Adaptive Scan-Correlation for Mobile Robot Localization in Unstructured Environments

by Chris Scrapper
B.S. in Computer Science and Applied Mathematics, American University

A Thesis submitted to
The Faculty of
School of Engineering and Applied Science
of The George Washington University
in partial fulfillment of the requirements
for the degree of Master of Science

May 15th, 2011

Thesis directed by
Dr. James Hahn
Professor of Computer Science

Abstract of Thesis

"Adaptive Scan-Correlation
for Mobile Robot Localization in Unstructured Environments"

Mobile robot localization is fundamental to the development of more proficient robots capable of operating in complex, unstructured environments. However, many environments in which mobile robots are operating may be devoid of the static landmark and/or lack geometric primitives required for feature-based localization techniques. For such environments, scan–correlation have been employed. These approaches rely on the temporal correlation of unprocessed data to measure the relative displacement, i.e. motion, between successive scans obtained from a laser range finder. This research provides a comprehensive analysis of an adaptive scan–correlation technique that leverages previous effort to address real-time computational constraints and data association issues for mobile robots localization in complex, unstructured environments. This analysis uses a two-pronged approach to identifying and testing performance singularities to identify errors a specific system is prone to, how these errors impact the overall performance of that system, and how performance of that system compares with competing approaches. This analysis will will lead to the discovering of three testing scenarios for characterizing the performance of mobile robot localization.

Contents

List of Figures

1. Introduction

Mobile robot localization is fundamental to the development of more proficient robots capable of operating in complex, unstructured environments. ***Mobile robot localization*** is defined here as *the ability of a system to sense the environment, create internal representations of its environment, and estimate pose (where pose consists of position and orientation) with respect to a fixed coordinate frame*. This core competency will enable mobile robots to autonomously navigate an unknown environment while avoiding obstacles and potential hazards present in the environment. It will enable robots to estimate of where they are and where they have been. However, the performance of these systems varies greatly depending on the sensors employed and assumptions about the operational environment. The success of an approach relies on its ability to identify systematic and non-systematic errors and to compensate for the errors, bounding the uncertainty in the pose estimate.

The performance of localization largely depends on its ability to reliably accomplish two fundamental tasks. First, it must measure its surrounding environment accurately. Second, it must determine valid correspondences between observations reliably e.g., associating an object in one observation with its counterpart in another. The type and conditions of the environment strongly influence the systems ability to accomplish either task. Furthermore, subtle differences in similar environments can have very different effects on overall system performance .

In an effort to mitigate performance issues, an objective evaluation frameworks must be employed to quantify the performance in repeatable and reproducible testing scenarios that isolate potential failure conditions in a controlled environment. Often the evaluation of mobile robot localization is based on qualitative approaches that do not take into account how specific environmental conditions impact the performance of the system. While this type of analysis provides some indication of the overall performance, it does not allow researchers to understand which errors a specific system is prone to, how these errors impact the overall performance of that system, and how performance of that system compares with competing approaches.

1.1. Motivation

The primary motivation of this research is to foster the development of a robust mobile robot localization solution to improve the proficiency of mobile robots operating in complex, unstructured environments and to provide a comprehensive analysis of mobile robot localization to support the development of the standardized test methods for for the ASTM International Standards Committee on Homeland Security Applications; Operational Equipment; Robots (E54.08.01) [1]. Although the analysis presented in this thesis serves as the basis for developing standard test methods for robotic mapping and localization, the actual development and validation of the test methods is outside of the scope of this thesis. For more information regarding the initiative to develop performance standards for response robots, please refer to the Department of Homeland Security Response Robot Performance Standards webpage [2].

1.2. Contribution

This research makes substantial contributions in the field of mobile robots localization by addressing the technical shortfalls through two major thrusts:

1). ***Adaptive Scan-Correlation:*** Scan–correlation provides an alternative approach to mobile robot localization that uses the direct correlation of data obtained from a laser range finder to measure relative motion of the vehicle. However, real-time computational constraints and data association issues have afflicted the progress of this approach. Adaptive–scan correlation introduces a viable option based on a variant of the Iterative Closest Point (ICP) algorithm that leverages previous efforts by Zhang [45], Bailey [10], and Nuechter [32].

2). ***Performance Singularity Identification and Testing:*** The development of standard test methods and objective evaluation methodologies is essential to benchmarking performance of a system and fostering innovative solution to rectify divergent behavior. Performance Singularity Identification and Testing provides a systematic approach to quantifying performance and identifying failure conditions. This approach will facilitate the inter-comparison of results and introduce three test scenarios for characterizing the performance of localization techniques for mobile robots operating in complex, unstructured environments.

1.3. Thesis Overview

The organization of this thesis is as follows:

Chapter 2 provides an overview of the emerging mobile robot market, approaches to mobile robot localization, and highlights the need objective evaluation methodologies.

Chapter 3 provides an overview of the complete body of work presented in this thesis.

Chapter 4 describes the adaptive scan–correlation as a variant of the Iterative Closest Point (ICP) algorithm. The chapter starts by recapitulating basic ICP algorithm, then describes the modification made to improve the performance of this approach for mobile robot localization in unstructured environments.

Chapter 5 discusses the performance evaluation adaptive scan–correlation as a viable alternative to dead–reckoning. This evaluation decomposes errors to reveal the presence of performance singularities and helps to diagnosis the cause and impact of the divergent behavior on the overall performance of the system.

Chapter 6 presents the performance analysis of the adaptive scan-correlation technique as compared to the basic scan–correlation. This analysis highlights the development of test scenarios developed to challenge scan–correlation in environments with varying degrees of complexity and provides a quantitative analysis of the convergence characteristics of the competing approaches.

Chapter 7 summarizes the development and analysis of the adaptive scan–correlation technique.

Appendix A discuses a high-fidelity robot simulation testbed used to support this research.

2. Background

Automated Guided Vehicles (AGVs) are one of the oldest established markets for mobile robotics, valued over 900 million dollars [25]. These systems have played a significant role in the service robot industry due to their ability to operate in hazardous environments and to outperform humans in repetitive or mundane tasks that require a high-level of accuracy and repeatability. For these reasons, AGVs have enabled US industries to stay competitive in the global market place by increasing productivity, decreasing production costs, and assuring the safety of the existing human workforce while maintaining the integrity of the high-quality goods to consumers [24].

Historically, AGVs have relied on a centralized control strategy that exploit highly-structured indoor environments that have been deliberatively engineered with reference markers and dedicated pathways to keep vehicles free of obstacles [25]. The dependence on infrastructure, the lack of autonomy, and the cost associated with the installation and modifications severely limit the flexibility and adaptability of these systems. This has led many industry experts to speculate the most important developments in AGV technology will be the ability to "encounter any plant layout and create its own map by exploration, autonomously navigate to any point within the plant, and avoid any obstacles along the way. If the plant layout changes day to day (reconfiguration) or even momentarily (a ladder in an aisle to change a light bulb), they adapt intelligently"[18]. This vision of the next generation AGV will provide a cost-effective solution that will enable these system to be rapidly integrated into the existing workforce while increasing their versatility for various tasks.

Response robots is an emerging market for unmanned ground vehicles intended to assist emergency response personnel in a variety of application domains; such as Urban Search and Rescue (USAR), Explosive Ordnance Disposal (EOD), and Intelligence, Surveillance, and Reconnaissance (ISR). These systems serve as an extension of the operator to improve remote situational awareness and to provide assistive capabilities that minimizes the risk to responders and maximize the effectiveness and efficiency of a response in a tactical environment.

An example where response robots have been been making an enormous impact is in the Urban Search And Rescue (USAR) community. USAR is primarily concerned with the extrication of victims trapped in man-made structures such as collapsed buildings. During the initial phase of a structural collapse rescue, it is imperative that the first responders "size-up" the situation and establish an Incident Control System that allows information to flow regarding the nature of the problem. During this size-up, reconnaissance teams are dispatched to assess the magnitude of the situation, identify hazards, and locate areas that have the lowest cost-benefit ratio of danger to rescuers versus live victims [20]. Commonly, these environments contain unstable structures, undulating terrain, and hazardous or toxic debris. Recent advancements in mobile robot localization have provided responders with an invaluable tool, enabling them to safely and efficiently assess hazardous environments from a remote location.

In the future, robots will play an increasingly vital role in assisting humans in a variety of domains ranging from assisting humans with household chores to force protection supporting our serviceman operating in foreign countries around the world [21]. As mobile robots become more ubiquitous, their utility will rely on the ability of the robotic system to safely operate in dynamic, unstructured environments. Central to the realization of this vision of mobile robots is the system's ability to develop localization techniques that will enable a system to autonomously navigate an unknown environment, locate obstacles and hazards, and provide an estimate of where they are and where they have been.

To date there has been many different forms of localization proposed or implemented, some with greater success than others. The capabilities and limitations of each of these approaches vary significantly based on the requirements of the end-user, the operational domain, and the limitations of the on-board sensor suite. Understanding the strengths and weaknesses for each of the different forms of localization is essential for developing or selecting solution to meet the operational requirements within a specific domain.

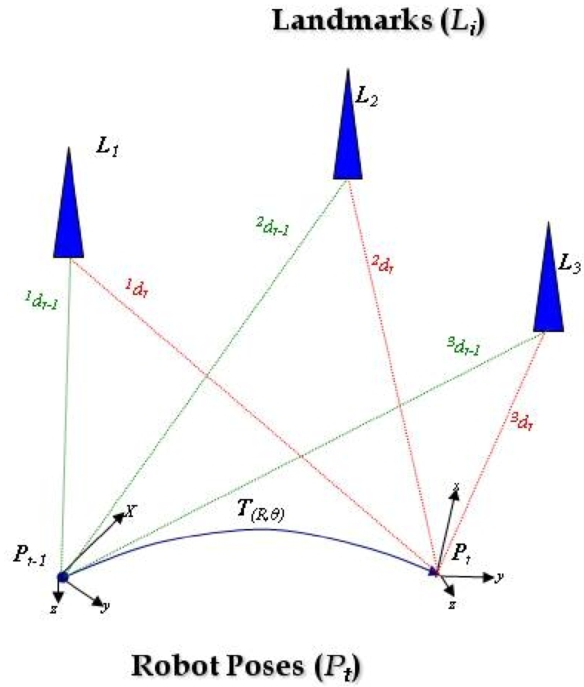

Figure 2.1: ***Mobile Robot Localization.*** In its most simple form mobile robot localization relies on mathematical principles and observation models for tracking landmarks to recursively compute a pose estimate consisting of the vehicles location and orientation.

Dead-reckoning is the most basic and common approach to localization. It is based on simple mathematical principles that 'advances" the pose estimate by recursively integrating motion to compute a new heading and the distance traveled. Dead-reckoning is favorable because it provides a simple, cost-effective solution that is self-contained. This encapsulation enables dead-reckoning solutions to generate pose estimates more efficiently and more frequently. The major drawback to dead-reckoning is two-fold [13]: 1). *sys-*

tematic and *non-systematic* dead-reckoning errors are hard to eliminate and 2). the recursive nature of the algorithms allows errors to propagate and accumulate in an unbounded manner, thus causing the pose estimate to diverge.

Another form of localization is feature-based approach [30, 19]. These approaches geometrically compute a pose estimate based on the recognition of distinct features, occurring naturally or artificially placed, in the environment. These approaches rely on the reliable acquisition and extraction of features from sensory data and its ability to exploit sensor data to accurately determine correspondences between perceived features with some navigational map [27]. While these methods, in general, provide an accurate pose estimate, they require either engineering the environment to provide an adequate set of features, or efficient recognition of features to use as landmarks [23]. In addition, these methods often rely on geometric primitives or models that are not guaranteed to exist in all environments.

In lieu of the feature–based approaches, scan-correlation techniques have been employed as another form of localization for unstructured environments. These approaches eliminate the need to decide what constitutes a feature by minimizing the discrepancies between the raw sensor data and a model of the environment. Using a maximum likelihood alignment to find the best fit between two sets of data points, scan-correlation is capable of providing a computationally efficient pose estimates in complex, unstructured environments. Examples of scan-correlation techniques in the literature are found in Lu and Milios [28], Censi [15], and Segal [37]

In the late 1980s, Smith *et al.* [38] introduced a new approach to localization and mapping that relied on the correlation of spatial relationships between the vehicle's pose and features in the environment. Later formalized by Leonard and Durant-White [27], Simultaneous Localization and Mapping (SLAM) uses statistical methods to fuse high-frequency predictions of vehicle maneuvers with low-frequency observation of the external environment to bound the errors in the pose estimate. Over the past decade, many implementations of SLAM have incorporated scan-correlation techniques as part of the SLAM framework to improve the versatility and accuracy these approaches [33, 31, 22].

Although recent advancements in mobile robot localization has improved the flexibility, utility, and survivability of overall system, in large these systems have failed to achieve a technology readiness level suitable for fielded systems deployed in their respective operational domains. Currently, there is no way to quantitatively measure the performance of localization and mapping techniques against user-defined requirements. Additionally, no consensus exists on what objective evaluation procedures need to be followed to deduce the performance of localization and mapping techniques in different domains. The lack of reproducible and repeatable test methods precludes researchers from working towards a common goal. It prevents the communication and comparison of the results, which prevents researchers from leveraging previous work and inhibits technology transfer from the "drawing board"

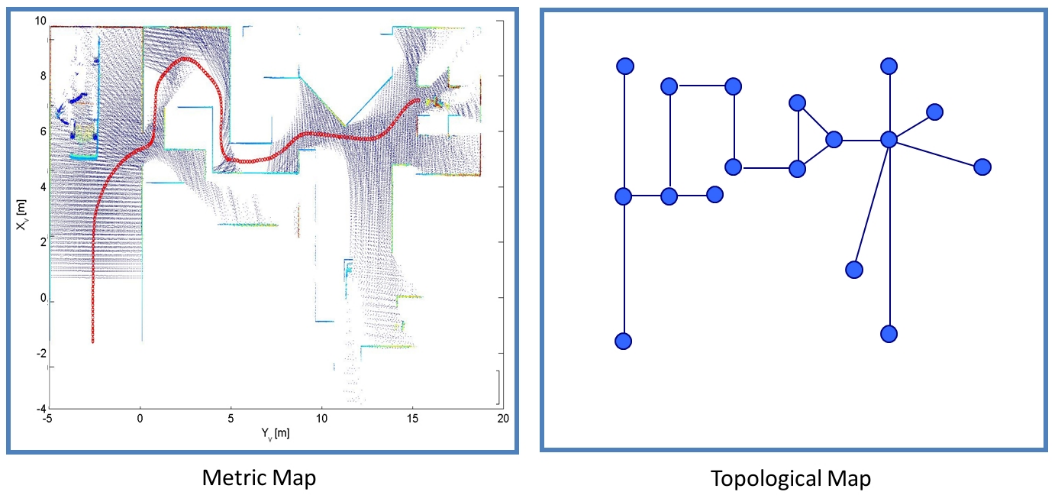

Figure 2.2: ***Robotic Mapping Paradigm*** The robotic mapping paradigm can be classified into two fundamental approaches; metric maps and topological maps. Metric maps provide a geometric representation where spatial relationships between objects are consistent with corresponding objects in the actual environment. Topological maps represent the environment as a connectivity graph of significant places or objects (nodes) and relationships between nodes (arcs) [40].

to the field.

Some researchers have recognized the need for the objective evaluation methodologies for assessing the performance of localization techniques for mobile robots. The common practice for characterizing the performance of these systems has been through the analysis of the map (or image) generated by these systems [42, 41]. Arguably, the most common mapping paradigm employed for robotic navigation is the metric mapping paradigm, shown in Figure 2.2. This intuitive mapping paradigm provides a representation where the spatial relationship between any two objects in the map is proportional to the spatial relationship of the corresponding objects in the actual environment [40]. Therefore, assessing the quality of metric maps is based on the spatial consistency of features, such as walls and hallways, between the map produced by the robot and the ground truth map of the actual environment.

While the analysis of global metric maps provide some indication of the overall performance, it does not allow researchers to identify problematic situations or how the propagation of errors impacts the performance of the overall system. Error propagation and sensitivity to performance singularities idiosyncratic to most localization techniques suggests the need to quantify the local (or regional) consistency of areas within the map, as well as the global consistency of the overall map. For example, some researchers [26, 16] have proposed a more introspective approach that examines the intrinsic relationships between the artifacts produced by the navigation solution, e.g. pose estimate and scans from the laser rangefinder, and models of the environment.

Although contributions by individual researchers are important steps to overcoming technological barriers impeding the development and fielding of localization and mapping solutions, a concerted effort among all interested parties is crucial. Test methods establishes a confident connection between developers and consumers regarding the expectations and performance objectives of robotic technologies. This is a cardinal step in fostering innovation and assessing the maturity of evolving technologies. They provide the basis for developers to understand the objective performance of a system and allows consumers to confidently select systems that will meet their requirements.

Test methods consist of well-defined testing apparatuses, procedures, and objective evaluation methodologies that isolate particular aspects of a system in known testing conditions [7]. The development of test methods start with a comprehensive analysis of the application domain to identify requirements with associated metrics and the range of performance, starting from a baseline threshold to the objective “best-case” performance. This analysis provides the basis for developing test methods and testing scenarios that are intentionally abstract so as to be repeatable across a statistically significant set of trials and reproducible by other interest parties.

The Department of Homeland Security (DHS) Science and Technology (S&T) Directorate has initiated an effort with the National Institute of Standards and Technology (NIST) to develop comprehensive set standard test methods and associated performance metrics to quantify key capabilities of emergency response robots as part of the ASTM International Standards Committee on Homeland Security Applications; Operational Equipment; Robots (E54.08.01) [2]. The set of test methods being developed focus on addressing responder-defined requirements for robot mobility, manipulation, sensors, energy, communications, mapping, human-robot interfaces, logistics and safety for response robots. The analysis conducted in the rest of this thesis is an effort to identify testing scenarios and metrics for developing test methods for robot localization and mapping.

3. Body Of Work

The primary focus of the research presented in this thesis is to foster the development of a robust mobile robot localization solution to improve the proficiency of mobile robots operating in complex, unstructured environments. This thesis attempts to address the technical shortcomings through two major thrusts. First, is the development of an adaptive scan-correlation technique, based on a variant to Iterative Closest Point (ICP) algorithm, to support mobile robot localization in unstructured environments. Second, is the development of three test scenarios for characterizing the performance of localization techniques using a two-pronged approach to identify errors a specific system is prone to, how these errors impact the overall performance of that system, and how performance of that system compares with competing approaches.

Scan-correlation is a popular form of localization due its ability to provide a cost-effective, computationally efficient solution that minimizes the dependence on feature models. Scan–correlation relies range image registration techniques to measure the relative displacement, i.e. motion, between successive scans obtained from a laser range finder. The major performance issue for scan–correlation is its ability to reliably determine valid correspondences, e.g. associating an object in one observation with its counterpart in another. In the field of mobile robot localization, there are two concepts of deployment for scan-correlation: 1). a sensor-based dead-reckoning solution commonly referred to *Scan-Matching* [28] or 2). the observation model within *Simultaneous Localization and Mapping* (SLAM) solutions [31].

Chapter 4 presents an adaptive scan–correlation technique that leverages previous efforts by Zhang [45], Bailey [10], and Nuechter [32] to address real-time computational constraints and data association issues endemic to the scan–correlation in complex, unstructured environments. The adaptive scan–correlation approach modifies the error metric [10], adds a reject phase for discarding invalid correspondences [45], and employs a new search strategy [32] to improve the convergence characteristics and performance of techniques based on the basic ICP algorithm. Since the type and conditions of the environment strongly influence the systems performance, it is essential to develop understand these implications and how they impact the overall performance of the mapping system.

Chapters 5 and 6, presents the comprehensive analysis for characterizing the performance of mobile robot localization. The premise behind this analysis is the identification of ***performance singularities***, or *the point where a system fails to be well-behaved.* This is essential for not only quantifying performance, but understanding how errors arise and speculating strategies for mitigating divergent behavior. This approach contains two distinct steps to first identify and then test performance singularities.

The first step in this process is the *performance evaluation*, which evaluates the performance of adaptive scan–correlation at the system-level. In Chapter 5, performance evaluation is used to assess the adaptive scan–correlation technique as a viable alternative to traditional dead–reckoning approaches. The evaluation of the performance will use ground truth information to decompose errors arising in the pose estimate of the competing approaches and will identify and diagnosis divergent behavior.

The second step is the *performance analysis*, which assesses adaptive scan–correlation at the algorithmic level. In Chapter 6, performance analysis is used to compare the performance of adaptive scan–correlation to a basic scan–correlation based on ICP. This analysis uses ground truth to compare the convergence characteristics of the two approaches in three test scenarios designed to challenge the systems ability to determine valid correspondences.

4. Adaptive Scan Correlation

Many environments in which mobile robots are currently operating may be devoid of the static landmark and/or lack geometric primitives required for feature-based localization techniques. For such environments, scan–correlation provides a viable alternative. One of the most common methods for developing scan–correlation is the use of a range image registration technique known as the Iterative Closest Point (ICP) algorithm. ICP enables scan–correlation to measure motion by computing the relative displacement that occurred between successive scans. However, error in the association of data can impede performance.

Adaptive scan–correlation is an ICP variant designed to address the data association problem and improve the convergence characteristics of the basic ICP. The remainder of this chapter is dedicated to the development of adaptive scan–correlation. It will start by recapitulating the ICP algorithm and explain how and why adaptive scan–correlation modified ICP to support mobile robot localization in complex, unstructured environments.

The Iterative Closest Point (ICP) algorithm refers to a class of fine range image registration techniques [35] that is widely recognized as the predominant method for the geometric alignment of 3D models [34]. Its popularity stems from its ability to use the direct correlation of unprocessed data, which eliminates the need to define feature models and avoids misclassification due to imperfect sensor models. This simplifies the registration process and provides a computationally efficient method for finding the maximum likelihood alignment between two sets of 3D point data, e.g. models.

$$\min_{\boldsymbol{\mathcal{T}}} \sum_i \|\boldsymbol{\mathcal{T}}\,\boldsymbol{p}_i - \hat{\boldsymbol{q}}_i\|^2 \tag{4.1}$$

The basic ICP algorithm computes the maximum likelihood alignment through an iterative process that minimizes the mean-square distance between two sets of data using the objective least-squares function shown in Equation 4.1. This iterative process starts with two sets of 3D point data, $\boldsymbol{\mathcal{S}}^{ref} : \{\boldsymbol{q}_j\}$ and $\boldsymbol{\mathcal{S}}^{obs} : \{\boldsymbol{p}_i\}$, and a transform that estimates the relative displacement between the two sets of data, $\boldsymbol{\mathcal{T}}$. At each iteration, $\boldsymbol{k}$, ICP refines $\boldsymbol{\mathcal{T}}$, based on correspondences found between the two sets of data at each iteration, until a termination condition has been met.

The basic ICP algorithm can be partitioned into three distinct phase; a *matching* phase, a *minimization* phase, and a *termination* phase. A formal description of each phase of the basic ICP algorithm is as follows:

1. ***Point-to-Point Matching***

   For each point $\boldsymbol{p}_i \in \boldsymbol{\mathcal{S}}^{obs}$, find its nearest neighbor, $\hat{\boldsymbol{q}}_i \in \boldsymbol{\mathcal{S}}^{ref}$ under the current transform $\boldsymbol{\mathcal{T}}_k$. Where $\hat{\boldsymbol{q}}_i$ is defined as:

$$\hat{\boldsymbol{q}}_i \stackrel{\triangle}{=} \arg\min_j \|\boldsymbol{\mathcal{T}}_k \boldsymbol{p}_i - \boldsymbol{q}_j\|^2 \tag{4.2}$$

2. ***Minimization using Singular Value Decomposition***

   Compute the incremental transform, $\boldsymbol{\mathcal{T}}$, that minimizes the mean-square distance between the correspondences found in Phase 1. using the Singular Value Decomposition method as follows [9] :

   - Compute the centroids of the correspondences.

$$\boldsymbol{p}_c = \frac{1}{N^{obs}} \sum_{i=1}^{N^{obs}} \boldsymbol{p}_i \tag{4.3}$$

$$\hat{\boldsymbol{q}}_c = \frac{1}{N^{obs}} \sum_{i=1}^{N^{obs}} \hat{\boldsymbol{q}}_i \tag{4.4}$$

   - Calculate the correlation matrix, $\boldsymbol{H}$

$$\boldsymbol{H} = \sum_{i=1}^{N^{obs}} (\boldsymbol{p}_i - \boldsymbol{p}_c)(\hat{\boldsymbol{q}}_i - \hat{\boldsymbol{q}}_c)^T \tag{4.5}$$

   - Find the Singular Value Decomposition (SVD) of $\boldsymbol{H}$

$$\boldsymbol{H} = \boldsymbol{U}\boldsymbol{\Omega}\boldsymbol{V}^T \tag{4.6}$$

   - Compute transform, $\boldsymbol{\mathcal{T}}$. Note $\boldsymbol{\mathcal{T}}$ is defined as $\boldsymbol{\mathcal{T}}\boldsymbol{p} \stackrel{\triangle}{=} \boldsymbol{R}\boldsymbol{p} + \boldsymbol{t}$, where $\boldsymbol{R}$, is a 3x3 rotational matrix and $\boldsymbol{t}$ translational vector.

$$\boldsymbol{R} = \boldsymbol{V}\boldsymbol{U}^T \tag{4.7}$$

$$\boldsymbol{t} = \hat{\boldsymbol{q}}_c - \boldsymbol{R}\boldsymbol{p}_c \tag{4.8}$$

3. ***Termination***

   If the magnitude of the translation recovered in Phase 2., e.g. $\|\boldsymbol{t}\|^2$, falls below a predefined threshold, say $\boldsymbol{\tau}$, or if the predefined limit to the number of iterations is exceeded then terminate, else iterate.

Although the basic ICP algorithm is guaranteed to monotonically converge to a local minimum, the convergence characteristics rely on the determination of valid correspondences. The association of data between the two data sets is complicated by spurious points, occlusions, and outliers. This can lead to correspondence errors and jeopardize the integrity of the algorithm. Although not addressed here, it is important to note that a good initial estimate of $\boldsymbol{\mathcal{T}}$, say $\boldsymbol{\mathcal{T}}_0$, is an important step to assure convergence to the global minimum.

Since first introduced by Besl and McKay [12] and Chen and Medioni [17], there has been many variants to the basic ICP algorithm proposed to improve the efficiency and the determination of valid correspondences [34]. New search strategies using kd-trees have reduced the time required for finding matches between the two sets of data, making it a viable candidate for real-

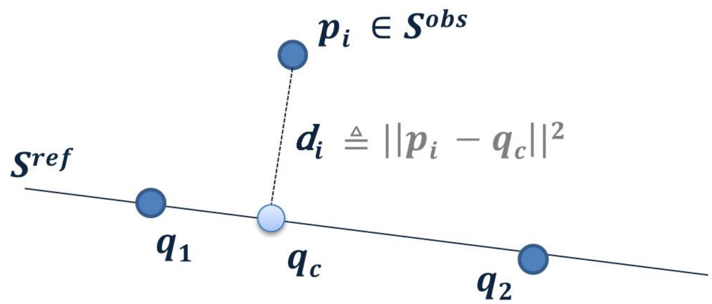

Figure 4.1: ***Point-to-Line Matching.*** The error metric modifies the traditional point-to-point error metric to approximate the real distance between a point and a line, represented as a series of points.

time applications [32]. Alternative error metrics [10, 15] and robust method for rejecting invalid correspondences have improved the convergence characteristics of the basic ICP algorithm [45].

The adaptive scan–correlation technique introduces an ICP variant leveraging previous efforts by Zhang [45], Bailey[10], and Nuechter [32] to address real-time computational constraints and data association issues to support mobile robots localization in complex, unstructured environments. The remaining two section in this chapter will discuss the modification made to the basic ICP algorithm and provides insight into how these modifications address the inherent shortcoming of ICP to rectify divergent behavior.

4.1. Point-to-Line Matching

Point-to-point matching employed by the basic ICP algorithm treats scan data as a set of discrete locations. It does not account for noise inherent to the sensor data or how slight perturbations between sensor reading can affect the locations of the hit points. It is obvious to see how this can lead to spurious point matching, which can in turn influence the convergence characteristics of the basic ICP algorithm.

$$\boldsymbol{q}_{1,2} \leftarrow \min_{j,l} \|\boldsymbol{\mathcal{T}}_k \boldsymbol{p}_i - \boldsymbol{q}_{j,l}\|^2 \tag{4.9}$$

$$\boldsymbol{q}_c = \boldsymbol{q}_1 + \frac{(\boldsymbol{p}_i - \boldsymbol{q}_1) \cdot (\boldsymbol{q}_2 - \boldsymbol{q}_1)}{\|\boldsymbol{q}_2 - q_1\|^2} (\boldsymbol{q}_2 - \boldsymbol{q}_1) \tag{4.10}$$

Adaptive scan–correlation addresses this shortcoming by assuming data obtained from a laser range finder represents a surfaces rather than a set of discrete locations. It modifies the error metric to approximate the real distance between a point in the observation model, $\boldsymbol{p}_i \in \boldsymbol{\mathcal{S}}^{obs}$, and a surface defined by $\boldsymbol{\mathcal{S}}^{ref}$. Figure 4.1 illustrates point-to-line matching, where $\boldsymbol{q}_c$ is the corresponding point in $\boldsymbol{\mathcal{S}}^{ref}$ for $\boldsymbol{p}_i \in \boldsymbol{\mathcal{S}}^{obs}$. Point-to-line matching first determines $\boldsymbol{p}_i$'s two nearest neighbors, $\boldsymbol{q}_{1,2} \in \boldsymbol{\mathcal{S}}^{ref}$, formally expressed in Equation 4.9. Using $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$ to approximate a line, $\boldsymbol{q}_c$ can be computed using the Equation 4.10.

4.2. Adaptive Thresholding

The least–square objective function used by the basic ICP algorithm (Equation 4.1) has no means to address uncertainties in the sensor data and to evaluate the validity of correspondences. This means all correspondences will be considered during the registration process, increasing the uncertainty in the pose estimate due to spurious points, occlusions, and outliers. This uncertainty can impact the performance of the algorithm and ultimately lead to the divergent behavior. Therefore, adaptive scan–correlation modifies the least–squares error function, shown in Equation 4.11, to address the uncertainty and correspondences errors that may arise.

$$\min_{\boldsymbol{\mathcal{T}}} \sum_{i} \boldsymbol{w}_i \, \|\boldsymbol{\mathcal{T}}\, \boldsymbol{p}_i \, - \, \hat{\boldsymbol{q}}_i\|^2 \tag{4.11}$$

$$\boldsymbol{w} = \begin{cases} 0 & \text{for } \boldsymbol{d}_i \, > \, \boldsymbol{\mathcal{D}}_k^{max} \\ 1 & \text{otherwise} \end{cases} \tag{4.12}$$

The modified least–squares error function (Equation 4.11) adds a mechanism, $\boldsymbol{w}_i$, to weight or reject correspondences determined at each iteration, $\boldsymbol{i}$. In this application, $\boldsymbol{w}$ is a binary weighting mechanism used to simply discard potential correspondence errors. The rejection criteria, shown in Equation 4.12, is based on an adaptive threshold, $\boldsymbol{\mathcal{D}}_k^{max}$ and Euclidean square distance between correspondences, denoted $\boldsymbol{d}_i \triangleq \|\boldsymbol{p}_i \, - \, \hat{\boldsymbol{q}}_i\|^2$.

$$\boldsymbol{\mu} = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{d}_i; \quad \boldsymbol{\sigma} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} (\boldsymbol{d}_i - \boldsymbol{\mu})^2} \tag{4.13}$$

$$\boldsymbol{\mathcal{D}}_{k+1}^{max} = \begin{cases} \boldsymbol{\mu} + 3\boldsymbol{\sigma} & \text{for } \boldsymbol{\mu} < \boldsymbol{\mathfrak{D}} \\ \boldsymbol{\mu} + 2\boldsymbol{\sigma} & \text{for } \boldsymbol{\mu} < 3\boldsymbol{\mathfrak{D}} \\ \boldsymbol{\mu} + \boldsymbol{\sigma} & \text{for } \boldsymbol{\mu} < 6\boldsymbol{\mathfrak{D}} \\ \boldsymbol{\mu} + \boldsymbol{\epsilon} & \text{otherwise} \end{cases} \tag{4.14}$$

The value for the adaptive threshold is determined through the statistical analysis of $\boldsymbol{d}_i$ of valid correspondences found in the current iteration. The analysis of $\boldsymbol{d}_i$, shown in Equation 4.13, is used to assess how the mean–square distance of the correspondence, $\boldsymbol{\mu}$, compares to the desired resolution of the registration process specified by the user, $\boldsymbol{\mathfrak{D}}$. Shown in Equation 4.14, this assessment is used to set the adaptive threshold for the next iteration, $\boldsymbol{\mathcal{D}}_{k+1}^{max}$. Note $\boldsymbol{\epsilon}$ is the median $\boldsymbol{d}_i$ for all valid correspondences.

As previously stated, $\boldsymbol{\mathfrak{D}}$ is a user-defined variable that represents the desired resolution or the expected mean–square distance between correspondences to be achieved during the registration process. It serves as the basis for characterizing the quality of the registration and computing a new value for $\boldsymbol{\mathcal{D}}^{max}$ accordingly. For instance in Equation 4.14, the first case, $\boldsymbol{\mu} < \boldsymbol{\mathfrak{D}}$, suggests the registration between the correspondences is good and the adaptive threshold is adjusted so more correspondences are considered. Conversely, the third case, $\boldsymbol{\mu} < 6\mathfrak{D}$, indicates the registration process is not achieving the desired accuracy so the adaptive threshold is set more conservatively to discard potential outliers or spurious points influencing the registration of the scans.

When defining $\mathfrak{D}$ there are two observation to consider:

1. If $\mathfrak{D}$ is too small, then valid correspondences may be discarded increasing the number of iterations required to converge

2. If $\mathfrak{D}$ is too big, then correspondences errors may be included increasing the probability for divergent behaviors.

5. Performance Evaluation

Performance evaluation is a vital first step in quantifying the performance of mobile robot localization. It uses the ground truth information about the location of the robot and the surrounding environment to decompose errors arising in the pose estimate. The juxtaposition of these errors provides the basis for the quantitative comparison of competing approaches. The decomposition of errors also reveals the presence of performance singularities and helps to diagnosis the cause and impact of the divergent behavior on the overall performance of the system.

This chapter will evaluate the performance of the adaptive scan–correlation as a stand-alone sensor-based dead-reckoning solution, commonly referred to as *scan-matching*. The performance of the scan-matching solution will be compared to two commonly used dead-reckoning solutions; Inertial Navigation System (INS) and encoder-based odometry using Unified System for Automation and Robot Simulation (USARSim) [6] and Mobility Open Architecture Simulation and Tools (MOAST) [4][1]. The remainder of this chapter will provide an overview of the dead–reckoning approaches and present ex-

[1] Appendix A provides an overview of the MOAST and USARSim as a high-fidelity robotic simulation testbed

perimental results on the performance evaluation of the three solutions.

5.1. Dead-Reckoning Approaches

This section provides an overview of the three dead-reckoning approaches being evaluated and describe the implementation of the virtual sensor models in USARSim developed for this research. As stated previously, dead-reckoning is the most basic form of localization, using simple mathematical principles measure relative motion and recursively “advances” the estimate of the pose. These solutions are beneficial because of their ability to supply short-term pose estimates at high-data rates, but lack mechanisms measure and bound uncertainty in the system.

5.1.1 Encoder-based Odometry

Encoder-based Odometry is the most commonly used dead-reckoning approach. This approach computes a 2D pose estimate based on a kinematic model of the vehicle and measurements of wheel rotation and steering angle. Discrepancies between the kinematic model and the actual vehicle, i.e. unequal wheel diameters, can lead to systematic errors that accumulate over time. While systematic errors impact the performance of these solutions, Encoder-based Odometry is more susceptible to non-systematic errors resulting from undulating terrain or wheel slippage.

The virtual Odometry model implemented in USARSim utilizes internal state information, available in the game engine, about the wheel velocities and diameter to compute the distance each of the tracks have travels on a skid–steering vehicle, $[\boldsymbol{U^L}, \boldsymbol{U^R}]$. This information is used to compute a 2D pose estimate, $[\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{\phi}}]$, in USARSim as follows:

Given a previous pose estimate, $[\hat{\boldsymbol{x}}_{\boldsymbol{t-1}}, \hat{\boldsymbol{y}}_{\boldsymbol{t-1}}, \hat{\boldsymbol{\phi}}_{\boldsymbol{t-1}}]$, and the distance traveled by the left and right tires/tracks over the past time step, $[\boldsymbol{U_t^L}, \boldsymbol{U_t^R}]$, compute a new pose estimate based on a skid-steered kinematic model shown in Equation 5.1, where $\boldsymbol{\ell}$ represents the wheel separation.

$$\begin{bmatrix} \hat{\boldsymbol{x}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{y}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{\phi}}_{\boldsymbol{t}} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{x}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{y}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{\phi}}_{\boldsymbol{t-1}} \end{bmatrix} + \begin{bmatrix} \frac{\boldsymbol{U_t^L}+\boldsymbol{U_t^R}}{2} \cos \hat{\boldsymbol{\phi}}_{\boldsymbol{t-1}} \\ \frac{\boldsymbol{U_t^L}+\boldsymbol{U_t^R}}{2} \sin \hat{\boldsymbol{\phi}}_{\boldsymbol{t-1}} \\ \arctan \frac{\boldsymbol{U_t^L}-\boldsymbol{U_t^R}}{\boldsymbol{\ell}} \end{bmatrix} \tag{5.1}$$

5.1.2 Inertial Navigation System

An *Inertial Navigation System* (INS) is another commonly used dead–reckoning solution based on Newton's laws of motion. It assumes that an object will remain in uniform motion unless it is acted on by an outside force. Forces acting on the system will produce accelerations in an inertial reference frame that can be measured and integrated over time to compute the relative motion of the vehicle. The change in position and orientation is accumulated to estimate its current pose with respect to a previously determined pose [39]. In many cases, inertial sensors are not able to dissociate accelerations due to external forces, such as gravity, from kinematic accelerations [29]. The failure of the systems to properly classify or compensate for non-kinematic noise will produce error in the inertial measurements. These errors coupled with numerical errors produced from the double integration of the accelerations will produce a gradual degradation of the navigation solution, commonly referred to as *drift*.

To support this research, a virtual model of an INS was developed for USARSim. This model utilized ground truth pose information obtained from the game engine to derive angular accelerations and the distance traveled to compute a pose estimate. The virtual model of the INS developed for USARSim computes a 3D pose estimate as follows:

1. Compute angular velocities, $[\boldsymbol{\omega_t^x}, \boldsymbol{\omega_t^y}, \boldsymbol{\omega_t^z}]$, of the robotic platform for current time step, $\boldsymbol{t}$, using the ground truth information obtained from the simulated world about the robots orientation,$[\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\psi}}, \hat{\boldsymbol{\phi}}]$:

$$\begin{bmatrix} \boldsymbol{\omega_t^x} \\ \boldsymbol{\omega_t^y} \\ \boldsymbol{\omega_t^z} \end{bmatrix} = \frac{1}{\boldsymbol{\Delta t}} \left(\begin{bmatrix} \boldsymbol{\theta_t} \\ \boldsymbol{\psi_t} \\ \boldsymbol{\phi_t} \end{bmatrix} - \begin{bmatrix} \boldsymbol{\theta_{t-1}} \\ \boldsymbol{\psi_{t-1}} \\ \boldsymbol{\phi_{t-1}} \end{bmatrix} \right) \tag{5.2}$$

2. Update the current orientation estimate using the angular velocities computed in Step 1. and the Gaussian noise model, $\mathbb{G}(\boldsymbol{\mu}, \boldsymbol{\sigma})$, based on

predefined mean, $\boldsymbol{\mu}$, and variance, $\boldsymbol{\sigma}$:

$$\begin{bmatrix} \hat{\boldsymbol{\theta}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{\psi}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{\phi}}_{\boldsymbol{t}} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{\theta}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{\psi}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{\phi}}_{\boldsymbol{t-1}} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\omega}_{\boldsymbol{t}}^{\boldsymbol{x}} \\ \boldsymbol{\omega}_{\boldsymbol{t}}^{\boldsymbol{y}} \\ \boldsymbol{\omega}_{\boldsymbol{t}}^{\boldsymbol{z}} \end{bmatrix} \; \mathbb{G}(\boldsymbol{\mu}, \boldsymbol{\sigma}) \; \boldsymbol{\Delta t} \tag{5.3}$$

3. Calculate an estimate of the Euclidean distance traveled over the past time step using ground truth and $\mathbb{G}(\boldsymbol{\mu}, \boldsymbol{\sigma})$:

$$\hat{\boldsymbol{V}}_{\boldsymbol{dist}} = \sqrt{\boldsymbol{\Delta x}^2 + \boldsymbol{\Delta y}^2 + \boldsymbol{\Delta z}^2} \;\; \mathbb{G}(\boldsymbol{\mu}, \boldsymbol{\sigma}) \tag{5.4}$$

4. Update the current position estimate using a 3D motion model [39] that takes into consideration the pitch of the vehicle, $\hat{\boldsymbol{\psi}}_{\boldsymbol{t}}$, and the yaw, $\hat{\boldsymbol{\phi}}_{\boldsymbol{t}}$, computed in Equations 5.3 and 5.4.

$$\begin{bmatrix} \hat{\boldsymbol{x}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{y}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{z}}_{\boldsymbol{t}} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{x}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{y}}_{\boldsymbol{t-1}} \\ \hat{\boldsymbol{z}}_{\boldsymbol{t-1}} \end{bmatrix} + \begin{bmatrix} \hat{\boldsymbol{V}}_{\boldsymbol{dist}} \cos \hat{\boldsymbol{\phi}}_{\boldsymbol{t}} \cos \hat{\boldsymbol{\psi}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{V}}_{\boldsymbol{dist}} \sin \hat{\boldsymbol{\phi}}_{\boldsymbol{t}} \cos \hat{\boldsymbol{\psi}}_{\boldsymbol{t}} \\ \hat{\boldsymbol{V}}_{\boldsymbol{dist}} \sin \hat{\boldsymbol{\psi}}_{\boldsymbol{t}} \end{bmatrix} \tag{5.5}$$

5.1.3 Scan-Matching

An alternative to traditional dead-reckoning is a sensor-based dead-reckoning approach known as *scan-matching.* Scan-matching measures the relative motion of a vehicle through the temporal correlation of consecutive scans obtained laser range finder. While scan-matching is also vulnerable to error propagation, the use of external observations minimizes the impact of systematic and non-systematic errors that plague other dead-reckoning solutions. As stated previous, the prominent source of errors arising in this solution is due to the data association problem, or the inability to determine valid correspondences during the registration process.

$$\min_{\boldsymbol{\mathcal{T}}} \sum_i \boldsymbol{w}_i \, \|\boldsymbol{\mathcal{T}}\, \boldsymbol{p}_i \, - \, \hat{\boldsymbol{q}}_i\|^2 \tag{5.6}$$

The scan-matching solution developed for this research is based on the adaptive scan–correlation algorithm discussed in Chapter 4. This approaches uses 2D scans obtained from a simulated laser range finder mounted on the vehicle in USARSim. Using the modified least–square objective function, shown in Equation 5.6, scan-matching measure the between the two scans to compute the a 2D pose estimate of the vehicle, $[\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{\phi}}]$.

5.2. Experimental Results

This experiment was conducted in USARSim using a simulated skid–steering vehicle shown in Figure 5.1. This skid–steered vehicle was configured with the simulated INS sensor (Section 5.1.2), a simulated Odometry sensor (Section 5.1.1), and a simulated 2D laser range finder configured to have a field of view of 180°with a beam separation of 1°, consisting of 181 returns. The data frequency for each the sensors was configured to serve data at 5 Hz. The simulated skid–steering vehicle was teleoperated in USARSim using MOAST

Figure 5.1: ***Skid–Steered Vehicle.*** A skid–steered vehicles refer to a tracked or wheeled ground vehicle that is maneuvered using differential control of the left and right tracks or wheels. The picture on the left shows the simulated skid–steered vehicle used for this research that was modeled in USARSim. The picture on the right shows the actual vehicle.

for approximately 4.5 minutes and traveled almost 26 meters at an average speed of 0.097 m/s. During the run, data streams were logged, processed, and compared to ground truth to quantify the performance of the three dead–reckoning solutions.

In this text, a *cumulative sensor map* refers to a composite map consisting of raw sensor data mapped into a relative coordinate frame using the pose estimate. This rudimentary map does not employ mapping facilities to improve the accuracy of the map, i.e. no filtering of data or pruning of the map. Close examination of the cumulative map produced by the scan-

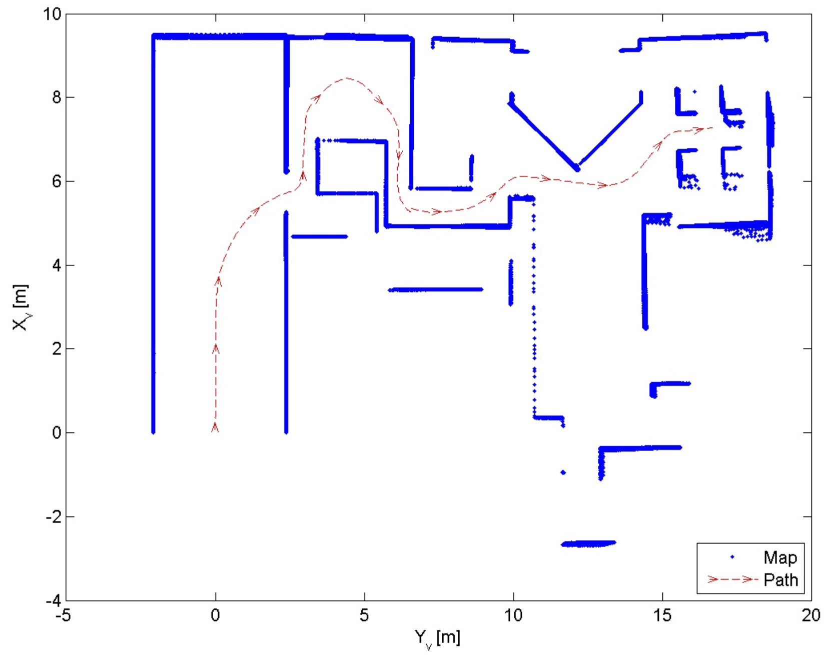

Figure 5.2: ***Scan-Matching Cumulative Sensor Map.*** This figure shows rudimentary map based on the scan–matching solution and raw sensor data.

matching solution, shown in Figure 5.2, illustrates the integrity and robust nature of an exteroceptive approach to formulating a pose estimate with only marginal errors being produced at the end of the run, visible in the top-right corner of the map.

Figure 5.3 shows a comparison of the position estimates produced by the dead–reckoning approaches, using the ground truth trajectory as a basis of comparison. This figure shows the scan-matching solution exhibits a more accurate representation of the actual path traveled by the vehicle. Although the paths of the odometry and INS solution appear to mimic the actual path,

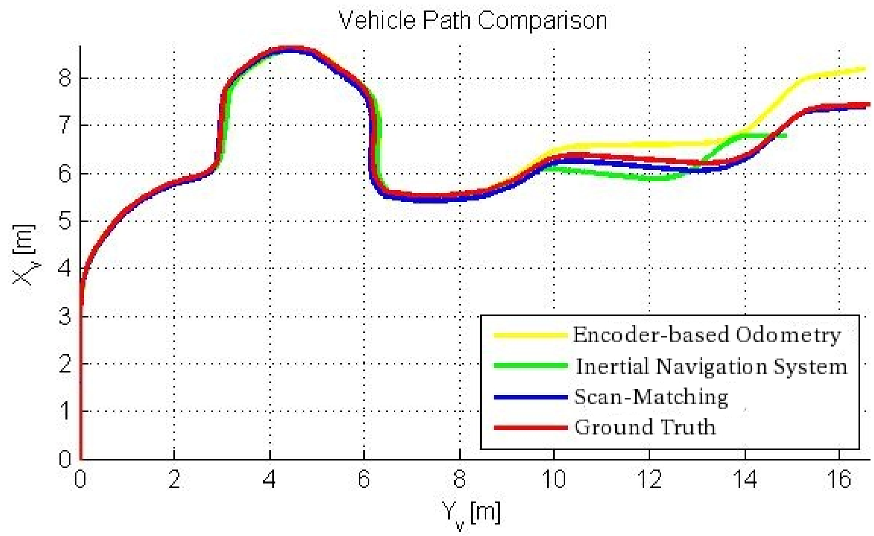

Figure 5.3: ***Pose Estimate Comparison.*** This figure compares the pose estimates of the three dead–reckoning approaches using ground truth information as the basis of comparison.

both have induced substantial errors spawning from a potential performance singularity occurring three-fourths of the way through the run, coordinates ($\boldsymbol{x} \approx 6,\ \boldsymbol{y} \approx 10$) in Figure 5.3.

Examining the positional errors in Figure 5.4 confirms the presence of a performance singularity occurring in both the odometry and INS between 150 and 200 seconds (approximately 180 seconds), causing the solutions to diverge. The simultaneous occurrence of the error in both the odometry and INS solutions anecdotally suggests the error can be attributed to a non-systematic error, i.e. hitting a wall. The divergent behavior emanating from this performance singularities exemplifies the how errors can propagate and grow unbounded in dead–reckoning approaches.

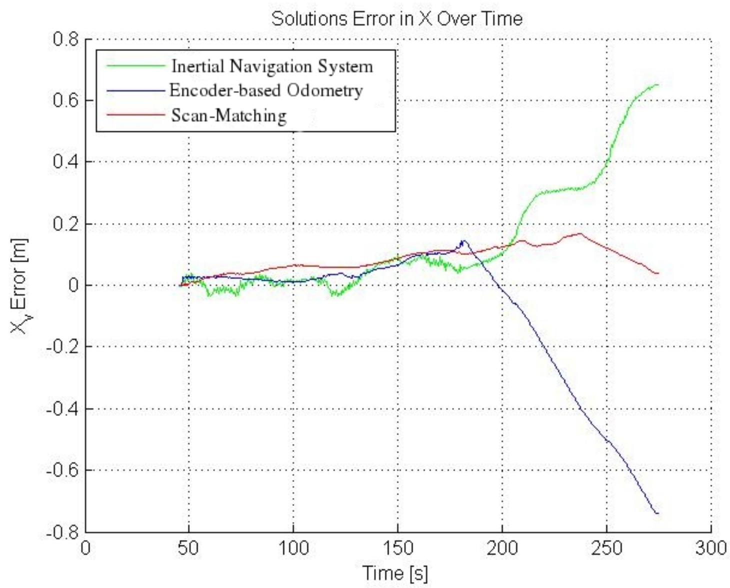

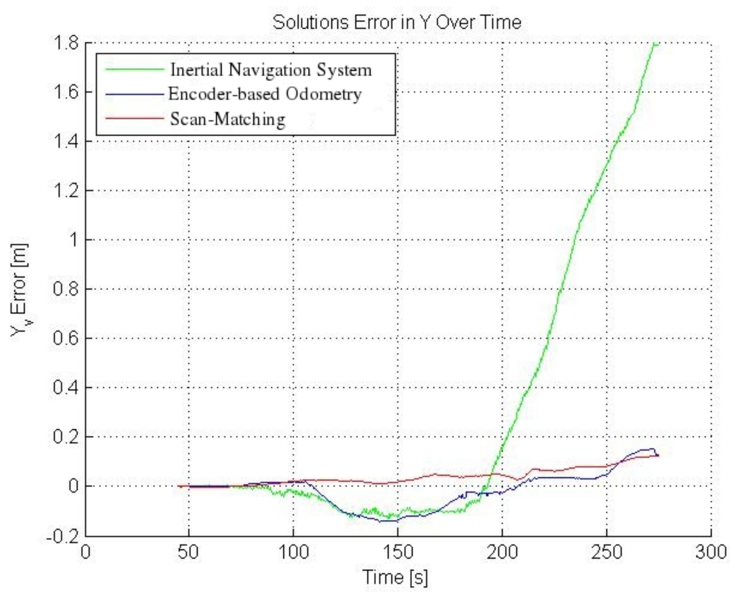

Figure 5.4: ***Evaluation of Positional Errors.*** This figure decomposes and compares the positional errors occurring in the dead–reckoning solutions with respect to time.

Comparing the positional errors depicted in Figure 5.4 affirms the scan–matching solution outperforms the other dead–reckoning approaches. The relative stability and accuracy exhibited in the scan–matching solution evinces its resilience to systematic and non-systematic errors that plague dead–reckoning approaches. In fact, the performance singularity causing divergent behavior in the odometry and INS solutions does not appear to influence the solution produced by the scan–matching approach. Using the error curves for the INS solution and Odometry as the basis for inferring trends, suggests errors accumulating in these solutions may have grown without bounds. In contrast, the error curves for the scan–matching solution indicates it provides a more robust approach to localization.

Looking at the orientational errors in Figure 5.5 shows a significant spike between 150 and 200 seconds in the odometry and INS solutions. This spike coincides with the divergent behavior seen in seen in Figure 5.4 for the odometry and INS solutions. This suggests the divergent behavior in these solutions is directly related oreintational errors. This also reinforces the assumption the noted performance singularity is due to non-systematic errors and not systematic errors biased to a specific approach.

In comparison, the heading errors for the scan–matching solutions shown in Figure 5.5 indicates the pose estimate produced by the scan–matching approach is more stable and accurate than odometry and INS. This stability is also reflected in the positional errors for seen in Figure 5.4. However, a sharp spike in the orientational errors of the scan–matching solution just after the

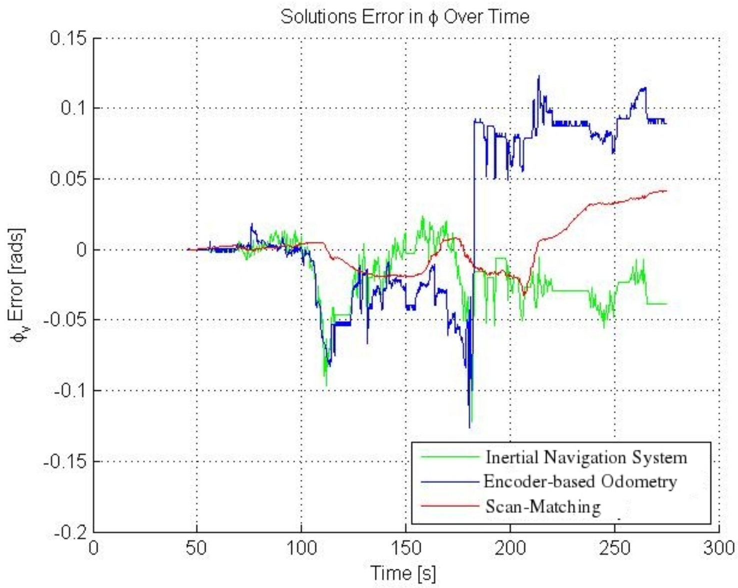

Figure 5.5: ***Evaluation of Orientational Errors.*** This figure compares the orientational errors occurring in the dead–reckoning solution with respect to time.

200 second mark indicates a potential systematic error has occurred. While this spike did not have significant impact on the positional errors, it could have contributed to the mapping errors in the upper-right hand corner of the cumulative sensor map shown in Figure 5.2. This systematic error ,which occurrs near the end of the run, is most likely a result of correspondence issues in the scan–matching algorithm. Furthermore, this identifies potential test scenarios for analyzing what factors contributed to systematic errors in scan–matching approach.

6. Performance Analysis

The primary focus of this chapter is to quantitatively analyze the performance of adaptive scan–correlation as compared to the basic scan–correlation discussed in Chapter 4. This analysis will introduce testing scenarios designed to challenge the ability to the system to determine valid correspondences in environments with different levels of complexities. The use of reference data sets within each test scenario will quantify the convergence characteristics of each solution and facilitate the inter-comparison of results. The ensuing subsections will first describe the test scenarios and how each of the scenarios will challenge the scan–correlation solutions. This will be followed by the performance analysis of the two scan-correlation variants in each of these scenarios with an emphasis on vehicle speed.

6.1. Testing Scenarios

In order to quantify the performance of the scan-correlation techniques, three testing scenarios were identified to challenge the system's ability determine valid correspondences in isolated, repeatable tests. Each test scenario is composed of a reference data set that simulates linear motion of the vehicle using Unified System for Automation and Robot Simulation (USARSim). Three dimensional scan data and the ground truth location of the vehicle is captured at 10 cm intervals along a straight line trajectory. The scan data has a horizontal field-of-view of 180° and vertical field-of-view of 20°, both with an angular resolution of 1°. This produces 3D scans that may contain 3801 hit points. Using different combinations of the reference scans will enable developers to test how linear displacement, a function of vehicle speed and the data rate of the sensor, affects the overall performance of scan-correlation algorithms.

6.1.1 Environments with Distinct Features

Environments with distinct features provides a scenario that limits environmental complexities to provide the best-case scenario, where scan–correlation solutions should perform optimally. As seen in Figure 6.1, this scenario uses a closed set of distinct mapping features and vertical walls that produces unique observations. This enables mapping systems to associate features and increases the likelihood of determining valid correspondences. Perpendicular surfaces, which allow for more accurate measurements of motion, appear in almost every scan. Limiting the environmental complexities allows developers to tune their systems and establishes baseline a for comparison.

6.1.2 Environments with Occluded Features

Environments with occluded features is designed to challenge the scan–correlation solution ability to determine valid correspondences as a result of occlusions,

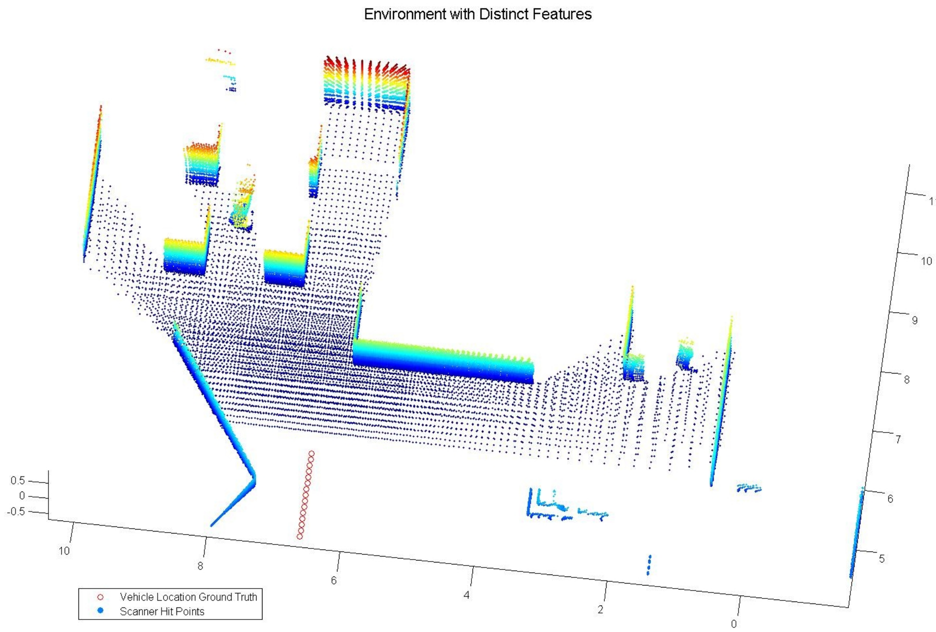

Figure 6.1: ***Environments with Distinct Features.*** This figure provides a 3D visualization of the testing scenario that provides an environment with distinct features. It uses the ground truth in the reference data set to compute the sensor hit points and plot the locations where the scans were logged. Scanner hit points are colorized based on height.

outliers, etc. In this test scenario, shown in Figure 6.2, nearby features may periodically occlude more distant features as the robot moves through the environment. This produces a situation where consecutive observations may not contain the same set of features, increasing the likelihood of correspondence errors. However, the nearby features, which are not obstructed, enable the system to make accurate measurements of the immediate vicinity and should help the system avoid catastrophic failures. This menagerie of features and occlusion in this scenario is indicative of unstructured en-

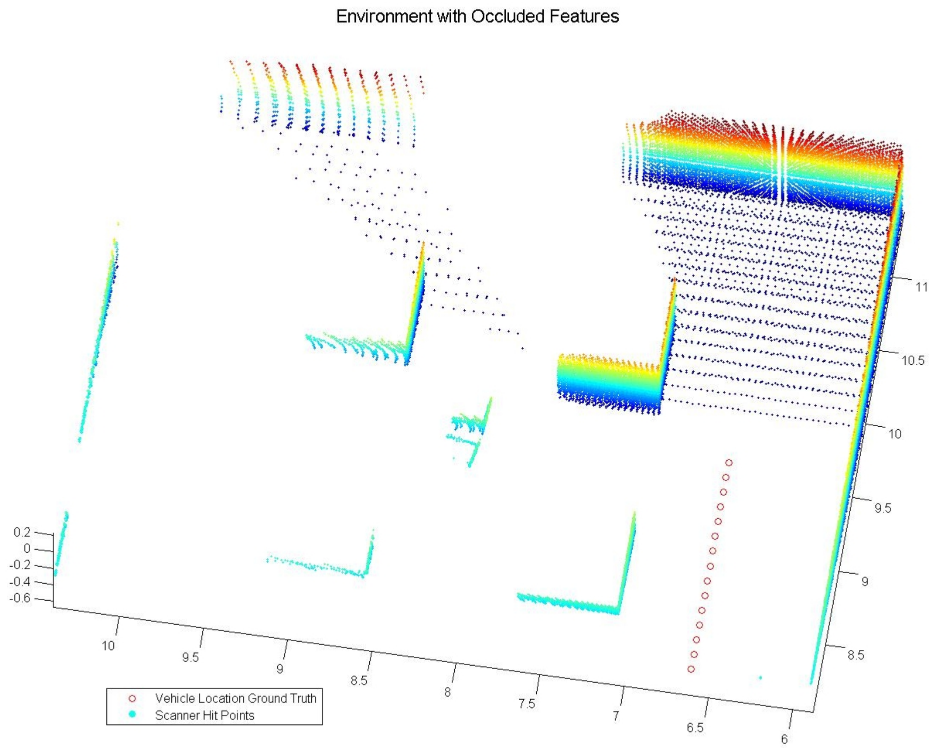

Figure 6.2: ***Environments with Occluded Features.*** This figure provides a 3D visualization of the testing scenario thatprovides an environment with occluded features. It uses the ground truth in the reference data set to compute the sensor hit points and plot the locations where the scans were logged. Scanner hit points are colorized based on height.

vironments; therefore, this is an essential test understanding how the scan-correlation approaches will perform in the real world.

6.1.3 Environments with Minimal Features

Environments with minimal features implements the degenerative case for scan-correlation techniques. As shown in Figure 6.3, this scenario presents a symmetric and featureless environment. This inhibits the system's ability

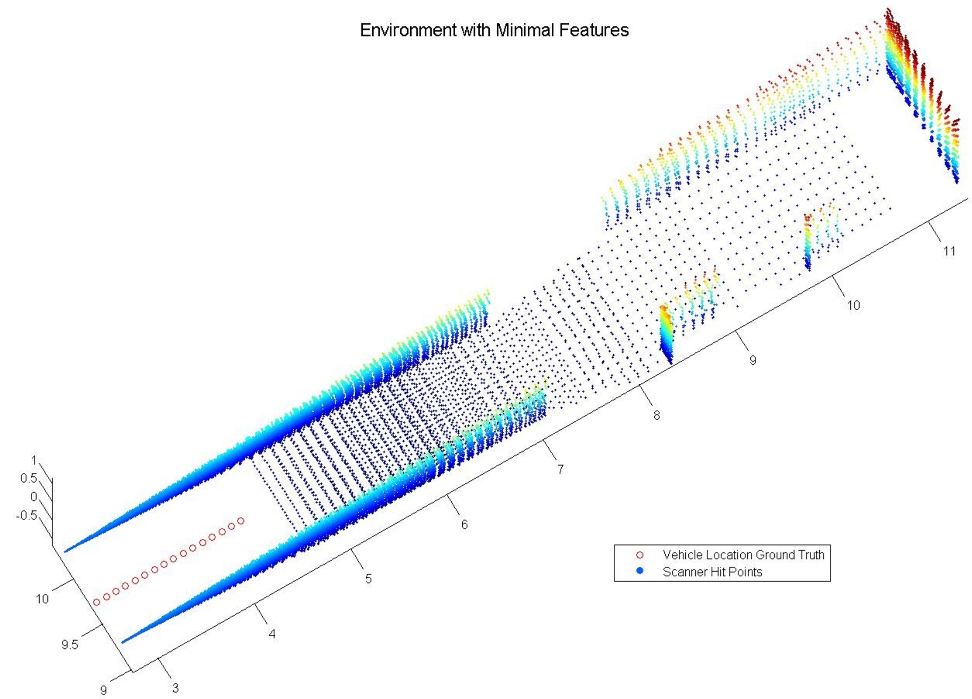

Figure 6.3: ***Environments with Minimal Features.*** This figure provides a 3D visualization of the testing scenario that provides an environment with minimal features. It uses the ground truth in the reference data set to compute the sensor hit points and plot the locations where the scans were logged. Scanner hit points are colorized based on height.

to make accurate measurements of its environments and it's ability to determine valid correspondences. The only distinct feature in the scenario is the far wall, which presents a perpendicular surface to the robot. The lack of distinct features increases the potential for catastrophic errors by preventing the convergence of the pose estimate in the scan-correlation techniques. While this situation does not occur commonly (except in culverts, sewers, and tunnels), this testing scenario is essential to understanding how the system fails.

6.2. Performance Analysis & Experimental Results

Performance analysis takes advantage of the ground truth in the reference data sets to measure the error in the pose estimate at each iteration. These errors are plotted to produce a convergence profile. The convergence profile not only shows how well the scan-correlation algorithm converges, it elucidates the convergence characteristics, such as the stability of the pose estimate. Other vital information can also be logged to help understand the performance characteristics found in the convergence profile. For instance, the correspondence profiles help to infer how the number of correspondences found at each iteration influences the performance of the system. In the case of the adaptive scan-correlation algorithm, an adaptive threshold profile is plotted to gain insight into how the value of $\boldsymbol{\mathcal{D}}^{max}$ reflects the quality of the registration and the stability/accuracy of the pose estimate. This analysis can also expose the presence of meta-level knowledge that may enable the system to recognize and diagnosis failure conditions arising in the system at run-time.

The performance analysis of the two scan-correlation techniques, discussed in Chapter 4, will be discussed using the basic scan-correlation technique as a baseline for analyzing the effects of the point-to-line matching and adaptive thresholding in the adaptive scan–correlation approach. In each testing scenario, described in Section 6.1, the algorithm will be subjected to several runs to measure how vehicle speed and/or data rate of the sensor

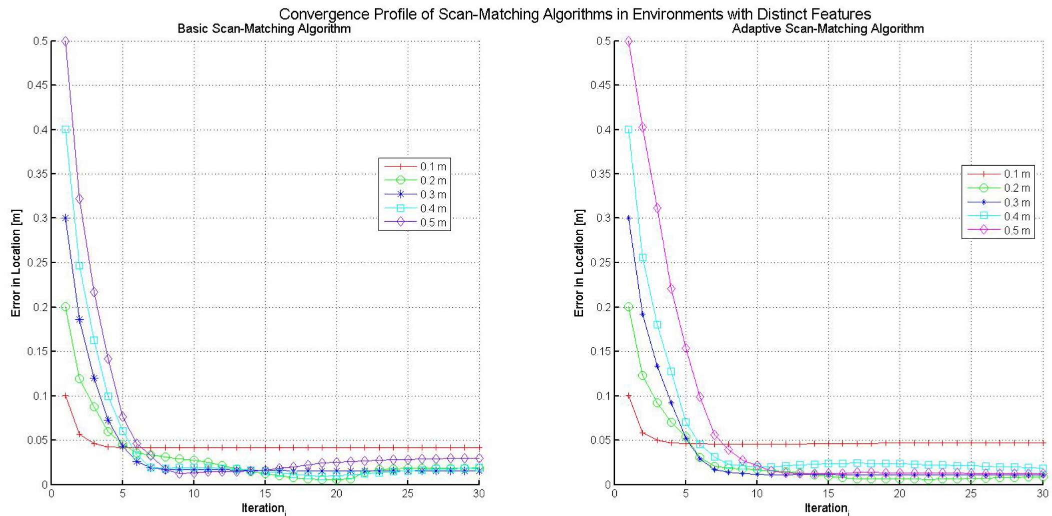

Figure 6.4: ***Correspondence Profiles: Distinct Features.*** The convergence profiles use the ground truth in the reference data set to compute the error in the pose estimate at each iteration of the scan-correlation algorithm. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

impacts the algorithm's ability to compute a single pose estimate in test scenarios with varying degrees of complexity. The linear motion of the vehicle is simulated by using different combination of scans in the referenced data set along a straight line trajectory. The first scan in each of the referenced data sets will serve as the reference scan, $\boldsymbol{\mathcal{S}}^{ref}$, in the registration process, or starting location. For each of the runs, the algorithm will iterate through the referenced data set, using each of the other scan location as the observation scan, $\boldsymbol{\mathcal{S}}^{obs}$, or the end location. Each of the combination will produce a test case simulating different linear displacements of 10 cm, 20 cm, 30 cm, 40 cm, and 50 cm along the straight line trajectory in each of the scenarios. Data will be logged for each test case and be used to analyze the performance of the competing approaches below.

In each of the test cases, the termination criteria for each of the scan–correlation algorithms was configured to run for 30 iterations. This was done in an effort to gain a better understanding the convergence characteristics for each of the scan–correlation techniques. Therefore, for each test cases, the scan–correlation techniques were required to register the two 3D data sets consisting of 3801 points. Although the computational complexity for each of the algorithms was not explored in this research, each of the algorithms were able to complete 30 iterations in under 200ms, or at 5Hz.

6.2.1 Analysis In Environments with Distinct Features

The convergence profiles of the two scan–correlation techniques in the test scenario with distinct features, shown in Figure 6.4, shows both approaches were able to rapidly converge to an accurate pose estimation (within 5 cm of ground truth) in under 10 iterations. However, it is noteworthy to point out three residual artifacts that are present in the convergence profile. First, in the 10 cm linear displacement run, the pose estimate computed by both algorithms is less accurate than their counterparts at larger displacements. Second, perturbations in the tail of the convergence profiles of the basic scan-correlation technique suggests the adaptive scan-correlation technique produces a more stable solution. Finally, the pose estimate in the basic scan-correlation algorithm appears to converge quicker than the adaptive

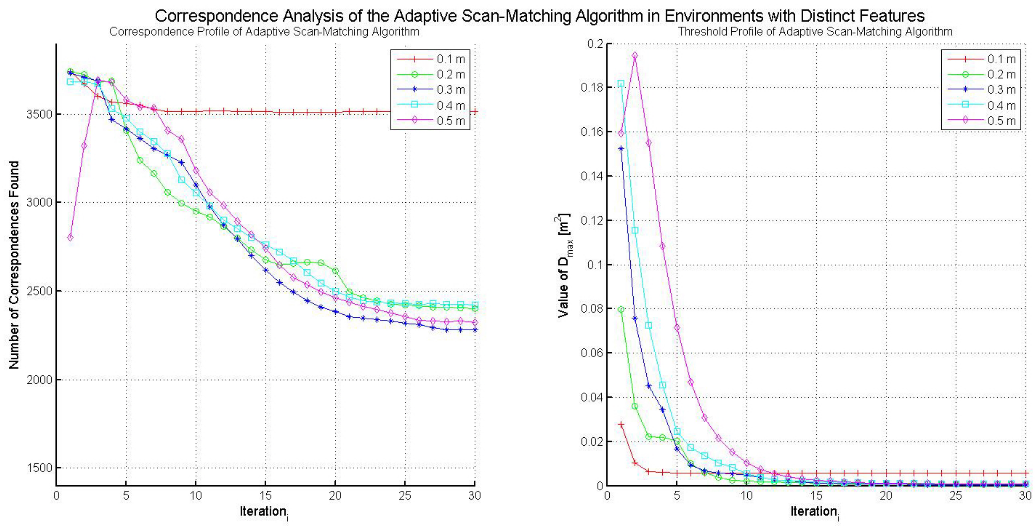

Figure 6.5: ***Adaptive Threshold Analysis: Distinct Features.*** The two figures here show the correspondence profile and the threshold profile of the adaptive scan-correlation algorithm in environments with distinct features. The correspondence profile plots the number of correspondences found at each iteration. The threshold profile shows the value of adaptive threshold, $\boldsymbol{\mathcal{D}}^{max}$, at each iteration. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

scan-correlation technique in the run with 50 cm linear displacement.

In order to gain insight into the nature of these artifacts, a close examination of the correspondence and threshold profiles of the adaptive scan-correlation technique is needed, as shown in Figure 6.5. First, the number of correspondences found in the 10 cm test case begins to plateau almost immediately and remains fairly constant for the remaining iterations. This differs drastically from the other runs where the number of correspondences is monotonically decreasing leading to more accurate pose estimates. Sec-

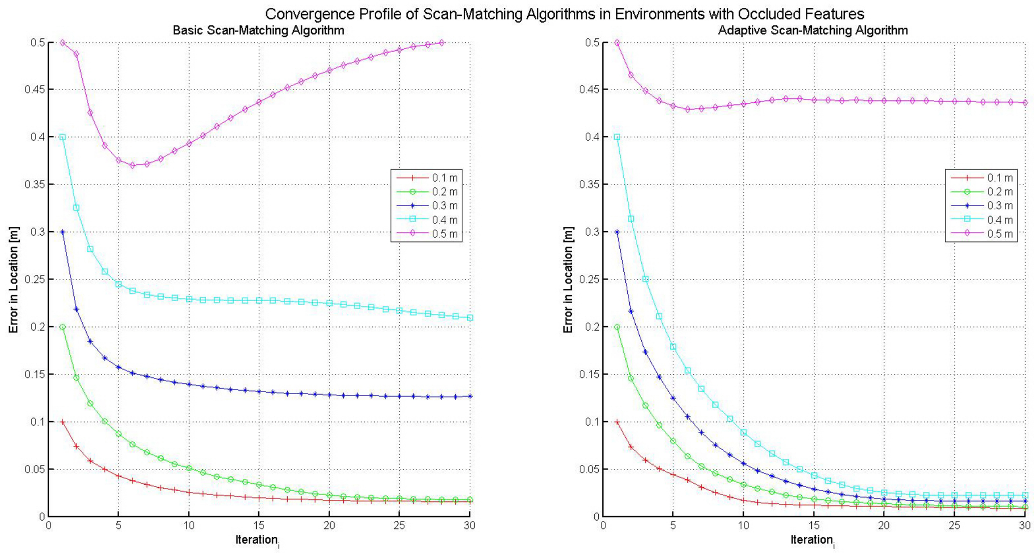

Figure 6.6: ***Correspondence Profiles: Occluded Features.*** The convergence profiles use the ground truth in the reference data set to compute the error in the pose estimate at each iteration of the scan-correlation algorithm. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

ond, the adaptive scan-correlation technique uses statistical analysis of the data to eliminate correspondence errors and improve registration between the data. Looking at the threshold profile, the value of $\mathcal{D}^{max}$ fluctuates until finally converging to a value close to zero. This indicates there is good registration between the points, making the pose estimate more stable. Finally, the adaptive scan-correlation uses a threshold based on distance that discards correspondences with large spatial relationships, preventing it from converging as quickly under ideal conditions with no occlusions.

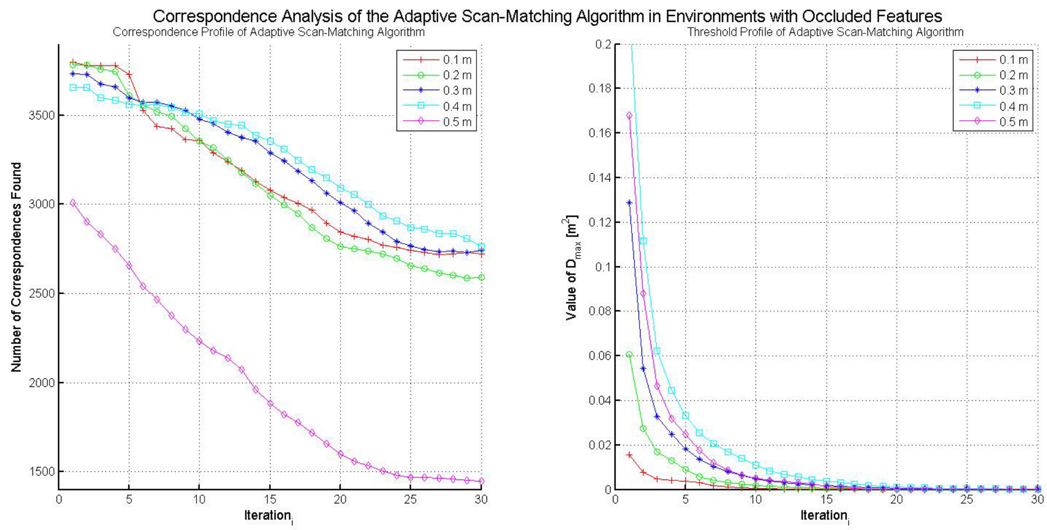

Figure 6.7: ***Adaptive Threshold Analysis: Occluded Features.*** The two figures here show the correspondence profile and the threshold profile of the adaptive scan-correlation algorithm in environments with occluded features. The correspondence profile plots the number of correspondences found at each iteration. The threshold profile shows the value of adaptive threshold, $\mathcal{D}^{max}$, at each iteration. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

6.2.2 Analysis In Environments with Occluded Features

In environments with occluded features the convergence profiles in Figure 6.6 shows the adaptive scan-correlation algorithm is more proficient and is able to outperform the basic scan-correlation algorithm. In the 50 cm test case, neither of the scan-correlation variants were able to converge to an accurate pose estimate. However, in the other test cases the adaptive scan-correlation algorithm was able to maintain stability, where the basic scan-correlation pose estimate diverges.

In order to understand why the adaptive scan–correlation algorithm was able to exhibit superior performance in this test scenarios, it is important to examine the correspondence and threshold profiles shown in Figure 6.7. Recall that each scan in the reference data set contains approximately 3801 hit points. The correspondence profile for the adaptive scan-correlation technique shows the number of correspondences found at each iteration is monotonically decreasing. This is driven by the value of $\mathcal{D}_{max}$, shown in the threshold profile. The statistical analysis of the data causes the value of $\mathcal{D}_{max}$ to converge as the registration between the data improves. It is also important to note two additional observation that indicate the presence meta-level knowledge that may help the adaptive scan-correlation technique recognize the stability of the system. First, the convergence of the threshold profile coincides with the convergence of the pose estimate. Second, the correspondence profile for the 50 cm test case is noticeably different than the profiles for the other run.

6.2.3 Analysis In Environments with Minimal Features

Figure 6.8 shows the effects of the degenerative case, environments that lack distinct features, on the scan-correlation algorithms. Even though the basic scan-correlation algorithm appears to converge slightly better, both scan-correlation variants fail to produce valid pose estimates. The lack of features

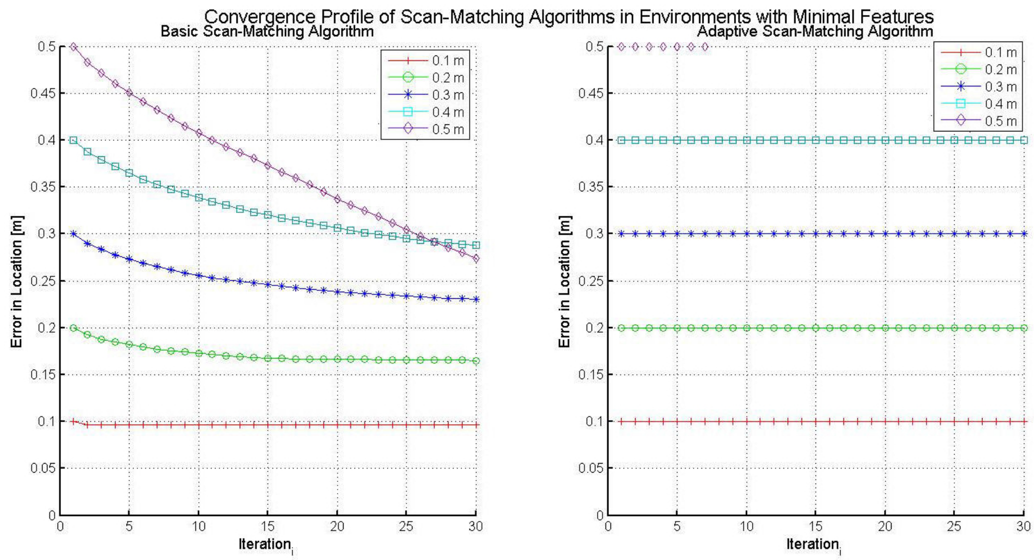

Figure 6.8: ***Correspondence Profiles: Minimal Features.*** The convergence profiles use the ground truth in the reference data set to compute the error in the pose estimate at each iteration of the scan-correlation algorithm. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

produces identical scan signatures where the points in each of the resulting scans are located in the same place from the perspective of the sensor. This undermines the ability of both techniques to determine valid correspondences, which impedes their ability to compute a valid pose estimate. Since the majority of the points in the resulting scans are aligned, the adaptive scan-correlation algorithm assumes good registration between the data and discards valid correspondences. The basic scan-correlation algorithm does not weight the correspondences, so all available information is used. This suggests valid correspondences are not discarded, which enables the pose

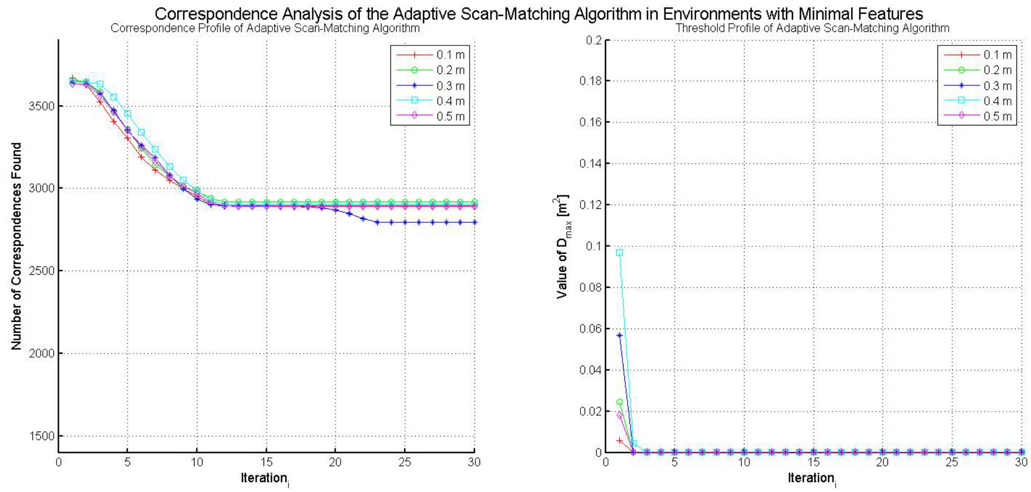

Figure 6.9: ***Adaptive Threshold Analysis: Minimal Features.*** The two figures here show the correspondence profile and the threshold profile of the adaptive scan-correlation algorithm in environments with minimal features. The correspondence profile plots the number of correspondences found at each iteration. The threshold profile shows the value of adaptive threshold, $\mathcal{D}_{max}$, at each iteration. Each of the figures shows the results of 5 runs of linear displacements along a straight line trajectory. The linear displacements for each of the runs is depicted in the legend.

estimate in the basic scan-correlation to converge slightly better.

In order to better understand the effects of this scenario on adaptive scan-correlation technique, it is essential to examine the correspondence and threshold profiles in Figure 6.9. First, it is important to note that the value of $\mathcal{D}_{max}$ in the threshold profile converges instantly. This indicates that the statistical analysis of the data incorrectly assumes there is good correspondence between the data sets. Moreover, the correspondence profile shows the number of correspondences are not monotonically decreasing and actually level out after 10 iterations. This differs drastically from the threshold and correspondence profiles seen the the other test scenarios where the adaptive scan-correlation algorithm was able to compute a valid pose estimate.

7. Conclusions and Further Discussions

The development of robust localization techniques is a core competency to improve the utility and proficiency of mobile robots operating in complex, unstructured environments. This research introduces an adaptive scan-correlation approach for mobile robot localization that minimizes the dependence on feature models and improve the convergence characteristics of scan-correlation. In an effort to quantify the performance and facilitate the inter–comparison of results, an evaluation methodology was developed to quantify performance based on the principle of identifying and testing performance singularities.

The performance evaluation, presented in Chapter 5, assesses the adaptive scan–correlation technique as a viable alternative for dead-reckoning. Evaluating the performance at the system level enabled the inter-comparison of three dead-reckoning approaches; Scan-Matching, Inertial Navigation system, and Encoder-based Odometry. Assessment at this level, explored the performance characteristics of each of the solutions and helped identify the capabilities and limitations for each solution.

The performance evaluation was able to show, through simulation, that the adaptive scan–correlation approach is more resilient to errors that plague the other dead-reckoning solutions and was able to bound uncertainty in the system. Although the adaptive scan–correlation is a computationally efficient form of localization and exhibits good performance in complex environments, it is vulnerable to failure due to the lack of redundancy and its inability to measure the amount of uncertainty in the system at any given time. Therefore, the concept of deployment should consider these limitations. For instance, a small unmanned ground vehicle conducting operations in a confined space may only need scan-matching solutions. However, for long duration missions the adaptive scan–correlation should be fused with other techniques to take into account errors and measure uncertainty.

The performance analysis, presented in Chapter 6, shows three testing scenarios developed to analyze the performance characteristics of the adaptive scan-correlation algorithms with an emphasis on linear displacement. These testing scenarios facilitated the analysis how correspondence determination can jeopardize the integrity of the techniques. It was illustrated how these methods facilitate the inter-comparison of experimental results between the basic scan–correlation approach and the adaptive scan–correlation approach. Testing at the algorithmic level also unearths meta-level knowledge that may help identify failure conditions, allowing mechanisms to help these solutions overcome and/or avoid performance singularities.

The performance analysis of the two scan-correlation techniques has shown that the adaptive scan-correlation technique is a more proficient solution. It provides a more robust data association technique that enables it to overcome performance issues in complex, unstructured environments. This analysis also shows the adaptive scan-correlation algorithm produces a more stable solution with meta-level knowledge that indicates when performance singularities may be occurring. Future research is necessary here to evaluate how to apply this knowledge and develop mechanisms to compensate for errors that may be introduced into the scan-correlation process.

Appendices

A. High-Fidelity Simulation Testbed

Robotic simulation systems, such as Micrsoft Robotics Studio [3] and The Player Project [5], are commonly used in the development of the autonomous systems and advanced robotic algorithms. They provide a cost-effective tool that enables developers to customize repeatable testing scenarios to challenge specific aspects of a system over potential failure conditions under the same environmental conditions [36]. In order to provide convincing arguments about a system's performance and reliability, the simulation systems must be capable of capturing the stochastic nature of a real world environment.

Unified System for Automation and Robot Simulation (USARSim) [6] is an open-source package that provides a high-resolution, physics-based simulation that solves many of the practical problems faced by robotic simulators. Initially developed to support development of robotic algorithms in the urban search and rescue environment, USARSim has expanded its core functionality to provide a general-purpose, multi-agent simulation system with a set of unique characteristics unmatched by other simulation systems [44].

MOAST [4] is an open-source, turn-key hierarchical control system that was originally developed to promote the research of advanced robotic algorithms [11]. Based on the *4-D Real-time Control System* (4D/RCS) architecture [8], MOAST provides a modularized hierarchical framework that allows for the transparent transference of data between a matrix of real and virtual components. This framework is glued together through well-defined interfaces and communications protocols, and detailed specifications on individual subsystem input/output (I/O) that allows developers to freely swap components. Internal tools provide developers with state-by-state, time-stamped snapshots that allow researchers to quantitatively measure and classify the performance characteristics of new algorithms and the means to analyze the overall impact on the system's performance by means of comparison.

Since the validity of the results obtained from such algorithms are directly related to the accuracy and realism of the underlying simulation models, it is important that the sensors provide realistic data. Significant efforts on the validation of simulated models in USARSim have resulted in close correspondence between simulated data extracted from USARSim and their real world counterparts [14, 43]. Therefore, integration of these high-fidelity models with MOAST allows researchers to develop advanced robotic algorithms, classify their performance characteristics, and evaluate the overall impact of the algorithms on a robotic system before implementation on real robotic hardware.

Bibliography

[1] "ASTM Committee E54.08.01 on Homeland Security Operations; Operational Equippment; Robots." [Online]. Available: http://www.astm.org/COMMIT/SUBCOMMIT/E5408.htm

[2] "Department of Homeland Security Response Robot Performance Standards." [Online]. Available: http://www.nist.gov/el/isd/ks/response_robot_test_methods.cfm

[3] "Microsoft Robotics Studio." [Online]. Available: http://msdn.microsoft.com/robotics

[4] "The MOAST Project." [Online]. Available: http://www.sourceforge.net/projects/moast

[5] "The Player Project." [Online]. Available: http://playerstage.sourceforge.net/

[6] "USARSim." [Online]. Available: http://www.sourceforge.net/projects/usarsim

[7] in *Form and style for ASTM Standards.* ASTM International, March 2007.

[8] J. Albus and A. Meystel, *Engineering of Mind: An Introduction to the Science of Intelligent Systems*, J. Albus, A. Meystel, and L. Zadeh, Eds. John Wiley and Sons, Inc 2001, 2001.

[9] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-Squares Fitting of Two 3-D Point Sets," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 5, pp. 698–700, 1987.

[10] T. Bailey, "Mobile Robot Localisation and Mapping in Extensive Outdoor Environments," Ph.D. dissertation, University of Sydney, 2001.

[11] S. Balakirsky, F. M. Proctor, C. J. Scrapper, and T. R. Kramer, "A Mobile Robot Control Framework: From Simulation to Reality," in *SIMPAR '08: Proceedings of the 1st International Conference on Simulation, Modeling, and Programming for Autonomous Robots.* Berlin, Heidelberg: Springer-Verlag, 2008, pp. 111–122.

[12] P. J. Besl and N. D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.

[13] J. Borenstein, H. Everett, and L. Feng, *Navigating Mobile Robots: Systems and Techniques.* A.K. Peters Ltd., 1996.

[14] S. Carpin, M. Lewis, J. Wang, S. Balakirsky, and C. Scrapper, “USARSim: a Robot Simulator for Research and Education,” in *Proceedings of the IEEE 2007 International Conference on Robotics and Automation*, October 2007.

[15] A. Censi, “An ICP Variant Using a Point-to-Line Metric,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Pasadena, CA, May 2009. [Online]. Available: http://purl.org/censi/2007/plicp

[16] Chandran-Ramesh, M. and Newman, P., “Assessing Map Quality using Conditional Random Fields,” in *Proceedings of the Field and Service Robotics*, July 2007.

[17] Y. Chen and G. Medioni, “Object Modelling by Registration of Multiple Range Images,” *Image Vision Comput.*, vol. 10, pp. 145–155, April 1992.

[18] B. Consulting, “AGV Industry Next-Generation Technology Priorities,” National Institute of Standards and Technology (NIST), Tech. Rep., April 2006.

[19] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-whyte, and M. Csorba, “A Solution to the Simultaneous Localization and Map Building (SLAM) Problem,” *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 229–241, 2001.

[20] *Rescue Systems I*, USFA/NFA-RS1-SM1993 ed., FEMA, 1993.

[21] B. Gates, "A Robot in Every Home," *Scientific America*, no. 58-65, 2007.

[22] G. Grisetti, G. Tipaldi, C. Stachniss, W. Burgard, and D. Nardi, "Fast and Accurate SLAM with Rao-Blackwellized Particle Filters," *Robotics and Autonomous Systems*, 2006.

[23] J. Gutmann, W. Burgard, D. Fox, and K. Konolige, "An Experimental Comparison of Localization Methods," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'98)*, 1998.

[24] E. Incorporated, "Robotics in Manufacturing Technology Raodmap," Tech. Rep., November 2006.

[25] A. Kelly, B. Nagy, D. Stager, and R. Unnikrishnan, "An Infrastructure-Free Automated Guided Vehicle Based on Computer Vision," *IEEE Robotics and Automation Magazine*, pp. 24–34, September 2007.

[26] R. Lakaemper and N. Adluru, "Using virtual scans for improved mapping and evaluation," *Auton. Robots*, vol. 27, pp. 431–448, November 2009.

[27] J. Leonard and H. Durrant-Whyte, "Simultaneous Map Building and Localization for an Autonomous Mobile Robot," in *IEEE International Workshop on Intelligent Robots and Systems*, 1991, pp. 1442–1447.

[28] F. Lu and E. Milios, "Robot Pose Estimation in Unknown Environments by Matching 2D Range Scans," *Journal of Intelligent and Robotic Systems*, vol. 18, pp. 249–275, 1994.

[29] I. Mandour and M. El-Dakiky, "Inertial Navigation System Synthesis Approach and Gravity-Induced Error Sensitivity," in *IEEE Transactions on Aerospace and Electronic Systems*, 1988.

[30] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem," in *Proceedings of the AAAI National Conference on Artificial Intelligence.* Edmonton, Canada: AAAI, 2002.

[31] J. Nieto, T. Bailey, and E. Nebot, "Recursive Scan-Matching SLAM," *Robotics and Autonomous Systems*, vol. 55, no. 1, pp. 39–49, 2007.

[32] A. Nuchter, K. Lingeman, J. Hertzberg, and H. Surmann, "6D SLAM with Approximate Data Association," in *Proc. Intl. Conf. on Adv. Robotics*, July 2005.

[33] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann, "6D SLAM—3D mapping outdoor environments: Research Articles," *Journal of Field Robotics*, vol. 24, no. 8-9, pp. 699–722, 2007.

[34] S. Rusinkiewicz and M. Levoy, "Efficient Variants of the ICP Algorithm," in *Third International Conference on 3D Digital Imaging and Modeling (3DIM)*, June 2001.

[35] J. Salvi, C. Matabosch, D. Fofi, and J. Forest, "A Review of Recent Range Image Registration Methods with Accuracy Evaluation," *Image Vision Comput.*, vol. 25, no. 5, pp. 578–596, 2007.

[36] C. Scrapper, S. Balakirsky, and E. Messina, "MOAST and USARSim: A Combined Framework for the Development and Testing of Autonomous Systems," in *The International Society for Optical Engineering (SPIE '06)*, vol. 6230, 1988.

[37] A. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Proceedings of Robotics: Science and Systems*, Seattle, USA, June 2009.

[38] R. Smith, M. Self, and P. Cheeseman, "Estimating Uncertain Spatial Relationships in Robotics," pp. 167–193, 1990.

[39] S. Sukkarieh, "Low Cost, High Integrity, Aided Inertial Navigation Systems for Autonomous Land Vehicles," Ph.D. dissertation, University of Sydney, 2000.

[40] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann, 2002, to appear.

[41] I. Varsadan, A. Birk, and M. Pfingsthorn, "Determining map quality through an image similarity metric," in *RoboCup 2008: Robot Soccer World Cup XII*, ser. Lecture Notes in Computer Science, L. Iocchi,

H. Matsubara, A. Weitzenfeld, and C. Zhou, Eds. Springer Berlin / Heidelberg, 2009, vol. 5399, pp. 355–365.

[42] A. I. Wagan, A. Godil, and X. Li, "Map quality assessment," in *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, ser. PerMIS '08. New York, NY, USA: ACM, 2008, pp. 278–282.

[43] J. Wang, M. Lewis, S. Hughes, M. Koes, and S. Carpin, "Validating USARSim for Use in HRI Research," in *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting(HFES05)*, 2005.

[44] M. Zaratti, M. Fratarcangeli, and L.Iocchi, "A 3D Simulator of Multiple Legged Robots based on USARSim," in *Robocup 2006: Robot Soccer World Cup X, Springer, LNAI*, 2006.

[45] Z. Zhang, "Iterative Point Matching for Registration of Free-Form Curves and Surfaces," *Int. J. Comput. Vision*, vol. 13, no. 2, pp. 119–152, 1994.