



Sensor Data Integrity and Mitigation of Perceptual Failures

FA2386-10-1-4153

Period of Performance: 05/2011 - 05/2012

PI: Dr. Thierry Peynot

Investigators: Marcos P. Gerardo Castro, Rowan McAllister,
Ken Ho, Dr. Robert Fitch, Prof. Salah Sukkarieh

**Australian Centre for Field Robotics
The University of Sydney**

Contact: Dr. Thierry Peynot

Address: ACFR, Rose St. Building J04

The University of Sydney, NSW 2006, Australia

Email: tpeynot@acfr.usyd.edu.au

Phone: +61 2 9036 9193

Fax: +61 2 9351 7474

August 2012

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 24 AUG 2012	2. REPORT TYPE Annual	3. DATES COVERED 07-09-2010 to 06-05-2011	
4. TITLE AND SUBTITLE Sensor Data Integrity and Mitigation of Perceptual Failures		5a. CONTRACT NUMBER FA23861014153	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Thierry Peynot		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The University of Sydney, Rose St. Building J04, Sydney 2006, Australia, NA, NA		8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD, UNIT 45002, APO, AP, 96338-5002		10. SPONSOR/MONITOR'S ACRONYM(S) AOARD	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-104153	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			
13. SUPPLEMENTARY NOTES			
14. ABSTRACT This report describes basic research to analyze and investigate methods for automatic detection and mitigation of sensor and perception data. Specifically, the PI will investigate 1) Metrics for sensor performance in context using multi-modal redundancy and information theoretic approaches; 2) Detection, localization and mitigation of faults and failures using hierarchical causal reasoning for diagnosis; 3) Develop methods for using low-level (e.g., sensor data fusion filter) versus high-level (autonomous inference) processes for automated selection of data to use for perception.			
15. SUBJECT TERMS Intelligent Systems, Sensing , Adaptive Robotics, Fault Diagnosis and Recovery			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	
			18. NUMBER OF PAGES 133
			19a. NAME OF RESPONSIBLE PERSON

Abstract

This document is a progress report for the second year of the research project “Sensor Data Integrity and Mitigation of Perceptual Failures”, sponsored by the AFOSR/AOARD under agreement number FA2386-10-1-4153. The objective of this research is to better understand and promote integrity and dependability of unmanned ground vehicles (UGVs), with a focus on their perceptual systems. This will provide UGVs with the ability to achieve long-term autonomous operations in off-road environments, including in challenging conditions. We investigate methods to mitigate, detect and/or recover from perceptual failures and failures due to perception. The document is divided into four main sections.

First, we further develop the concept of laser-to-radar sensing redundancy for resilient perception in adverse environmental conditions (e.g. presence of smoke, airborne dust or heavy rain). We show that we can separate laser points due to dust or smoke clouds from the points corresponding to actual dense obstacles using a consistency test with data acquired by a mm-wave radar. A direct benefit is that the UGV does not misinterpret dust/smoke clouds or heavy rain particles as actual obstacles. A more general benefit is that using this technique, a UGV can keep building accurate environment models in clear environmental conditions, while maintaining resilient perception in adverse environmental conditions.

Second, we analyse the influence of different terrain geometry representations (in particular state-of-the-art techniques based on Gaussian Process regression) on the estimation of traversability of a ground vehicle. In particular, we discuss how these terrain geometry modelling techniques may whether mitigate or generate errors in the estimation of traversability. This preliminary study will drive further investigation into traversability estimation techniques that are more accurate and more robust to occlusions and sensing errors.

In the third part, we further develop the concept of resilient navigation through probabilistic modality reconfiguration that we introduced in the previous progress report in 2011. The update includes new experiments, and a comparison with a simpler thresholding technique to decide the most appropriate navigation modality. The benefit of this navigation modality reconfiguration technique is the online mitigation of, or recovery from, unpredictable errors such as control deviations, map failures and localisation faults.

Finally, in the last part of this document pushes the study to a higher level in the architecture

of a UGV system. We show that the platform's safety can be increased and the robot made more robust to control and localisation errors by achieving path planning with stochastic control, i.e. by *anticipating* possible errors at the planning stage. In the proposed method, the outcomes of desired control actions are learned from experience and represented statistically using Gaussian process regression models. We provide an experimental validation of this approach on a planetary rover.

Acknowledgments

This material is based on research sponsored by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

This work was also supported in part by the Australian Centre for Field Robotics (ACFR) and the NSW State Government.

The authors would like to thank Alen Alempijevic at the University of Technology Sydney (UTS) for providing the indoor robot used for the experiments in Chapter 4 and for his contribution to these experiments.

Disclaimer

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

Contents

Contents	i
1 Introduction	1
1.1 Laser-to-radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions [1]	2
1.2 Analysis of Terrain Geometry Representations for Traversability of a UGV [2]	3
1.3 Resilient Navigation through Probabilistic Modality Reconfiguration [3] . .	5
1.4 Motion Planning and Stochastic Control with Experimental Validation on a Rover [4]	6
1.5 New Multimodal Sensing Datasets	7
1.6 Publications and Communication of this Research	8
1.7 Related References	11
1.8 Outline	11
2 Laser-to-Radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions [1]	12
3 Analysis of Terrain Geometry Representations for Traversability of a Mars Rover [2]	18
4 Resilient Navigation through Probabilistic Modality Reconfiguration [3]	33
5 Motion Planning and Stochastic Control with Experimental Validation on a Planetary Rover [4]	49

Chapter 1

Introduction

The general objective of this research is to better understand and promote integrity and dependability of perceptual systems for unmanned ground vehicles (UGVs), and UGV systems in general, to provide them with the ability to achieve long-term autonomous operations in off-road environments, including (and in particular) in challenging conditions.

Aspects considered in this project are:

- Sensor Data Integrity: what (combination of) sensors for what challenging environments?
- Characterisation of Perceptual Failures and Failures due to Perception in a UGV system,
- Detection or Mitigation of Perceptual Failures (mostly using Multimodal Sensing Redundancy),
- Mitigation of Failures due to Perception (through Modality Reconfiguration).

This document is mostly articulated around the publications that resulted from this project in the 2011-2012 period of performance (see Section 1.6). The following sections provide a summary of the update on the research items that were initially introduced in the 2011 report of this project¹, as well as new research items introduced this year.

¹labelled (Project Report, 2011) in the rest of this document (see Sec. 1.7)

In (Project Report, 2011) we considered the problem of diagnosis for outdoor robotics, characterising failures and their possible causes and consequences within a UGV system, represented by the generic diagram in Fig. 1². In this report, we further consider the

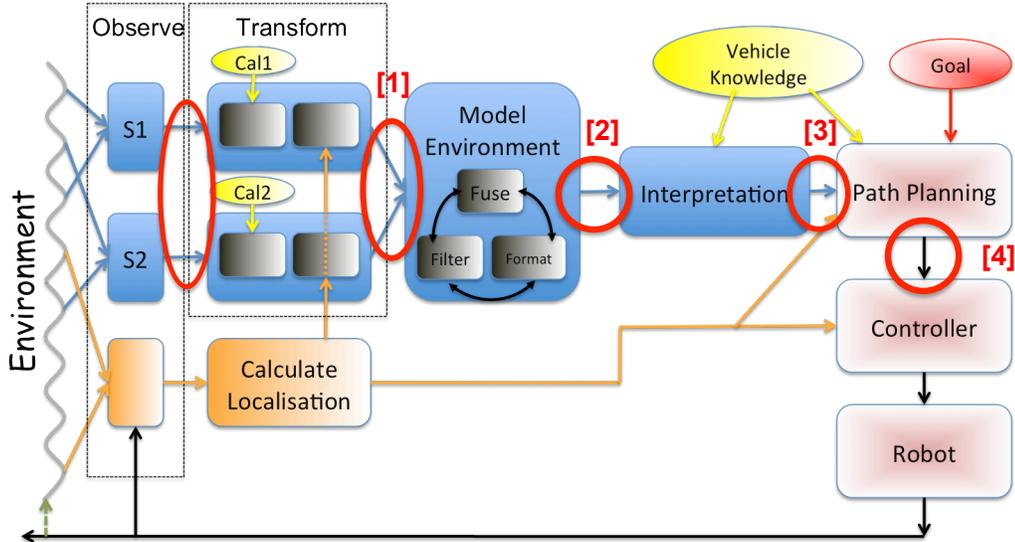


Figure 1.1: Functional components of a typical UGV System. The red ellipses indicate the parts of the system that are considered (separately) in this document. The associated labels (e.g. [2]) refer to the corresponding publications (see Sec. 1.6).

problems of perceptual failures and failures due to perception in this representative system, with experimental validation on three different types of platforms (which further illustrates the generality of the diagram).

1.1 Laser-to-radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions [1]

In (Project Report, 2011), we proposed a preliminary exploration of multi-modal redundancy for the mitigation and/or detection of perceptual failures, specifically between laser and radar. This initial study pointed out the benefit for a UGV, which could safely navigate in the presence of airborne dust, smoke or heavy rain. [1] further develops this study. It proposes a method to separate 3D data points that are considered as *consistent* between

²In this diagram only 2 (different) sensors are represented for simplicity, however, the number and types of sensors do not suffer from such limitation in reality.

observations of a laser range finder and a mm-wave radar, from data points that are *inconsistent*. Consistent observations mean that both types of sensors detect the same target (at the same geographical location), while inconsistent points are typically found when environmental elements such as airborne dust or smoke particles are detected by the laser (i.e. they strongly attenuate laser EM waves) and not by the radar. Experiments conducted using the *Marulan datasets* (SDI Report, 2009) show how this method can separate dust and smoke clouds from raw laser data, preventing traditional perception systems from considering them as obstacles. Exploiting the best of laser and radar sensors and their combination opens the door to resilient navigation of UGVs in challenging environmental conditions. The model of the environment built by the UGV can still be as accurate as models typically built from laser data in clear conditions (i.e. without dust, smoke or rain), while safe navigation can still be maintained in adverse conditions (i.e. in the presence of dust, smoke) thanks to the radar.

Future Work

Ongoing and future work on this subject include the use of a better sensor model for the radar. Research is limited in this area, and a better understanding of radar data (and the differences with laser data) would not only improve radar-based perception that is needed in adverse environmental conditions, it would also allow for a more appropriate and trustful comparison between laser and radar data. Future work will also look at learning how to separate consistent and inconsistent data acquired by different sensing modalities³ automatically.

1.2 Analysis of Terrain Geometry Representations for Traversability of a UGV [2]

In this study, the task of interpretation of the sensing data consists in analysing the terrain traversability (i.e. predicted vehicle response on the terrain, or predicted attitude and configuration of the platform). This terrain traversability analysis allows the UGV to provide a map of difficulty of the unstructured terrain (where a piece of terrain with highest difficulty will be considered as an obstacle), which is necessary to the path planner (see Fig. 1.2). We

³i.e. using a distinct physical process, for example operating at different electromagnetic frequencies.

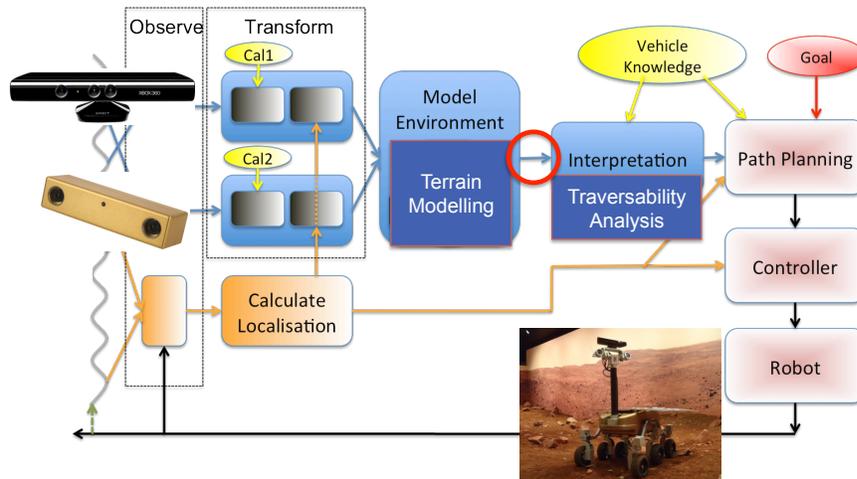


Figure 1.2: In this example of a particular UGV System architecture, the robot, a planetary rover, needs to build a terrain model to achieve the terrain traversability analysis required for its motion planning.

consider and discuss different methods for representing the geometry of the terrain from sensing data (typically 3D point clouds provided by a range sensor, e.g. stereovision or RGB-D camera). We discuss how these terrain modelling techniques may mitigate, or on the contrary generate, errors in the traversability analysis due to inaccurate or incomplete representations of the terrain.

This research was published in [2]. The abstract of the paper follows.

Abstract

For a planetary rover to successfully traverse across unstructured terrain autonomously, one of the major challenges is to assess its local traversability such that it can plan a trajectory to achieve its mission goals efficiently while minimising risk to the vehicle itself. This paper aims to provide a comparative study on different approaches for representing the geometry of Martian terrain for the purpose of evaluating terrain traversability. An accurate representation of the geometric properties of the terrain is essential as it can directly affect the determination of traversability for a ground vehicle. We explore current state-of-the-art techniques for terrain estimation, in particular Gaussian Processes (GP) in various forms, and discuss the suitability of each technique in the context of an unstructured Martian terrain. Furthermore, we present the limitations of regression techniques in terms

of spatial correlation and continuity assumptions, and the impact on traversability analysis of a planetary rover across unstructured terrain. The analysis was performed on datasets of the Mars Yard at the Powerhouse Museum in Sydney, obtained using the onboard RGB-D camera.

Ongoing and Future Work

Our ongoing work in this area focuses on directly learning vehicle response on the terrain from experience (as opposed to learning the terrain geometric model and predicting the vehicle response on this model). This will improve the traversability estimation, in particular in areas of poor or missing sensing data.

Future work will use machine learning to better anticipate (and accommodate for) situations that are difficult to predict from direct perception, in particular terrain deformations due to the interaction with the rover. These situations would be typically regarded as errors in the terrain model (inconsistencies between observation and terrain model built *a priori*) when monitoring the status of the rover, and are rarely accounted for.

1.3 Resilient Navigation through Probabilistic Modality Reconfiguration [3]

In (Project Report, 2011) we also proposed an investigation into techniques of reconfiguration to recover from possible failures (or mitigate them), using a multi-modality approach for navigation. The concept was implemented and experimented on an indoor robot, although it is equally applicable to UGVs (as shown by prior work from the PI). In 2012, new series of experiments were conducted, with a more accurate localisation system, allowing for a better proof of concept that could be focussed on each source of error independently. A comparison was also made with a simpler approach using thresholds on distances between the robot and closest obstacles. It was shown that the latter approach was undesirable, due to the occurrence of oscillations in the modality recommendation.

This research, including these 2012 updates, was published in [3]. The abstract of the paper follows.

Abstract

This paper proposes an approach to achieve resilient navigation for indoor mobile robots. Resilient navigation seeks to mitigate the impact of control, localisation, or map errors on the safety of the platform while enforcing the robot’s ability to achieve its goal. We show that resilience to *unpredictable* errors can be achieved by combining the benefits of independent and complementary algorithmic approaches to navigation, or *modalities*, each tuned to a particular type of environment or situation. In this paper, the modalities comprise a path planning method and a reactive motion strategy. While the robot navigates, a Hidden Markov Model continually estimates the most appropriate modality based on two types of information: context (information known *a priori*) and monitoring (evaluating unpredictable aspects of the current situation). The robot then uses the recommended modality, switching between one and another dynamically. Experimental validation with a SegwayRMP-based platform in an office environment shows that our approach enables failure mitigation while maintaining the safety of the platform. The robot is shown to reach its goal in the presence of: 1) unpredicted control errors, 2) unexpected map errors and 3) a large injected localisation fault.

1.4 Motion Planning and Stochastic Control with Experimental Validation on a Rover [4]

This study concentrates efforts at a level of the UGV system further away from the input sensing data. We argue the need for considering control uncertainty at the stage of path planning, since the outcome of an executed action is not deterministic, and therefore cannot be predicted consistently and accurately. This is often due to imperfections of the robot controller. However, unexpected deviations from desired control actions can also be caused (or increased) by an erroneous or incomplete terrain model, and/or localisation errors during the execution of control actions. In the method proposed and implemented in [4], although we do not identify the causes of control action errors, we do mitigate for some localisation errors in our system (as well as “pure” control errors) by learning the outcome of the available control actions from experience, and constructing policies for navigation that account for the stochastic nature of the control actions.

This research will be published in October in the proceedings of the IEEE/RSJ IROS conference [4]. The abstract of the paper follows.

Abstract

Motion planning for planetary rovers must consider control uncertainty in order to maintain the safety of the platform during navigation. Modelling such control uncertainty is difficult due to the complex interaction between the platform and its environment. In this paper, we propose a motion planning approach whereby the outcome of control actions is learned from experience and represented statistically using a Gaussian process regression model. This model is used to construct a control policy for navigation to a goal region in a terrain map built using an on-board RGB-D camera. The terrain includes flat ground, small rocks, and non-traversable rocks. We report the results of 200 simulated and 35 experimental trials that validate the approach and demonstrate the value of considering control uncertainty in maintaining platform safety.

1.5 New Multimodal Sensing Datasets

To begin to address the issues of sensor data integrity, in a technical effort preliminary to this project in 2008 that was sponsored by AFRL, synchronised data were gathered from a representative UGV platform using a wide variety of sensing modalities (see (SDI Report, 2009) in Sec. 1.7). These included multiple 2D laser range finders, a visual camera, an infrared camera, and a mm-wave radar, in addition to a dGPS/INS unit for accurate localisation. These large volumes of data were made available to the community to evaluate the performance of perception algorithms. This was later published as a “data paper” in one of the top-end robotics journals: the International Journal of Robotics Research (see (IJRR, 2010)).

As this research was being developed, it became clear that we should extend this library of datasets, both in terms of sensing modalities and challenging conditions. Therefore, we recently acquired new datasets using an extended suite of sensors. In addition to the same sensors mentioned above, we used stereovision, a 3D laser (Velodyne HDL-64E), a 360° visual camera (Point Grey Ladybug 2). We also used two UGV platforms to collect the

data. The challenging conditions were also extended to sunset, sudden switch from artificial lighting to complete or partial darkness, and fire (or presence of strong heat waves in the air, that can constitute an obscurant for an IR camera as opposed to a visual camera).

At the end of this period of performance, these datasets are still being processed and analysed for accuracy and completeness. Once this analysis is completed and the datasets are thoroughly documented, we plan to make this data available to AFRL, other partners and the research community in the near future (within the next period of performance). However, some of these data can already be made available to AFRL scientists upon request.

1.6 Publications and Communication of this Research

This section lists the publications that resulted from full or partial funding from this grant, as well as presentations at conferences and invited presentations, during the current period of performance.

Conference Publications

- [2] K. Ho, T. Peynot and S. Sukkarieh, “Analysis of Terrain Geometry Representations for Traversability of a Mars Rover”, *11th NCSS/NSSA Australian Space Science Conference*, Canberra, Australia, September 2011.
- [3] T. Peynot, R. Fitch, R. McAllister and A. Alempijevic, “Resilient Navigation through Probabilistic Modality Reconfiguration”, *12th International Conference on Intelligent Autonomous Systems (IAS)*, Jeju Island, Korea, June 2012.
- [4] R. McAllister, T. Peynot, R. Fitch and S. Sukkarieh, “Motion Planning and Stochastic Control with Experimental Validation on a Planetary Rover”, to appear in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vilamoura, Portugal, October 2012. (Accepted 2 July 2012).

Workshop Publication (Peer-Reviewed)

- [1] M. P. Gerardo Castro and T. Peynot, “Laser-to-Radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions”, *Beyond Laser and Vi-*

sion: Alternative Sensing Techniques for Robotics Perception, Workshop, Robotics: Science and Systems (RSS), Sydney, Australia, 11-12 July 2012.

Conference Presentations

- “Laser-to-Radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions”, presented by M. P. Gerardo Castro at the Workshop *Beyond Laser and Vision: Alternative Sensing Techniques for Robotics Perception, Robotics: Science and Systems (RSS)*, Sydney, Australia, 11-12 July 2012.
- “Resilient Navigation through Probabilistic Modality Reconfiguration”, presented by T. Peynot at the *12th International Conference on Intelligent Autonomous Systems (IAS)*, Jeju Island, Korea, 29 June 2012.
- “Analysis of Terrain Geometry Representations for Traversability of a Mars Rover”, presented by K. Ho at the *11th NCSS/NSSA Australian Space Science Conference*, Canberra, Australia, September 2011.
- “Persistent Perception for Long-term Autonomy of Ground Vehicles”, presented by T. Peynot at *Workshop on Long-term Autonomy, IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, 9 May 2011.
- “Autonomous Reconfiguration of a Multi-Modal Mobile Robot”, presented by T. Peynot at *Workshop on Automated Diagnosis, Repair and Re-Configuration of Robot Systems*, IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China, 9 May 2011.

Invited Presentations

The following lists invited presentations given by the PI in the current period of performance that included some research developed in the context of this project.

- “Sensor Data Integrity and Mitigation of Perceptual Failures”, Air Force Office of Scientific Research (AFOSR) Program Reviews, Arlington, VA, USA, 25 January 2012.

- “Dependable Autonomy of Planetary Rovers”, Congreso Internacional en Aeronautica: Avances en desarrollo e innovacion tecnologica, Universidad Militar Nueva Granada, Bogota, Colombia, 12 October 2011.
- “Perception Integrity and Dependable Autonomy for Mobile Robots”, Universidad Militar Nueva Granada, Bogota, Colombia, 5 October 2011.
- “Perception Integrity and Dependable Autonomy for Mobile Robots”, Model-based Embedded and Robotic Systems group (MERS), Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, 9 June 2011.
- “Sensor Data Integrity and Mitigation of Perceptual Failures”, Air Force Office of Scientific Research (AFOSR) Robust Computational Intelligence Program Review, Arlington, VA, USA, 7 June 2011.

Collaborations and other Interactions

Collaboration with the Centre for Autonomous Systems at the University of Technology (UTS), Sydney, made the experimental validation of [4] possible.

Interesting research questions and discussions came out of the two AFOSR program reviews the PI attended during this period of performance, in June 2011 and January 2012, in Arlington, VA.

At the Robust Computational Intelligence (RCI) program review in June 2011, key people I had the pleasure to exchange ideas with were: Tom Russell, director, Peter Friedland (host and interim project manager) and David Atkinson (former RCI program manager), in addition to my colleagues PIs and co-PIs of the RCI program.

At the extended AFOSR Cognition, Decision, and Computational Intelligence program review in January 2012, key AFRL people I had the pleasure to discuss this research and other related research with included: Jay Myung (new program manager), Peter Friedland (AFRL Advisor), Michael A. Vidulich (AFRL) and Kevin Gluck (AFRL), in addition to my colleagues PIs and co-PIs of the different programs.

1.7 Related References

These references are reports from previous AFRL/AOARD grants related this project, or prior progress reports.

- **(Project Report, 2011)** T. Peynot et al., “Sensor Data Integrity and Mitigation of Perceptual Failures - Progress Report”, Technical Report for AFRL/AFOSR/AOARD, Australian Centre for Field Robotics, The University of Sydney, May 2011.
- **(SDI Report, 2009)** T. Peynot, S. Terho and S. Scheduling, “Sensor Data Integrity: Multi-Sensor Perception for Unmanned Ground Vehicle”, Technical Report ACFR-TR-2009-002, Australian Centre for Field Robotics, The University of Sydney, February 2009.
- **(IJRR, 2010)** T. Peynot, S. Scheduling, and S. Terho, “The Marulan Data Sets: Multi-Sensor Perception in Natural Environment with Challenging Conditions”, International Journal of Robotics Research, vol. 29, no. 13, pp. 1602-1607, November 2010.

1.8 Outline

The following chapters contain the publications summarised above, and listed in Sec. 1.6.

Chapter 2

Laser-to-Radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions [1]

by M. P. Gerardo Castro and T. Peynot,

in *Workshop Beyond Laser and Vision: Alternative Sensing Techniques for Robotics Perception, Robotics: Science and Systems (RSS)*,

Sydney, Australia, 11-12 July 2012.

Laser-to-Radar Sensing Redundancy for Resilient Perception in Adverse Environmental Conditions

Marcos P. Gerardo Castro, Thierry Peynot,
Australian Centre for Field Robotics (ACFR)
The University of Sydney, NSW 2006, Australia
m.castro@acfr.usyd.edu.au, tpeynot@acfr.usyd.edu.au

I. INTRODUCTION

This work focuses on the development of reliable perception systems for outdoor unmanned ground vehicles (UGV), in particular in adverse environmental conditions (e.g. presence of airborne dust, smoke, thick fog or rain). The problem of modelling and mitigating *systematic* errors in perception models, such as sensor measurement errors or sensor misalignment, has been extensively studied by robotics researchers and thorough solutions have been proposed (e.g. [1] for range sensors such as laser range finders (LRF) or radars). However, the main remaining challenge lies in *interpretation* errors. These errors can be very random, are difficult to predict, and can often be orders of magnitude larger than the systematic errors mentioned above. Arguably, a reliable perception system should use different sensor modalities [3], [2], especially for outdoor operations. As these modalities sense the environment using different physical processes, they also respond differently to environmental conditions. For example, a mm-wave radar has excellent properties of penetration through heavy dust and smoke contrary to a laser, and an infrared camera can see through smoke, contrary to a visual camera. Therefore, a more reliable perception system can be obtained by intelligently combining the data provided by such different sensing modalities [4].

While the fusion of observations made by a laser and a radar in clear conditions, e.g. without the presence of challenging conditions such as dust or smoke, is straightforward when a good sensor error model is available [1], it relies on the assumption (or *pre-condition*) that the two sensors actually detect the same targets in the environment. If, for example, a LRF does not see through a heavy dust cloud while a radar does, this assumption does not hold any more. Therefore, in such a situation data fusion should not be executed, at least not in its traditional form. Consequently, to be robust to adverse environmental conditions, the perception system should have the ability to verify this assumption of data consistency prior to fusion. Another advantage to this ability is that the data provided by a LRF can be conveniently filtered, separating points returned because of dust or smoke that a radar would hardly be affected by. The radar could then ensure that detection of actual obstacles and terrain modelling stay operational, albeit less accurate (since the radar accuracy is typically not comparable to the laser's as described in Table I).

Recently, laser range finders capable of returning multiple echoes for each emitted pulse have been introduced commercially (e.g. the Sick LMS5xx series [5] or LD-MRS [6] for automotive applications). Although this ability has made such laser sensors more robust to adverse environmental conditions (e.g. compared to the LMS2xx series), they cannot provide a full solution of the problem. Because of the level of attenuation of the laser signal, a mm-wave radar will still be able to penetrate better through obscurants such as heavy dust that would eventually block laser signals [7], [8]. Moreover, an analysis of pre-conditions for laser-radar fusion and for separating dense objects from such obscurants would still be required to obtain a resilient navigation of the UGV.

The idea of using laser-radar data comparison for perception in the presence of airborne dust was introduced in previous work by the authors. However, if [9] delivered a proof of concept with promising initial results, this work had several limitations: 1) the two sensors were considered perfectly aligned, allowing for a direct comparison of the measured ranges they provide for each bearing angle, 2) the laser-radar data comparison was specifically designed and used as an airborne dust filter, 3) this filter was demonstrated on only one particular dataset. In practice, not only is the alignment assumption a strong constraint on the system, but such alignment is practically nearly impossible unless the two sensors use the same mirror and scanning mechanism. The technique proposed in this paper does not require that the sensors are perfectly aligned, instead it uses a 6-DOF calibration allowing to correct the mis-alignment of the sensors. The comparison of the data can then be realised in a coordinate frame related to the body of the vehicle (instead of one of the sensor frames as in [9]), accounting for the extrinsic calibration of the sensors. In this paper we also exploit more information from the spectrum provided by the radar, allowing for a closer comparison between the two types of data. Finally, if the technique can also be used as a dust filter, it is not designed as such specifically, so that other causes of inconsistencies can be detected as well. Some of these causes will be discussed below.

The paper is organized as follows. In Sec. II, we discuss the methodology to perform the laser-radar consistency test. Sec. III presents an experimental study to characterise the laser-radar distance. In Sec. IV, we describe results measuring the consistency test in scenarios with the presence of airborne

dust, smoke or none of the above.

II. LASER-RADAR REDUNDANCY

In order to compensate for the mis-alignment of the laser and radar sensors, we need to perform an extrinsic calibration of the relative transformation between the two sensors (or the transformation between each sensor and a frame linked to the body of the vehicle, which we will call the *body frame*). In this paper we use the calibration technique described in [1], which can achieve a joint extrinsic calibration of multiple exteroceptive range-based sensors such as lasers and radar. Since the configurations of the sensors are different, only a (common) part of a synchronised pair of laser-radar scans contain points that can be considered *consistent*¹. This part can be seen as a “common footprint” (or “footprint overlap”) of the two scans and can be conveniently expressed as a range of bearing angles for each type of scan. Hereafter, all comparison of laser and radar points is made within this common part of the scans. Another important thing to consider during this comparison is the range resolution of the two sensors. As described in Table I, radar resolution is much bigger than the laser resolution.

The rest of the process can happen systematically on-line. Sec. II-A describes how target data points are extracted from the radar raw data (i.e. noise removal). Then, Sec. II-B shows how radar and laser points are effectively compared after their transformation into the body frame.

TABLE I
RANGE SENSOR SPECIFICATIONS

Sensor (model)	Maximum range	Range resolution	Horizontal FOV	Angular resolution	Scanning rate
Horizontal Laser (Sick LMS291)	80 m	0.01 m	180°	0.25°	≈18 Hz
Radar (Custom built at ACFR)	40 m	0.2 m	360°	≈1.90°	≈3 Hz

A. The Radar Data

For each bearing angle the radar provides an FFT (Fast Fourier Transform) spectrum. Using the “radar equation” [7] this spectrum can be mapped to a function of intensity vs. range. Most robotics applications only use the highest peak of that spectrum as a range value provided by the radar (such as in our prior work in [9]). However, this leads to the loss of a significant amount of useful information contained in the rest of the spectrum. As an example, [10] exploited the shape of this spectrum to make a more accurate estimation of the ground. The resulting ground estimation was significantly more accurate than when using the highest peak of the spectrum only. However, this particular technique can only be used if a model of the spectrum profile obtained for a given target (such as a roughly flat piece of ground) is known a priori. In order to make a “fair” comparison of the radar points

¹The adjective consistent will be used to refer to the local agreement between laser and radar observations.

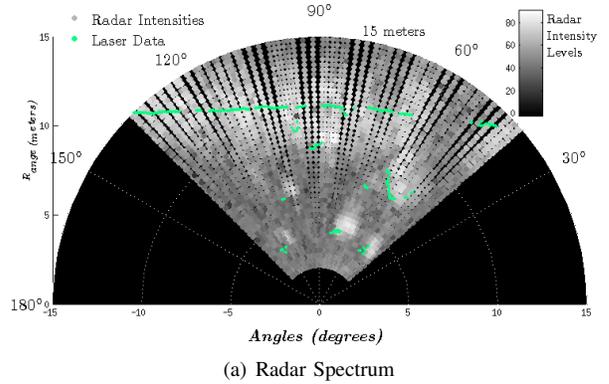


Fig. 1. (a) Radar Spectrum, coloured by intensities from black to white. The corresponding laser points are showed in green. (b) Laser Projected in a visual image from the the platform at the same area

with observed laser points, in this paper we extract other peaks (local maxima) from the spectrum, in addition to the highest peaks (the global maxima), see Fig. 1(a). This will provide us with a better resolution in the discrimination of laser-radar data. First, for each radar bearing angle, all intensity peaks above a lower threshold of intensity are extracted from the radar spectrum. The lower threshold of intensity was defined in order to minimise the radar noise. Then, given that: a) the laser provides much more accurate data than the radar, b) we know that generally both sensors detect the same targets in clear conditions, c) we have an accurate calibration of the sensors and a very accurate localisation on our robot, we used the laser data as a reference in large datasets of a rural *static* environment to determine a relevant criteria for an automatic extraction of the peaks from the noise in the radar data. For our radar, extracting peaks that have an intensity above 55% of the intensity of the highest peak was found to be appropriate. From the points extracted using the threshold criteria we then defined the radar *candidate peaks*. A radar point can be considered as a candidate peak if a laser point is closer (in terms of 3D euclidean distances) to the radar point/peak than the highest peak of the current radar spectrum.

B. Laser-Radar Comparison (Consistency Test)

The actual comparison between laser points and candidate radar peaks relies on the computation of the 3D euclidean distance between each laser point and the closest radar target

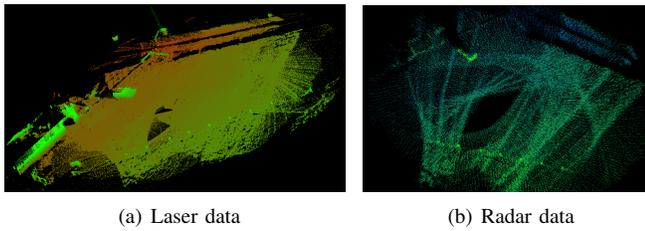


Fig. 2. View from the top of the scene observed in clear conditions by the four lasers on the Argo (a) and the radar (b). Points are coloured by elevation. We can see the posts of a fence at the bottom and a shed on the left of (a). The area is about $56 \times 55 m^2$.

point found in the synchronised scan, which will be called the *laser-radar distance*. A model of the laser-radar distance based on 3D distance comparison (which will be described below) between laser and radar points was used to decide whether the laser and radar are observing the same target. The main reasons for not observing the same target (i.e. laser-radar measurement discrepancy) are the following:

- the laser actually detects dust, smoke or rain particles that the radar waves penetrate through,
- the perception is inconsistent because of the material the target is made of (e.g. the radar may detect the presence of a window that the laser sees through and therefore does not detect),
- the relative extrinsic calibration between the laser and the radar is wrong,
- the echo returned by the sensor is the result of a multi-path effect (see [7], [8]).

To determine an appropriate threshold on the 3D distance between comparable laser and radar points, we used a dataset in clear conditions in a rural environment (see Fig. 2), limiting the risks of multi-path or distinct reaction of the radar and the laser to particular materials. Since in these conditions a close match should always be found, the dataset (containing about 1.7 million laser points) could be used as a reference.

Fig. II-B shows the number of *inconsistent* points for a varying value of distance threshold δ (i.e. number of laser points for which the closest radar peak was at a distance superior to δ). A distance threshold of $\delta = 0.8m$ was found to be appropriate. With this threshold in the static environment used as reference, only about 0.5% of the points were inconsistent.

Section III and IV show an experimental study to characterise the laser-radar distance and different examples and applications of the laser-radar comparison.

III. EXPERIMENTAL SETUP

The experiments were conducted with the Argo UGV, an 8 wheel skid-steering platform (see Fig. 4) equipped with a reliable navigation system composed of a Novatel SPAN (Synchronised Position Attitude & Navigation) System and a Honeywell Inertial Measurement Unit. This unit usually provides a 2-cm accuracy localisation, with a constant update of the estimated uncertainties on this solution.

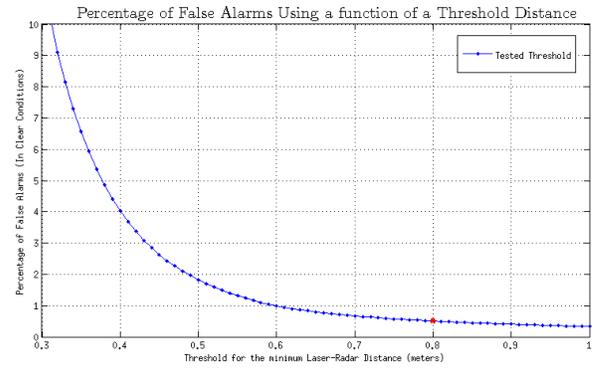


Fig. 3. Percentage of *inconsistent* points vs. threshold on the laser-radar distance (in metres). An 0.5% error was found with a threshold at $0.8m$ (blue cross).



Fig. 4. The Argo UGV and its sensors.

The following exteroceptive sensors were mounted on the vehicle (Fig. 4):

- 4 Sick LMS291/221 laser range scanners, with 180° field of view (FOV), 0.25° angular resolution, and a range resolution of $0.01m$.
- a $94GHz$ Frequency Modulated Continuous Wave (FMCW) Radar, custom built at ACFR for environment imaging, with 360° FOV, 2° angular resolution and a range resolution of $0.2m$,
- a visual camera and an infrared camera.

The Laser indicated in Fig. 4 was only roughly aligned with the Radar to have a similar perspective of the environment, therefore this laser was chosen to provide the data to be compared with the radar data (recall that only a rough physical alignment is sufficient, as mentioned earlier, as long as an extrinsic calibration between the two sensors is available). Fig. 5 shows an example of scans provided by these two sensors.

The experiments were conducted with the *Marulan Datasets* described in [11]. We used various datasets with the vehicle driven around two different areas. Each dataset featured the presence of airborne dust (Fig. 8), smoke (Fig. 9), rain, or

none of the above (i.e. *clear* conditions). The environment was not known by the vehicle a priori.

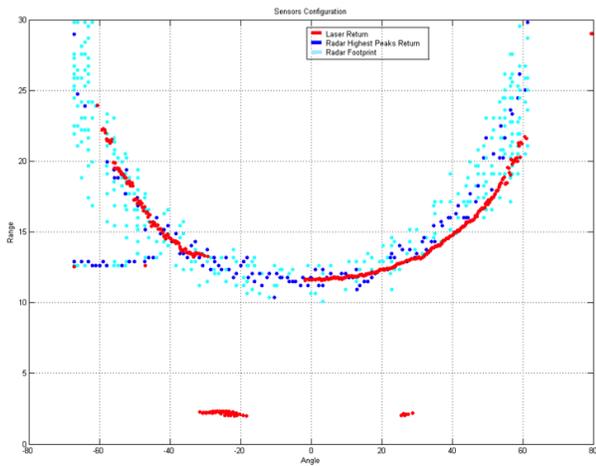


Fig. 5. Example of laser and radar scans displayed as range vs. bearing angle. Red points are laser returns while blue points are radar peaks (the highest peaks for each bearing angle are shown in dark blue). Note the laser returns due to dust at shorter range, which are clearly inconsistent with the radar measurements.



Fig. 6. Experiments with adverse environmental conditions including presence of airborne dust.

IV. RESULTS

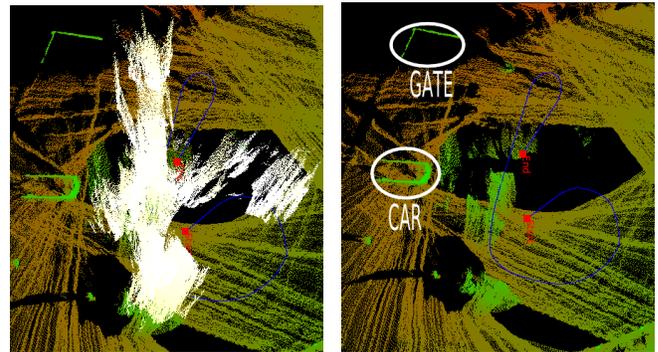
In these experiments, synchronised pairs of laser and radar scans are compared to separate consistent and inconsistent points. In practice, since the laser scanner has a higher scanning rate than the radar scanner (see Table I), for each laser scan the closest radar data available in time is used for the comparison and the consistency check.

Fig. 8 shows an experiment realised in the same area as in Fig. 2 but with presence of heavy airborne dust. We can see that most dust points in the laser data have been well cleaned out from the dataset, after being found inconsistent with the radar data. However, some dust points returned by the laser have remained, as they were too close to the ground, which was still seen by the radar, to be called inconsistent.

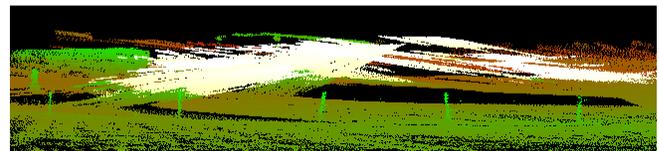
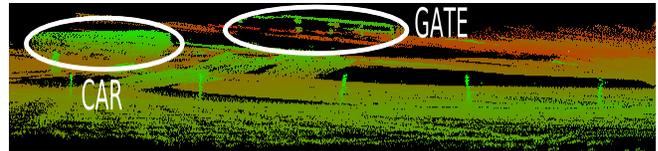
Fig. 9 shows another experiment, conducted in a different area (a more natural and unstructured environment with surrounding trees), with presence of smoke. It shows how smoke also significantly affects the laser data and how the consistency test with the radar data allows for an effective separation of the smoke cloud.



Fig. 7. Visual Image from the platform perspective, where results from Fig. 8 is shown.

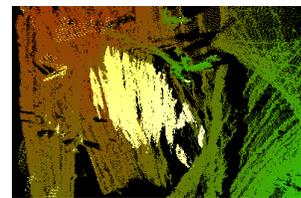


(a) All Laser points, bird's eye view. (b) Without the inconsistent points

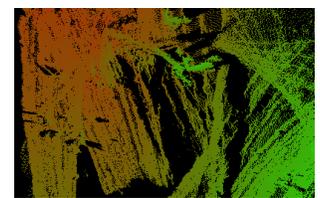


(c) Side view. Top: all points. Bottom: *consistent* points only.

Fig. 8. Experiment with heavy airborne dust (see Fig. III). Points are coloured by elevation. The laser points found to be *consistent* were coloured from green to red, while *inconsistent* points were coloured from yellow to white. The blue line shows the path followed by the platform while collecting this dataset.



(a) All Laser points.



(b) *Consistent* points only.

Fig. 9. Experiment with smoke, bird's eye view. Points are coloured by elevation. The laser points found to be *consistent* were coloured from green to red, while *inconsistent* points were coloured from yellow to white.

V. DISCUSSION

The method presented in this paper enables to maintain the safe operation of a UGV in the presence of adverse environmental elements such as airborne dust or smoke, which are strong obscurants for common robotic sensing modalities such as a laser or a visual camera. When dust or smoke are present and block the laser perception, the UGV may still go through the obscurant cloud, with the radar allowing for a persistent obstacle detection. On the other hand when no obscurant cloud is present, laser perception will be preferred since it is more accurate compared with the radar data.

In the experiments presented in this paper we have observed that some dust/smoke points may not be labelled as inconsistent when they are too close to dense obstacles, as their discrimination is limited by the resolution and the noise of the radar data.

Another situation that this method may not be able to identify is when airborne dust or smoke particles are detected by the laser in the immediate proximity of radar returns due to multi-path effect. This is because in such situation the system will consider these radar returns as a confirmation that the target detected by the laser is in fact a dense object (therefore a potential obstacle for the UGV). To overcome this situation another modality such as visual or infrared can be used.

The proposed method relies on the availability of an accurate exteroceptive calibration between the laser and the radar. If the calibration is jeopardised during a mission of the UGV (for example one of the sensors is accidentally put out of place), the consistency test might reject a large part of the laser data even in clear conditions. Consequently, the UGV would have to rely entirely and systematically on the radar data (which is typically less accurate). However, such situation could be recognised over time since the inconsistency between the laser and radar data would then be very stable and geometrically constant. This could let the system distinguish this case from the presence of dust or smoke for example.

A sensor model that accounts for uncertainties will be introduced in future work. Uncertainties in the comparison test will also be analysed (see Fig. II-B) by considering context information, e.g. facing scenarios where dust is close to the ground.

ACKNOWLEDGEMENTS

This work was supported in part by the Australian Centre for Field Robotics (ACFR) and the NSW State Government. This material is based on research sponsored by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

REFERENCES

- [1] J. P. Underwood, A. Hill, T. Peynot, and S. J. Scheduling, "Error modeling and calibration of exteroceptive sensors for accurate mapping applications," *Journal of Field Robotics, Special Issue: Three-Dimensional Mapping, Part 3*, vol. 27, no. 1, pp. 2–20, January/February 2010.
- [2] C. Urmson et al., "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.
- [3] C. Thorpe et al., "Dependable perception for robots," in *Proceedings of International Advanced Robotics Programme IEEE*. Seoul, Korea: Robotics and Automation Society, May 2001.
- [4] A. Kelly et al., "Toward reliable road autonomous vehicles operating in challenging environments," *International Journal of Robotics Research*, vol. 25, pp. 449–483, May/June 2006.
- [5] SICK Inc., "SICK LMS5xx Laser Measurement Technology," <https://www.mysick.com/partnerPortal/ProductCatalog/DataSheet.aspx?ProductID=75420#>, 2012.
- [6] —, "SICK LD-MRS Laser Measurement Technology," <https://www.mysick.com/partnerPortal/ProductCatalog/DataSheet.aspx?ProductID=34057#>, 2012.
- [7] G. Brooker, *Sensors for Ranging and Imaging*. SciTech Publishing, Inc., 2009.
- [8] J. Ryde and N. Hillier, "Performance of laser and radar ranging devices in adverse environmental conditions," *Journal of Field Robotics*, vol. 26, no. 9, pp. 712–727, September 2009.
- [9] T. Peynot, J. Underwood, and S. Scheduling, "Towards reliable perception for unmanned ground vehicles in challenging conditions," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009.
- [10] G. Reina, J. Underwood, G. Brooker, and H. Durrant-Whyte, "Radar-based perception for autonomous outdoor vehicles," *Journal of Field Robotics*, vol. 28, no. 6, pp. 894–913, November/December 2011.
- [11] T. Peynot, S. Scheduling, and S. Terho, "The Marulan Data Sets: Multi-Sensor Perception in Natural Environment with Challenging Conditions," *International Journal of Robotics Research*, vol. 29, no. 13, pp. 1602–1607, November 2010.

Chapter 3

Analysis of Terrain Geometry Representations for Traversability of a Mars Rover [2]

by K. Ho, T. Peynot and S. Sukkariéh,
in *11th NCSS/NSSA Australian Space Science Conference*,
Canberra, Australia, September 2011.

Analysis of Terrain Geometry Representations For Traversability of a Mars Rover

Ken Ho*, Thierry Peynot* and Salah Sukkarieh*

* *Australian Centre for Field Robotics, J04, University of Sydney, NSW, Australia, 2006*

Summary: For a planetary rover to successfully traverse across unstructured terrain autonomously, one of the major challenges is to assess its local traversability such that it can plan a trajectory to achieve its mission goals efficiently while minimising risk to the vehicle itself. This paper aims to provide a comparative study on different approaches for representing the geometry of Martian terrain for the purpose of evaluating terrain traversability. An accurate representation of the geometric properties of the terrain is essential as it can directly affect the determination of traversability for a ground vehicle. We explore current state-of-the-art techniques for terrain estimation, in particular Gaussian Processes (GP) in various forms, and discuss the suitability of each technique in the context of an unstructured Martian terrain. Furthermore, we present the limitations of regression techniques in terms of spatial correlation and continuity assumptions, and the impact on traversability analysis of a planetary rover across unstructured terrain. The analysis was performed on datasets of the Mars Yard at the Powerhouse Museum in Sydney, obtained using the onboard RGB-D camera.

I. Introduction

Robotic missions have been utilised to explore various scientific aspects of the Mars surface, including surface geology and the possibility of life. Teleoperating robots from Earth is proving difficult due to the communication delay between Earth and Mars which can be as long as 22 minutes. Early Lunar exploration rovers such as the Lunokhod required a five-man team to operate by sending driving commands from Earth in real time. Despite having a much smaller communication of 3 seconds, the team experienced many challenges to manoeuvre the Lunokhod on the Lunar surface. To perform the robotic mission more efficiently, low order or time-critical tasks such as obstacle avoidance and motor control can be handled autonomously, while high order mission tasks such as “explore area A” or “travel to rock B” can be handled by ground operators on Earth. By incorporating autonomous or semi-autonomous capabilities to the rover, operations from Earth can be more focused towards high level mission goals.

To achieve autonomy for high order tasks, planetary rovers need to be capable of traversing across the terrain in an efficient and safe manner. The level of autonomy of the rover is related to its capability to sense, represent and interpret the surrounding environment. An environment such as the Mars surface involves a great diversity in terrain features, including highly uneven geometry which is difficult to model, therefore accurate and reliable techniques are required to represent the terrain surface.

Many recent advances have been made in the area of terrain modelling to better estimate terrain geometry in areas with little or no data, such as techniques to preserve discontinuities in terrain models [7] and incorporating visibility constraints [9] to improve the accuracy of the estimated terrain geometry. However, these techniques have not yet been applied in the area of space exploration to construct accurate terrain maps in unstructured environments.

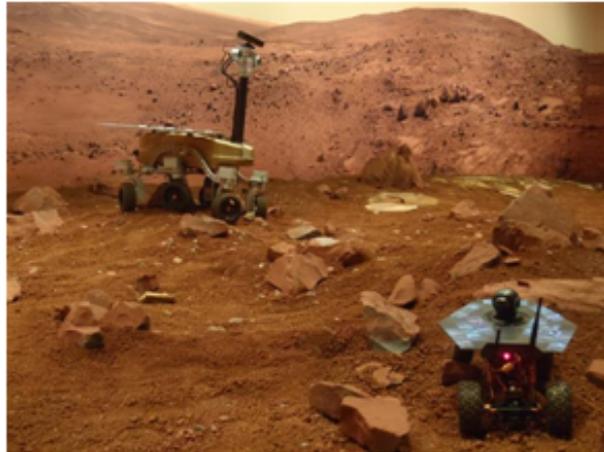


Fig. 1: Mars Rover (Mawson) at the back and Scout Rover at the front in the Sydney Powerhouse Museum Marscape

Once the environment is modelled, the rover needs to be able to interpret the data and assess the associated risks or difficulties of traversing across the terrain. Traversability analysis provides a metric for evaluating planning and control strategies to avoid hazardous areas, and thus provide efficiency and safety for rover operation. Numerous techniques for evaluating traversability metrics have been implemented in existing rover platforms with varying degrees of success, such as the systems implemented on the NASA Mars Exploration Rovers [1] and the LAAS Marsokhod Rover [2]. However, with advances in terrain modelling and terrain traversability, we need to explicitly draw the connection between the two fields, i.e. perform terrain modelling purely for the purpose of traversability, to promote synergy in the system.

In this paper we compare state-of-the-art techniques for terrain estimation and discuss the suitability of each technique in the context of an unstructured Martian terrain. By linking previous work in the area of terrain modelling and traversability analysis, we investigate the effects of terrain geometry models on terrain traversability analysis for planetary rovers, in particular the effects on estimation of vehicle attitude and configuration on terrain models constructed using terrain estimation techniques. We also reconsider state-of-the-art terrain model estimation techniques based on experimental data, and present limitations of current terrain estimation methods in the application of terrain traversability estimation.

In Section II, we outline previous work in the area of terrain modelling and traversability analysis. Section III reviews the theory behind some terrain representation, in particular Digital Elevation Maps (DEM) and Gaussian Processes (GPs), along with the limitations of each technique. We describe the steps taken to evaluate the traversability metric using experimental data in Section IV. Section V outlines the experimental setup of the rover and the Mars

Environment. In Section VI we show initial results of traversability analysis using different terrain modelling techniques, and discuss the effects of linking terrain representation and traversability analysis. Section VII summarises our conclusions and future work in this area.

II. Related Work

The area of terrain model estimation and terrain traversability analysis have been well explored in each of their respective fields. Research in the area of terrain model estimation aims to improve the accuracy and reliability of the predicted terrain geometry using available sensor data, while the work in terrain traversability aims to best estimate vehicle behaviour over the terrain.

A. Terrain Modelling

In the area of terrain model estimation, Digital Elevation Maps (DEMs) have been used to create a discrete geometric representation of the terrain. Much work has been performed to improve on Digital Elevation Maps (DEMs) to create a more complete model of the terrain, i.e. to estimate elevation in regions of little or no data. Lang et al. proposed the use of adaptive non-stationary kernel regression in Gaussian Processes (GPs) to deal with varying data densities and to preserve discontinuities in terrain models [7]. Vasudevan et al. compared the performance of different covariance functions for large scale terrain modelling, and introduced multi-output GPs to incorporate the RGB and the elevation values in the training data [8]. Hadsell et al. extended the traditional kernel-based learning approaches for estimating continuous surfaces by providing upper and lower bounds on the surface [9]. This was done by exploiting visibility constraints of the sensor to the terrain surface and applying kernel-based regression techniques to improve the precision of the terrain geometry estimate.

B. Terrain Traversability

The development of the Grid-based Estimation of Surface Traversability Applied to Local Terrain (GESTALT) system by Goldberg et. Al. has been successfully implemented on the Mars Exploration Rovers (MER) Spirit and Opportunity [1]. It is based on Carnegie Mellon's Morphin algorithm [3,4] and is a local planner which uses stereo cameras to evaluate terrain safety and avoid obstacles. The system uses stereo vision to calculate a disparity image, which is mapped to a 3D Cartesian location using camera geometry to produce an elevation map. Once the local elevation map is obtained, GESTALT determines the next best direction for the rover to reach its goal safely. The traversability of each cell is determined by merging the moment statistics of the set of Cartesian points on each grid cell to find the best fit plane, and then using the plane statistics to calculate hazard measures [1]. Finally, hazard and waypoint arc votes are used to select the set of arcs for the rover to follow until the desired waypoint is reached. While GESTALT provides a computationally efficient method of calculating terrain traversability, plane fitting methods may not provide accurate results specific to the vehicle.

Lacroix et al. explored the possibility of long range autonomous navigation with a 6-wheel Marsokhod chassis [2]. On rough terrains, the chassis internal configuration is calculated from the digital elevation map (DEM) [5] and a path is selected to maximize the interest/cost ratio. The DEM was preferred over other methods because critical constraints to traversing over rough terrain are stability, collision and configuration constraints, in order for the rover to overcome terrain irregularities. The proposed technique of short-range path planning using elevation map considers rover mobility over the terrain and also reflects the capability of the vehicle. However, this technique can become computationally expensive as it relies on simulation to determine vehicle configuration.

More recently, Helmick et al. presented the Terrain Adaptive Navigation (TANav) system [6], designed to enable planetary rovers to operate more robustly over a terrain of varying slip. The system encompasses the areas of goodness map generation, terrain triage, terrain classification, remote slip prediction, path planning, high fidelity traversability analysis, and slip compensated path following. The goodness map generated is based on classification of known classes, such as rocks, sand, gravel, with predefined properties. The TANav system is able to efficiently determine terrain traversability but is limited to the defined terrain classes which may be limiting on the Mars surface with different terrain properties.

It can be seen that there has been significant progress in the area of terrain modelling and traversability analysis, but little effort has been made to link the two areas of research to develop an accurate vehicle specific traversability model.

III. Terrain Representation

To accurately predict rover response on the terrain, terrain representation need to be performed in a manner that best represents the geometry and characteristics of the terrain, as well as the associated uncertainties to determine the "quality" of the prediction. In this section we will be exploring representation using DEM from raw data, and using Gaussian process regression.

A. Digital Elevation Map from Raw Data

Digital Elevation Maps are often used to model terrain surfaces. By representing the terrain as an elevation map, the amount of stored data can be scaled using the grid size, which is favorable in applications where memory and computational resources are limited. Using raw data from sensors such as stereo cameras or laser range finders, a DEM can be constructed by taking the mean elevation of the data points at each grid cell. Figure 2 shows a DEM produced using raw data from a single instance, and it can be seen that there are areas (shown in white) which are occluded by rocks from the sensor field of view.

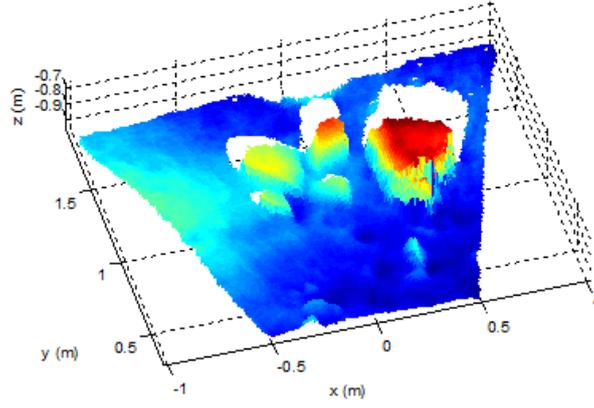


Fig. 2: Digital Elevation Map Produced from Raw Data

B. Gaussian Process Regression

Even with the use of modern sensors, there always exists occlusions and areas with lower density of data. In areas of little or no data, interpolation techniques can be used to estimate elevation. Gaussian process (GP) regression provides a means of learning the underlying model of spatially correlated data with uncertainty. As such, it has been the proposed method for estimating missing information in incomplete datasets in applications such as mapping or system identification. In our problem, we will be estimating the elevation (z) using the (x, y) coordinates of the data point. Gaussian approaches can be thought of as a normally distributed probability density function characterized by a mean $m(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$

$$m(\mathbf{x}) = E[f(\mathbf{x})]$$

$$k(\mathbf{x}, \mathbf{x}') = E[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$$

where $\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$, denoting our input variable.

The covariance function, also referred to as kernel, defines the correlation between the random variables in the training data. A popular kernel is the squared-exponential kernel, which can be given as

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T \Sigma (\mathbf{x} - \mathbf{x}')\right)$$

where $\Sigma = \begin{bmatrix} l_x & 0 \\ 0 & l_y \end{bmatrix}$ is the length scale matrix and measures the rate at which the modelled function changes in the x and y direction (in our case, for a 2D grid); σ_f^2 is the variance of the modelled function.

An example of a DEM produced using GP regression with SE Kernel with mean affine function can be seen in figure 3.

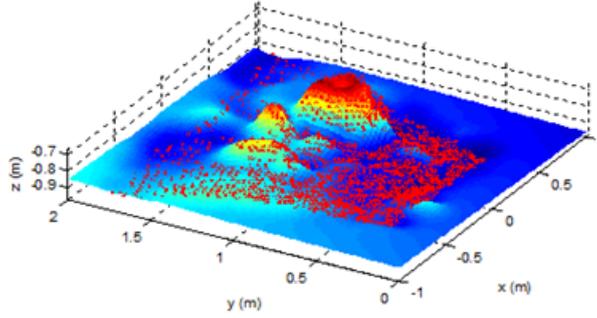


Fig. 3: Digital Elevation Map Produced using GP regression and Squared Exponential Kernel with Mean Affine Function

The neural network kernel is another commonly used kernel and can be given as

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \arcsin \left(\frac{\beta + 2\mathbf{x}^T \Sigma \mathbf{x}'}{\sqrt{(1 + \beta + 2\mathbf{x}^T \Sigma \mathbf{x})(1 + \beta + 2\mathbf{x}'^T \Sigma \mathbf{x}')}} \right)$$

where β is the bias factor and Σ is the length scale matrix.

The squared-exponential kernel function is stationary, and has a smoothing effect on the data by nature of the shape of the kernel function. Vasudevan's work [8] showed that the neural network kernel was more effective than the squared-exponential kernel function at handling discontinuous data which is common in data sets containing unstructured terrain.

To learn the model using a training data set, a kernel needs to be chosen and the relevant hyperparameters for the kernel need to be optimised. This is commonly performed by formulating the problem in a log marginal likelihood framework, then solving as a non-convex optimization problem.

Defining X and z to be the inputs and outputs from the training data respectively, the log marginal likelihood of the training outputs z given training inputs X and hyperparameters θ is given by

$$X = [(x_i, y_i)]_{i=1}^n$$

$$z = [z_i]_{i=1}^n$$

The log marginal likelihood has three terms - the first describes data fit, the second penalizes model complexity, and the third is a normalization constant for the number of data points. By minimizing the log marginal likelihood, the optimal set of hyperparameters which fit the data set is found. In this work, the Polack-Ribiere flavor of conjugate gradients was first

used to compute search directions [10]. A line search using quadratic and cubic polynomial approximations, and the Wolfe-Powell stopping criteria together with slope ratio method were used to estimate the initial step sizes for gradient based optimisation.

Once the GP model is learned, it can be applied across a grid to estimate the elevation information. This process is commonly known as Gaussian process regression.

Since the joint distribution of any finite number of random variations of a GP is Gaussian, the joint distribution of the training inputs z and test outputs f^* can be given as

$$\begin{bmatrix} z \\ f^* \end{bmatrix} \sim N \left(0 \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X, X_*) \\ K(X, X_*) & K(X_*, X_*) \end{bmatrix} \right)$$

The posterior or expected value can be given as

$$\bar{f}_* = K(X_*, X)[K(X, X) + \sigma_n^2 I]^{-1}z$$

and the covariance or uncertainty can be given as

$$cov(f_*) = K(X_*, K_*) - K(X_*, X)[K(X, X) + \sigma_n^2 I]^{-1}K(X, X_*)$$

For n training points and n_* test points, $K(X, X_*)$ represents the $n \times n_*$ covariance matrix evaluated at all the pairs of training and test points. This framework was used to estimate the elevation information at any point in the grid given the incomplete data set shown in Figure 2.

While GPs provide a framework for estimating elevation information with uncertainty at areas where there is sparse or no data, it has a few limitations which may render it unsuitable for terrain geometry estimation in unstructured outdoor environments. Firstly, GPs are implicitly continuous and assume single output value, i.e. the GP is expected to calculate a single elevation value at each grid cell. However, this problem still exists when producing DEMs using raw data and is a general problem to all elevation representations where $z = f(x, y)$. This may lead to misrepresentations of overhanging terrain features. The second limitation is the problem of spatial correlation, which is a common limitation among interpolation methods. This has the effect of smoothing out terrain features and thus reduces the accuracy of the estimation of terrain geometry.

IV. Traversability Metric

To evaluate the traversability of the vehicle over the terrain, we performed a forward propagation of discrete vehicle states over each grid cell in the DEM using a simplified model of the Mawson Rover (Figure 4). Mawson is a six wheeled vehicle with individual steering servo motors on each wheel. As such, the vehicle can be treated as a holonomic vehicle. Mawson's chassis is designed as a rocker-bogie system, which is designed to reduce motion of the main body. By lowering the vehicle and placing it at each grid cell, the configuration of the Rocker-Bogie suspension and vehicle attitude are simulated. A similar technique was employed by Peynot in [5] to satisfy the configuration constraints in the articulated chassis of the Marsokhod rover.

The simplified model of the rover can be seen in Figure 4.

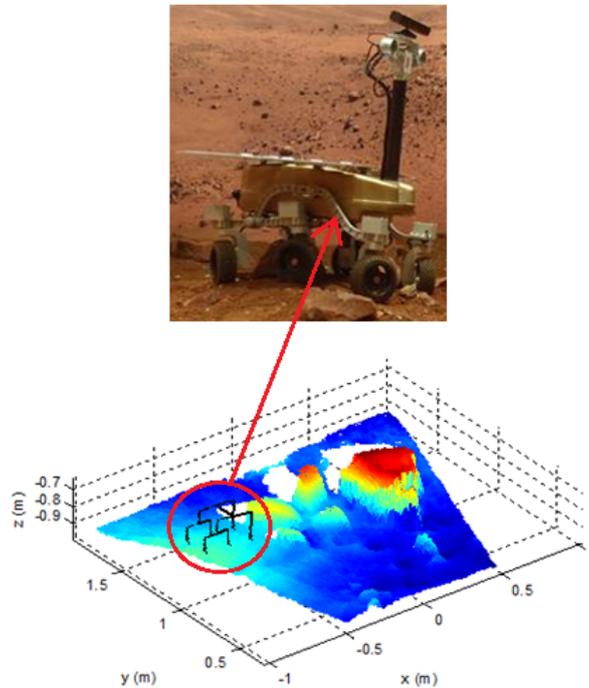


Fig. 4: Simplified rover model. The black frame represents the Rocker-Bogie suspension of the Mawson rover.

To determine the configuration of the Rocker-Bogie suspension and vehicle attitude, the vehicle attitude is first initialised as zero in pitch, roll and yaw, and the altitude is initialised such that one wheel is in contact with the terrain. While keeping the Rocker-Bogie joint angles at zero, a set of heuristics was used to find the vehicle attitude to minimise the total distance from each wheel to the ground. A similar set of heuristics was then used to find the Rocker-Bogie joint angles.

The simulation only accounts for only static scenarios and does not yet consider the transition of vehicle states from one cell to the next. The wheel-terrain interaction and friction

in the rotating joints of the Rocker-Bogie suspension, as well as the mass distribution of the vehicle and payload are ignored in this simulation. It should be noted that the simulation requires a specified yaw angle of the vehicle and determines the resulting pitch and roll angles based on terrain geometry only.

V. Experimental Setup

The experiments were conducted with Mawson, a planetary rover named after an Australian Antarctic explorer, which was developed at the Australian Centre for Field Robotics (ACFR). The rover footprint is approximately 0.5 m by 0.3 m. More details about the design and development of the rover can be found in [11].



Fig. 5: The Mawson Rover

Mawson carries an array of sensors onboard, including

- RGB-D Camera (Microsoft Kinect),
- Colour Cameras,
- Encoders on Rocker-Bogie Suspension.

For the purposes of this work, the sensor data from the RGB-D camera will be primarily used. Using a IR emitter and IR camera, the RGB-D camera measures the time of flight of the emitted IR beams of each pixel and builds a 2.5D map of the environment. The sensor provides a maximum resolution of 640 by 480 pixels at 30 frames per second, has a range of approximately 8 m, and is mounted at 1 m above the ground. It should be noted that although the RGB-D camera may not be an appropriate sensor for some/all outdoor operations, the geometric point cloud can also be obtained using other sensors such as stereo vision. The conclusions of this study do not depend on the type of sensor used to acquire the geometric point cloud.

The experiments were conducted at the Marscape at the Sydney Powerhouse Museum, which contains rocks and inclines with varying degrees of slip and cohesion (Figure 1). Other elements in the Mars environment, such as lighting, terrain geometry and composition, can also be controlled and adjusted if necessary.

VI. Results and Evaluation

The following experimental evaluation was conducted on point clouds acquired with a single sensor snapshot, representing an area of $2 \times 1.75m$, formatted in a grid with a cell size of $0.05 \times 0.05m$, the vehicle attitude and Rocker-Bogie joint angles were calculated at each grid cell. The traversability metrics were determined as the Root-Sum-Squared (RSS) of the predicted vehicle attitude and Rocker-Bogie joint angles (radians).

Figure 6 shows the estimated terrain geometry using the Squared Exponential (SE), Squared Exponential with Mean Affine Function, and Neural Network (NN) kernel in the GP framework. The data points are denoted as red dots and the estimated terrain height denoted as the colour coded surface. Comparing the data points to the estimated terrain height using GP regression, it can be seen that GP regression has a smoothing effect in its estimation, especially in areas that do not have a lot of data from the sensor. This affects vehicle attitude and Rocker-Bogie angles and can cause the vehicle to be overconfident in its estimation of terrain traversability. It can also be seen that the use of different kernels affects the predicted terrain in particular areas that are not well observed. This is because the nature of the kernel function has a much bigger effect on the prediction in areas with little or no observations, as the assumptions about the shape of the terrain is implied in the kernel function.

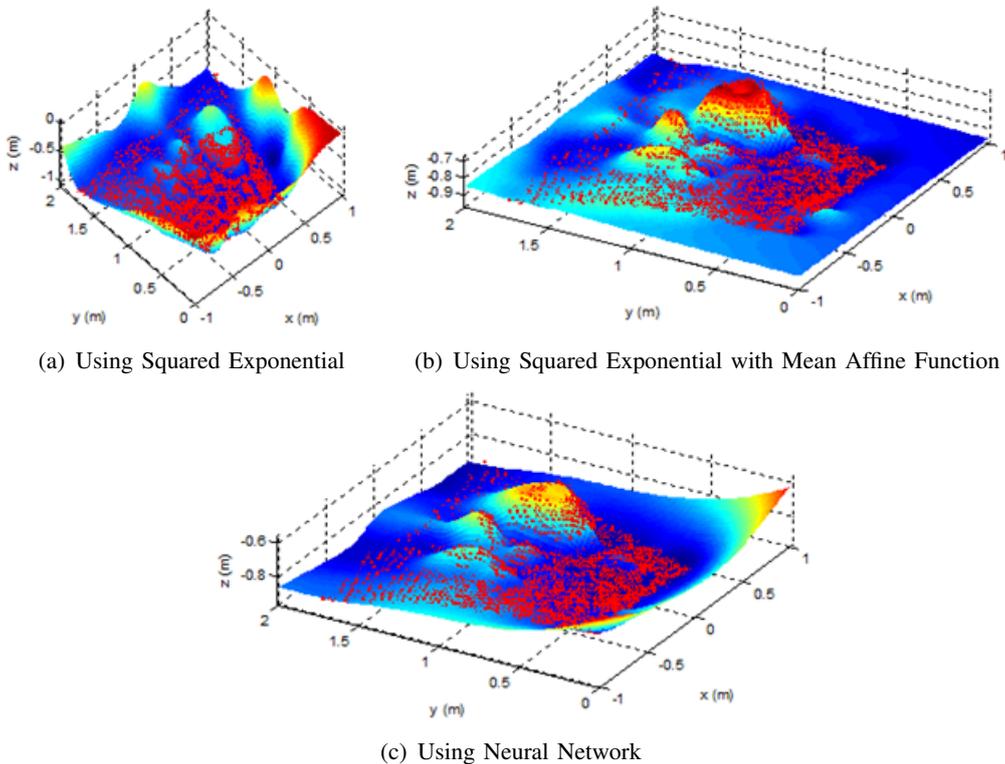


Fig. 6: Terrain Geometry using SE, SE with Mean Affine Function, and NN Kernels

In the case where traversability is determined using the DEM produced using raw sensor data (Figure 2), occluded areas are declared as untraversable, i.e. if any of the 6 wheels comes into contact with an area with no elevation data, the cell which the rover is on is declared untraversable. It can be seen in figure 6 that there is a large area that is declared

untraversable in this strategy which may limit path planning options, but at the same time is very conservative for the rover in terms of making a decision about the risks involved in going over an area that it has no information on. It can also be seen that the affine mean function improves the accuracy of the terrain geometry estimation by assuming an average plane of elevation throughout the grid cells. Performing a similar traversability analysis using the GP-generated terrain, it can be seen in Figure 6 that the smoother terrain geometry results in a smoother change in value of traversability between cells. In rough terrains, this would underestimate traversability, especially in areas with little or no data where terrain geometry estimation are made using assumptions from the selection of the GP kernel.

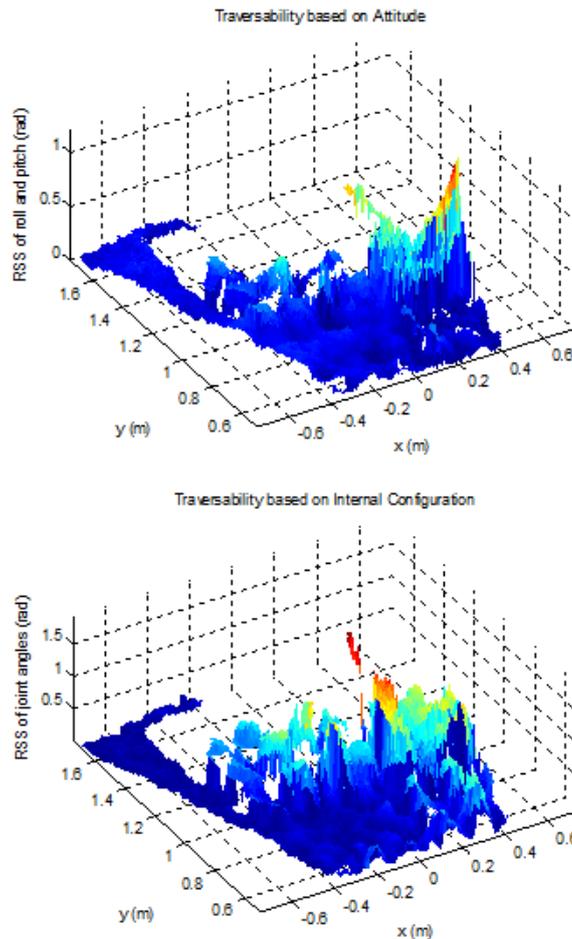


Fig. 7: RSS of vehicle pitch and roll (top) using raw sensor data, RSS of Rocker-Bogie joint angles (bottom) using raw sensor data

Comparing the estimated traversability using raw data (shown in Figure 7) and GPs (shown in Figure 8) at the region from $y = [0.6, 1]$, it can be seen that the RSS of vehicle pitch and Rocker-Bogie joint angles in the DEM generated by raw data is visibly higher, has higher fidelity, and less smooth than the DEM generated using GPs. This is a direct effect of the smoothed terrain produced by GP regression resulting in underestimating traversability in an area. Due to the continuity assumption, the terrain geometry estimated using GPs results in smooth transition between each estimated elevation point.

In occluded areas, traversability estimated using GP methods is largely dependent on the kernel behaviour in GP methods, as seen in the higher vehicle attitude and joint angles using the NN kernel (Figure 9) compared with using the SE kernel with mean affine function (Figure 8). Within these areas, the elevation estimation is highly uncertain as there are no data points contributing towards the dataset used to train the GP, and the resulting vehicle attitude and joint angles varies greatly based on the kernel used. Since the terrain geometry is largely affected by the shape of the kernel in these areas, the resulting traversability estimate will be a smooth surface with variations in elevation conditioned on data points in areas which are visible to the sensor. This can result in underestimation of the traversability in rough terrains, and overestimation of traversability in flat terrains. On the other hand, the same area is simply declared untraversable in the DEM constructed from raw data, which is a conservative approach but does not possess the same variation and uncertainty of the GP approach.

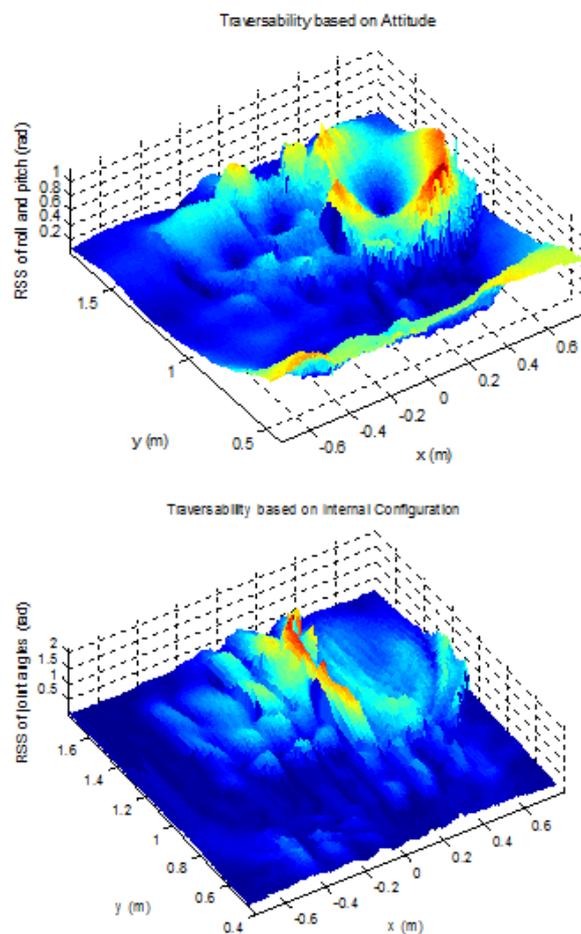


Fig. 8: RSS of vehicle pitch and roll (top) using GP regression with SE, RSS of Rocker-Bogie joint angles (bottom) using GP regression with SE

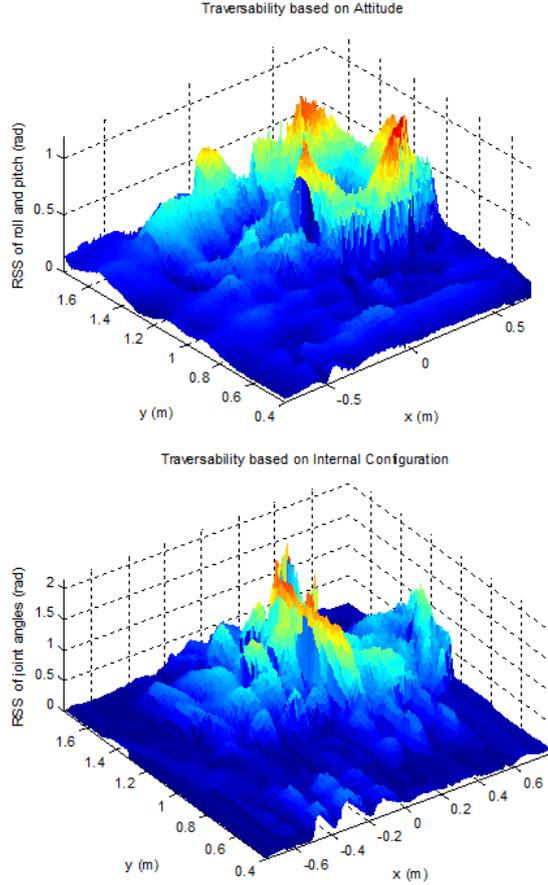


Fig. 9: RSS of vehicle pitch and roll (top) using GP regression with NN, RSS of Rocker-Bogie joint angles (bottom) using GP regression with NN

VII. Conclusion/Future Work

From the results presented in section VI, it can be seen that the GP representation of terrain geometry is inherently continuous and its smoothing nature may cause the vehicle to become overconfident (i.e. assessing the terrain to higher traversability) in its assessment of stability. On the other hand, building an elevation map from raw data preserves terrain geometry in unstructured terrain. However, it is more affected by occlusion from terrain features and lacks an uncertainty estimate of the resulting elevation geometry. The current method is to consider all occluded terrain to be untraversable, which led to large sections of the map to be classified as untraversable. This can lead to overly conservative estimates of terrain traversability resulting in no possible solutions to reach desired waypoints in challenging terrain.

As there are shortcomings and limitations to both raw data and regression based techniques, an area of future work is to consider a terrain modelling technique purely for the purposes of traversability analysis which would explicitly consider the vehicle-terrain interaction. To account for sensor noise, the current approach is to assume the noise from the sensor data to be non-coloured within the sensor operating range, and the variance in the noise is captured to some extent in the uncertainty estimation of the GP. However, to explicitly incorporate

uncertainties in the elevation obtained from raw data, a sensor error model would need to be developed.

Acknowledgment

The authors would like to acknowledge the work of the technical staff at ACFR - Muhammad Esa Attia, Jeremy Randle, Taufik Yunahar, and Victor Chan, for their support and assistance in the assembly and setup of the Mawson Rover, and to the staff at the Sydney Powerhouse Museum for their work in setting up in the Mars Yard which made this work possible.

The work was supported by the Australian Research Council, the Australian Space Research Program Stream A, Pathways to Space: Empowering the Internet Generation, and the Australian Centre for Field Robotics (ACFR). This material is also based on research partially sponsored by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

References

- [1] S. B. Goldberg, M. W. Maimone and L. Matthies, *Stereo Vision and Rover Navigation Software for Planetary Exploration*, IEEE Aerospace Conference Proceedings, Big Sky, Montana, USA, March 2002.
- [2] S. Lacroix, A. Mallet, D. Bonnafous, G. Bauzil, S. Fleury, M. Herrb, R. Chatila, *Autonomous Rover Navigation on Unknown Terrains Functions and Integration*, IEEE International Symposium for Experimental Robotics (ISER), Sant'Angelo d'Ischia, Italy, July 2002.
- [3] R. Simmons, L. Henriksen, L. Chrisman, G. Whelan, *Obstacle Avoiding and Safeguarding for a Lunar Rover*, AIAA Forum on Advanced Developments in Space Robotics, Madison, WI, USA, August 1996.
- [4] S. Singh, K. Schwehr, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, *Recent Progress in Local and Global Traversability for Planetary Rovers*, IEEE International Conference on Robotics and Automation (ICRA), San Francisco, CA, USA, April 2000.
- [5] T. Peynot and S. Lacroix, *Selection and Monitoring of Navigation modes for an Autonomous Rover*, Experimental Robotics: The 10th International Symposium on Experimental Robotics, Vol. 39, Springer-Verlag, Berlin, pp. 121-130, 2008.
- [6] D. Helmick, A. Angelova, L. Matthies, *Terrain Adaptive Navigation for Planetary Rovers*, Journal of Field Robotics - Special Issue on Space Robotics, Part II, Volume 26 Issue 4, April 2009.
- [7] T. Lang, C. Plagemann, W. Burgard, *Adaptive Non-Stationary Kernel Regression for Terrain Modeling*, Robotics: Science and Systems (RSS), Seattle, WA, USA, June 2009.
- [8] S. Vasudevan, F. Ramos, E. Nettleton, H. Durrant-Whyte, *Large-scale Terrain Modeling from Multiple Sensors With Dependent Gaussian Processes*, International Conference on Intelligent Robots and Systems (IROS), Taipei, Taiwan, October 2010.
- [9] R. Hadsell, J. A. Bagnell, D. Huber, M. Hebert, *Accurate Rough Terrain Estimation with Space-Carving Kernels*, Robotics: Science and Systems (RSS), Seattle, WA, USA, June 2009.
- [10] C. E. Rasmussen, C. K. I. Williams, *Gaussian Processes for Machine Learning*, Cambridge, Massachusetts, USA: The MIT Press, 2006.
- [11] M. E. Attia, A. H. Goktogan, S. Sukkarieh, *Martian Rover 'Pathways To space' Technical Paper*, Currently Under Review, Australian Space Science Conference (ASSC), Canberra, ACT, Australia, September 2011.

Chapter 4

Resilient Navigation through Probabilistic Modality Reconfiguration [3]

by T. Peynot, R. Fitch, R. McAllister and A. Alempijevic,
in *12th International Conference on Intelligent Autonomous Systems (IAS)*,
Jeju Island, Korea, June 2012.

Resilient Navigation through Probabilistic Modality Reconfiguration

Thierry Peynot, Robert Fitch, Rowan McAllister and Alen Alempijevic

Abstract This paper proposes an approach to achieve resilient navigation for indoor mobile robots. Resilient navigation seeks to mitigate the impact of control, localisation, or map errors on the safety of the platform while enforcing the robot's ability to achieve its goal. We show that resilience to *unpredictable* errors can be achieved by combining the benefits of independent and complementary algorithmic approaches to navigation, or *modalities*, each tuned to a particular type of environment or situation. In this paper, the modalities comprise a path planning method and a reactive motion strategy. While the robot navigates, a Hidden Markov Model continually estimates the most appropriate modality based on two types of information: context (information known *a priori*) and monitoring (evaluating unpredictable aspects of the current situation). The robot then uses the recommended modality, switching between one and another dynamically. Experimental validation with a SegwayRMP-based platform in an office environment shows that our approach enables failure mitigation while maintaining the safety of the platform. The robot is shown to reach its goal in the presence of: 1) unpredicted control errors, 2) unexpected map errors and 3) a large injected localisation fault.

1 Introduction

Motion planning and control of a mobile robot necessarily involves multiple sources of uncertain information such as control uncertainty, localisation uncertainty, and mapping errors. Current research seeks to address these sources of uncertainty by

Thierry Peynot, Robert Fitch, Rowan McAllister
Australian Centre for Field Robotics, The University of Sydney, NSW 2006, Australia, e-mail: {tpeynot, rfitch, r.mcallister}@acfr.usyd.edu.au

Alen Alempijevic
Mechatronics and Intelligent Systems Group, University of Technology Sydney, NSW 2007, Australia, e-mail: Alen.Alempijevic@uts.edu.au

modelling them in the context of the planning problem [1, 5]. However, problems arising during execution of a plan are not always predictable (and hence able to be modelled). For example, it is difficult to predict localisation errors ahead of time, or to anticipate which map locations actually contain large errors. We are interested in mitigating the impact of such unpredictable errors on robot performance and safety. We introduce the term *resilience* to refer to this goal. Resilient navigation seeks to mitigate the impact of control, localisation, and map errors on the safety of the platform while enforcing the robot’s ability to achieve its goal.

In this paper, we study resilient navigation in the context of indoor mobile robots. We believe resilience is best achieved by combining the benefits of multiple independent algorithmic approaches, or *modalities*, each tuned to a particular type of environment or situation. The idea is to develop a set of modalities that covers the range of possible situations, and then to reconfigure the system dynamically in response to unpredicted errors. The key challenges are: 1) how to choose a suitable set of modalities, 2) how to represent information that describes the robot’s context, and 3) how to decide which modality is most appropriate at any given time. Because we are dealing with uncertain information, these challenges require solutions in probabilistic form.

Our approach in choosing a set of modalities is to include a motion planning strategy that requires global information, and a reactive strategy that requires only local information. These two modalities are complementary. If the navigation goal is within the field of view (FOV) of the robot, a reactive obstacle avoidance approach (e.g. [10, 8]) can be feasible. However, reactive approaches have known limitations. They can become trapped in dead-ends or U-shape obstacles, and it is difficult to obtain smooth trajectories. If the goal is located outside of the robot’s FOV, the recommended strategy is to use a motion planning algorithm that reasons more globally, especially if some prior knowledge of the environment is available. In addition, smoother and more efficient paths can be obtained (see Fig. 1). However, in cluttered environments, such a strategy can only be effective if sufficiently accurate map and global localisation are available. In addition, the control of the platform needs to be robust and precise enough to follow the planned trajectory.

An alternative is to combine the two strategies to obtain a *hybrid* system [4]. Typically, a motion planning algorithm computes a global plan, generating a list of waypoints along the computed trajectory which are passed to a reactive motion method. A drawback of these hybrid techniques is that even if the motion planner can produce smooth trajectories (or trajectories respecting some pre-defined constraints), the execution of such types of trajectories cannot be enforced. Another inconvenience is that events that provoke failure of one of the components will often provoke failure of the combination, whereas this can be mitigated by using the appropriate method at the right time. A comparison of the different strategies discussed in this paper is shown in Table 1.

Instead, we propose a modality-switching algorithm based on a *hidden-markov model* (HMM) that considers *context* and *monitoring* information. If the system is aware that path execution cannot safely handle a difficult situation such as a narrow doorway, it is appropriate to switch to a reactive strategy. This situation can be

Table 1 Comparison of navigation strategies

Strategy:	Planning	Reactive	Hybrid	Our approach
Robust to dead-ends	✓	✗	✓	✓
Robust to dynamic obstacles	✗	✓	✓	✓
Robust to errors in localisation or map	✗	✓	✓	✓
Optimised paths (when possible)	✓	✗	✗	✓

evaluated using the localisation of the robot in a map, and detecting the presence of this narrow passage. However, reasoning only on this context information will not be sufficient to handle situations where the error/uncertainty of global localisation is high, where elements of the map have moved, or where a dynamic obstacle has appeared. Fast local replanning integrating map updates can partially address this problem but is computationally expensive and can lead to instabilities in control. Therefore, we propose to choose a modality based on context information *and* monitoring information (such as proximity to obstacles observed from laser data).

We evaluate our approach through hardware experiments with an indoor mobile robot in an office environment. We show that failures can be mitigated in challenging situations while maintaining the safety and liveness of the platform. The situations we consider include: control errors, localisation errors, map errors (unexpected obstacles), and presence of an “aggressive” human dynamic obstacle.

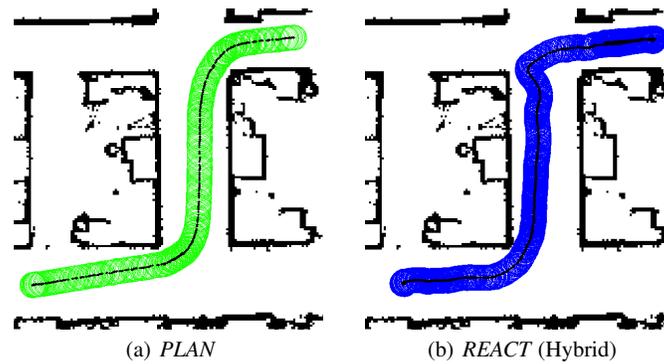


Fig. 1 Trajectory obtained using a planner (a) and a *hybrid* approach (b). Obstacles in the map are in black. Circles represent the radius of the robot. Approximate size of the area shown: $6.5m \times 7.5m$.

2 Related Work

Previous work has considered multi-modal systems for the navigation of an indoor mobile robot. [11] proposed the Robel system: a robot controller that learns different ways of combining sensory-motor functions to achieve a navigation task. Robel uses a Markov decision process (MDP) to provide a policy. However, MDPs are generally computationally expensive and policies often have to be computed off-line or at low frequencies. Our system was designed to be efficient enough for the robot to be reactive: modality switching can happen quickly when needed. Besides, an MDP requires the states to be fully observable. [15] proposed a system based on a partially observable MDP (POMDP) that can be used to detect, diagnose and recover from faults. However, the policy is computed off-line and the robot does not have a real alternative navigation modality when the path planning strategy fails. Our approach does not require explicit detection and identification of specific faults such as a localisation error; it focuses on mitigating failures that could occur in consequence, finding alternatives to obtain robustness while maintaining safety.

Motion planning and obstacle avoidance are well-studied problems in the literature. See [6] for a comprehensive review up to 2005. More recently, researchers have sought to address motion planning under uncertainty in control [1], localisation [5] or sensing and environment map [9]. However, these studies typically require the ability to *predict* possible errors, as they need to model the uncertainty in the context of the planning problem. In this paper, we are interested in mitigating the impact of *unpredictable* errors.

3 Probabilistic Modality Reconfiguration

The approach we propose is a probabilistic framework for an indoor robot endowed with two main navigation modalities: 1) a global planner (*PLAN*), and 2) a reactive motion approach (*REACT*). In addition, a *STOP* modality is included for emergency and safety. This method builds on our previous work for an outdoor mobile robot with modalities appropriate to flat terrain and rough terrain [13].

Our approach is to estimate the likelihood of each modality being most suitable using an HMM. The HMM is appropriate since states are not directly observable and it provides a time integration that prevents jitter (too frequent modality changes, see Fig. 5). Crucially, the probabilistic approach allows the system to handle uncertainty in the different sources of information.

The goal of the HMM is to provide a *modality recommendation*. The HMM is constructed such that the number of states is equal to the number of available modalities. Fig. 2 provides a graphical representation of our three-state HMM, where each state x_k corresponds to the proposition: “modality m_k is the appropriate modality to apply at this point in time.”

Two categories of information are input to the HMM: 1) *context* information is global environmental knowledge known *a priori*, and 2) *monitoring* information is

online execution knowledge of the observed immediate situation. The framework is designed as a Markov *conditional estimation* system [2]. It estimates the conditional state $x_{k,t}$ at time t , knowing context observation until time t , $O_{1:t}$, and online monitoring information $M_{1:t}$. If the robot is endowed with N different modalities, the probability that m_k is the appropriate modality to apply at time t can be written, $\forall k \in \llbracket 1, N \rrbracket$,

$$P(x_{k,t} | O_{1:t}, M_{1:t}) \propto P(O_t | x_{k,t}) \sum_{i=1}^N P(x_{k,t} | x_{i,t-1}, M_t) P(x_{i,t-1} | O_{1:t-1}, M_{1:t-1}) \quad (1)$$

where $P(O_t | x_{k,t})$ is an observation probability (computed using the *context* information), and $P(x_{k,t} | x_{i,t-1}, M_t)$ is the conditional probability of transition from state x_i to state x_k , knowing the *monitoring* data M_t at time t . The following sub-sections describe more specifically the different modalities of the robot used in this paper and the nature of the context and monitoring information.

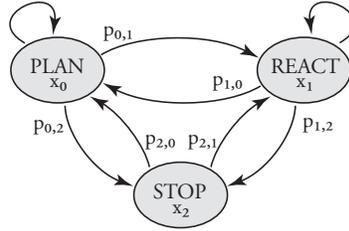


Fig. 2 Graphical representation of the 3-state HMM

3.1 Context

The context information relates to the distance d from the robot boundaries to the closest obstacles as seen on an *a priori* global map. This information is used to predict the likely modality at a given map location. We determined experimentally that the *PLAN* modality is likely to fail in situations where the robot is too close to obstacles, i.e. closer than a *security distance* $d_s = 0.15m$, equal to half the radius of the robot. Therefore, intuitively, the *a priori* recommendation based on context information is to use *PLAN* in areas sufficiently clear from obstacles ($d > d_s$), *REACT* in areas that are close to an obstacle on the global map ($d \leq d_s$), and *STOP* in places immediately proximal to obstacles ($d < d_c$, *critical distance*).

d (the observation O_t) is calculated online using the current localisation of the robot in the map. To integrate this observation in the system (HMM), probability density functions (pdf) are used to take into account uncertainties. The main sources of uncertainties are the (x, y) localisation of the robot in the map and the location of

the obstacles in the map itself. The map is an occupancy grid that was built using the laser of the robot assuming perfect localisation. Therefore, the map uncertainty can be expressed as $\sigma_{map} = \sigma_{laser}$, the standard deviation of the range measurements of the laser scanner. $\sigma_{laser} = 0.03 \text{ m}$ for the Hokuyo laser on our robot.

The uncertainty in the *a priori* map is independent of the uncertainty of the current localisation, as the map was built beforehand, using a different localisation. Therefore, the standard deviation on the observation of d can be expressed as the sum of the uncertainties: $\sigma = \sigma_{map} + \sigma_{loc}$, where σ_{loc} represents the localisation uncertainty provided by the algorithm mentioned in Sec. 4.

3.1.1 Modality STOP

We define the distribution of $p(O_t|x_2)$, or $p(d|STOP)$, as an inverse sigmoid centred on the critical distance d_c (see Fig. 3 in red):

$$P(d|STOP) = 1 - \frac{1 - \alpha}{1 + e^{-(d-d_c)/\sigma}} \quad (2)$$

where σ partly defines the curvature of the sigmoid. σ (similar to the standard deviation of a Gaussian) corresponds to the uncertainty in d , and $d_c = 0$ is the critical distance.

The distribution $p(d|STOP)$ represents the likelihood that observation d is made, knowing that the robot should stop. The inverse sigmoid accounts for the uncertainty in the observation and in the knowledge of this threshold value. The limit of this sigmoid, when d tends to infinity, is superior to zero (see Fig. 3). This accounts for the fact that the map does not capture all information in the world, in particular dynamic obstacles. The value of this limit represents the chance of having to stop the robot while infinitely away from map obstacles, i.e. the chance of having a dynamic object appearing within less than d_c of the robot, α . It is crucial to account for the possibility of this event sufficiently so that the system maintains a chance of capturing it [7]. Thus, this value is set to a value higher than the actual probability of occurrence as would be determined statistically. In our implementation we set $\alpha = 0.05$.

3.1.2 Modality REACT

The distribution of $p(O_t|x_1)$, or $p(d|REACT)$, is defined as a sigmoid centred on d_c (see Fig. 3 in blue):

$$P(d|REACT) = \frac{1 - (\alpha + \gamma)}{1 + e^{-(d-d_c)/\sigma}} \quad (3)$$

where $\sigma = \sigma_{map} + \sigma_{loc}$, as defined earlier. To guarantee safety, the main restriction for this modality is that it cannot be used too close to obstacles ($d < d_c$), hence the sigmoid.

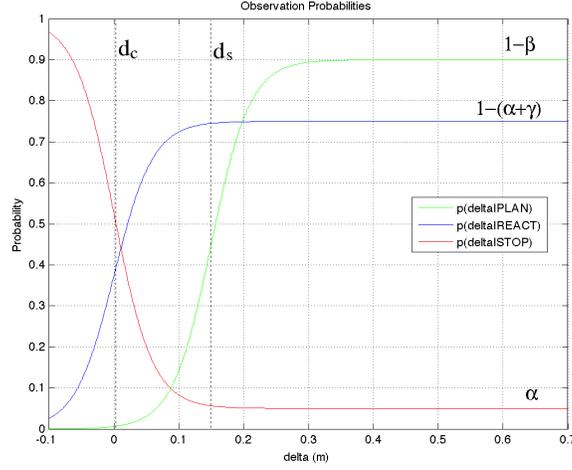


Fig. 3 Probability density functions for the context information (shown before normalisation, with $\sigma_{loc} = 0$). These are functions of d , representing the distance to the closest obstacles to the robot, as seen in the map.

There are two secondary restrictions. One consideration is the chance of a dynamic object appearing within a distance d_c to the robot, i.e. $\alpha = 0.05$. The other is the *a priori* chance of failure of *REACT* in general, even in an open map (recall that this modality is subject to local minima). This chance of failure highly depends on the environment, which we capture with the probability: $\gamma = 0.20$. Considering the events represented by α and γ as independent, the limit of the sigmoid $p(d|REACT)$ when d tends to infinity is set to $1 - (\alpha + \gamma) = 0.75$.

3.1.3 Modality PLAN

The distribution of $p(O_t|x_0)$, or $p(d|PLAN)$, is also defined as a sigmoid, centred on the security distance d_s (see Fig. 3 in green):

$$P(d|PLAN) = \frac{1 - \beta}{1 + e^{-(d-d_s)/\sigma}} \quad (4)$$

Note that once again the limit of the sigmoid $p(O_t|m_0)$ when d tends to infinity is lower than 1. This accounts for the chance of having dynamic objects appearing within d_s of the robot bounds. We consider the prior probability of this event to be $\beta = 0.10$ ($\beta > \alpha$), therefore the limit of the sigmoid distribution is $1 - \beta = 0.9$. For high values of d it is important to set the chance of success of *PLAN* higher than *REACT* (if the goal is far, it is known that *PLAN* is more likely to succeed), i.e. $\beta < \alpha + \gamma$. Finally, note that these distributions need to be normalised before integration in the HMM.

3.2 Monitoring

Contrary to context information, the purpose of monitoring information is to check the actual “appropriateness” of the current situation, with regard to the possible modalities, using data only observable during execution. The online monitoring uses δ , the distance from the robot bounds to the closest obstacle detected in laser measurements. Recall that the online monitoring contributes to the computation of the transition probabilities of the HMM. If the robot gets too close to obstacles seen in current laser scans while operating in *PLAN*, it should switch to *REACT*. In this way, if global localisation is temporarily inaccurate, or if obstacle points are in a different location than on the (static) global map, this situation can be handled by *REACT*, contrary to *PLAN*.

More specifically, the intuitive rules of transitions (given here without uncertainty, for convenience) are the following. The corresponding transition probabilities used in the HMM are given in parenthesis, in both full (e.g. $P(x_2|x_1, \delta)$) and equivalent reduced form (e.g. $p_{1,2}$). First, let us consider the output transitions of x_0 (i.e. *PLAN*).

- The transition $P(x_1|x_0, \delta) = p_{0,1}$ (*PLAN* to *REACT*) is likely if $d_c < \delta < d_s$, i.e. an obstacle is detected by the laser in the *intermediate* proximity of the robot.
- The transition $P(x_2|x_0, \delta) = p_{0,2}$ (*PLAN* to *STOP*) is likely if $\delta < d_c$, i.e. an obstacle is detected by the laser in the *immediate* proximity of the robot.
- $P(x_0|x_0, \delta) = p_{0,0}$ (*PLAN* to *PLAN*) is likely if $\delta > d_s$, i.e. the robot is clear from obstacles.

The other transitions can be defined similarly, using the same short notations as above: $p_{1,0} = p_{2,0} = p_{0,0}$, $p_{1,1} = p_{2,1} = p_{0,1}$, $p_{1,2} = p_{2,2} = p_{0,2}$.

Because of the uncertainty in δ (the laser measurements), these rules are defined probabilistically using sigmoid distributions similar to those defined in Sec. 3.1 and shown in Fig. 3. In this case the main source of uncertainty is the relative inaccuracy of the laser measurements, therefore the σ of the sigmoids is: $\sigma = \sigma_{laser}$. The output transition probabilities from each state are normalised, as their sum must equal 1.

4 Implementation

Our experimental platform consists of the Segway RMP100 base with onboard PCs and various sensors, including a Hokuyo UTM-30LX laser range-finder and encoders in the mobile base for odometry [12]. Localisation is computed using the Monte Carlo Localisation (MCL) algorithm [14]. The robot’s belief is represented by a set of weighted hypotheses which approximate the posterior under a common Bayesian formulation of the localisation problem. We update this distribution using data from odometry, the laser range-finder, and a predefined map of the environment.

The test area is an office environment occupied by over 25 people and consisting mainly of student workstations and fixed and movable furniture. This area is thus well-suited for evaluating real-world applicability.

4.1 Available Modalities

Modality m_0 is *PLAN*. We implemented the well-known Latombe Grid-Search algorithm [3] for nonholonomic planning, customised to find paths with minimum change in curvature. Although the name may seem to imply a discrete search space, the algorithm does use continuous coordinates. A detailed summary can be found in [6]. The planner is complete with respect to the resolution of its given proximity grid and time interval of the path set [3]. Because this proof is not constructive, we do not have a method for determining parameter values analytically. We hand-tuned them empirically and found a reasonable grid resolution of $0.2m \times 0.2m \times \frac{\pi}{8} rad$ and path set time intervals of 2 or 4 seconds. Our path set has angular velocities chosen from $\{-\frac{\pi}{4}, -\frac{\pi}{8}, -\frac{\pi}{16}, -\frac{\pi}{32}, 0, \frac{\pi}{32}, \frac{\pi}{16}, \frac{\pi}{8}, \frac{\pi}{4}\}$ and linear velocities from $\{0.2m/s, 0.1m/s\}$. Our priority queue uses a cost function that combines minimum distance to goal with minimum change in curvature.

Modality m_1 is *REACT*. This is a reactive collision avoidance method that avoids sensed obstacles. We implemented a potential field method derived from a model of human navigation [8]. This method directly controls angular acceleration and produces smooth paths. We chose this method because the robot operates in an office-like environment amenable to human-like paths. Because the laser cannot scan all 360° around the robot, the perception data that *REACT* uses is a local fusion of laser scans. Odometry is used for localisation in order to avoid the influence of errors in global localisation.

Modality m_2 is *STOP*. This is the safety modality; it applies when the robot has come too close to an obstacle and the only reasonable option is to halt. If *STOP* was provoked by monitoring information, the robot can only resume when the obstacle responsible for the stop is dynamic and has moved away. To account for this, our system continues to evaluate the HMM recommendation even though the robot is stationary.

4.2 Modality Switching

PLAN is the default starting modality, as it has the highest prior probability. When switching from *PLAN* to *REACT*, a goal waypoint must be chosen. We initially choose the next waypoint on the last path computed by *PLAN*. However, because of localisation or map errors, this waypoint may intersect an obstacle. In this case, the next waypoint of the plan that is confirmed as clear from obstacles becomes the new

goal for *REACT*. It is then preferable to switch back to *PLAN* quickly to avoid the risk of *REACT* falling into local minima.

5 Experiments

The experimental validation in this section illustrates the resilience of our probabilistic reconfiguration approach, which allows the mitigation of unpredictable failures. Examples of causes of such failures are: errors of the controller while executing a planned path, errors in the map (i.e. presence of objects that could not be integrated in the map early enough) and large localisation uncertainty or error. We also compare our method to simple threshold rules for modality switching. Results were obtained using the platform described in Sec. 4.

The illustrations show the estimated robot trajectory during each test, using coloured points to represent the modality used at the time. The selected modality corresponds to the highest output probability $P(x_{k,t}|O_{1:t},M_{1:t})$ at each time step t . The HMM and the modality selection were updated at $10Hz$.

5.1 Modality Reconfiguration in Static Environment

5.1.1 Unpredicted Control Error

The experiment in Fig. 4(a) illustrates how our framework allowed the system to maintain the robot’s safety in the presence of unpredicted errors of the controller during execution of the planned trajectory. The robot started executing a planned path similar to the one in Fig. 1(a), which was successful using *PLAN* only. However, at the (expected) end of the turn around the first corner, the controller “overshot”, risking the safety of the platform. This event was detected by our system, which switched to *REACT* to recover. When safe, the robot returned to the *PLAN* modality to complete its mission.

5.1.2 Going Through a Narrow Doorway

We also tested the robot’s ability to follow a corridor and then pass through a $0.85m$ -wide doorway (the robot’s diameter is $0.6m$). As the corridor is reasonably large (about $1.7m$ in average), the robot first used *PLAN* and only switched to *REACT* to negotiate the doorway passage.

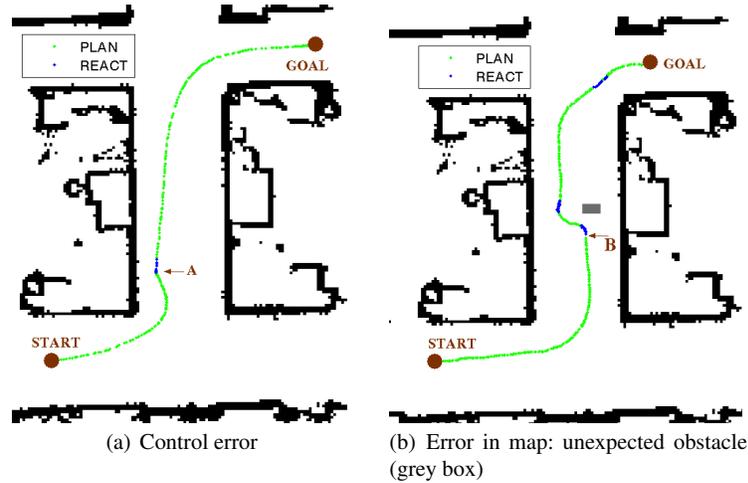


Fig. 4 Examples of robot trajectory executed with modality switching. Known obstacles in the map are shown in black, while the colour points show the (estimated) positions of the centre of the robot. Green means the recommended modality is *PLAN*, while blue means the recommended modality is *REACT*. Approximate size of the area shown: $6.5m \times 7.5m$.

5.2 Unexpected Map Error and Comparison to Simple Threshold Rules

Fig. 4(b) shows another example of modality switching to negotiate an unexpected situation safely: an unpredictable large error in the map. This situation is caused by the presence of an unexpected obstacle. This simulates a map error. In order to avoid the box, the robot switches to *REACT*, then returns to *PLAN* once the situation is safe and the map is more consistent with the current observation. A likely collision was thus avoided.

Fig. 5 illustrates a similar test using a recommendation based on simple logical rules comparing d and δ to “hard” thresholds equal to d_c and d_s . It can be seen that such strategy can provoke frequent undesirable modality switches, contrary to the HMM of our approach.

5.3 Presence of Dynamic Obstacles

We validated that the robot is resilient to the presence of highly dynamic obstacles. In the test shown in Fig. 6(a), a pedestrian coming from the top left of the scene walked quickly towards the robot. On approach, the robot first switched to *REACT* and then tried to evade (event C). Once the human had left the vicinity, the robot

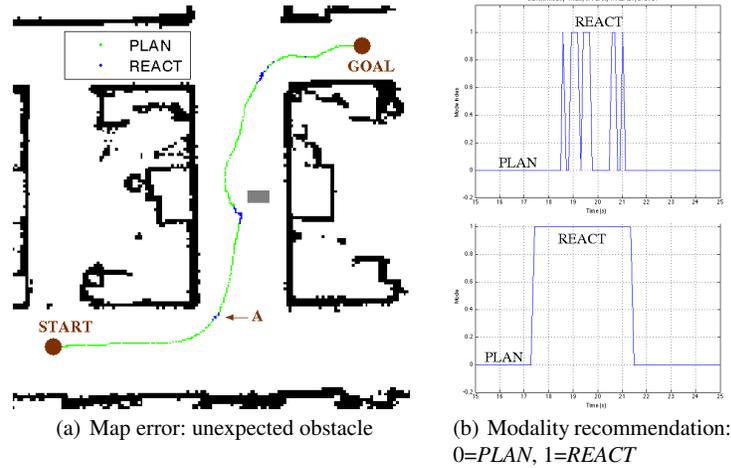


Fig. 5 Comparison with a simple threshold strategy. The robot encounters an unexpected obstacle in (a). (b) shows the chosen modality over time for the 10s surrounding event A in (a). Modality switching with fixed thresholds results in unacceptable oscillation (top) compared to our method (bottom).

could resume its mission. A similar situation occurred again later in the test, with an even more sudden appearance of the human in the FOV of the robot. This event was again safely handled by the robot. This test shows that, although the robot can nominally execute optimised trajectories, it can also safely react to dynamic obstacles, comparably to a pure reactive motion strategy.

5.4 Injected Localisation Fault

In this experiment, a significant localisation fault was artificially created by introducing a sudden and unexpected offset of $1m$ to the output of the localisation estimator. Fig. 6(b) shows the clear offset to the right between the estimated position and the reality. However, the robot was still able to safely achieve its mission by switching to *REACT* when appropriate.

The context information is shown in Fig. 7 and the corresponding partial probabilities of modality recommendation are shown in Fig. 8. The moment of the localisation fault injection is clearly visible at $t = 35s$ (event D) on both figures. Fig. 7 indicates that with context information alone the robot should definitely *STOP*. If the robot could rely only on its current global localisation estimate and map, it would not have been safe to continue, since according to its estimated position the robot appears to be *on* the location of known obstacles in the map (see Fig. 6(b)). However, our system recommended a better alternative: a prudent switch to *REACT*. The

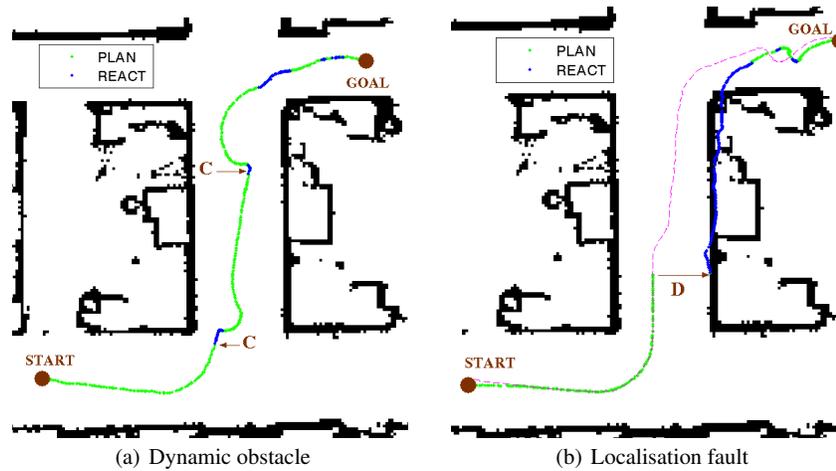


Fig. 6 (a): Presence of highly dynamic obstacles (events C). (b): Modality switching with injected localisation fault (sudden offset to the right of 1m, see D). As a result, the estimated positions of the robot appear to be *on* the right wall (blue line). The magenta dashed line shows the reference localisation.

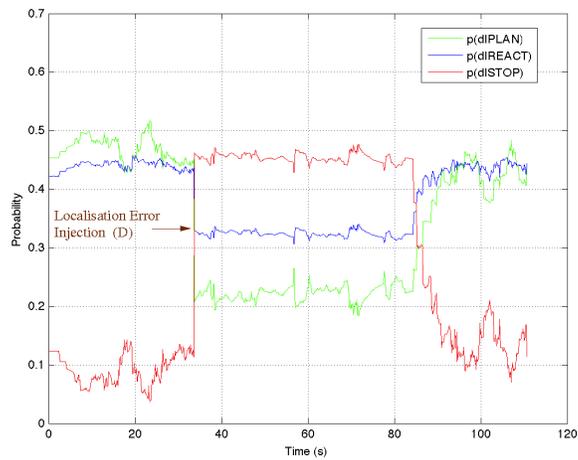


Fig. 7 Partial probabilities for context information, corresponding to the test in Fig. 6(b).

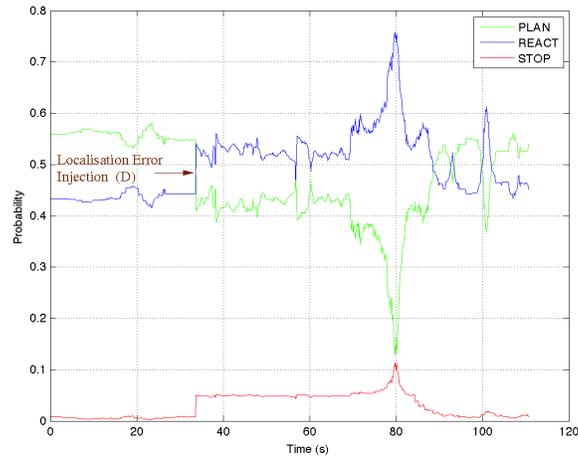


Fig. 8 Partial probabilities for all three states (i.e. modality recommendation), corresponding to the test in Fig. 6(b).

localisation algorithm then progressively corrected its estimation up to a point when the situation became comfortable enough again to use *PLAN* to finish the mission. This shows our system to be resilient to large unpredicted localisation or map errors.

6 Conclusion

We have shown that a multiple modality strategy for resilient navigation, based on a probabilistic framework, can be applied to an indoor mobile robot to combine the advantages of navigation modalities while maintaining the safety of the platform. In our implementation, the robot is able to plan and execute smooth paths (minimising change in curvature) when possible, while being very reactive when needed. The system was applied online and shown to be reliable and robust in the presence of map errors and large localisation uncertainties or offsets. The concept is demonstrated with one choice of planning and reactive modality, however, these planning and reactive methods are easily interchangeable.

Future work will exploit the monitoring and context information for diagnosis and recovery of particular components of the system. Currently, both types of information are only exploited to compute a modality recommendation. However, we saw in Sec. 5.4 that the discrepancy between context and monitoring indicates a likely error in the map or global localisation. This could be used to actively repair these components, while the robot uses a reactive modality.

7 Acknowledgments

This work was supported in part by the Australian Centre for Field Robotics (ACFR) and the NSW State Government. This material is based on research sponsored by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

References

1. Alterovitz, R., Simeon, T., Goldberg, K.: The stochastic motion roadmap: A sampling framework for planning with markov motion uncertainty. In: *Robotics: Science and Systems II* (2007)
2. Arnaud, E., Memin, E., Cernuschi-Frias, B.: Conditional filters for image sequence based tracking - application to point tracking. *IEEE Transactions on Image Processing* **14**(1), 63–79 (2005)
3. Barraquand, J., Latombe, J.: Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica* **10**, 121–155 (1993)
4. Brock, O., Khatib, O.: High-speed navigation using the global dynamic window approach. In: *Proc. of IEEE International Conference on Robotics and Automation* (1999)
5. Bryand, A., Roy, N.: Rapidly-exploring random belief trees for motion planning under uncertainty. In: *Proc. of IEEE International Conference on Robotics and Automation* (2011)
6. Choset, H., Lynch, K.M., Hutchinson, S., Kantor, G.A., Burgard, W., Kavraki, L.E., Thrun, S.: *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, Cambridge, MA (2005)
7. Dearden, R., Clancy, D.: Particle filters for real-time fault detection in planetary rovers. In: *12th International Workshop on Principles of Diagnosis* (2002)
8. Huang, W.H., Fajen, B.R., Fink, J.R., Warren, W.H.: Visual navigation and obstacle avoidance using a steering potential function. *Robotics and Autonomous Systems* **54**(4), 288–299 (2006)
9. Kurniawati, H., Bandyopadhyay, T., Patrikalakis, N.M.: Global motion planning under uncertain motion, sensing, and environment map. In: *Robotics: Science and Systems VII* (2011)
10. Minguez, J., Montano, L.: Nearness diagram navigation (ND): Collision avoidance in troublesome scenarios. *IEEE Transactions on Robotics and Automation* **20**(1), 45–59 (2004)
11. Morisset, B., Ghallab, M.: Learning how to combine sensory-motor functions into a robust behavior. *Artificial Intelligence* **172**, 392–412 (2008)
12. N. Kirchner et al.: RobotAssist - a Platform for Human Robot Interaction Research. In: *Proc. of ARAA Australasian Conference on Robotics and Automation* (2010)
13. Peynot, T., Lacroix, S.: Selection and monitoring of navigation modes for an autonomous rover. In: O. Khatib, V. Kumar, D. Rus (eds.) *Experimental Robotics: The 10th International Symposium on Experimental Robotics*, vol. 39, pp. 121–130. Springer-Verlag, Berlin (2008)
14. Thrun, S., Fox, D., Burgard, W., Dellaert, F.: Robust monte carlo localization for mobile robots. *Artificial Intelligence* **128**(1–2), 99–141 (2000)
15. V. Verma R. Simmons, J.F.: Probabilistic models for monitoring and fault diagnosis. In: *2nd IARP and IEEE/RAS Joint Workshop on Technical Challenges for Dependable Robots in Human* (2002)

Chapter 5

Motion Planning and Stochastic Control with Experimental Validation on a Planetary Rover [4]

R. McAllister, T. Peynot, R. Fitch and S. Sukkarieh,
to appear in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*,
Vilamoura, Portugal, October 2012. (Accepted 2 July 2012).

Motion Planning and Stochastic Control with Experimental Validation on a Planetary Rover

Rowan McAllister, Thierry Peynot, Robert Fitch and Salah Sukkarieh

Abstract—Motion planning for planetary rovers must consider control uncertainty in order to maintain the safety of the platform during navigation. Modelling such control uncertainty is difficult due to the complex interaction between the platform and its environment. In this paper, we propose a motion planning approach whereby the outcome of control actions is learned from experience and represented statistically using a Gaussian process regression model. This model is used to construct a control policy for navigation to a goal region in a terrain map built using an on-board RGB-D camera. The terrain includes flat ground, small rocks, and non-traversable rocks. We report the results of 200 simulated and 35 experimental trials that validate the approach and demonstrate the value of considering control uncertainty in maintaining platform safety.

I. INTRODUCTION

Motion planning for mobile robots in unstructured environments must consider various forms of uncertainty. One significant source of uncertainty in outdoor terrain is *control uncertainty*. Robots such as planetary rovers are designed for mobility in challenging environments, but understanding the associated control uncertainty for the purpose of motion planning is difficult due to the complexity of this type of environment. It is critical to consider control uncertainty in motion planning, particularly in environments that expose the robot to the risk of serious mechanical damage. We are interested in this problem in the context of planetary rovers [1]. Our goal is to navigate while maintaining the safety of the platform in potentially dangerous terrain.

The goal of classical geometric motion planning is to minimise time or distance while avoiding obstacles [2]. The conceptual distinction between free space and obstacles for planetary rovers, however, is less clear. It is important to avoid obstacles, but it is also desirable to avoid free space where, due to control uncertainty, the robot has high likelihood of encountering an obstacle during execution. This situation cannot be modelled by simple distance thresholds surrounding obstacles because risk varies across free space, and is probabilistic.

Accurately predicting executed behaviour in response to a given control input is difficult for planetary rovers due to complex terramechanics [3]. For previously unobserved terrain, prior models of terrain properties may not be available. It is thus important to model control uncertainty with a



Fig. 1. Planetary rover “Mawson” used for experimental validation, shown in the Mars yard at the Powerhouse Museum in Sydney, Australia.

method that can be feasibly executed online during operation of the robot, and to validate such a model experimentally.

Our approach is to build a statistical model directly from observed behaviour, represented as a Gaussian process (GP). We consider uncertainty in the heading of the platform and in distance travelled. We use this GP model to build a stochastic transition function for use in motion planning. The planning goal is to compute a policy that allows the robot to reach a given goal location while maintaining the safety of the platform. We assume that a map of the environment is available, represented as a digital elevation map. Platform safety is represented by a cost function over this terrain map, which is constructed *a priori* using on-board sensors. We compute the policy using dynamic programming (DP), where the resolution of discretised geometric states is equal to that provided in the elevation map.

In this paper, we present the details of our approach and its implementation for the planetary rover shown in Fig. 1. The environment consists of flat terrain, traversable rocks, and non-traversable rocks. We learn GP models for rock traversal that map environment features to a distribution of resulting rover configurations (in state space) for two types of control actions. The cost map is constructed from data collected by an on-board RGB-D camera. We report results from 200 simulated and 35 experimental trials that evaluate the rover’s ability to traverse flat terrain and small rocks while avoiding non-traversable rocks. We compare rover performance in executing policies constructed with and without control uncertainty. Our results show empirically that planning with control uncertainty improves the rover’s ability to navigate while avoiding non-traversable areas, and demonstrate the value of planning under uncertainty for planetary rovers using a real platform in a realistic environment.

The work was supported by the Australian Research Council, the Australian Space Research Program *Pathways to Space: Empowering the Internet Generation*, the Australian Centre for Field Robotics, and the Air Force Research Laboratory under agreement FA2386-10-1-4153.

R. McAllister, T. Peynot, R. Fitch and S. Sukkarieh are with the Australian Centre for Field Robotics, The University of Sydney, Australia {r.mcallister, tpeynot, rfitch}@acfr.usyd.edu.au

II. RELATED WORK

A common approach for considering control uncertainty in motion planning is to express the uncertainty as a cost and then to plan a path that minimises this cost assuming deterministic control [4], [5]. Another family of approaches plans a path using a sampling-based algorithm, and then evaluates the control uncertainty along the path selected [6], [7]. In classical motion planning, the desired path is provided to a feedback controller for execution. Various forms of control strategies (such as LQG) can be used to model potential deviations from the path and hence to select a path with least risk in terms of platform safety [8], [9].

For non-deterministic systems Markov decision processes (MDPs) are commonly used to formulate problems in motion planning with uncertainty [2], [10]. Control uncertainty is represented as a stochastic transition function, and a policy can be computed using dynamic programming [11]. The partially-observable Markov decision process (POMDP) is another common formulation [11]. However, these techniques are most often evaluated in simulation only and there is a critical need for further validation using real robots.

Recent work by Brooks and Iagnemma [12] models control uncertainty as a function of terrain in a self-supervised learning framework. This approach uses visual features to classify terrain types and learn associated proprioceptive mechanical properties.

Finally, physics-based approaches that study terramechanics provide detailed mobility models by considering features such as soil cohesion and density [13]. Statistical mobility prediction using terramechanics has been proposed that generates a Gaussian distribution over predicted future states on homogeneous terrain [3]. However, it is difficult to precisely model non-homogeneous terrain that includes rocks of different sizes and shapes that may move in reaction to the force exerted by a rover wheel.

In our work we directly search for a path with low risk of entering a non-traversable area, but our model of stochastic actions is tied to observed environmental features that vary across the terrain and learned through experience. We furthermore consider risk at the level of primitive actions and construct a policy that is executed directly. Our approach uses statistical regression techniques in performing the inference and direct learning is applied showing meaningful improvements to motion planning are possible without a complex terramechanics model.

III. MOTION PLANNING AND LEARNED CONTROL UNCERTAINTY

With the aim of reaching a given goal region *safely* while optimising the total cost of traversal over the *executed* trajectory, our approach is to take into account the stochasticity of the control of the robot at the planning stage. This requires modelling the control uncertainty, which we achieve by experience, using machine learning. This section first describes the planning algorithm used in our approach (Sec. III-A), followed by the presentation of the learning technique used to model the control uncertainty (Sec. III-B).

A. Planning Algorithm

We compute a control policy for the robot using dynamic programming (DP). DP computes an optimal policy with respect to a discrete set of (stochastic) primitive motions and given resolution of the state space [2]. DP is a feasible method in low dimensional state spaces; in our problem the state s can be defined using two lateral dimensions x and y and one dimension for orientation ψ , i.e., $s = \{x, y, \psi\}$.

This formulation treats the motion planning problem as a Markov Decision Process, assuming accurate localisation and stochastic control. The optimal motion policy is computed using the Bellman optimality equation iteratively:

$$V^*(s) = \max_a \left\{ \sum_{s'} P(s'|s, a) (R(s'|s, a) + \gamma V(s')) \right\}, \quad (1)$$

where s is the robot state (discretised into uniform cells), a is an action (or motion primitive) from the action set A , and $\gamma = 1$ is the discount factor. The transition function $P(s'|s, a)$ describes the probability that a state-action pair (s, a) transitions to state s' . The optimal policy is given by:

$$\pi^*(s) = \arg \max_a \left\{ \sum_{s'} P(s'|s, a) (R(s'|s, a) + \gamma V(s')) \right\}, \quad (2)$$

where the reward $R(s'|s, a)$ is computed from a cost map that represents the difficulty of the terrain. $P(s'|s, a)$ is not known *a priori*. Therefore, these state transitions are learned from experience, as described below.

B. Learning-based Mobility Prediction

$P(s'|s, a)$ can be expressed using a PDF of the relative transition between states, $p(\Delta s|s, a)$, where $\Delta s \equiv s' - s$:

$$P(s'|s, a) = \int p(\Delta s|s, a) f(s + \Delta s, s') d\Delta s, \quad (3)$$

where $f(s_1, s_2) = 1$ if the discretisation of s_1 corresponds to s_2 and $f(s_1, s_2) = 0$ otherwise. In unstructured terrains, Δs may strongly depend on factors such as the terrain profile along the executed path and also the action executed. Terrain profiles are represented by a vector of features (or characteristics): $\lambda(s, a)$. Therefore, $p(\Delta s|s, a)$ is learned as a function of $\lambda(s, a)$ and a :

$$p(\Delta s|s, a) = p(\Delta s|\lambda(s, a), a). \quad (4)$$

The estimation of relative change in state is achieved using Gaussian Process (GP) regression. GP is a standard framework to learn a model of spatially correlated data and provides estimations with uncertainty. The GP framework is especially effective in cases where the input data are sparsely populated.

We use the GP formulation from [14]. The input vector \mathbf{x} is a function of the terrain features, shifted such that the input has zero mean, i.e., $\mathbf{x} = \lambda - \bar{\lambda}_{train}$, where $\bar{\lambda}_{train}$ is the mean of each terrain feature in the training data. We define $\Delta s_{i,a}$ as the i^{th} single-valued component of the change of state Δs resulting from executing action a . We define a GP for each $\Delta s_{i,a}$, $i \in \llbracket 1, N \rrbracket$, $a \in A$, where $N = Dim(s)$. The output value $y_{i,a}$ of one of these

GPs is defined as $y_{i,a} = \Delta s_{i,a} - \bar{\Delta s}_{i,a}$, where $\bar{\Delta s}_{i,a}$ is the mean value of Δs_i across all executions of action a in the training data. $y_{i,a} = f_{i,a} + w_{i,a}$ has a zero mean and variance σ_n^2 equal to the variance of the additive noise $w_{i,a}$, i.e., $p(y_{i,a}|f_{i,a}) = \mathcal{N}(0, \sigma_n^2)$.

The covariance function used to describe the spatial correlation between two input vectors is the squared exponential:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^\top \Lambda^{-2}(\mathbf{x} - \mathbf{x}')\right) + \sigma_n^2 I, \quad (5)$$

where σ_f^2 is the input variance and Λ is a length scale matrix of diagonal elements that describes the smoothness of the input data. The predictive distribution is given by a Gaussian,

$$p(\Delta s_{i,a} | \lambda(s, a)) - \bar{\Delta s}_{i,a} = p(f_* | X, Y, \mathbf{x}_*) \sim \mathcal{N}(\mu_*, \Sigma_*), \quad (6)$$

with predictive mean

$$\mu_* = K(\mathbf{x}_*, X)[K(X, X) + \sigma_n^2 I]^{-1} Y,$$

and variance

$$\begin{aligned} \Sigma_* &= K(\mathbf{x}_*, \mathbf{x}_*) \\ &\quad - K(\mathbf{x}_*, X)[K(X, X) + \sigma_n^2 I]^{-1} K(X, \mathbf{x}_*), \end{aligned}$$

where X is the $n \times m$ matrix of all n training input vectors, Y is the $n \times 1$ vector of all training output values, and \mathbf{x}_* is the test input vector. $K(X, \mathbf{x}_*)$ is a covariance matrix which stores the covariance of each training input value against the test input values.

Thus, $p(\Delta s_{i,a})$ can be computed for untraversed terrain profiles using Eq. (6), where the test input vector \mathbf{x}_* has not been observed directly but rather estimated by employing a kinematics model of the robot over the series of states that would be traversed by executing action a from state s .

Finally, our planner considers the uncertainty in each component Δs_i separately by using the full distribution learnt from Δs_i and the expectation of the other components. Representing $\lambda(s, a)$ as λ , Eq. (4) is calculated as:

$$\begin{aligned} p(\Delta s | \lambda, a) &= p(\{\Delta s_1, \dots, \Delta s_i, \dots, \Delta s_N\} | \lambda, a) \approx \\ &= \{\mathbb{E}(\Delta s_1 | \lambda, a), \dots, p(\Delta s_i | \lambda, a), \dots, \mathbb{E}(\Delta s_N | \lambda, a)\}. \end{aligned} \quad (7)$$

C. System Outline

Fig. 2 shows an outline of the system. Once an elevation map of the robot's local environment has been computed, a kinematics model is used to estimate both the terrain profile characteristics and cost of traversal ($R(s'|s, a)$) at each state-action pair. Characteristics of the terrain profile, as observed by the Inertial Measurement Unit (IMU) and the localisation system, are used as input data to train the GP offline. The GP can then be used to estimate the stochastic transition function $P(s'|s, a)$ on terrain similar to that encountered during training. Using both $P(s'|s, a)$ and $R(s'|s, a)$, the motion planner computes the value function (Eq. (1)) over the observed state space to follow greedily as per the policy given in Eq. (2).

IV. IMPLEMENTATION

This section describes the implementation of the proposed approach on our experimental Mars rover shown in Fig. 3(a).

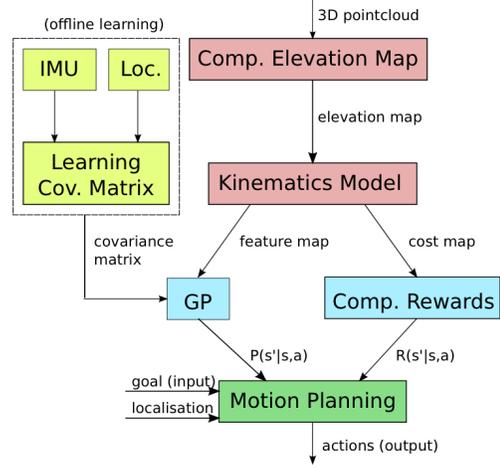


Fig. 2. System Outline. Colours indicate perception (red), offline learning (yellow), estimation (blue) and planning (green).

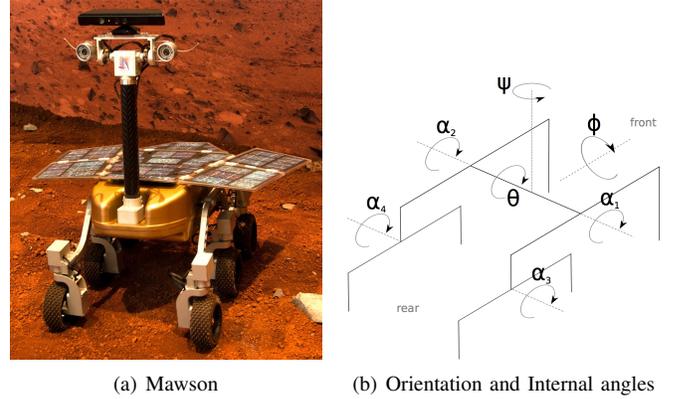


Fig. 3. The Mawson Rover (a) and its chassis configuration (b).

A. The Robot

Mawson is a six-wheeled rover with individual steering servo motors on each wheel and a Rocker-bogie chassis. The platform is equipped with:

- an RGB-D camera (Microsoft Kinect) mounted on a mast, tilted down 14° , used for terrain modelling and localisation,
- a 6-DOF IMU used to measure the roll (ϕ) and pitch (θ) of the robot,
- three potentiometers to observe the configuration of the chassis by measuring both bogie angles and the rocker differential (α_i in Fig. 3(b)).

For localisation and terrain modelling we use the RGB-D SLAM algorithm [15], implemented in the Robot Operating System (ROS) [16], which uses data from the RGB-D camera to perform simultaneous localisation and mapping (SLAM) online. An elevation map is generated from the point clouds supplied by the RGB-D camera by distributing elevation points in a regular Cartesian grid. The grid resolution is $0.025m \times 0.025m$.

B. Kinematics Model

To predict the attitude angles $\{\phi, \theta\}$ and chassis configuration $\{\alpha_2 - \alpha_1, \alpha_3, \alpha_4\}$ of the rover at given positions on the elevation map, we use a method similar to [17]. Although it does not take into account the dynamics of the platform, this simplified model is a sufficient approximation since the rover operates at low speeds. This kinematics model is used to: 1) estimate terrain profile characteristics for (s, a) pairs the planner queries, and 2) compute a cost map of the observed terrain (see Fig. 2).

1) *Feature Map*: From the estimates of the configurations of the rover on the map, a feature map is built for each (s, a) pair using the GP model to predict $P(\Delta s)$ for each (s, a) pair the planner considers.

2) *Cost Map*: The cost function chosen to generate the cost map from the elevation map penalises large absolute values of roll, pitch, and configuration angles of the chassis at a given position $s = \{x, y, \psi\}$:

$$\text{cost}_{\text{terrain}}(s) = (\text{cost}_{\phi\theta}(s) + \text{cost}_{\alpha}(s))^2, \quad (8)$$

where

$$\text{cost}_{\phi\theta}(s) = (\phi^2 + \theta^2), \quad (9)$$

$$\text{cost}_{\alpha}(s) = (\alpha_2 - \alpha_1)^2 + \alpha_3^2 + \alpha_4^2. \quad (10)$$

Since the configuration of the robot at a given position on the elevation map depends on its orientation, a 2D (x, y) cost map needs to be generated for each discretised orientation. The result is a 3-dimensional (x, y, ψ) cost map.

C. Planning

1) *State Space*: As mentioned in Sec. III-A, the rover's state s is defined as its position and orientation:

$$s \triangleq \{x, y, \psi\} \in \mathbb{R}^3. \quad (11)$$

This definition specifies all other orientations and internal angles (ϕ, θ, α_i) at each state using the kinematics model. State resolution was required to be smaller than the uncertainty bounds of resultant positions of actions in order for uncertainty to be considered by the DP. A discretisation of $0.025m \times 0.025m \times \frac{\pi}{32}rad$ is sufficient.

2) *Action Set*: We define two action types for the rover: *crabbing* and *rotation*. Crabbing corresponds to executing a straight line translation in the xy -plane by a given distance and heading, with no change in ψ , and constant linear velocity ($0.11m/s$). The rover is able to crab in any direction. Rotation is a spin-on-the-spot motion primitive at constant angular velocity ($0.24rad/s$). It changes ψ by a given magnitude. In total, the action set A is composed of 2 rotation and 8 crabbing motion primitives:

$$A \triangleq \{\text{rotate}(\pi/4), \text{rotate}(-\pi/4), \text{crab}(0.3m, n\pi/4) \forall n \in [-3, 4]\}. \quad (12)$$

Restricting the actions to this set was the result of a trade-off between rover dexterity and algorithm complexity.

3) *Reward Function*: The reward $R(s'|s, a)$ of an action is defined as the negative of the average cost of states that lie on a linear interpolation between a start state s and a resultant state s' :

$$R(s'|s, a) = -\xi - \frac{1}{K} \sum_{i=0}^K \text{Cost}\left(s_x + \frac{i}{K}(s'_x - s_x), s_y + \frac{i}{K}(s'_y - s_y), s_\psi + \frac{i}{K}(s'_\psi - s_\psi)\right), \quad (13)$$

where $\xi = 0.003$ is a small penalty used to deter excessive motions on flat terrain and K is the sampling resolution.

D. Learning and GP

To achieve mobility prediction from training data, we used the GP implementation from [14]. These training data were obtained through experimental runs of the rover executing each action a from A multiple times, while logging discretised sets of $\{\phi, \theta, \text{time}, s, a\}$ continuously. The training data collected comprise the rover's attitude angles, provided by the on-board IMU, and the deviation from the expected motion when executing a given action, measured using the localisation system. Due to the left-right symmetry of the platform, training was only required on 6 of the 10 actions from A (Fig. 4 shows which set of training data can be combined between symmetric actions).

Due to slow localisation updates, Δs could only be measured at the end of each action primitive. However, multiple values of $\{\phi, \theta\}$ were recorded during the execution of an action primitive, resulting in a vector of these values $\{\phi, \theta\}$, representing the terrain profile. Terrain profiles were encoded more compactly by extracting features of the evolution of $\{\phi, \theta\}$ to avoid overfitting. A combination of terrain profile features were tested with GP regression using cross validation. The features (vector of functions λ) which produced the lowest Root Mean Square error between the GP mean estimations and test datum output were chosen:

$$\lambda \triangleq \{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}, \quad (14)$$

where

$$\lambda_1(\phi, \theta) = \max(\phi_j - \phi_i), \forall i, j : i < j \quad (15)$$

$$\lambda_2(\phi, \theta) = \min(\phi_j - \phi_i), \forall i, j : i < j \quad (16)$$

$$\lambda_3(\phi, \theta) = \max(\theta_j - \theta_i), \forall i, j : i < j \quad (17)$$

$$\lambda_4(\phi, \theta) = \min(\theta_j - \theta_i), \forall i, j : i < j, \quad (18)$$

i.e., largest increase of ϕ , largest decrease of ϕ , largest increase of θ and largest decrease of θ during an action primitive.

When planning over previously untraversed states, ϕ and θ are estimated using the kinematics model from a state-action pair: $\hat{\phi}(s, a)$ and $\hat{\theta}(s, a)$. The λ functions are then applied to these estimates, which are input to the GP. Thus, the expression

$$p(\Delta s_i | \lambda(s, a), a) = p(\Delta s_i | \lambda(\hat{\phi}(s, a), \hat{\theta}(s, a), a))$$

is substituted in Eq. (7).

The components Δs_i of Δs (see Sec. III-B) were defined according to radial coordinates, as per the control space of crabbing and rotations, opposed to the Cartesian representation of the state space $\{x, y, \phi\}$ itself. The Δs_i components we consider are heading and distance travelled for crabbing actions

$$\Delta s_1 = \Delta s_{head} = \text{atan2}(\Delta y, \Delta x) \quad (19)$$

$$\Delta s_2 = \Delta s_{dist} = \sqrt{(\Delta x)^2 + (\Delta y)^2}, \quad (20)$$

and yaw for rotation actions

$$\Delta s_3 = \Delta s_{yaw} = \Delta \psi. \quad (21)$$

Therefore, Δs is represented by the tuple

$$\Delta s \triangleq \{\Delta s_{head}, \Delta s_{dist}, \Delta s_{yaw}\}. \quad (22)$$

In the rest of the paper, we define control errors (δs) as the difference between observed change in state (Δs) and “ideal” change in state ($\tilde{\Delta s}$) (i.e., if the controller followed the action-command perfectly); i.e., $\delta s = \Delta s - \tilde{\Delta s}$, a tuple analogous to Eq. (22).

V. EXPERIMENTAL RESULTS

We evaluated our approach both in simulation and through experiments with a planetary rover robot. Our experimental methodology consisted of two phases. First, we learned statistics of control error through empirical trials, described in Sec. V-A. Then, we performed navigation experiments using these learned data to build the motion planner’s stochastic transition function. Sec. V-B describes experiments in simulation, and Secs. V-C and V-D describe experiments using the robot.

Experiments were conducted in the Mars Yard environment shown earlier in Fig. 1. This environment is $117m^2$ in area and was designed to reproduce typical Martian terrain. It contains rocks of various sizes, small craters, and various grades of sand, dirt and gravel.

A. Training on Flat and Rough Terrain

Training was conducted for two cases: flat-terrain traversal and rough-terrain traversal. For the flat terrain case, control errors were learned by executing multiple runs for each action. Rough-terrain training additionally involved traversal of various rocks (one at a time).

In flat-terrain traversal, variations of values of $\{\phi, \theta\}$ were negligible, therefore, motion errors were learnt with respect to action only. Fig. 4 shows example heading errors for each action, and Table I lists learned statistical values. Although the terrain was mostly flat, the variance in motion primitive error is significant, which validates the need to take uncertainty into account in the planning. We found that the distributions could reasonably be approximated by a Gaussian.

During rough-terrain traversal, various rocks were traversed. Note that in some cases some rocks shifted under the weight of the rover, slightly sinking into the sand or rolling over. These types of situations, which are extremely

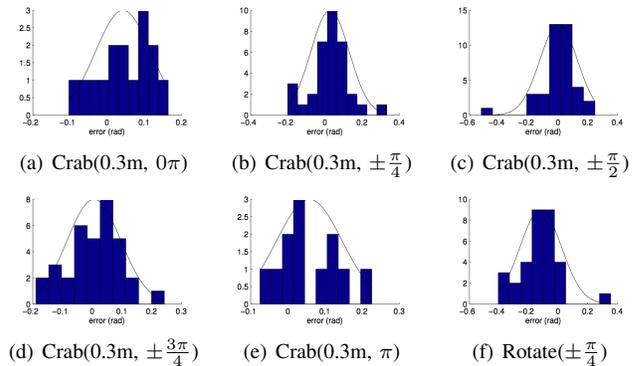


Fig. 4. Mobility prediction by action on flat terrain.

TABLE I
MOBILITY PREDICTION BY ACTION, GP FEATURES NOT INCLUDED

Action:	crab 0π	crab $\pm\pi/4$	crab $\pm\pi/2$	crab $\pm3\pi/4$	crab π	rotate $\pm\pi/4$
Error:	δs_{head}	δs_{head}	δs_{head}	δs_{head}	δs_{head}	δs_{yaw}
Flat Terrain						
mean (rad)	0.043	0.028	0.004	0.006	0.058	-0.117
std. (rad)	0.074	0.103	0.127	0.091	0.088	0.140
# samples	15	34	39	34	12	32
Rough Terrain - marginalised by Action						
mean (rad)	0.044	0.010	0.060	0.037	0.037	-0.063
std. (rad)	0.081	0.115	0.158	0.117	0.089	0.119
# samples	43	58	58	48	35	54

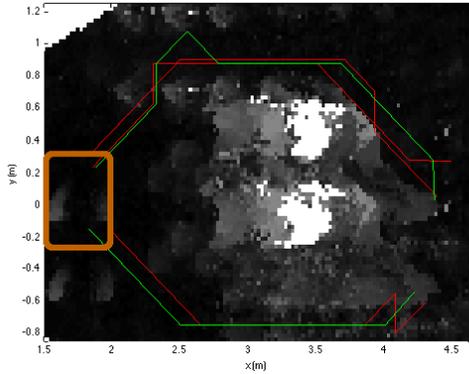
difficult to predict by modelling, were therefore captured in our learning data. As expected, the control errors were more significant in rough terrain data. Table II shows the GP hyperparameters obtained.

B. Simulation of Flat Terrain Traversal

We simulated the robot traversing flat terrain. Control uncertainty was simulated using the learned data described in Sec. V-A. The robot’s environment was simulated using a point cloud acquired by the robot’s RGB-D camera. This environment is a roughly flat area with a cluster of rocks, shown in Fig. 5(b). Trials consisted of placing the robot randomly around the cluster of rocks and directed to a unique goal region opposite the rock cluster. We used a cost function where any rock on the terrain appears as an obstacle.

TABLE II
GP HYPERPARAMETERS TRAINED FROM TRAVERSALS IN ROUGH TERRAIN

Action	Error	Λ_{11}^{-2}	Λ_{22}^{-2}	Λ_{33}^{-2}	Λ_{44}^{-2}	σ_f	σ_n
crab 0π	δs_{head}	0.056	0.052	10.59	10.56	0.251	0.029
crab $\pm\pi/4$	δs_{head}	0.068	0.010	0.063	0.087	0.183	0.032
crab $\pm\pi/2$	δs_{head}	0.644	0.448	0.015	0.024	0.000	0.111
crab $\pm3\pi/4$	δs_{head}	1.89	0.058	0.068	1.78	0.284	0.040
crab π	δs_{head}	0.055	0.042	0.151	2.47	0.163	0.032
rotate $\pm\pi/4$	δs_{yaw}	0.007	0.246	0.036	0.014	0.135	0.042



(a) Simulated trajectories



(b) The real terrain

Fig. 5. (a) shows a few samples of simulated trajectories to navigate around a cluster of rocks (shown in (b)). The cost on the map is shown as levels of grey, with white indicating the highest cost, for a single orientation value: the rover facing left. The brown rectangle on the left indicates the common goal region. Red trajectories were computed without considering uncertainty, while green trajectories considered uncertainty.

Planning without uncertainty (whereby the expectation of the transition function $\mathbb{E}(P(\Delta s | \lambda(s, a), a))$ learned is used instead of the full distribution) and planning considering uncertainty were each tested 100 times. When planning considering uncertainty, the flat-terrain learning data were used, whereby Δs_{head} was considered during crab actions, and Δs_{yaw} was considered during rotate actions. Resultant trajectories were assessed in light of the known Δs_{head} and Δs_{yaw} distributions to determine the probability of colliding with a rock as well as the expectation of accumulated cost for each trajectory. Results obtained for both methods can be compared using the statistics on all *executed* trajectories in Table III. These statistics represent: the averages of total cost accumulated from start position to goal, the probability of hitting an obstacle summed over the entire trajectory and the minimum distance to an obstacle over the trajectory. Fig. 5(a) shows a few examples of trajectories executed.

Results in Table III highlight two of the major consequences of planning without uncertainty: a platform will be more likely to collide with an obstacle and will, on average, accumulate more cost in traversing to the goal region. In fact, when planning without uncertainty, the projected accumulated cost will always be underestimated when the robot deviates at least once from the lowest cost path it computes,

TABLE III
SIMULATED PLANNING AROUND A CLUSTER OF ROCKS

No Uncertainty	Cost	Prob.	Min. Dist. (m)
mean	1.132	0.028	0.154
std.	0.605	0.101	0.131
max.	6.879	0.68	0.530
min.	0.727	0	0.015
Head. Uncertainty	Cost	Prob.	Min. Dist. (m)
mean	1.080	0.005	0.161
std.	0.138	0.035	0.113
max.	1.449	0.34	0.505
min.	0.733	0.0	0.025

which is frequently the case with imperfect control. Thus planning without uncertainty cannot provide any guarantee of total cost accumulated to reach a goal region, which is important when a decision needs to be made regarding if a goal region is worth visiting up to a certain cost of traversing there. The safety issue of rock collisions occurs when the planner follows paths very close to an obstacle without considering consequences of slight deviations from its path.

C. Experiments of Flat Terrain Traversal

We conducted a series of experiments using the physical robot traversing flat terrain. The robot navigated from a starting position to a goal region whilst avoiding large rocks. The terrain was flat sand and gravel, however due to the imperfect control on the loose terrain, the control uncertainties were significant, as shown earlier in the training data in Sec. V-A.

Trajectories were planned and executed on the robot 10 times for planning with uncertainty (in heading during crabbing actions and yaw during rotation actions). For comparison, 10 trajectories were also planned and executed without accounting for uncertainty. In both cases, two different starting points were used, while the goal region was the same.

Fig. 6 shows a few examples of the actual trajectories executed by the rover. Occlusions shown in the cost map are due to “shadows” of rocks when they were observed by the robot.

The results of planning with uncertainty of heading during crabbing actions and yaw during rotation actions are compared statistically against results of planning without uncertainty in Table IV. Results indicate that the robot would, on average, plan wider berths around rocks with uncertainty considered and execute safer paths in practice.

D. Experiments on Unstructured Terrain: Traversing Rocks

We also conducted a series of experiments where the robot traverses rough terrain. Trained GP models were used to predict control uncertainty as described in Sec. V-A. As in the training phase, the experiments were conducted in unstructured and rough terrain because of the presence of rocks that the robot sometimes has to travel across, but they were limited to areas with negligible terrain slopes.

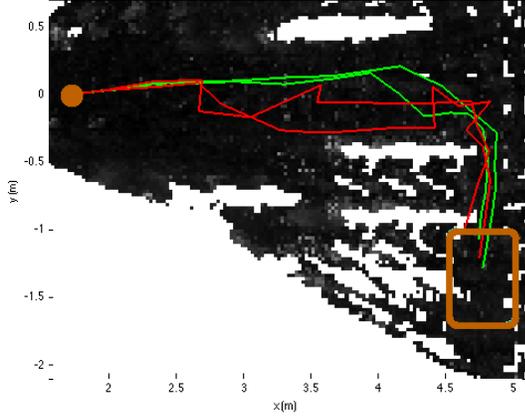


Fig. 6. Example of trajectories taken to avoid several rocks on otherwise flat terrain. Red: without uncertainty considered. Green: with uncertainty in heading. The starting position is indicated by the brown circle on the left and the goal region is shown as the brown rectangle at the bottom right. Planning with uncertainty generates paths that are less sensitive to control error.

The learned rough-terrain models were limited to traversal of one rock at a time, and so a rock field (Fig. 7) was set up accordingly. However, the layout of rocks was dense enough to cause the robot to traverse multiple rocks while navigating to the goal region.

In addition to Δs_{yaw} during rotation actions, two types of uncertainty were considered (separately) with crab actions: Δs_{head} and Δs_{dist} . Fig. 8 shows an example policy computed by our planning algorithm.

Resultant trajectories in Fig. 9 show that whilst each planning method attempted to navigate between rocks, the method of planning without uncertainty (red) would tend to “zig zag” more severely in attempts to attain the least possible cost path, finely navigating every rock. Methods considering uncertainty (green, cyan), by contrast, appear smoother. The “zig zag” action sequences can actually result in more rock traversals in total, due to a greater distance travelled in the rock field.

Results shown in Table V indicate that policies chosen by the planning with uncertainty case were favourable, with less accumulated cost. In these tests, considering heading uncertainty was most significant to overall cost. In this table, “stuck states” refers to situations where the rover dug one of its wheels in next to a rock during the test and could not initially mount the rock as the result. This occurred 20% of

TABLE IV
FLAT TRAVERSAL: PROBABILITY ASSESSMENT

No Uncertainty			Cost	Prob. Hit	Min. Dist. (m)
# trials:	10	Mean:	1.295	0.380	0.095
# collisions:	1	Std.:	0.385	0.492	0.084
Uncertainty			Cost	Prob. Hit	Min. Dist. (m)
# Trials:	10	Mean:	1.177	0.016	0.165
# Collisions:	2	Std.:	0.262	0.005	0.088

TABLE V
PLANNING WITH TRAVERSAL OVER ROCKS

Uncertainty considered	# Trials	# Stuck States	Mean Cost	Std. Cost
none	5	1	1.46	0.050
$\Delta s_{dist}, \Delta s_{yaw}$	5	1	1.31	0.083
$\Delta s_{head}, \Delta s_{yaw}$	5	0	1.19	0.061

the time when no uncertainty was considered. When considering uncertainty in distance travelled, a stuck state still occurred. However, planning with heading uncertainty again made more impact, with all stuck states avoided. Although the sample size is limited, this trend in the data supports the claim that our approach to consider the control uncertainty can significantly improve the safety of the platform at the execution of planned trajectories.

Note that in these experiments we considered two sources of uncertainty independently. Considering them in combination should then have an even stronger impact on the platform safety. This will be tackled in future work.

VI. CONCLUSION

Since the motion of any real mobile robot is stochastic to some degree, considering control uncertainty at the planning stage enables us to significantly enhance the safety of the platform (e.g. mitigating chances of collisions) and lower the cost accumulated on average during the execution of planned trajectories in real environments. These claims were demonstrated in this paper using a holonomic planetary rover. A model of uncertainty was built using learning data and Gaussian processes to predict motions over flat and unstructured terrain in a Mars yard setup. The trained model was then exploited to plan policies using dynamic programming, and to execute paths following the planned policies. Experimental validation was achieved both in simulation and in real experiments, in a variety of situations, taking into account uncertainties in heading and distance travelled. The experimental validation demonstrates the results obtained on the actual executed trajectories compared to trajectories executed when planning without considering uncertainty (de-



Fig. 7. Rock traversal experiment setup: rover is shown at its starting position. It must traverse over a “rock field” to reach a goal region to the right (out of the picture).

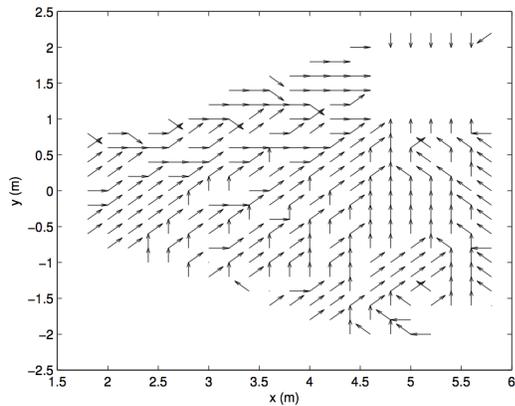


Fig. 8. Example policy obtained for rough terrain experiment with Δs_{head} considered, projected onto the $x - y$ plane. The goal is the empty square in the upper right corner of the figure. Arrows indicate the preferred crab actions at each state. Dots indicate rotations.

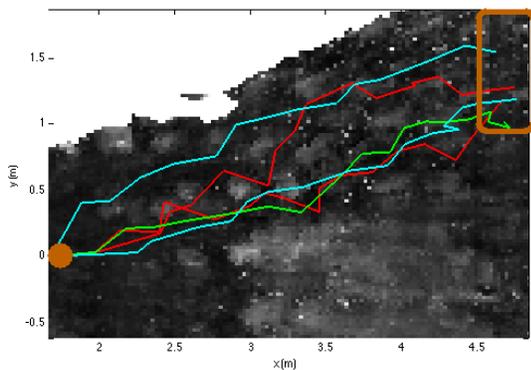


Fig. 9. Examples of trajectories during the rock traversal experiments, comparing planning when considering different uncertainty sources. Red: without uncertainty considered. Green: with uncertainty in heading. Cyan: with uncertainty in distance. Cost map shown in grayscale, from dark to light as cost increases. The starting position is indicated by the brown disk on the left and the goal region is shown as the brown rectangle at the top right corner. The planner generally attempts to navigate between most of the rocks which are high cost.

terministic control). These results show the improvement in safety and accumulated cost when accounting for uncertainty.

In future work, we will consider more complex terrain with larger slopes, denser collections of small rocks. We are particularly interested in studying the ability of our approach to learn and predict the deviations of control

actions due to loose rocks that shift during traversal. We will also investigate an extension of the GP learning to include observations collected during navigation.

REFERENCES

- [1] P. S. Schenker, T. L. Huntsberger, P. Pirjanian, E. T. Baumgartner, and E. Tunstel, "Planetary rover developments supporting mars exploration, sample return and future human-robotic colonization," *Autonomous Robots*, vol. 14, pp. 103–126, 2003.
- [2] S. LaValle, *Planning Algorithms*. Cambridge Univ. Press, 2006.
- [3] G. Ishigami, G. Kewlani, and K. Iagnemma, "Statistical Mobility Prediction for Planetary Surface Exploration Rovers in Uncertain Terrain," in *IEEE International Conference on Robotics and Automation*, 2010.
- [4] S. Lacroix, A. Mallet, D. Bonnafous, G. Bauzil, S. Fleury, M. Herrb, and R. Chatila, "Autonomous Rover Navigation on Unknown Terrains: Functions and Integration," *The International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 917–942, 2002.
- [5] D. Helmick, A. Angelova, and L. Matthies, "Terrain Adaptive Navigation for planetary rovers," *Journal of Field Robotics*, vol. 26, no. 4, pp. 391–410, 2009.
- [6] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *IEEE International Conference on Robotics and Automation*, 2011.
- [7] G. Ishigami, K. Nagatani, and K. Yoshida, "Path planning for planetary exploration rovers and its evaluation based on wheel slip dynamics," in *IEEE International Conference on Robotics and Automation*, 2007.
- [8] S. Patil, J. Berg, and R. Alterovitz, "Motion planning under uncertainty in highly deformable environments," in *Robotics: Science and Systems VII*, 2011.
- [9] J. Berb, P. Abbeel, and K. Goldberg, "LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information," in *Robotics: Science and Systems VI*, 2010.
- [10] S. M. LaValle and S. A. Hutchinson, "An Objective-Based Framework for Motion Planning under Sensing and Control Uncertainties," *The International Journal of Robotics Research*, vol. 17, no. 1, pp. 19–42, 1998.
- [11] H. Kurniawati, T. Bandyopadhyay, and N. Patrikalakis, "Global motion planning under uncertain motion, sensing, and environment map," in *Robotics: Science and Systems VII*, 2011.
- [12] C. A. Brooks and K. Iagnemma, "Self-Supervised Terrain Classification for Planetary Surface Exploration Rovers," *Journal of Field Robotics, Special Issue on Space Robotics, Part I*, vol. 29, no. 3, pp. 445–468, 2012.
- [13] K. Iagnemma, H. Shibly, A. Rzepniewski, and S. Dubowsky, "Planning and control algorithms for enhanced rough-terrain rover mobility," in *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2001.
- [14] C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [15] F. Endres, J. Hess, N. Engelhard, J. Sturm, D. Cremers, and W. Burgard, "An evaluation of the rgb-d slam system," in *IEEE International Conference on Robotics and Automation*, 2012.
- [16] M. Quigley et al., "Ros: an open-source robot operating system," in *Open-Source Software Workshop, IEEE International Conference on Robotics and Automation*, 2009.
- [17] D. Bonnafous, S. Lacroix, and T. Simeon, "Motion generation for a rover on rough terrains," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2001.