Planning and Perception for Robotic Manipulators in Semi-structured Environments

by Fouad Sukkar

Thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

under the supervision of Prof. Robert Fitch

University of Technology Sydney
Faculty of Engineering and Information Technology

February 2022

Certificate of Original Authorship

I, Fouad Sukkar declare that this thesis, is submitted in fulfilment of the requirements for the award of Doctor of Philosophy, in the School of Mechanical and Mechatronic Engineering at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research is supported by the Australian Government Research Training Program.

Production Note:
Signature removed prior to publication.

Fouad Sukkar

February 2022

Planning and Perception for Robotic Manipulators in Semi-structured Environments

by

Fouad Sukkar

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

Abstract

New applications of robotic manipulators require a far greater degree of autonomy than traditionally has been needed. Industrial robot manipulators such as collaborative robots are designed for advanced manufacturing applications where they should operate safely in dynamic work environments shared with humans and should be able to adapt quickly to perform a variety of tasks. Manipulators for fruit harvesting in agricultural robotics need to operate quickly in semi-structured outdoor environments in order to perform at a commercially viable level.

These applications require many components to work together robustly and efficiently in order to be viable for real world use. For example, robotic harvesting involves task and motion planning, task space control, active perception, object segmentation and grasping. Extending this to multiple robots requires coordination among agents which is non-trivial and further increases problem complexity. Existing methods for solving such problems in real-time often resort to making coarse approximations to one or many parts of the problem, for example a simplified motion or sensor model. This can lead to undesirable behaviour particularly for highly non-linear systems such as robotic arms or complex sensor models such as those for RGB-D cameras. Furthermore, when searching for objects it remains a challenge to effectively balance between exploration of a scene and targeted observation of key regions.

This thesis is concerned with developing component algorithms that solve such problems effectively and exhibit computational behaviour that is well-understood analytically. Key to achieving this is the exploitation of known a priori structure in the problem, typical of semi-structured environments. First, a novel planning framework called Hausdorff approximation planner (HAP) is described which utilises prior knowledge for building a library of trajectories which are guaranteed to be of bounded length and approximately correspond to geodesics in the task space. These trajectories can be concatenated reliably over long task horizons, adapted quickly online and queried for accurate costs in constant time. Then, a novel targeted information gathering algorithm is presented which is integrated with HAP and other components into a non-myopic active perception framework suitable for systems with complex motion and sensor models and effectively balances between exploration and targeted observations of the scene. This is then generalised into a decentralised coordination framework suitable for flexible multi-robot perception problems that facilitate downstream processes, such as manipulation, object segmentation, and collaborative tasks. To demonstrate the flexibility and effectiveness of these frameworks they are evaluated extensively in various problem settings with several different robotic platforms both in simulation and hardware, including robotic pruning, robotic harvesting, and searching for objects in clutter using multiple robotic manipulators with eye-in-hand sensors.

Acknowledgements

This thesis would not have been possible without the help and support of the many people around me and whom I have met throughout my studies, both locally and abroad. First and foremost, thank you to Prof. Robert Fitch. You have put up with me as a supervisor for a majority of my academic career, spanning multiple degrees, and have never failed me in your advice and guidance. Your ability to communicate technical ideas in a clear and inspiring way never ceases to amaze me and I have learnt to become a better communicator as a consequence. Thank you to my co-supervisors Dr Alen Alempijevic and Dr Chanyeol Yoo for your words of encouragement and invaluable advice.

I am deeply grateful to my fellow researchers from the robotics group at UTS, especially Felix Kong, Giovanni D'urso, Brian Lee, Cadmus To, James Lee, Jennifer Wakulicz and Jayant Khatkar; you have all inspired creativity and helped make my PHD an unforgettable experience. I am thankful for the opportunity to have collaborated and spent time with amazing groups from Oregon State University and Washington State University - Graeme Best, Joseph R. Davidson, Alexander You and Manoj Karkee - my thesis has greatly benefited from your help and advice. To the reading groups I've attended over the past years at UTS and Usyd, I appreciate the invaluable discussions and feedback which helped me refine the ideas in this thesis and produce high quality publications.

Last but not least, thanks to my Mum, Dad, sisters Jessica and Clarissa, and friends and family for your unending support. It has been a challenging time and you have helped me through the best and worst, this thesis is dedicated to you.

Contents

Declaration of Authorship iii

Abstract v

Acknowledgements vii

Contents ix

List of Figures xv

List of Tables xv

List of Algorithms xvii

List of Theorems and Definitions xxi

Nomenclature xxi

1 Introduction 1
- 1.1 Planning In Semi-Structured Environments . . . 3
- 1.2 Manipulator Planning for Perception . . . 5
- 1.3 Limitations of Current Manipulator Planners . . . 6
- 1.4 Planning with Hausdorff Approximations (HAs) . . . 8
- 1.5 Scope . . . 11
  - 1.5.1 Manipulator Task and Motion Planning . . . 11
  - 1.5.2 Information Gathering and Active Perception . . . 12
  - 1.5.3 Task-Space Control and Manipulation . . . 12
  - 1.5.4 Multi-Robot Active Perception . . . 13
  - 1.5.5 Problem Settings . . . 14
- 1.6 Principal Contributions . . . 17
- 1.7 Thesis Outline . . . 19

2 Related Work 21
- 2.1 Planning with Domain Knowledge . . . 21

2.2 General Motion Planners . . . 24
2.2.1 Grid-Based Planning . . . 24
2.2.2 Sampling-Based Planning . . . 25
2.2.3 Trajectory Optimisation . . . 28
2.3 Task and Motion Planning . . . 31
2.4 Coverage Planning . . . 32
2.5 Next-Best-View Planning . . . 32
2.6 Information Gathering . . . 35
2.7 Multi-Robot Active Perception . . . 36

3 Hausdorff Approximation Planner (HAP) 37
3.1 Overview . . . 37
3.1.1 Chapter Outline . . . 40
3.2 Problem Formulation . . . 40
3.2.1 Notation . . . 40
3.2.2 Task and Motion Planning Problem . . . 40
3.2.3 Problem Reduction . . . 41
3.3 Online Motion Planning with HAs . . . 43
3.3.1 Retrieval of Trajectories . . . 43
3.3.2 Modifying Trajectories to Satisfy Tasks . . . 44
3.3.3 Task Sequencing . . . 45
3.3.4 Online Adaptation of Trajectories . . . 45
3.4 Trajectory Library Construction . . . 46
3.4.1 Algorithm Overview . . . 46
3.4.2 Modified Dijkstra Algorithm for Finding HAs . . . 48
3.4.3 Practical Considerations . . . 49
3.5 Motivation for Multiple HAs . . . 55
3.6 Analysis . . . 58
3.6.1 Complexity Analysis . . . 58
3.6.2 Bounded Suboptimality Analysis . . . 59
3.7 HAP Evaluation . . . 62
3.7.1 Experimental Setup . . . 62
3.7.2 Implementation Details . . . 65
3.7.3 Bookshelf Environment with 6D Pose Tasks . . . 66
3.7.4 Task-Space Control . . . 70
3.8 Robotic Pruning . . . 76
3.8.1 Experimental Setup . . . 77
3.8.2 Hardware Results . . . 81
3.9 Summary . . . 82

4 Targeted Information Gathering 83
4.1 Overview . . . 83
4.1.1 Chapter Outline . . . 85
4.2 Problem Formulation . . . 85

4.2.1 Robot and Environment 85
4.2.2 Observation Model 86
4.3 Problem Statement 87
4.4 Viewpoint Evaluation 88
4.4.1 Surface Reconstruction 88
4.4.2 Volumetric Information Formulation 88
4.5 Active Perception Framework 91
4.5.1 Monte Carlo Tree Search 92
4.5.2 Reward Function 93
4.5.3 Computing Realistic Trajectory Costs for Budgeting 95
4.6 Apple Harvesting 96
4.6.1 Experimental Setup 96
4.6.2 Implementation Details 98
4.6.3 Hardware Results 102
4.7 Summary 110

5 Decentralised Coordination Framework 111
5.1 Overview 111
5.1.1 Chapter Outline 112
5.2 Robot Team 112
5.3 Problem Statement 113
5.4 Decentralised Coordinated Planning 113
5.4.1 Decentralised Monte Carlo Tree Search 114
5.4.2 Integration with Active Perception Framework 115
5.4.3 System Execution 116
5.4.4 Problem Complexity Analysis 117
5.5 Searching for Objects with Multiple Robot Manipulators 119
5.5.1 Experimental Setup 120
5.5.2 Simulation Results 121
5.5.3 Hardware Results 122
5.6 Unoccluding Objects Via Manipulation 124
5.6.1 Experimental Setup 125
5.6.2 Hardware Demonstration 128
5.7 Summary 129

6 Discussion 135
6.1 Importance of Good Trajectory Initialisation 135
6.2 Analysing the Behaviour of HAP 137
6.3 Using Easy-To-Find Information to Bootstrap Hard-To-Find Information . 141
6.4 Manipulation for Perception 142

7 Conclusions and Future Work 145
7.1 Thesis Summary 146
7.1.1 Hausdorff Approximation Planner (HAP) (Chapter 3) 146
7.1.2 Targeted Information Gathering (Chapter 4) 147

7.1.3 Decentralised Coordination Framework (Chapter 5) 147
7.1.4 Discussion (Chapter 6) 148
7.2 Summary of Contributions 148
7.2.1 HAP Framework 148
7.2.2 Targeted Information Gathering VI Metric 149
7.2.3 Active Perception Framework 149
7.2.4 Decentralised Coordination Framework 150
7.2.5 Theoretical Analysis 150
7.2.6 Empirical Results 150
7.3 Future Work 151
7.3.1 HAP Framework 151
7.3.2 Targeted Information Gathering Framework 153
7.3.3 Decentralised Coordination Framework 154
7.4 Outlook 155

Bibliography **159**

Appendices **184**
A Experiment Supporting Figures 185
A.1 Naive Task Region Base Allocations for HAP Experiments 185
A.2 Apple Trellis Configurations for Harvesting Experiments 186
B Supplementary Algorithmic and Implementation Details 188
B.1 Null-Space Control for Avoiding Joint Limits 188
B.2 Upper Confidence Bound 1 Applied to Trees (UCT) 189
B.3 Discounted UCT (D-UCT) 189

List of Figures

1.1 Examples of robots operating in semi-structured environments. 4
1.2 Simple illustration of Hausdorff Approximation concept. 8
1.3 Illustration of Hausdorff Approximation mapping from task space to configuration space. 9
1.4 Two robot arms coordinate their actions to discover and reconstruct apples on a trellis. 16

3.1 Example bookshelf environment illustrating graph construction. 42
3.2 A side-by-side comparison of HAs with largest coverage. 55
3.3 Visualisation of HA mappings for UR5. 56
3.4 Visualisation of HA mappings for UR5 with mobile base. 57
3.5 Bookshelf object configurations used for the simulation experiments. 63
3.6 Planning success rate and times for 7-DOF Sawyer experiments with the fixed-base model. 67
3.7 Planning success rate and times for 7-DOF UR5 with the mobile-base model. 68
3.8 Task-space control experiment setup for HAP evaluation. 71
3.9 HAs computed for UR5 with static base for task-space control experiments. 73
3.10 Subspaces computed for HAP-Mobile. 74
3.11 UR5 task-space control experiment results. 75
3.12 Our autonomous pruning system with a custom cutter mechanism. 76
3.13 Our pruning robot setup, consisting of a UR5e arm, an Intel RealSense D435 RGB-D camera, and custom pneumatically-actuated shears. 77
3.14 The test setup used for experimental evaluation. 77
3.15 Simulation visualisation of cutpoint generation. 79
3.16 Robot pruning a branch with simulation environment and in-hand camera view. 81

4.1 Searching for objects of interest, here apples, that are sparsely located and in some cases heavily occluded. 84
4.2 Ray tracing visualisation. 87
4.3 Visualisation of view-point evaluation function. 90
4.4 Two separate octrees (shown here as quadtrees for simplicity) used for the exploration and targeted information VIs. 94
4.5 Experimental setup for apple harvesting experiments. 97
4.6 Apple harvesting execution pipeline and subsystem layout. 98

4.7 Discretised viewpoints computed by HAP that are used by the active perception framework to search over. . . . 101
4.8 An example sequence of views computed by the active perception component in the apple harvesting experiment. . . . 105
4.9 Resulting reconstruction and segmented apples. . . . 106
4.10 An example grasp sequence from the apple harvesting experiment. . . . 108

5.1 High-level overview of decentralised planning framework. . . . 114
5.2 Experimental setup for multi-robot active perception. . . . 118
5.3 Multi-robot active perception roadmap construction. . . . 119
5.4 Apple trellis configurations for simulated and hardware experiments. . . . 120
5.5 Samples viewpoints from dataset used for simluated experiments. . . . 122
5.6 Apple detection rates. . . . 123
5.7 Simulated experimental results. . . . 124
5.8 Hardware experimental results. . . . 125
5.9 Examples of obstructions in a real apple orchard. . . . 126
5.10 Experimental setup for unoccluding objects via manipulation. . . . 127
5.11 Poke obstruction example. . . . 128
5.12 Push obstruction example. . . . 129
5.13 Hardware demonstration of manipulator unoccluding objects for active perception. . . . 132
5.14 Final reconstructed point cloud with segmented apples. . . . 133

6.1 Side view of bookshelf environment showing comparison of naive planner's initial trajectory seed and HAP's. . . . 136
6.2 Example subset of a Naive-Full trajectory sequence for Sawyer bookshelf experiment. . . . 137
6.3 Example HAP-Full trajectory sequence for fixed-base model bookshelf experiment. . . . 138
6.4 Example subset of a Naive-Mobile trajectory sequence for mobile-base model bookshelf experiment. . . . 139
6.5 Example HAP-Mobile trajectory sequence for mobile-base model bookshelf experiment. . . . 140
6.6 Detected ROIs overlaid on apples trellis. . . . 141

7.1 Unseen bookshelf environment. . . . 152
A.1 Subspaces and base positions (manually defined) for Naive-Mobile in single- and batch-query experiments. . . . 185
A.2 Subspaces and base positions (manually defined) for the Naive-Mobile method in task-space control experiments. Design intention as in previous figure. . . . 186
A.3 Trellis configurations used for apple harvesting experiments. . . . 187

List of Tables

3.1 Variants used for the Sawyer experiments. 64
3.2 Variants used for the UR5 experiments. 64
3.3 Parameters used for experiments. 66
3.4 HAP validation results. 70
3.5 Results from pruning experiments. 80

4.1 Experimental results for apple grasping. 109
4.2 Classification of hardware experiment failures. 109

List of Algorithms

3.1 Online Planner 43
3.2 Generate Trajectory Library 49
3.3 Modified Dijkstra 50
3.4 Get Candidate 50
3.5 Update 51
4.1 MCTS rollout with targeted information gathering VI 95
5.1 Dec-MCTS rollout with targeted information gathering VI 115
5.2 Decentralised Coordination Framework 116
5.3 Dec-MCTS Overview 117

List of Theorems and Definitions

3.1 Problem (Task and motion planning problem) 41
3.2 Problem (Motion planning problem) 42
3.3 Problem (Task sequencing problem) 42
3.1 Definition (ϵ-Gromov-Hausdorff approximation (ϵ-GHA)) 47
3.1 Lemma (Runtime of *Get Candidate* (Algorithm 3.4)) 58
3.1 Theorem (Runtime of *Modified Dijkstra* (Algorithm 3.3)) 59
3.1 Remark (Runtime dependence on number of mappings) 59
3.2 Remark (Runtime with mobile base) 59
3.2 Theorem (θ is an ϵ-GHA) 60
3.2 Definition (Geodesics) 61
3.3 Theorem (ϵ-GHAs preserve geodesics) 61

4.1 Problem (Viewpoint plannning) 87

5.1 Problem (Decentralised viewpoint planning) 113
5.1 Theorem (Complexity of decentralised viewpoint planning problem) 117

Nomenclature

Acronyms & Abbreviations

AGV	autonomous ground vehicle
AUV	autonomous underwater vehicles
BIT*	Batch Informed Trees
Dec-MCTS	decentralised Monte Carlo tree search
DOF	degrees of freedom
D-UCT	discounted-UCT
ϵ-GHA	ϵ-Gromov-Hausdorff approximation
FOV	field of view
GPIS	Gaussian process implicit surface
HA	Hausdorff approximation
HAP	Hausdorff approximation planner
IG	information gain
IK	inverse kinematics
MCTS	Monte Carlo tree search
NBV	next best view
PBVS	position-based visual servoing
PRM	probabilistic roadmap
ROI	region of interest
ROS	robot operating system
RRT	rapidly-exploring random tree

TAMP	task and motion planning
TSP	travelling salesman problem
UCB	upper confidence bound
UCT	UCB applied to trees
UR	Universal Robots
UTS	University of Technology, Sydney
VI	volumetric information

List of Publications

Conference Papers

1. **F. Sukkar**, G. Best, C. Yoo and R. Fitch, "Multi-Robot Region-of-Interest Reconstruction with Dec-MCTS," *Proc. IEEE Int. Conf. on Robotics and Automation*, May, 2019. **Winner of the IEEE Best Paper Award in Service Robotics.**

2. A, You, **F. Sukkar**, R. Fitch, M. Karkee and J. R. Davidson "An Efficient Planning and Control Framework for Pruning Fruit Trees," *Proc. IEEE Int. Conf. on Robotics and Automation*, May, 2020.

Journal Papers

1. **F. Sukkar**, J. Wakulicz, K. M. B. Lee and R. Fitch, "Motion planning in task space with Gromov-Hausdorff approximations", *The International Journal of Robotics Research.* ***submitted to**

Awards

Conferences

1. IEEE ICRA 2019 Best Paper Award in Service Robotics (Nominated for Service Robotics and Robotic Vision Awards)

University

1. UTS Faculty of Engineering and Information Technology 2019 Research Showcase - First Prize (Oral Presentation)
2. School of Mechanical and Mechatronic Engineering 2019 Research Showcase - First Prize (Oral Presentation)