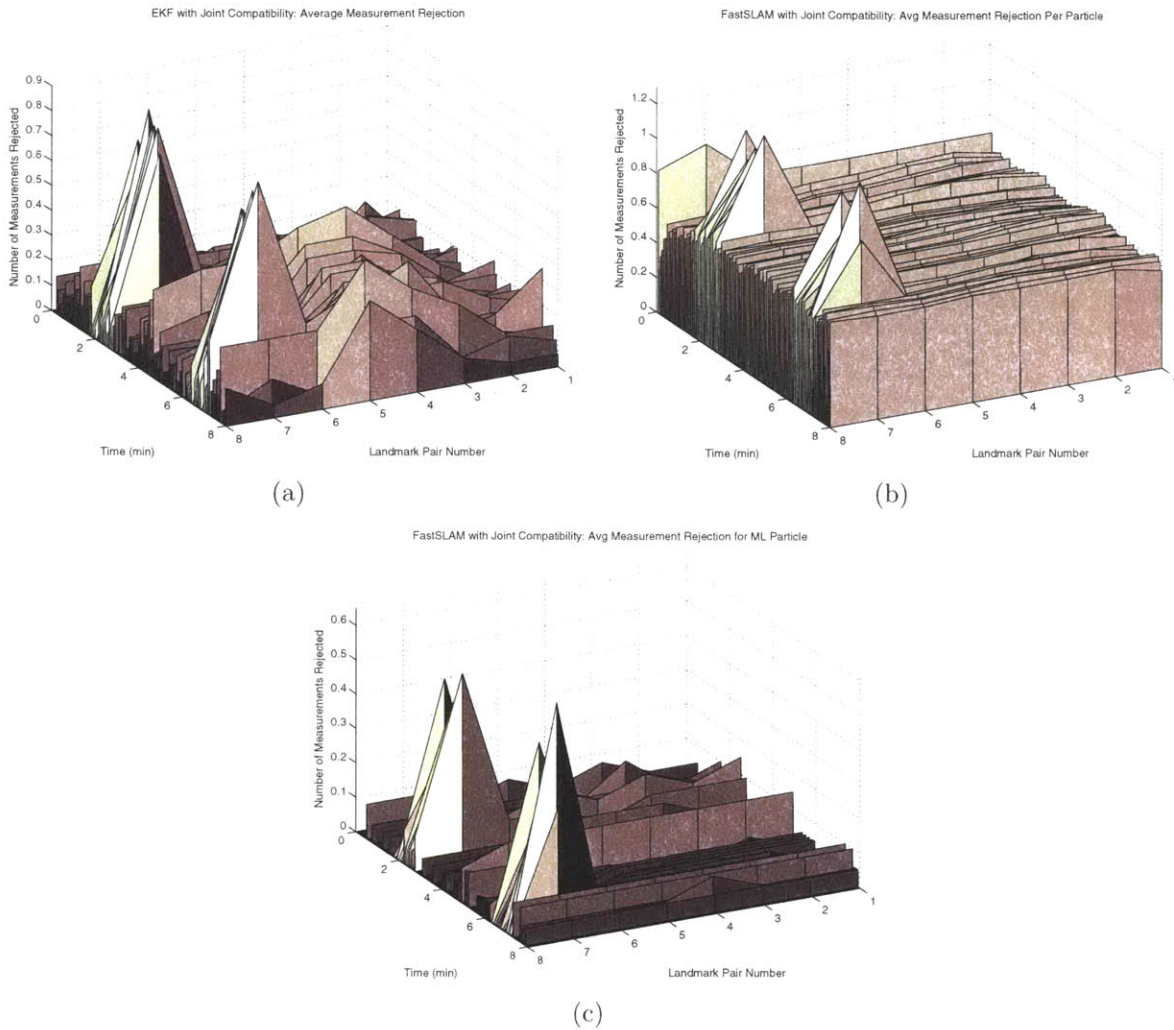


Figure 4-8: Joint Compatibility data association measurement rejection characteristics for the case where the landmarks are located outside of the trajectory. The average number of measurements rejected versus time and landmark separation for EKF (a) and FastSLAM, where (b) shows the average per-particle rejection and (c) shows results for just the maximum likelihood particle.. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

that this pairing has both good and poor results. The most obvious demonstration of the poor results comes from the large increase in RMS position errors for the 4 meter landmark separation case. Additionally, there is some decrease in performance, via landmark error, for some larger landmark separations. However, those increases are not as significant.

For the FastSLAM filter, the performance with JC has not changed substantially from what was seen with SCNN. The most significant change comes in the increase in agent

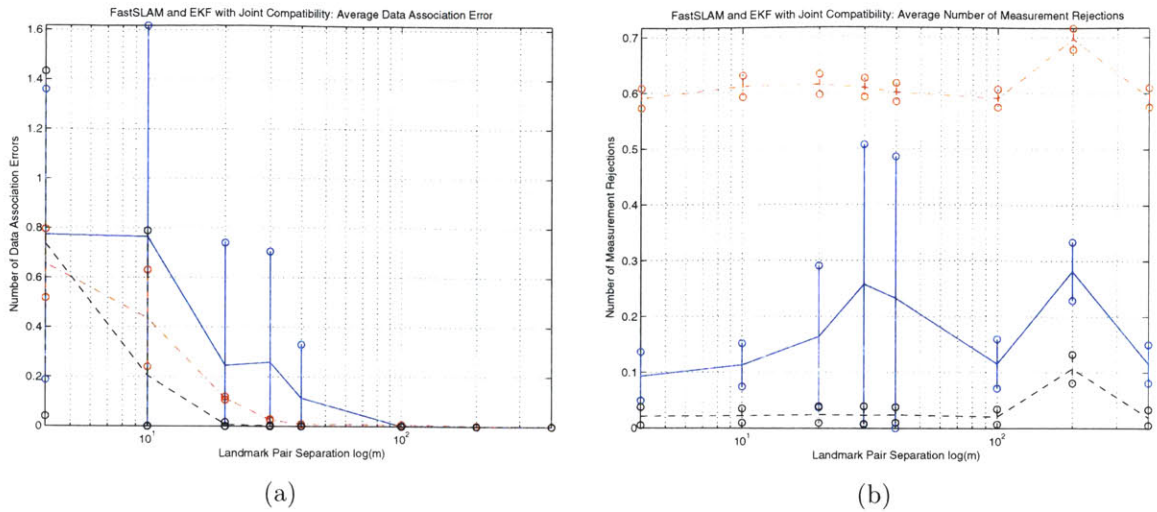


Figure 4-9: Average measurement rejection, (a), and data association errors, (b), for Joint Compatibility data association for both EKF and FastSLAM where the landmarks are outside of the trajectory. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

and landmark position error for the 200 meter separation case. This is again attributed to the geometry of this case, with regards to the positioning of landmarks along the line of motion of the agent. The SCNN method achieves different results from JC in this situation because of the measurement rejections it is making, this is reflected in almost twice as many measurement rejections than JC on average.

4.2.5 Results for Maximum Likelihood Data Association

As can be seen in figure 4-11 the Joint Maximum Likelihood (JML) method of data association produces the most consistent results, in terms of data association errors, across time when paired with both the EKF and FastSLAM, when compared with the Joint Compatibility and Sequential Compatibility Nearest Neighbor methods. For the EKF the number of data association errors being made, for small landmark separations of 4 meters and 10 meters, is comparable to Joint Compatibility, however slightly worse than what was seen for SCNN. At intermediate separations of 20 and 30 meters Joint Maximum Likelihood continues to have similar performance, in terms of data association errors made, to that of JC using an EKF. However, when compared to the

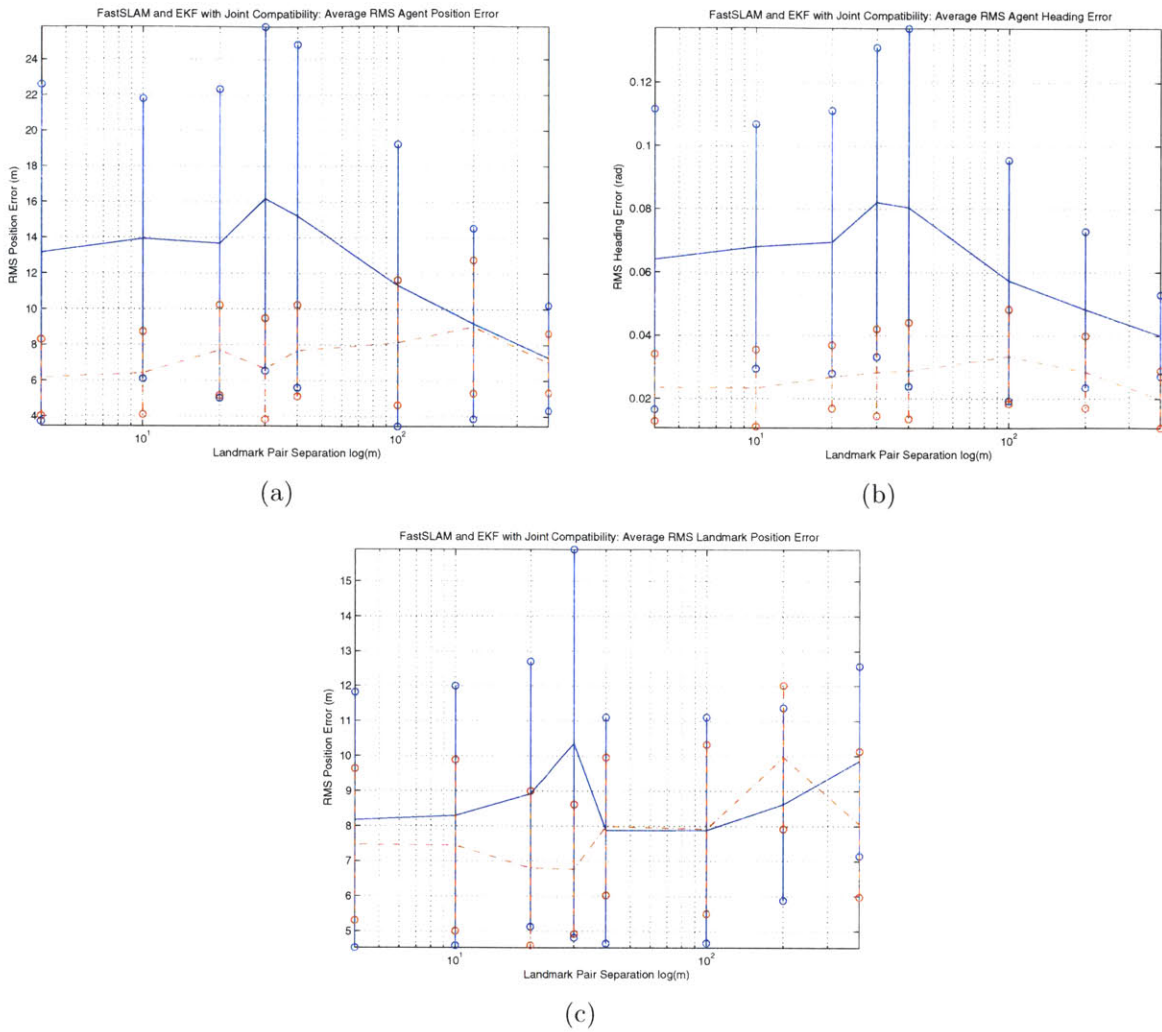


Figure 4-10: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Joint Compatibility data association where the landmarks are located outside of the trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

SCNN method, when paired with an EKF, the performance is similar to 30 meter separations but at the smaller separation of 20 meters SCNN performs worse than the other two methods.

When compared to the previous two methods studied for usage by the EKF the JML method sets itself apart for the cases of landmark separation greater than 30 meters by making no data association errors after that distance.

The pairing of FastSLAM with Joint Maximum Likelihood method appears to have slightly worse performance than the other two methods for the smallest landmark

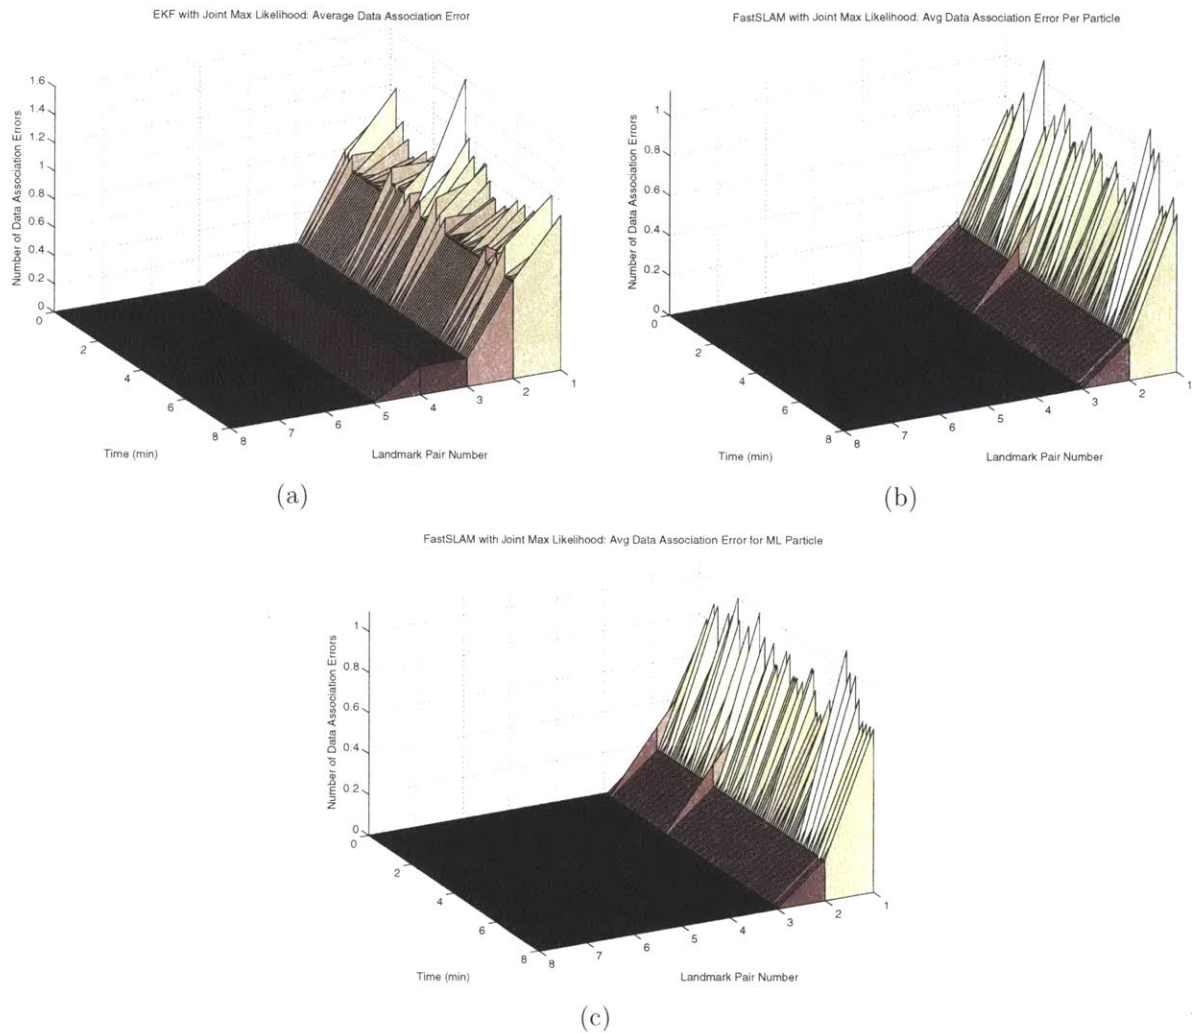


Figure 4-11: Joint Maximum Likelihood data association performance variation over time and landmark separation for EKF (a) and FastSLAM, where (b) shows the average per-particle error and (c) shows just the maximum likelihood particle error. The average number of data association errors versus time and landmark separation. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

separation considered, 4 meters. However, for separations of 10 meters and greater the JML method paired with FastSLAM appears to perform far better than any of the combinations studied. For this pairing no data association errors occur for landmark separations of greater than 10 meters. As can be seen in subfigures 4-11(b) and 4-11(c) the performance of the average particle and the maximum likelihood particle in making data association choices correctly are very similar. This is very interesting as it has implications for the performance of the this filter-data association pairing when applied

to the problem where the possibility of new, previously not-seen landmarks are allowed. This behavior indicates that when using JML, when faced with such a problem, FastSLAM would have a set of particles that would be fairly similar in their concept of the map of the environment.

The plots of the measurement rejections made using JML can be seen in figure 4-12. Examining these plots, it is obvious that the loss of observability effect caused by walking directly towards a landmark from which measurements are being received is once again showing up. Here it affects the data association algorithm for both EKF and FastSLAM. It shows up in measurement rejections during the portion of the trajectory where the agent moves either directly toward, or away from, a landmark because the lack of observability causes an increase in the error of the landmark estimate. This in turn is reflected in the data association algorithm rejecting measurements because for some period of time measurements are more likely to not fit the current state estimate.

The average number of data association errors made and measurements rejected for the pairing of Joint Maximum Likelihood with both the EKF and FastSLAM filter can be seen in figure 4-13. The most notable features here are that the average data association error for both the mean per-particle and maximum likelihood particle are essentially equivalent for FastSLAM. Additionally, the spike in measurement rejections for the 200 meter landmark separation can be seen for both EKF and FastSLAM in subfigure 4-13(b).

Figure 4-14 shows the mapping and localization performance obtained for marrying JML with EKF and the FastSLAM filter. These results are very encouraging since for most of the cases shown demonstrate performance that is very comparable to the perfect data association case. For the EKF the only large deviation from this case is the error in landmark localization when the separation the landmarks is less than 30 meters or less. The FastSLAM filter when paired with Joint Maximum Likelihood seems to perform as good or better than in the perfect data association case. The only exception to this is the error in landmark location the 100 meter landmark separation, however even here the result is comparable to that obtained with perfect data association.

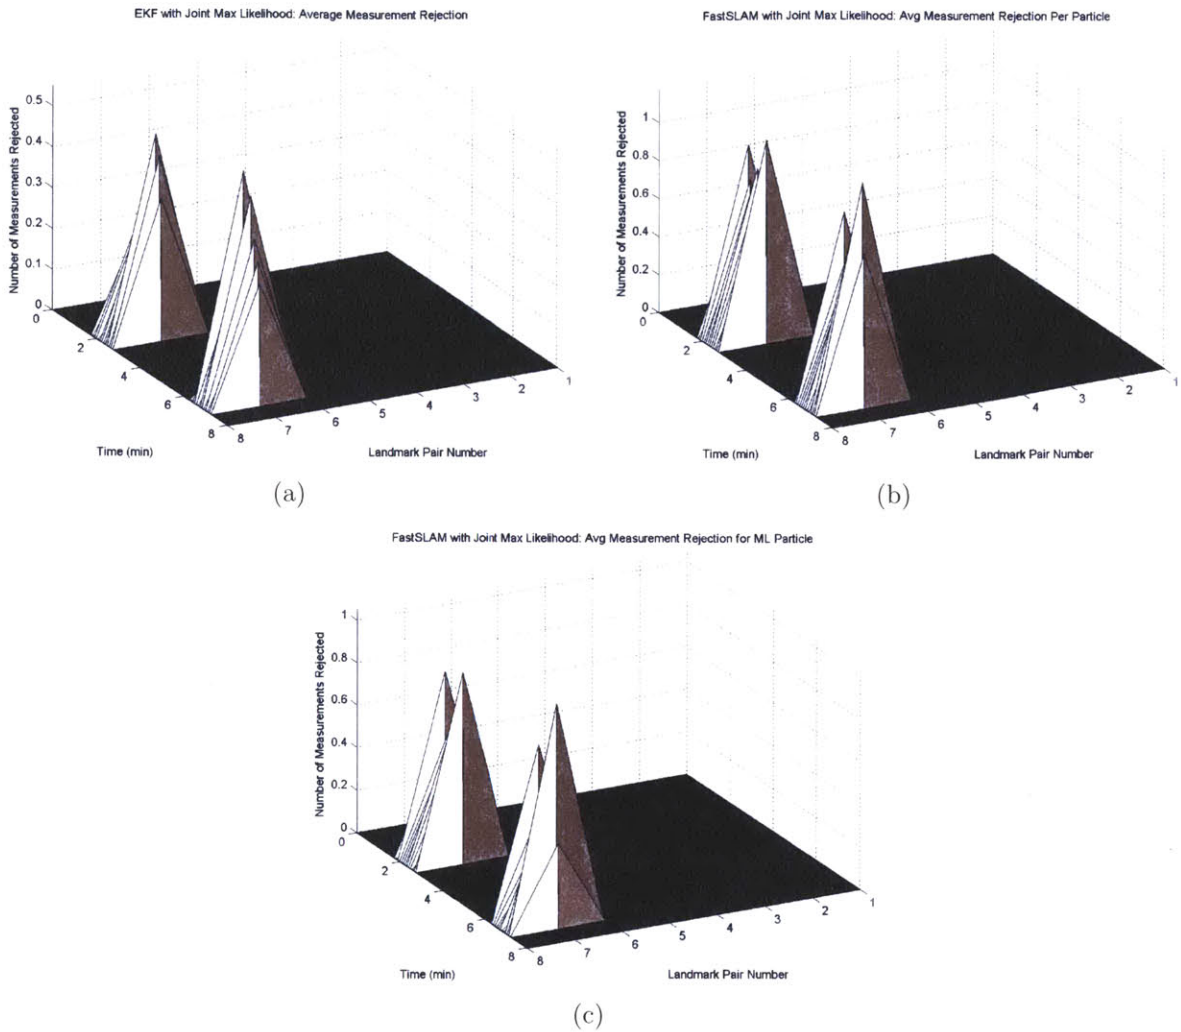


Figure 4-12: Joint Maximum Likelihood data association measurement rejection characteristics for the case where the landmarks are exterior to the trajectory. The average number of measurements rejected versus time and landmarks separation for EKF (a) and FastSLAM, where (b) shows average per-particle rejection and (c) shows results for just the maximum likelihood particle. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

4.2.6 A Quick Performance Comparison of the Various Filter-Data Association Marriages

Examining table 4.3 it is apparent that for this case, where the landmarks are located well outside of the agent trajectory, FastSLAM has better performance than the EKF in comparing perfect data association as well as all data association methods considered. For the EKF based marriages used, under uncertain data association, it appears as

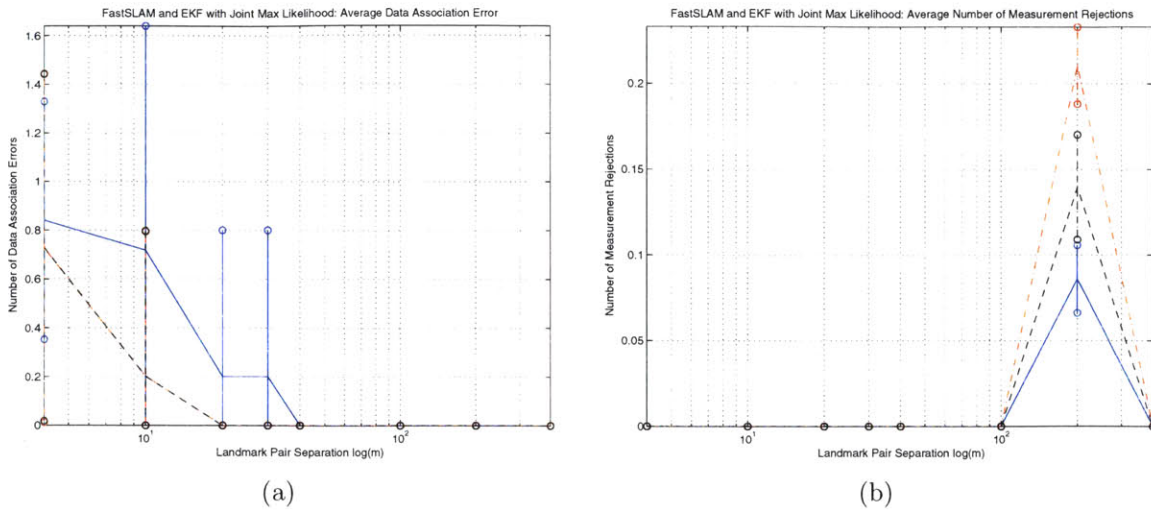


Figure 4-13: Average measurement, (a), and data association errors, (b), for Joint Maximum Likelihood data association for both EKF and FastSLAM. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

though EKF-JML performs the best in terms of minimum error for agent localization. In comparing the FastSLAM-data association combinations, no one method appears to really stand out as the best performing for agent localization, although JML may have a slight edge over the other two in terms of the smallest maximum average RMS error. A summary of the performance of the various filter-data association combinations for landmark localization is given in table 4.4. This data also indicates that the EKF-JML combination gives slightly better performance in landmark localization when compared with the other two EKF based methods. For the FastSLAM based methods it is once again difficult to discern which of the three data association methods leads to the best filter performance in terms of minimal landmark position error. It does appear as though the JML and SCNN produces similarly good results with a slight edge over the JC method. An important item to note from the information given in this table is that the average landmark position errors across EKF and FastSLAM based methods are very comparable, with FastSLAM maintaining a much smaller edge for this metric. The average data association error information for JML, JC, and SCNN when paired with EKF and FastSLAM are shown in table 4.5 for the outside of trajectory landmark situation. The most notable trends for both EKF and FastSLAM are how the SCNN

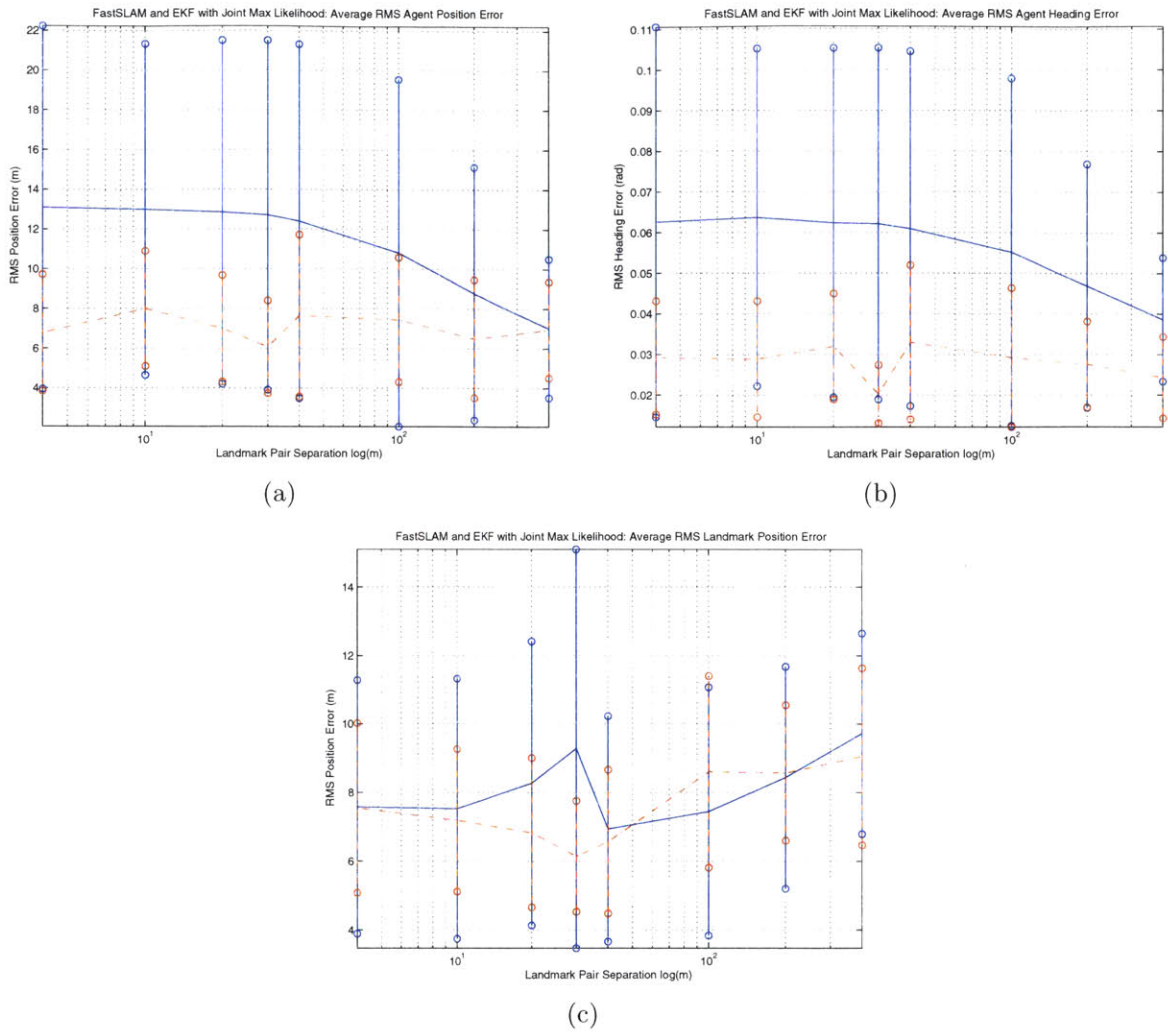


Figure 4-14: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Joint Maximum Likelihood data association for the case where the landmarks are exterior to the trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

data association method has the lowest average number of mistakes at small landmark separations, but also has a long tail making mistakes out to 40 meter landmark separations. At the same time the JML method seems to have the opposite property, making large numbers of mistakes for the very close together pair, 4 meters, but then dropping off very quickly with increasing landmark separation.

One other observation that can be made examining all three of these tables 4.3, 4.4, and 4.5 is the effect of using data association algorithms even in situations where no

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
Perfect Data Association - EKF (m)	12.3	12.3	12.2	12.1	12.0	10.2	8.3
$\sigma_{Perfect-EKF}$ (m)	± 7.9	± 7.8	± 7.6	± 7.6	± 7.6	± 8.0	± 6.5
JC - EKF (m)	13.2	13.9	13.7	16.2	15.2	11.4	9.2
σ_{JC-EKF} (m)	± 9.4	± 7.8	± 8.6	± 9.7	± 9.6	± 7.9	± 5.3
SCNN - EKF (m)	11.3	16.3	14.9	15.6	16.2	11.2	9.4
$\sigma_{SCNN-EKF}$ (m)	± 6.2	± 7.9	± 8.4	± 8.5	± 11.3	± 8.5	± 5.9
JML - EKF (m)	13.1	13.0	12.9	12.7	12.4	10.8	8.8
$\sigma_{JML-EKF}$ (m)	± 9.1	± 8.3	± 8.7	± 8.8	± 8.9	± 8.7	± 6.4
Perfect Data Association - FS (m)	7.7	6.8	7.1	8.3	6.7	7.6	8.4
$\sigma_{Perfect-FS}$ (m)	± 2.3	± 1.9	± 2.2	± 3.3	± 2.7	± 3.6	± 4.0
JC - FS (m)	6.2	6.4	7.7	6.7	7.7	8.2	9.0
σ_{JC-FS} (m)	± 2.1	± 2.3	± 2.5	± 2.8	± 2.6	± 3.5	± 3.7
SCNN - FS (m)	5.5	7.5	6.2	8.5	7.0	7.2	7.0
$\sigma_{SCNN-FS}$ (m)	± 2.3	± 2.9	± 3.0	± 4.2	± 2.3	± 3.4	± 2.0
JML - FS (m)	6.8	8.0	7.0	6.1	7.7	7.4	6.5
σ_{JML-FS} (m)	± 2.9	± 2.9	± 2.7	± 2.3	± 4.1	± 3.1	± 3.0

Table 4.3: Performance comparison of filter-data association marriages for agent position error using an average RMS error metric. In this case the landmarks are located outside of the agents trajectory.

incorrect data associations are being made. Namely, that the performance of the filters are not identical to the performance under the perfect data association. Obviously, if no incorrect data associations are being made in these instances the change in performance must be due to measurement rejections that are occurring. In essence all data association methods have the ability to act as a pre-filter for measurements coming in, and in fact will even in situations where data association appears to be trivial. In some cases this has a pronounced positive effect, in other cases it does not, nonetheless the effect is there and it is important to be aware of.

4.3 Landmarks Inside of Trajectory

An additional set of scenarios were studied for the data association-SLAM filter problem using simulated data. For these scenarios the landmark separation is varied over the same intervals as before, however the locations vary drastically from the outside of

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
Perfect Data Association - EKF (m)	7.3	7.2	7.1	7.2	7.2	7.3	7.8
$\sigma_{Perfect-EKF}$ (m)	± 2.9	± 3.0	± 3.0	± 3.0	± 3.1	± 3.6	± 3.8
JC - EKF (m)	8.2	8.3	8.9	10.4	7.9	7.9	8.6
σ_{JC-EKF} (m)	± 3.6	± 3.7	± 3.8	± 5.5	± 3.2	± 3.2	± 2.7
SCNN - EKF (m)	6.4	9.1	9.1	8.9	8.3	8.4	9.1
$\sigma_{SCNN-EKF}$ (m)	± 2.0	± 4.4	± 6.1	± 6.8	± 3.0	± 3.3	± 2.9
JML - EKF (m)	7.6	7.5	8.3	9.3	7.0	7.5	8.4
$\sigma_{JML-EKF}$ (m)	± 3.7	± 3.8	± 4.1	± 5.8	± 3.3	± 3.6	± 3.2
Perfect Data Association - FS (m)	7.9	7.0	7.4	7.3	6.5	6.9	8.5
$\sigma_{Perfect-FS}$ (m)	± 2.0	± 1.7	± 2.0	± 2.4	± 1.8	± 2.1	± 2.5
JC - FS (m)	7.5	7.4	6.8	6.8	8.0	7.9	10.0
σ_{JC-FS} (m)	± 2.2	± 2.4	± 2.2	± 1.8	± 2.0	± 2.4	± 2.0
SCNN - FS (m)	6.3	7.6	7.8	7.4	6.0	8.2	6.8
$\sigma_{SCNN-FS}$ (m)	± 2.9	± 2.2	± 2.7	± 2.1	± 1.1	± 2.4	± 1.7
JML - FS (m)	7.6	7.2	6.8	6.2	6.6	8.6	8.6
σ_{JML-FS} (m)	± 2.5	± 2.1	± 2.2	± 1.6	± 2.1	± 2.8	± 2.0

Table 4.4: Performance comparison of filter-data association marriages for landmark position error using an average RMS error metric. In this case the landmarks are located outside of the agents trajectory.

trajectory case. The trajectory used along with the landmark pairings is shown in figure 4-15.

Here the rightmost landmark, denoted with XX is positioned at 85 meters East and 0 meters North. This landmark remains in the same location for all scenarios and the westerly landmark is moved further west based on the given landmark separation.

4.3.1 Results for Perfect Data Association

The results for this set of landmark pair positions given perfect data association is shown in 4-16. As can be seen here the results for these scenarios vary drastically from those seen in the landmarks outside of trajectory case. One example of this is the dramatic decrease in performance of both filters for almost all landmark separations. This is particularly true for the FastSLAM filter which, in the majority of scenarios, performs worse than the EKF.

While these differences in performance are significant, of greater importance is the fact

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
JC - EKF	0.78	0.76	0.25	0.25	0.11	0	0
σ_{JC-EKF}	± 0.59	± 0.85	± 0.5	± 0.45	± 0.22	0	0
SCNN - EKF	0.43	0.39	0.32	0.16	0.08	0	0
$\sigma_{SCNN-EKF}$	± 0.11	± 0.26	± 0.32	± 0.22	± 0.12	0	0
JML - EKF	0.84	0.72	0.20	0.20	0	0	0
$\sigma_{JML-EKF}$	± 0.49	± 0.92	± 0.6	± 0.6	0	0	0
JC - FS	0.66	0.44	0.11	0.026	0.008	0.007	0.001
σ_{JC-FS}	± 0.14	± 0.2	± 0.01	± 0.003	± 0.001	± 0.001	± 0.001
JC - FS for ML Particle	0.73	0.20	0.007	0.001	0	0	0
$\sigma_{JC-FSw/ML}$	± 0.70	± 0.58	± 0.009	± 0.002	0	0	0
SCNN - FS	0.58	0.5	0.2	0.007	0.004	0.003	0
$\sigma_{SCNN-FS}$	± 0.04	± 0.12	± 0.16	± 0.003	± 0.001	± 0.001	0
SCNN - FS for ML Particle	0.97	0.69	0.38	0.01	0	0	0
$\sigma_{SCNN-FSw/ML}$	± 0.32	± 0.39	± 0.38	± 0.02	0	0	0
JML - FS	0.73	0.2	0.0004	0.0002	0.0001	0	0
σ_{JML-FS}	± 0.72	± 0.6	± 0.001	± 0.0004	± 0.0002	0	0
JML - FS for ML Particle	0.73	0.2	0	0	0	0	0
$\sigma_{JML-FSw/ML}$	± 0.71	± 0.6	0	0	0	0	0

Table 4.5: Performance comparison of filter-data association marriages for data association errors made by the combination used. In this case the landmarks are located outside of the agents trajectory.

that the overall trends continue to make physical sense. The agent localization errors decrease with increasing landmark spacing, while the landmark position errors increase. These trends are essentially the same as what was seen in the former scenario.

4.3.2 Description of Data Format Used for Filter-Data Association Algorithm Marriages

For purposes of brevity, a more limited set of figures will be examined for the landmarks inside of trajectory scenarios. In this case two sets of plots will be examined for each data association method considered. The first set of plots will show average RMS error in: agent position, agent heading, and landmark position. The second set of plots demonstrate the average number of data association errors and measurement rejections made by the EKF filter, average particle and the maximum likelihood particle for

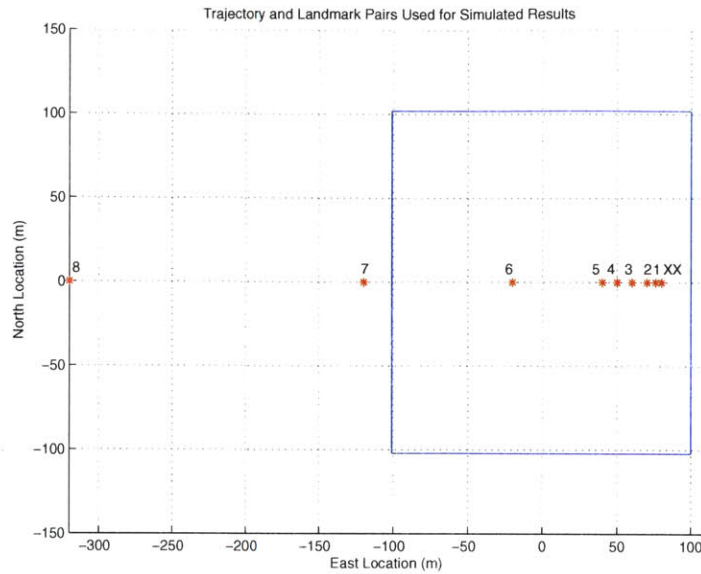


Figure 4-15: The trajectory and landmark pairings used in computing the simulated results for the inside of trajectory case. Landmark Pair Numbers [1,...,8] correspond to separations of [4m,10m,20m,30m,40m,100m,200m,400m].

FastSLAM. After these plots are used to examine each data association algorithm individually, a set of three tables will be used to examine the performance across all filter- data association marriages.

4.3.3 Results for Sequential Nearest Neighbor Data Association

Examining subfigure 4-17(a), which shows the average data association errors for the EKF and FastSLAM filters when paired with SCNN for the inside of trajectory landmark case, it appears that the characteristics seen here are similar to those seen in the outside of trajectory landmark case.

One of the characteristics that remains the same is that the SCNN algorithm performs better at smaller landmark separations when paired with the EKF than it does with FastSLAM. Additionally, for FastSLAM the average particle and the maximum likelihood particle have similar data association error characteristics to what was seen previously. The main difference appears to be an effect whereby the EKF-SCNN pair makes more data association errors for 20 meter landmark separation than it does for 10

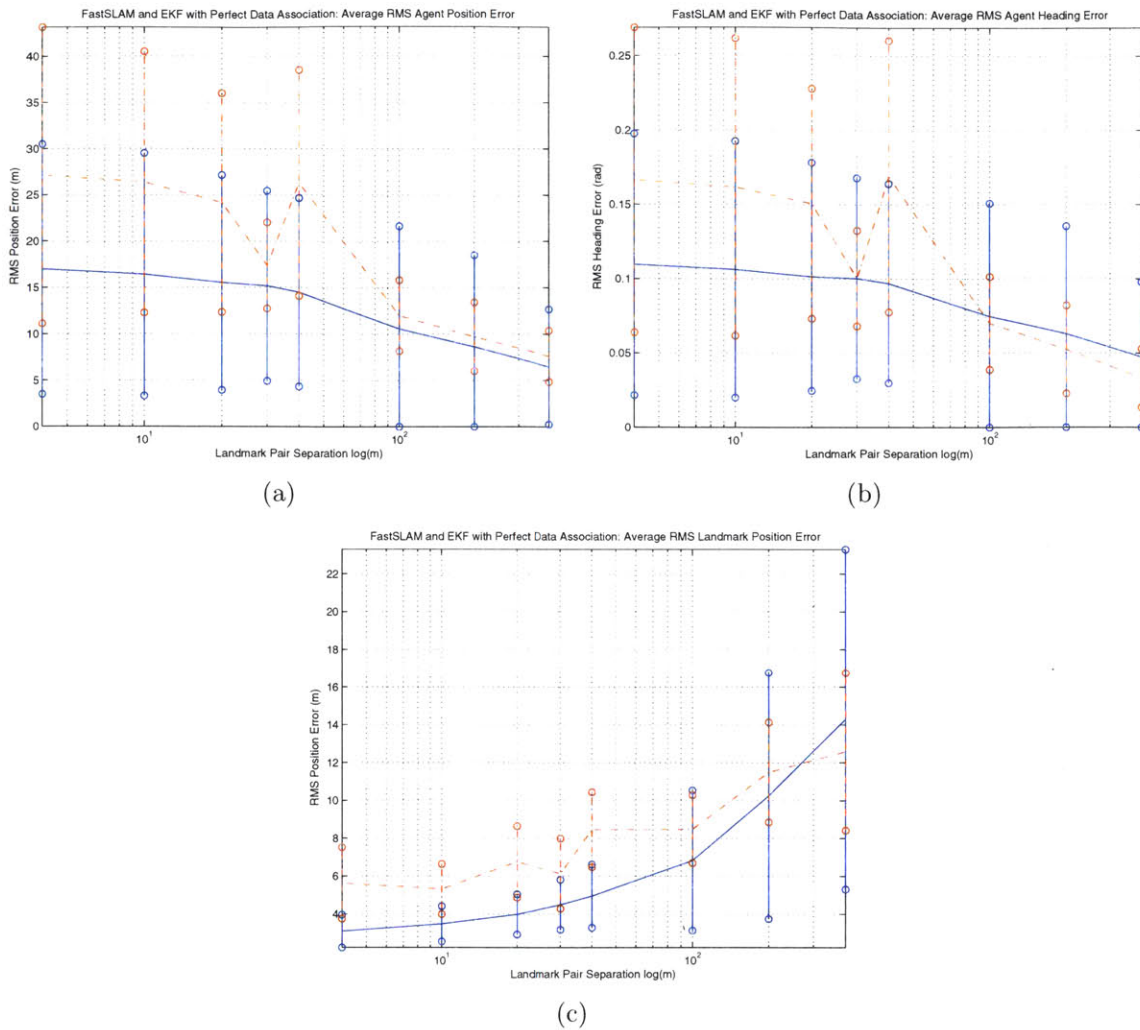


Figure 4-16: Performance of the EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, filters for the perfect data association case with Landmarks inside. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

meter separation. This is an unusual result as the relationship between data association errors and landmark separation has for the majority of cases been shown to be a monotonically decreasing one. The measurement rejection characteristics seen in subfigure 4-17(b), are very similar to previous landmark configuration with the only significant difference being the absence of the spikes in rejection that had been seen for 200 meter landmark separations.

Figure 4-18 shows the filter performance obtained using SCNN for this configuration of landmark locations. While these results are difficult to compare with those seen of the

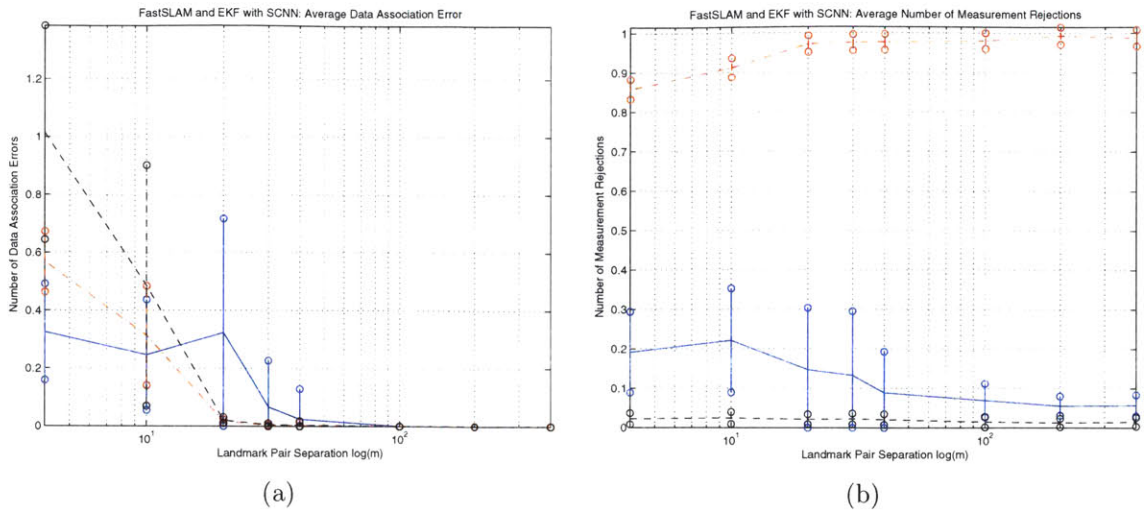


Figure 4-17: Average measurement rejection, (a), and data association errors, (b), for Sequential Nearest Neighbor data association for both EKF and FastSLAM when the landmarks considered are inside of the agent trajectory. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

outside of landmark case, they are still interesting unto themselves. One interesting result can be seen in plots 4-18(a) and 4-18(b), which demonstrate that the FastSLAM filter performs significantly better at agent localization than the EKF at landmark separations of 4 and 10 meters. This is particularly interesting because it was shown in the previous plots that for these separations fewer data association errors are being made by the EKF-SCNN combination.

Two other interesting results can be seen in subfigure 4-18(c): the first is the large spike in RMS landmark position error for the EKF with a landmark separation of 20 meters and the second is the increase in performance over the perfect data association case for FastSLAM in landmark localization. The spike in position error for the EKF is attributed to the increase in data association errors for this same condition, of course these two results have the potential to interact, that is, an increase in the number of data association errors leads to greater errors in localization error, at the same time greater localization error leads to an increase in data association ambiguity and a greater probability of increased data association errors.

It is unclear what leads to the increase in performance of FastSLAM in landmark

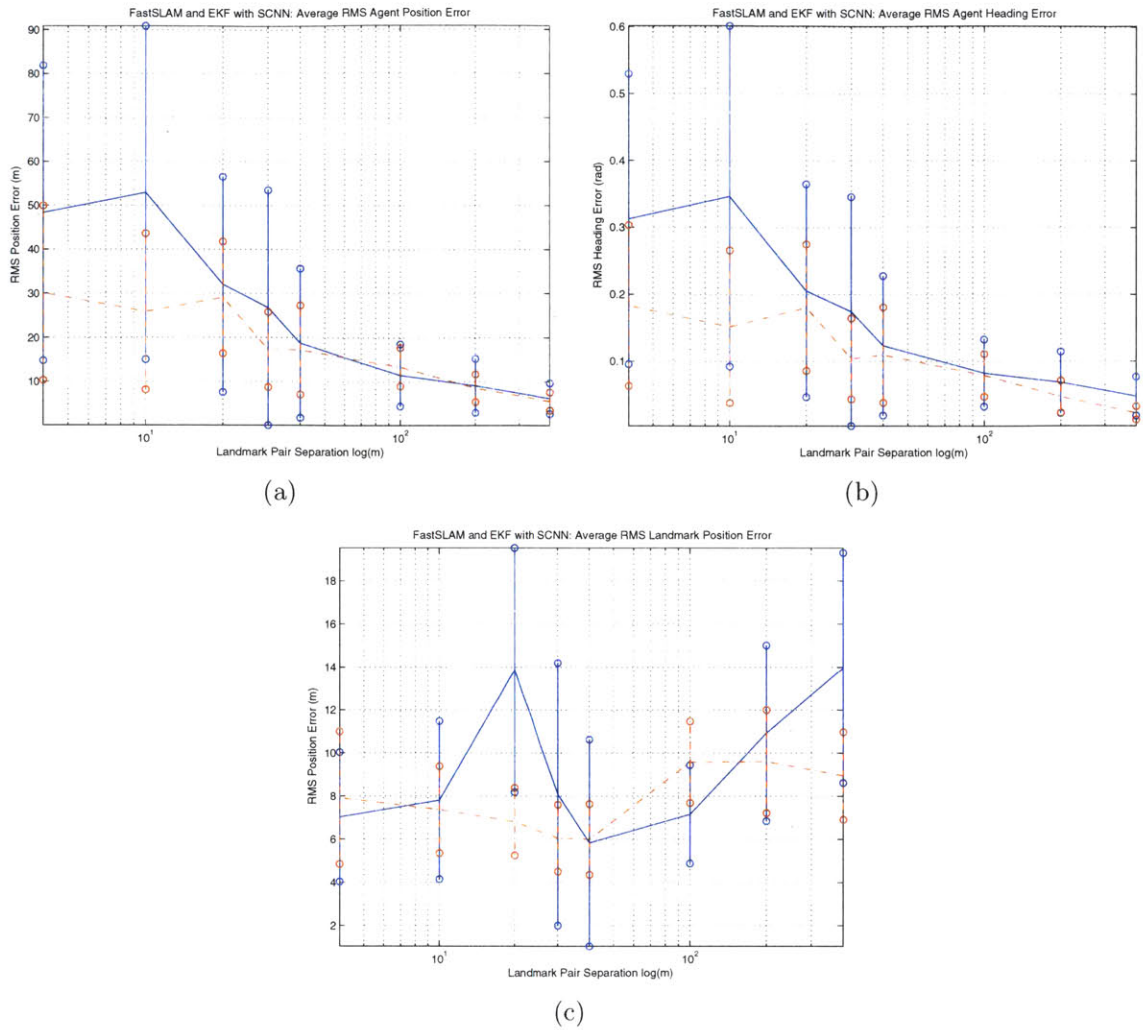


Figure 4-18: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Sequential Nearest Neighbor data association for the case where the landmarks are within the trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

localization, but it is likely related to the large number of measurement rejections occurring for these scenarios.

4.3.4 Results for Joint Compatibility Data Association

Figure 4-19 shows the data association error and measurement rejection characteristics for Joint Compatibility data association when paired with the EKF and FastSLAM filters for the situation where the landmarks are located inside of the agent's trajectory.

Examining subfigure 4-19(b) it is apparent that while there are some small changes in the measurement rejection characteristics of the filter-JC combination when compared with the *landmarks outside of trajectory* situation, these changes do not appear to be significant.

Comparing the results shown here for data association errors, in plot 4-19(a) to those seen previously in plot 4-9(a) for the *landmarks outside of trajectory* case, there is a drastic change in the average number of errors made for both filters. One of the most significant changes is with the EKF, where there are far fewer errors made in this case, no errors for landmark separations greater than 30 meters, and the number of errors decreases monotonically with landmark separation. Most importantly is the result that in this case Joint Compatibility, when paired with the EKF, makes fewer data association errors than either the average particle or the maximum likelihood particle when JC is paired with FastSLAM for landmark separations of 4 and 10 meters. This is in contrast to what was seen in the case where the landmarks were located outside of the agents trajectory.

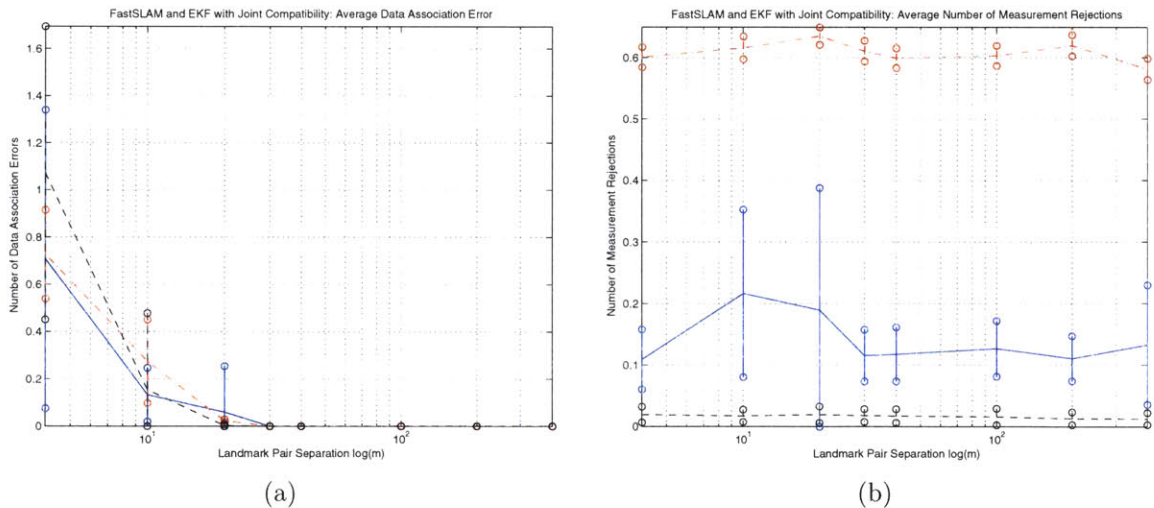


Figure 4-19: Average measurement rejection, (a), and data association errors, (b), for Joint Compatibility data association for both EKF and FastSLAM when the landmarks are within the trajectory. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

It is also notable that the FastSLAM maximum likelihood particle experiences average

data association errors of greater than one for the landmark separation of 4 meters. Such a large number of data association errors being made by the maximum likelihood particle would lead one to believe that the FastSLAM filter would perform quite poorly under this condition. However, by examining the filter performance plots in 4-20 it does not appear that these data association errors are having a negative effect, since the RMS state errors for the 4 meter landmark separation are in fact better than in the perfect data association case 4-16(c). This behavior is the result of incorrect data associations having a minimal effect when these errors occur for landmarks that are close together.

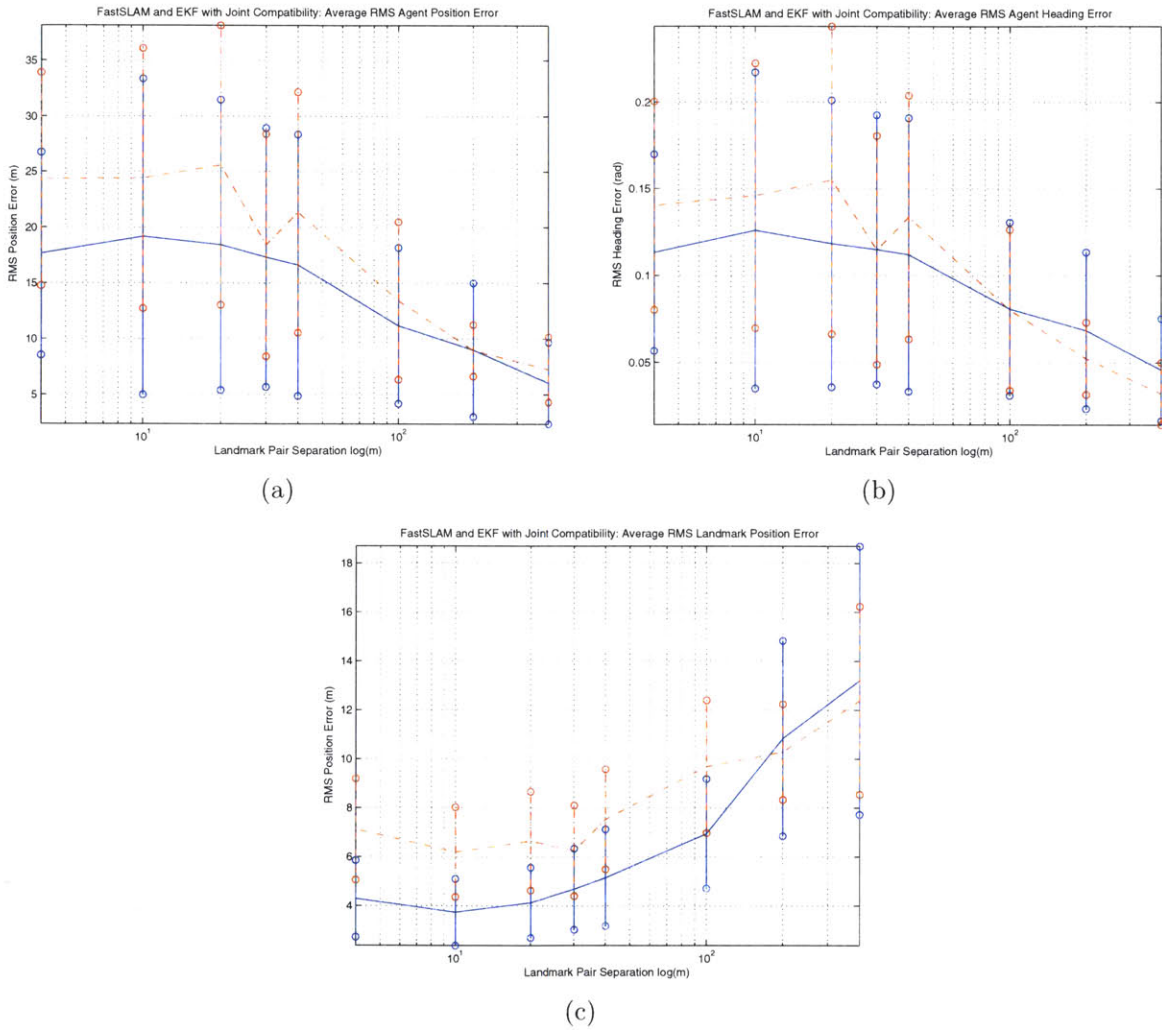


Figure 4-20: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Joint Compatibility data association when the landmarks are located inside of the trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

One other interesting behavior to notice in figure 4-19 is the increase in the RMS error that FastSLAM experiences in the change in landmark separation from 10 meters to 20 meters. While this may seem a bit out of the ordinary, it appears to be the result of the existence of data association errors at this landmark separation, which can be seen in subfigure 4-20(a). As very few data association errors occurred at this landmark separation for FastSLAM, however there is a noticeable effect on the performance of the filter. This indicates that in this situation there is an unhappy medium for the FastSLAM-JC marriage whereby there is a reasonable probability of a data association error occurring and such an error will cause a significant negative effect on the filter by inducing larger state errors for landmark separations near 20 meters.

4.3.5 Results for Maximum Likelihood Data Association

Examination of subfigure 4-21 indicates some interesting results for the Joint Maximum likelihood method of data association. First, in comparison to the results obtained using the JML method for the *outside of trajectory* configuration, a couple of differences can be noted. First, the performance when paired with the EKF is very good at determining the correct data associations, especially when compared with the EKF in the other landmark configuration situation and FastSLAM in this situation. A second difference to note is the substantial increase in the average data association error experienced by FastSLAM for the 4 meter landmark separation scenario, for the current situation. The measurement rejection plot, 4-21(b), is also informative, as it indicates that the average number of measurements rejected by this method is very minimal. This plot indicates that no measurements are rejected by the EKF or the maximum likelihood particle in FastSLAM. There are a few measurements rejected by the average particle in FastSLAM, but at less than 10^{-7} average rejections any effects are almost negligible. The most interesting part of this result is found when considering what was seen in figure 4-13(b), which showed the measurement rejection when using JML for the *outside of trajectory* landmark case. In that case there were minimal measurement rejections except for a single special case when the landmarks were separated by 200 meters. The fact that this is not seen here presents indirect evidence that the

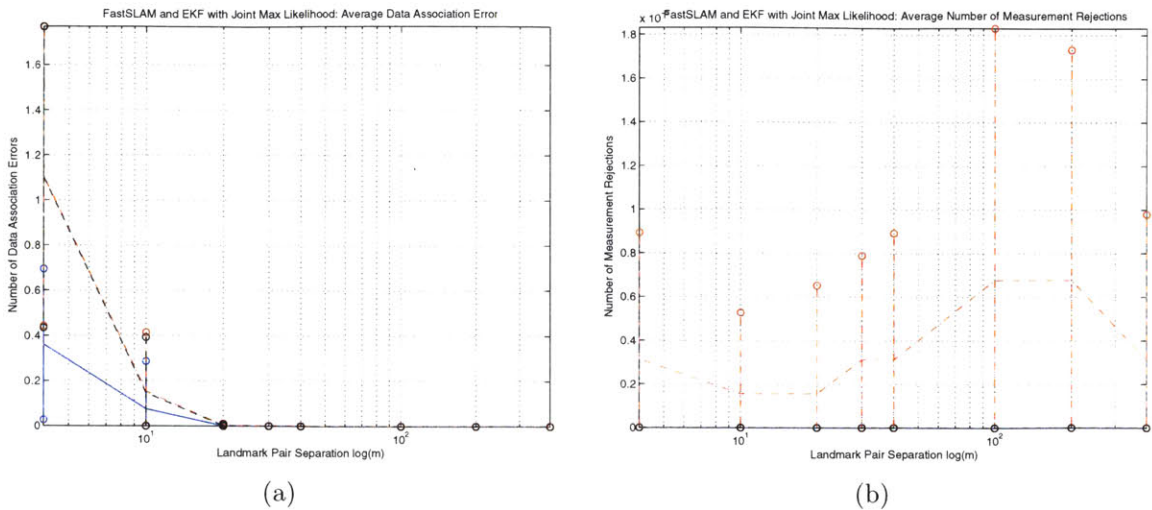


Figure 4-21: Average measurement, (a), and data association errors, (b), for Joint Maximum Likelihood data association for both EKF and FastSLAM when the landmarks are located inside of the trajectory. Here the EKF data is denoted using solid lines, and FastSLAM data is denoted with dash-dot lines for average particle behavior and with dashed lines for maximum likelihood particle behavior.

explanation presented to explain this phenomena was valid.

Figure 4-22 shows plots of the average RMS state errors for the JML-filter marriage case. The most interesting characteristic about these plots is how well behaved the EKF implementation is here, particularly in terms of its monotonic behavior with changing landmark separation. This behavior is very similar to what was seen in the perfect data association case for this situation. While the average RMS state errors are still larger than those in the perfect data association case they are very close to these values. These characteristics are an indication that the EKF and JML may be a good marriage.

The results for the FastSLAM filter and JML are also good, as they remain close to what was seen in the perfect data association case. In fact for certain situations the average RMS errors are smaller for the unknown data association case when JML is used. This is the result of JML acting as a pre-filter, which in effect alters the weighting of particles for the re-sample step in FastSLAM. In this case the effect of this turns out to be positive for reducing RMS state errors.

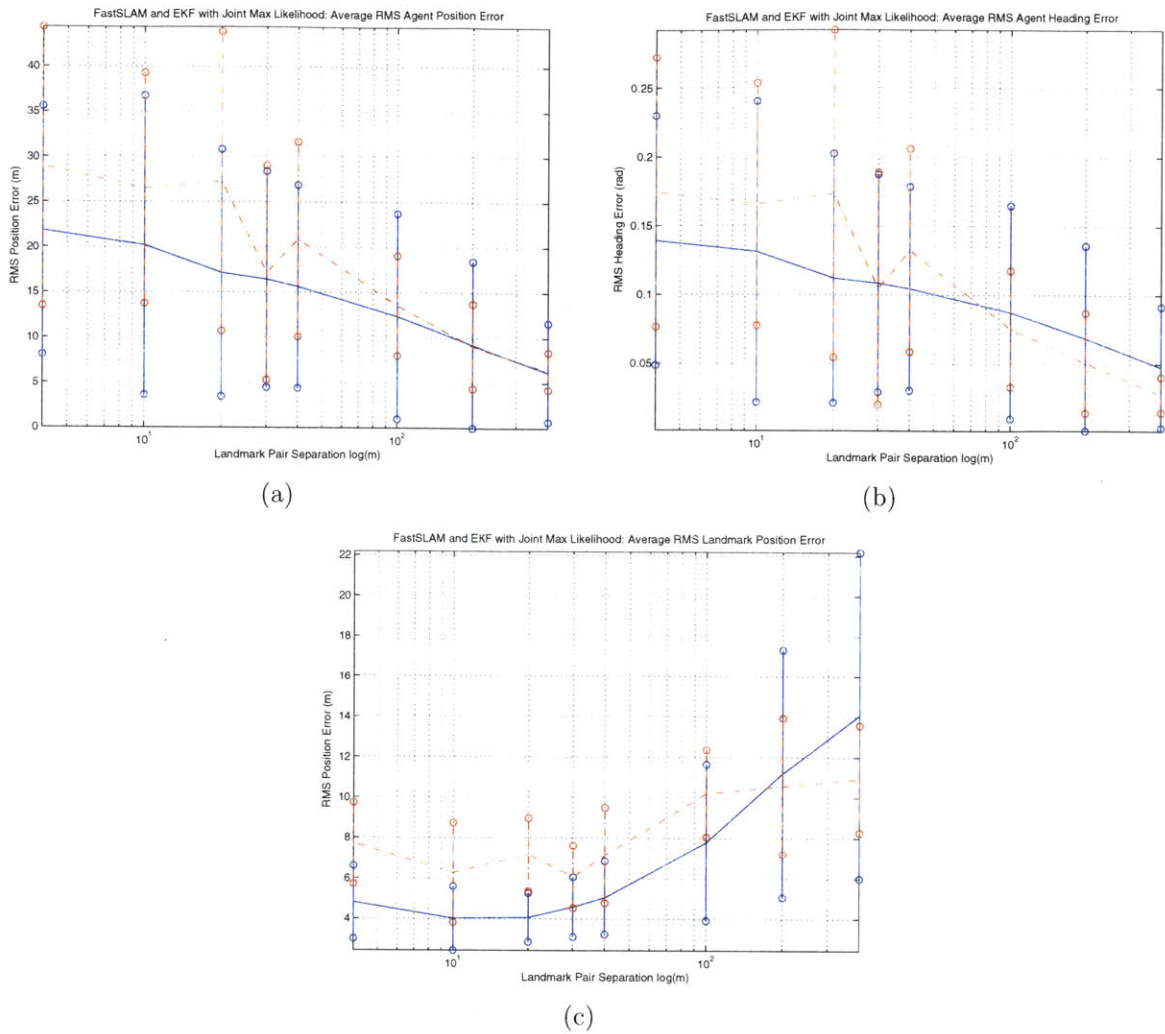


Figure 4-22: EKF, denoted using solid lines, and FastSLAM, denoted using dash-dot lines, performance using Joint Maximum Likelihood data association when the landmarks are located inside of the agents trajectory. Subfigures (a) and (b) show the agents RMS error in position and heading, while (c) shows the average RMS error in the estimated landmark positions.

4.3.6 A Quick Performance Comparison of the Various Filter-Data Association Marriages

Table 4.6 shows the average agent position error results for all filter-data association combinations, as well as the perfect data association case, for the landmarks inside of trajectory case. One of the most interesting results seen here is that the best performing cases are EKF based, which is in contrast to what was seen in the previous section 4.3. However, it should also be noted that the worst performing case here is also EKF based,

that is EKF-SCNN.

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
Perfect Data Association - EKF (m)	17.0	16.4	15.5	15.2	14.5	10.6	8.6
$\sigma_{Perfect-EKF}$ (m)	± 13.5	± 13.1	± 11.6	± 10.3	± 10.2	± 11.1	± 9.9
JC - EKF (m)	17.7	19.2	18.4	17.3	16.6	11.2	9.0
σ_{JC-EKF} (m)	± 9.1	± 14.2	± 13.0	± 11.6	± 11.8	± 7.0	± 6.0
SCNN - EKF (m)	48.3	52.9	32.1	26.8	18.7	11.5	9.1
$\sigma_{SCNN-EKF}$ (m)	± 33.6	± 38.0	± 24.5	± 26.7	± 16.9	± 7.1	± 6.1
JML - EKF (m)	21.8	20.2	17.1	16.4	15.6	12.3	9.2
$\sigma_{JML-EKF}$ (m)	± 13.8	± 16.6	± 13.7	± 12.0	± 11.2	± 11.4	± 9.3
Perfect Data Association - FS (m)	27.1	26.4	24.2	17.4	26.3	12	9.7
$\sigma_{Perfect-FS}$ (m)	± 16.0	± 14.1	± 11.8	± 4.6	± 12.2	± 3.8	± 3.7
JC - FS (m)	24.3	24.4	25.6	18.4	21.4	13.4	9.0
σ_{JC-FS} (m)	± 9.6	± 11.7	± 12.5	± 10.0	± 10.8	± 7.1	± 2.3
SCNN - FS (m)	30.1	25.9	29.1	17.3	17.2	13.4	8.6
$\sigma_{SCNN-FS}$ (m)	± 19.8	± 17.7	± 12.7	± 8.5	± 10.1	± 4.4	± 3.1
JML - FS (m)	28.9	26.5	27.3	17.1	20.9	13.5	9.1
σ_{JML-FS} (m)	± 15.5	± 12.8	± 16.6	± 11.9	± 10.8	± 5.5	± 4.7

Table 4.6: Performance comparison of filter-data association marriages for agent position error using an average RMS error metric, where the data is given in meters. In this case the landmarks are located inside of the agents trajectory.

It is notable that for this situation the EKF-JC is one of the best performing marriages, with the EKF-JML having similar performance. For the FastSLAM algorithm the JC method appears to give the best results, in terms of agent localization at landmark separations of 20 meters or less. For landmark separations of greater than 30 meters the SCNN method when paired with FastSLAM does the best. This is an indication of dependence of landmark separation on filter-data association marriage performance for FastSLAM filters.

Examination of table 4.7, which gives the average RMS landmark position errors, indicates that both EKF-JC and EKF-JML have similar performance for this metric, with Joint Compatibility having a slight edge. It is also notable that in this situation SCNN when paired with an EKF is clearly the worst performer. Interestingly enough a similar behavior to what was seen in the agent localization results are also seen here for the FastSLAM marriages, where JC gives the best results up until 30 meter landmark

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
Perfect Data Association - EKF (m)	3.1	3.5	4.0	4.5	5.0	6.8	10.3
$\sigma_{Perfect-EKF}$ (m)	± 0.86	± 0.93	± 1.1	± 1.3	± 1.7	± 3.7	± 6.5
JC - EKF (m)	4.3	3.7	4.1	4.7	5.2	7.0	10.9
σ_{JC-EKF} (m)	± 1.6	± 1.4	± 1.4	± 1.7	± 2.0	± 2.2	± 4.0
SCNN - EKF (m)	7.0	7.8	13.9	8.1	5.8	7.2	10.9
$\sigma_{SCNN-EKF}$ (m)	± 3.0	± 3.7	± 5.7	± 6.1	± 4.8	± 2.3	± 4.1
JML - EKF (m)	4.8	4.0	4.1	4.6	5.1	7.8	11.2
$\sigma_{JML-EKF}$ (m)	± 1.8	± 1.6	± 1.2	± 1.5	± 1.8	± 3.9	± 6.1
Perfect Data Association - FS (m)	5.6	5.3	6.8	6.1	8.5	8.5	11.5
$\sigma_{Perfect-FS}$ (m)	± 1.9	± 1.3	± 1.9	± 1.9	± 2.0	± 1.8	± 2.6
JC - FS (m)	7.1	6.2	6.6	6.3	7.5	9.7	10.3
σ_{JC-FS} (m)	± 2.1	± 1.8	± 2.0	± 1.8	± 2.0	± 2.7	± 2.0
SCNN - FS (m)	7.9	7.4	6.8	6.1	6.0	9.6	9.6
$\sigma_{SCNN-FS}$ (m)	± 3.1	± 2.0	± 1.6	± 1.6	± 1.6	± 1.9	± 2.4
JML - FS (m)	7.7	6.3	7.2	6.1	7.2	10.2	10.6
σ_{JML-FS} (m)	± 2.0	± 2.5	± 1.8	± 1.5	± 2.4	± 2.2	± 3.4

Table 4.7: Performance comparison of filter-data association marriages for landmark position error using an average RMS error metric, where the data is given in meters. In this case the landmarks are located inside of the agents trajectory.

separation after which SCNN clearly performs the best. This is further evidence of the dependence on landmark separation for which methodology works the best.

Similar trends in the average number of data association errors made (to what was seen for the first set of simulations) are seen for the landmarks inside of trajectory case which are shown in table 4.8. One of these trends is the relatively small number of errors made by SCNN for landmark pairs that are close together and a tail for which the algorithm continues to make ever-decreasing errors out to landmark separations of 40 meters. One noticeable difference in the results for this situation are the significantly smaller number of errors made by JML when paired with the EKF.

4.4 Summary of Results

The motivation for performing the sets of simulated experiments described above is multi-faceted. In part the motivation is to develop some level of intuition about the

	Landmark Pair Separation (meters)						
	4	10	20	30	40	100	200
JC - EKF	0.71	0.13	0.06	0	0	0	0
σ_{JC-EKF}	± 0.63	± 0.11	± 0.19	0	0	0	0
SCNN - EKF	0.33	0.25	0.32	0.07	0.02	0	0
$\sigma_{SCNN-EKF}$	± 0.17	± 0.19	± 0.39	± 0.16	± 0.1	0	0
JML - EKF	0.36	0.075	0.0005	0	0	0	0
$\sigma_{JML-EKF}$ (m)	± 0.33	± 0.21	± 0.002	0	0	0	0
JC - FS	0.73	0.27	0.023	0.0005	0.0002	0	0
σ_{JC-FS}	± 0.19	± 0.18	± 0.006	± 0.0003	± 0.0001	0	0
JC - FS for ML Particle	1.1	0.15	0.002	0	0	0	0
$\sigma_{JC-FSw/ML}$	± 0.62	± 0.32	± 0.005	0	0	0	0
SCNN - FS	0.57	0.31	0.017	0.003	0.003	0	0
$\sigma_{SCNN-FS}$	± 0.10	± 0.17	± 0.006	± 0.001	± 0.01	0	0
SCNN - FS for ML Particle	1.01	0.49	0.02	0.005	0.004	0	0
$\sigma_{SCNN-FSw/ML}$	± 0.37	± 0.42	± 0.04	± 0.005	± 0.01	0	0
JML - FS	1.1	0.15	0.003	0.0006	0.0004	0.0001	0
σ_{JML-FS}	± 0.66	± 0.26	± 0.007	± 0.0004	± 0.0003	± 0.0001	0
JML - FS for ML Particle	1.1	0.15	0.002	0	0	0	0
σ_{JML-FS}	± 0.67	± 0.25	± 0.006	0	0	0	0

Table 4.8: Performance comparison of filter-data association marriages in terms of data associations, here the metric used is the average number of data associations made by the EKF, the average FastSLAM particle and the maximum likelihood FastSLAM particle. In this case the landmarks are located inside of the agents trajectory.

operation of the various filter-data association algorithm marriages, as well as possibly determine if particular marriages appear to work better than others by analyzing the basic simulations considered. Additionally, it would be advantageous to be able to use this simple data to get a handle on what conditions may lead to poor performance on the part of both the data association algorithms and the filter-data association combinations.

4.4.1 Interpretation of Data Association Results

Based partially on the intuition developed in section 1.4, which helped to clarify those elements which can lead to data association ambiguity, as well as on the results that were presented in this chapter two sets of effects, first and second order, are presented here. The first order effects are those which appear to have the greatest effect on the

operation of the data association-filter marriages. The second order effects also have an effect on the operation of the data association algorithms, and therefore the SLAM filters, however the effects are less influential in overall performance and may not affect all methods in the same manner.

First Order Effects:

1. Filter performance under correct data association varies with the landmark location. In this case that effect is translated into a dependence on both total landmark separation and distance of landmarks from the agent.
2. All of the data association algorithms studied demonstrate a variation in performance with landmark pair separation, worse at smaller landmark separations and better at larger separations (with one exception).
3. The effect of an incorrect data association on the SLAM filter is also dependent on landmark pair separation, for the landmarks for which the data association occurs. For example, a data association error made between two landmarks that are very close together can have a very minimal effect on the filter. However, such a data association error made between two landmarks that are separated by a significant distance will cause significant errors in the filter.

Second Order Effects:

1. For data association performance, in terms of errors made, there is a small but noticeable dependence on the location of the agent along the trajectory for some of the data association algorithms.
2. Again for data association performance there appears to also be a small change in performance over time that the filter is in operation, which would be dependent on the correlations that are built-up in the filter over time and the information gain. This effect appears to be specific to the EKF-Joint Compatibility marriage, and cannot be an effect of agent location as it is not cyclical.

The first order effects discussed above become coupled in the implementation of a SLAM filter-data association algorithm marriage. The result of this is that at small

landmark separations many data association errors are made, however at these landmark separations data association errors do not have a significant negative effect on the operation of the filter. At the same time the filter does not naturally perform well, in a relative manner, at agent localization when the landmarks that are seen are close together. As the separation of landmarks becomes larger both the data association algorithms and the SLAM filter increase in performance. However, with this increase in performance the negative effect of an incorrect data association also increases. In the end this indicates that one would expect there to be some intermediate landmark separation where there is a non-negligible probability of a data association error occurring for which the effect of a data association error will be significant. It may be that the performance of the various marriages will be dependent on this effect. In the case discussed here it appears that this problematic landmark separation would appear to be in the 20 to 40 meter range.

4.4.2 Interpretation of Measurement Rejection Results

For the problem as set-up here it is an important result to determine whether or not significant measurement rejections have a negative effect on the SLAM filter. In essence, the question can be posed as, "Are the data association algorithms rejecting good measurements due to poor performance or are they in effect acting as a pre-filter by only allowing the best measurements to be processed by the SLAM filter?"

Further measurement rejection can be significant for real-world applications whereby the filter must recognize new, previously unseen, landmarks in the environment and initialize these landmarks as states in the SLAM filter. Measurement rejection is significant because for these applications the decision process used here to reject a measurement is the same as that used to initialize a new landmark. In essence significant measurement rejection could pose problems when the same method is applied to real-world applications.

4.4.3 Filter Dependent Effects

One additional consideration is how data association errors and measurement rejections are treated for the two SLAM filters, EKF and FastSLAM. For EKF this is relatively straightforward and well understood, as a measurement rejection or data association error experienced by the EKF are experienced by the filter as a whole. However, these choices are made on a per-particle basis for FastSLAM, and after these choices are made the particles are then sampled with replacement. Because of this process and the fact that both of these events affect particle weighting differently it is not straightforward how these occurrences affect the filter as a whole.

Chapter 5

Results from Application of Filter-Data Association Marriages to Experimental Data

In order to get a sense of how the various filter-data association marriages might operate in a real-world application, an attempt is made here to use these marriages to compute SLAM solutions on experimental data. The data used in this effort was publicly available via the website of Professor Eduardo Nebot at the University of Sydney and is titled the *Car Park* data set [1]. This data consists of time-stamped inertial measurements, GPS measurements, and range and bearing measurements of a set of beacons.

The results from this process will show that in this particular situation the Joint Maximum Likelihood and Sequential Nearest Neighbor methods of data association produce equivalent results when paired with an EKF implementation. Along with this, some interesting results will be presented for the implementation of Joint Compatibility data association for use with the EKF. Additionally, there is a short discussion of the problems this particular data set poses for implementations of the FastSLAM 1.0 algorithm and how the limitations of the FastSLAM 1.0 algorithm limit the use of the type of data association algorithms considered here.

5.1 Experimental Set-up

The *Car Park* data set was obtained by driving a small truck around the top-level of a flat, open parking ramp. This truck was equipped with a kinematic GPS system that gives quoted accuracy of 2 cm Circular Error Probable (CEP) [34], which will be used in establishing ground truth. For motion measurements the truck was equipped with inertial measurement units that have a one-sigma accuracy of 7 degrees steering and 0.7 meters-per-second in velocity [1]. Additionally, the truck had a SICK laser measurement unit mounted on the front of it, which supplied range and bearing measurements of the environment. The landmarks considered in this experiment were a set of steel poles of 6 cm diameter. These poles were also covered in reflection tape. This set-up made the problem of landmark extraction trivial, such that the range and bearing measurement of landmarks were quite accurate. In addition, the location of the steel poles was obtained using GPS giving a ground truth for the environment as well.

5.2 Limitations of FastSLAM 1.0 for the Car Park Data Set

There are a number of factors which may have led to the difficulties that were faced when using a FastSLAM 1.0 implementation to solve the SLAM problem posed by the *Car Park* data set. A leading candidate for causing problems in any FastSLAM 1.0 implementation is particle depletion. The problem of particle depletion is the process of limiting particle diversity, which can eventually lead to the situation where every particle in existence is an ancestor of a single particle. In the study performed here only 200 particles were used in the FastSLAM filter, and it is possible that increasing this number of particles would alleviate some of the problems associated with particle depletion. However, successful implementations of FastSLAM for processing experimental data have been reported using an equivalent number of particles [34]. Additionally, it should be noted that the particle depletion problem is one that has been addressed in the more recent version of this algorithm, FastSLAM 2.0 [30, 29].

As documented in one of the original papers on this filtering method, the leading cause of particle depletion is the use of highly accurate range and bearing measurements [30]. The *Car Park* data set gives one sigma range accuracies on the order of 0.1 meters and bearing accuracies on the order of 1 degree, which would indicate that the particle depletion problem might be a likely suspect for problems experienced.

However, the fact that other authors have successfully used this data set on FastSLAM implementations for the situation where data associations were known *a priori* and for the case of unknown data associations using a particle splitting multiple hypothesis method, [34], suggests that if particle depletion is part of the problem it is only part of the story. Additionally, the successful implementation of FastSLAM using the SCNN data association presented here suggests that some other negative synergy exists for the implementation of FastSLAM paired with Joint Maximum Likelihood and Joint Compatibility, considering that neither of these methods were successful when married to FastSLAM and applied to the *Car Park* data set.

This negative synergy appears to come from the combination of the very accurate SICK laser, range and bearing, measurements and the way in which FastSLAM 1.0 carries around its notion of state uncertainty. This becomes problematic in the data association methods discussed here because the data association choices are made on a per-particle basis and all depend on the measurement likelihood.

As shown in equation (2.40), the measurement likelihood is a function of uncertainty in the states involved in the measurement, as well as the uncertainty in the measurement. However, in FastSLAM the pose states for a given particle are considered to be known perfectly in the likelihood function. Additionally, the covariance matrix for each landmark known to a given particle, which is a measure of the uncertainty of the landmark position, is driven down very fast due to the accurate measurements and perfect pose information at the particle level. These operational peculiarities along with the fact that an individual particle is in fact going to have an inaccurate pose mean that the likelihood of many incoming measurements goes to zero.

Due to the effects of this, the use of Joint Maximum Likelihood or Joint Compatibility data association in this situation, leads to numerous non-existent landmarks being

initialized into the FastSLAM filter. This behavior in effect destroys both the quality of the map and causes the path estimate to diverge. It will be shown that a reasonable solution to the problem posed by the *Car Park* data set can still be found using FastSLAM 1.0 with the SCNN data association method.

5.3 Landmark Localization

The results of both map building and agent localization can be seen in figure 5-1, for SCNN 5-1(a), JML 5-1(b), and JC 5-1(c) data association methods when paired with an EKF implementation and for SCNN when paired with FastSLAM 5-1(d). In order to allow the filters to deal with both spurious measurements and new landmark initialization, an *ad hoc* method had to be employed. This method was accomplished by counting the number of times (or hits), each landmark in the filter is observed, as decided by the data association method. Only those features in the filter which are seen three or more times are included in the map the agent has built. All other features are considered extraneous and to be the result of spurious measurements. This method is reflected in the map information shown here.

By visual inspection of these plots it can be seen that the results obtained using an EKF when paired with SCNN and JML data association produce essentially identical results. Due to the fact that the range and bearing measurements being considered are so accurate with respect to landmark separations, along with the fact that this implementation of JML makes use of the individual compatibility requirement, as does SCNN in a similar manner, this is not an incredibly surprising result.

Examining the estimated landmark locations in subfigure 5-1(c) indicates that in this situation the EKF-JC combination has improved landmark estimates over the use of EKF with SCNN and JML. It can also be seen from this plot that this method has missed one of the known features. Something that is more difficult to observe is that the EKF-JC method has two estimates of a single landmark. These two situations are likely related.

Finally, the results for the FastSLAM-SCNN marriage can be seen in subfigure 5-1(d).

This plot indicates that the map built by this combination of SLAM algorithm and data association method is noticeably worse than the previously mentioned EKF based combinations.

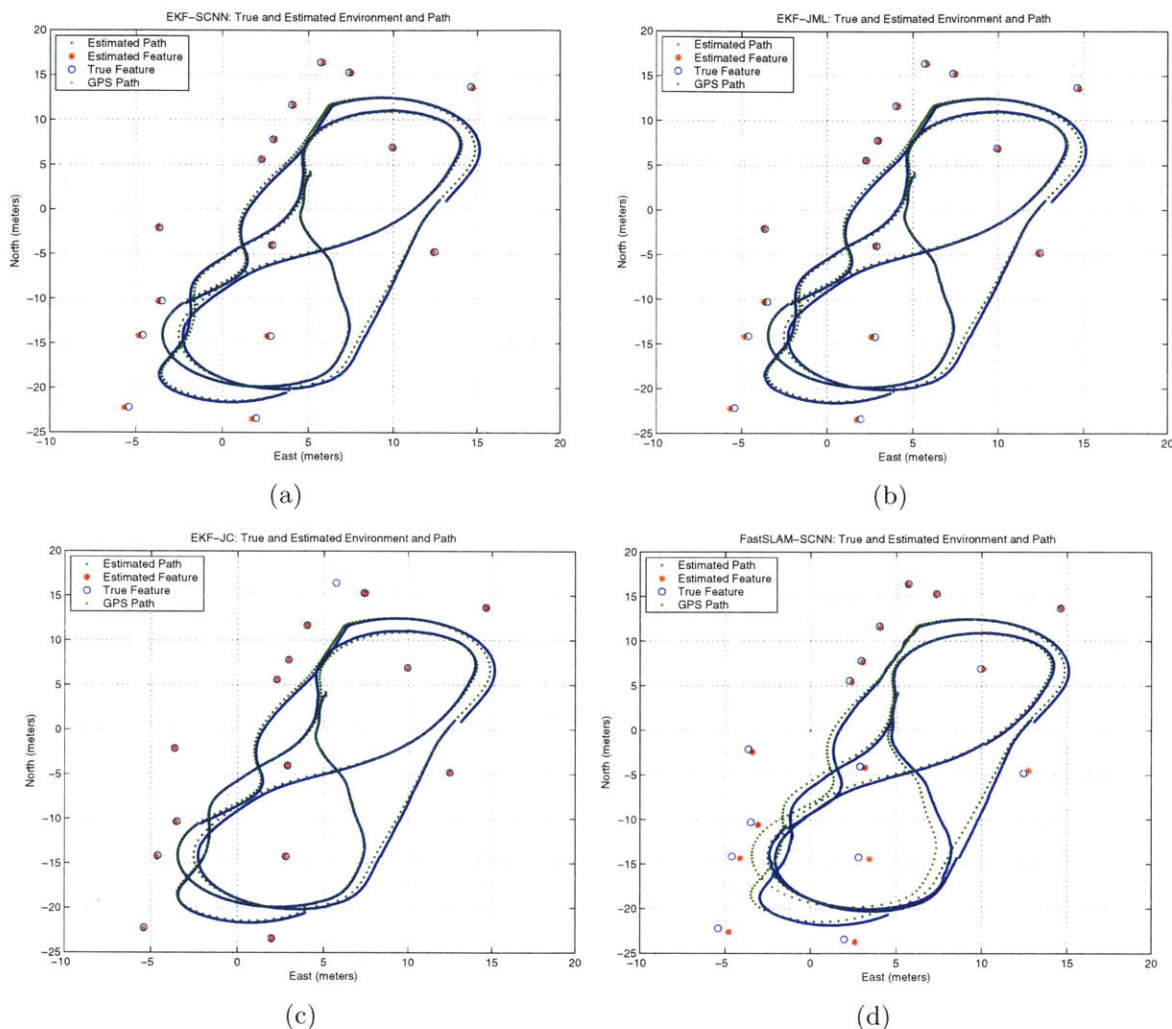


Figure 5-1: Plots of the path and map produced by the SLAM-Data Association combinations superimposed on the path and map given by GPS data. The results for each method are given by the following sub-figures (a) shows EKF-SCNN, (b) shows EKF-JML, (c) shows EKF-JC, and (d) shows FastSLAM-SCNN.

In order to quantify how well each of the four SLAM filter-data association algorithms perform map building, a set of statistics was generated and are shown in table 5.1. These statistics aim to describe the quality of the map generated by each. The statistics given are: N_F the number of features recognized, \bar{d}_F the average error in feature position estimate, σ_{d_F} the standard deviation of the error in the feature position

estimates, d_F^{max} the maximum feature position estimate error, and N_{extra} the number of features included in the filter but not in the map, i.e. those with fewer than 3 hits. Examination of table 5.1 verifies what was seen in subfigures 5-1(a) and 5-1(b), which indicated that the EKF paired with SCNN and JML have identical performance in map-building. One of the most interesting items to note from this table is the performance of the EKF-JC combination. The average landmark position estimate for this method is a two fold improvement over EKF with SCNN and JML, additionally there is a much smaller standard deviation in error and a noticeably smaller maximum error.

	N_F	\bar{d}_F (m)	σ_{d_F} (m)	d_F^{max} (m)	N_{extra}
SCNN - EKF (m)	15	0.1493	0.2661	0.0771	4
JML - EKF (m)	15	0.1493	0.2661	0.0771	4
JC - EKF (m)	14	0.0615	0.1640	0.0517	58
SCNN - FS (m)	15	0.3507	0.747	0.2329	4

Table 5.1: Map building characteristics of the SLAM Filter-Data Association combinations considered here. The descriptive variables used are as follows: N_F is the number of features recognized by the 3 hit criteria, \bar{d}_F is the average error in all of the feature position estimates, σ_{d_F} is the standard deviation of the error in all feature position estimates, d_F^{max} is the maximum error over all estimates, and N_{extra} is the number of extraneous landmarks included in the filter that do not meet the 3 hit criteria.

While these statistics indicate exemplary performance by EKF-JC, the values given by N_F and N_{extra} may be of concern. The fact that this method only recognizes 14 unique features is likely connected with it also having two estimates of one feature. This indicates that some combination of mistakes made by the JC data association algorithm and the EKF caused either multiple initializations of a single landmark, or the more likely scenario is that some portion of the measurements that came from the two estimate landmark were processed on a landmark that was initialized to be the true feature with a missing estimate. The second scenario is more likely given the evidence. Concern with regard to N_{extra} being 58, may be well founded, as such a large number of extraneous landmarks will slow down the operation of the SLAM filter, particularly traditional EKF implementations. This would be of particular concern for situations with a longer run-time and potentially larger number of extraneous features being added

to the filter. To alleviate this problem some systematic method would have to be employed, which could possibly be as simple as an enhanced method of what was used here, i.e. the 3 hit method.

In the end, this larger number of extraneous features created by the JC data association method is an indication that these methods can act not only to resolve data association ambiguity, but also as a pre-filter, where only the best measurements are accepted. This was also reflected in the simulated results as a large number of measurement rejections for the JC method.

The results for the FastSLAM-SCNN combination are interesting as they indicate similar performance to the EKF with SCNN and JML for feature recognition, having the same number of features recognized, 15, and the same number of extraneous features, 4. However, the performance results in terms of \bar{d}_F , σ_{d_F} , and d_F^{max} were the worst of the 4 methods. This is somewhat of a surprising result, but is likely a remnant of particle depletion issues discussed in section 5.2 that plague FastSLAM 1.0 for use with this data set. The mapping results seen here for FS-SCNN appear to be worse than those shown in [34], which has FastSLAM paired with the multiple hypothesis method. Unfortunately these statistics are not available for a comparison.

5.4 Agent Localization

The results from the four SLAM filter-data association methods considered are shown in figures 5-2 and 5-3. The first figure shows four plots, for the four methods considered, which depict North and East agent position errors with respect to the GPS estimated path versus time. The second figure shows the same information in a single plot, but here it is total error, or root-square of the North and East Error components, versus time.

The plots 5-2(a) and 5-2(b) are consistent with the results seen for landmark localization, where EKF-SCNN and EKF-JML produce essentially identical results. Additionally, the agent localization results for the EKF-JC combination are nearly identical to those of the other two EKF based marriages, which can be seen in subfigure

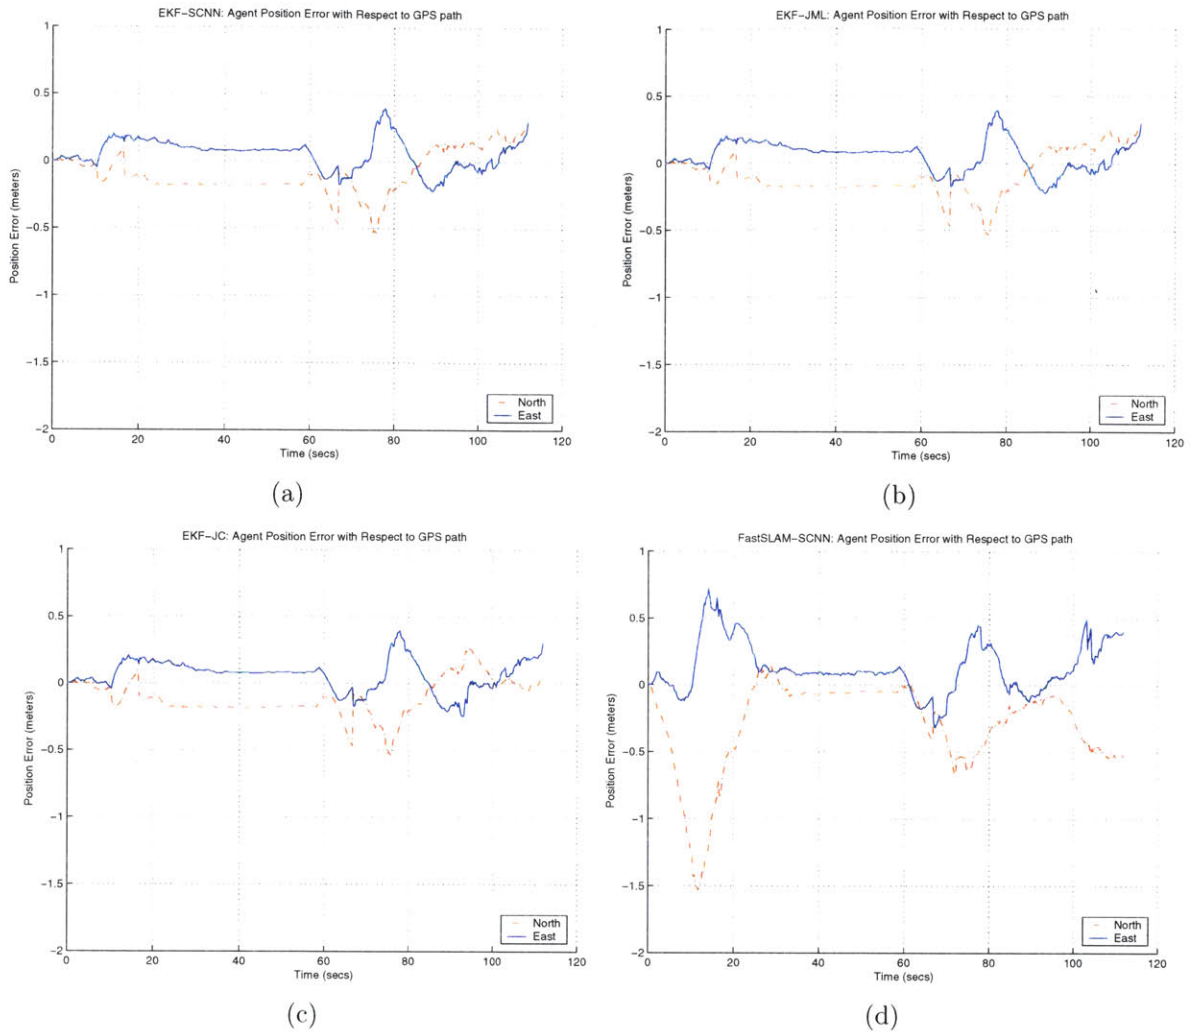


Figure 5-2: Plots of North and East agent position errors with respect to the GPS indicated location for the various SLAM Filter-Data Association marriages. The results for each method are given by the following sub-figures (a) shows EKF-SCNN, (b) shows EKF-JML, (c) shows EKF-JC, and (d) shows FastSLAM-SCNN.

5-2(c) and figure 5-3. In comparing these three methods it is difficult to say if one performs substantially better than the other two for agent localization.

Finally, an examination of subfigure 5-2(d) and figure 5-3 depict the agent localization performance for the FastSLAM-SCNN combination. It is clear that this method performs worse than the three EKF based ones. At the same time the localization error does remain well-bounded which is a positive note. In comparing these results to those presented for the FastSLAM multiple hypothesis method shown in [34] that method does appear to perform better at agent localization.

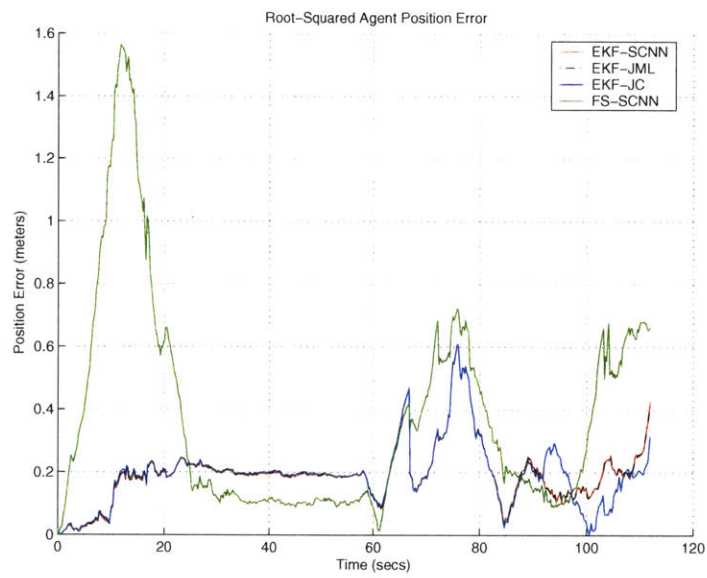


Figure 5-3: Comparison plot of the root-squared agent position error with respect to the GPS indicated location for the various SLAM Filter-Data Association marriages.

[This page intentionally left blank.]

Chapter 6

Conclusions and Future Work

One of the primary results of the work done here was that the data association methods considered demonstrate heavy dependence on landmark separation. In particular each of the three data association methods considered, JC, JML, and SCNN, appear to perform differently for a given landmark separation. For example, in the simulated results section it was demonstrated that, on average, the SCNN has the characteristic of making fewer data association errors when considering landmarks of small separation, in comparison to the results obtained using JC or JML. The SCNN method also had the property that the data association errors tapered off much more slowly at larger landmark separations than the other two methods. Additionally, it was generally found that both JC and JML have the opposite property where they make a relatively large number of errors at small landmark separations, but then drop off very quickly making few errors even at intermediate landmark separations. These results generally held true for both sets of simulations considered.

An additional observation that can be drawn from the simulated data, is that the occurrence of data association errors do not in and of themselves facilitate filter divergence for either the EKF or FastSLAM. The effects of these errors on filter performance is heavily dependent on the separation of the landmarks for which these errors occur. Additionally, the evidence suggests that even when data association errors occur for landmarks of various separations using either EKF or FastSLAM, they will not necessarily diverge. Both of these methods also appear to have some capability to

recover from data association errors, at least for landmarks that are separated by a reasonable distance.

In analyzing the filter performance for the simulated results it is difficult to conclude that a single filter-data association combination gives the best results. Essentially, there appears to be a great deal of variability. First, there is variation between the two simulation environments considered. In the case where the landmarks are located outside of the trajectory FastSLAM based combinations dominate the performance. However, in the case where the landmarks are located inside of the trajectory EKF based methods dominate. At the same time there is a large amount of variability of performance within the EKF and FastSLAM based methods, where for a given situation different data association methods give better results for the various landmark separations considered. One of the most intriguing results in the data presented in the simulated results chapter is the importance of measurement rejection properties of the data association algorithms discussed here. It can be seen in this data that for several cases where no data association errors are being made, the filter performance varies from that obtained under perfect data association. For these simulated experiments this variation in performance must be due to measurement rejections. The most important aspect of this observation is that in a number of instances the rejection characteristics appear to lead to significant improvements in filter performance.

In many ways the results presented for the experimental *Car Park* data set complement the simulation results presented. These results indicate that in situations, such as the one here, where highly accurate measurements are used along with good initial pose estimates, that the SCNN and JML (with individual measurement compatibility requirement) data association methods produce identical results when paired with an EKF SLAM solution. Another interesting result was the performance of the EKF-JC marriage as this method gave a map which was more accurate than that built by any of the other methods considered. This accuracy was compromised by the fact that it failed to map one of the fifteen objects in the environment and mapped one feature twice. One of the most intriguing aspects of the EKF-JC result is its apparent connection to measurement rejections, as the filter contained 58 extraneous objects, i.e., those features

initialized into the filter that were viewed fewer than three times. This is further evidence of the importance of the measurement characteristics of these methods, and that allowing these methods to reject measurements can lead to increased filter performance.

The results obtained using FastSLAM 1.0 with the experimental data were somewhat disappointing, however not completely unexpected as the particle depletion issues that potentially go along with this method are well known. The problem of particle depletion would not be unlikely in this situation because of the highly accurate instruments being used. At the same time the results that were obtained indicate that the problems faced were not just a result of particle depletion, but also that in the situation considered here two of the three data association methods may not have been effective when used in conjunction with FastSLAM 1.0. It may be that by employing some *ad hoc* methodologies, these data association algorithms could be retrofitted to work in this situation with FastSLAM.

6.1 Future Work

A number of issues that should be addressed in the area of data association and SLAM, as well as possible work that stems from the results presented here, are as follows:

1. Perform a similar comparison to what was done here using FastSLAM 2.0. The results from applying this implementation to the *Car Park* data set would be particularly interesting as this method claims to resolve particle depletion problems faced by the 1.0 version. This would also be interesting because it would indicate if the problems with using this data set were more a function of particle depletion or more a function of the data association algorithms being used.
2. Develop a better understanding of how the measurement rejection property of the data association algorithms interacts with the operation of the SLAM filters, i.e., if it improves or degrades performance.
3. Develop a more rigorous method for removing extraneous landmarks from the

EKF implementation, particularly something that could operate on-line. Such a tool may employ methodologies similar to what can be done using negative information in FastSLAM.

4. Additionally, while the application of the SLAM filter-data association marriages to the *Car Park* data set are useful, in many ways this data set does not offer an incredibly challenging data association problem due to the high level of accuracy given by the SICK laser. This being the case it would be very informative to apply these same marriages to a data set that supplies a more challenging SLAM-data association problem.

Bibliography

- [1] ACFR, experimental outdoor dataset. In *http://www.acfr.usyd.edu.au/homepages/academic/enebot/dataset.htm*.
- [2] D. Anguelov, R. Biswas, D. Koller, B. Limketkai, S. Sanner, and S. Thrun. Learning hierarchical object maps of non-stationary environments with mobile robots. In *Proceedings of the 17th Annual Conference on Uncertainty in AI (UAI)*, 2002.
- [3] T. Bailey, E.M. Nebot, J.K. Rosenblatt, and H.F. Durrant-Whyte. Data association for mobile robot navigation: A graph theoretic approach. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, 2000.
- [4] Tim Bailey. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, The University of Sydney, 2002.
- [5] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, Inc, Orlando, FL, 1988.
- [6] Yaakov Bar-Shalom and Thiagalingam Kirubarajan. Probabilistic data association techniques for target tracking in clutter. *Proceedings of the IEEE*, 92(3):536–557, 2004.
- [7] Michael Bosse, Paul Newman, John Leonard, and Seth Teller. An Atlas for scalable mapping. In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, 2003.

- [8] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [9] G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba. A computationally efficient solution to the simultaneous localisation and map building (slam) problem. In *Working Notes of ICRA'2000 Workshop W4: Mobile Robot Navigation and Mapping*, April 2000.
- [10] Gamini Dissanayake, Salah Sukkarieh, Eduardo Nebot, and Hugh Durrant-Whyte. The aiding of a low-cost strapdown inertial measurement unit using vehicle model constraints for land vehicle applications. *IEEE Transactions on Robotics and Automation*, 17(5):731–747, 2001.
- [11] A. Doucet, N. de Freitas, and editors N.J. Gordon, editors. *Sequential Monte Carlo Methods In Practice*. Springer, 2001.
- [12] Arnaud Doucet, Nand de Freitas, Kevin Murphy, and Stuart Russell. Rao-blackwellised particle filters for dynamic bayes networks. In *Proceedings of the 16th Annual Conference on Uncertainty in AI UAI*, 2000.
- [13] H. Durrant-Whyte, S. Majunder ans S. Thrun, M. de Battista, and S. Scheduling. A bayesian algorithm for simultaneous localization and map building. In *Proceedings of the 10th International Symposium of Robotics Research (ISRR '01)*, 2001.
- [14] Austin Eliazar and Ronald Parr. Dp-slam: Fast, robust simultaneous localization and mapping without predetermined landmarks. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003. IJCAI.
- [15] Austin Eliazar and Ronald Parr. Dp-slam 2.0. In *IEEE International Conference on Robotics and Automation (ICRA)*, New Orleans, LA, 2004. ICRA.

- [16] Martin Fisher and Robert Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [17] N. Freitas. Rao-blackwellized particle filtering for fault diagnosis. *IEEE Aerospace*, 2002.
- [18] Arthur Gelb, editor. *Applied Optimal Estimation*. The MIT Press, 1984.
- [19] José Guivant and Eduardo Nebot. Optimization of simultaneous localization and map-building algorithm for real-time implementation. *IEEE Transactions on Robotics and Automation*, 17(3):242–257, 2001.
- [20] José Guivant and Eduardo Nebot. Solving computational and memory requirements of feature-based simultaneous localization and mapping algorithms. *IEEE Transactions on Robotics and Automation*, 19(4):749–754, 2003.
- [21] D. Hähnel, W. Burgard, B. Wegbreit, and S. Thrun. Towards lazy data association in SLAM. In *Proceedings of the 11th International Symposium of Robotics Research (ISRR'03)*, Sienna, Italy, 2003. Springer.
- [22] Andrew H. Jazwinski. *Stochastic Processes and Filtering Theory*, volume 64 of *Mathematics in Science and Engineering*. Academic Press, Inc, New York, New York, 1970.
- [23] Z. Khan, T. Balch, and F. Dellaert. A rao-blackwellised particle filter for eigentracking. In *Proceedings of the Computer Society on Computer Vision and Pattern Recognition CVPR '04*, 2004.
- [24] Genshiro Kitagawa. Monte carlo filter and smoother for non-gaussian non-linear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.
- [25] J.J. Leonard and H.J.S. Feder. A computationally efficient method for large-scale concurrent mapping and localization.

- [26] J.S. Liu and R. Chen. Sequential monte carlo methods for dynamical systems. *Journal of the American Statistical Association*, 93(443):1032–1044, 1998.
- [27] Y. Liu and S. Thrun. Results for outdoor-SLAM using sparse extended information filters. In *IEEE International Conference on Robotics and Automation (ICRA)*, Taipei, Taiwan, 2003. ICRA.
- [28] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *IEEE International Conference on Robotics and Automation (ICRA)*, Taipei, Taiwan, 2003. ICRA.
- [29] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003. IJCAI.
- [30] Michael Montemerlo. *FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem with Unknown Data Association*. PhD thesis, Carnegie Mellon University, 2003.
- [31] J. Neira and J.D. Tardos. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics and Automation*, 17(6):890–897, 2001.
- [32] Paul Newman, John Leonard, Richard Rikoski, and Juan Tardós an José Neira. Towards robust data association and feature modeling for concurrent mapping and localization. In *10th International Symposium of Robotics Research (ISRR'2001)*, 2001.
- [33] P.M. Newman, J.J. Leonard, and R.J.Rikoski. Towards constant-time slam on an autonomous underwater vehicle using synthetic aperture sonar. In *Proceedings of the 11th International Symposium of Robotics Research (ISRR'03)*, Sienna, Italy, 2003. Springer.

- [34] J. Nieto, J. Guivant, E. Nebot, and S. Thrun. Real time data association for FastSLAM. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.
- [35] D.B. Rubin. *Bayesian Statistics 3*. Oxford University Press, 1988.
- [36] Stuart Russell and Peter Norvig, editors. *Artificial Intelligence: A Modern Approach*. Pearson Education, Inc, Upper Saddle River, NJ, 2004.
- [37] R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. Technical Report TR 4760 & 7239, SRI, 1985.
- [38] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In *Autonomous Robotic Vehicles*. Springer-Verlag, Orlando, FL.
- [39] S. Thrun. Particle filters in robotics. In *Proceedings of the 17th Annual Conference on Uncertainty in AI (UAI)*, 2002.
- [40] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002. to appear.
- [41] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 1998.
- [42] S. Thrun, D. Ferguson, D. Hähnel, M. Montemerlo, W. Burgard, and R. Triebel. A system for volumetric robotic mapping of abandoned mines. In *IEEE International Conference on Robotics and Automation (ICRA '03)*, Taipei, Taiwan, 2003. ICRA.
- [43] S. Thrun, D. Koller, Z. Ghahramani, H. Durrant-Whyte, and A.Y. Ng. Simultaneous mapping and localization with sparse extended information filters: Theory and initial results. *The International Journal of Robotics Research*, 23(7-8):693–716, 2004.
- [44] S. Thrun, C. Martin, Y. Liu, D. Hähnel, R. Emery-Montemerlo, D. Chakrabarti, and W. Burgard. A real-time expectation maximization algorithm for acquiring

multi-planar maps of indoor environments with mobile robots. *IEEE Transactions on Robotics and Automation*, 20(3):433–442, 2004.

- [45] S. Thrun, M. Montemerlo, D. Koller, B. Wegbreit, J. Nieto, and E. Nebot. Fastslam: An efficient solution to the simultaneous localization and mapping problem with unknown data association. *Journal of Machine Learning Research*, 2004. To appear.
- [46] Brian Williams and Nick Roy. Solving *Constraint Satisfaction Problems*. In *Course Notes from 16.413 @MIT*. 2004.