DISS. ETH NO. 26179

Environmental Mapping and Informative Path Planning

for UAV-based Active Sensing

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Marija Popović

MEng. Integrated Mechanical & Electrical Engineering,
University of Bath

born on May 14, 1993

citizen of Belgrade, Serbia

accepted on the recommendation of

Prof. Dr. Roland Y. Siegwart, Examiner
Dr. Teresa Vidal-Calleja, Co-examiner

2019

Institute for Robotics and Intelligent Systems, Autonomous Systems Lab
ETH Zurich
Switzerland

© 2019 Marija Popović. All rights reserved.

Abstract

Robotic platforms represent a new frontier for data acquisition in a wide range of monitoring and exploration missions, including agriculture, surveillance, post-disaster assessment, and environmental sensing. A key challenge to realize their full potential is deciding how an autonomous agent should act to gather the most useful data about an uncertain environment within the set of its resource constraints. To address this, this thesis investigates the problem of *active sensing*: in such scenarios, how can the actions of the agent be planned in order to maximize the efficiency of the data collection process? It proposes new methods for environmental mapping and informative path planning as crucial elements towards achieving this goal. The strategies are developed in the context of two distinct applications: (a) terrain monitoring, and (b) active sensing under localization uncertainty, with a focus on Unmanned Aerial Vehicle (UAV) systems.

The main contribution of this thesis is an informative planning framework that is applicable for general active sensing tasks. The overall approach integrates individual contributions on three different fronts, which aim to address the challenges associated with information gathering in uncertain 3-D environments with computationally limited systems. Firstly, a new method for environmental field mapping is presented for terrain monitoring scenarios. The strategy exploits a Gaussian Process (GP) model as a prior for recursive Bayesian data fusion with probabilistic, variable-resolution sensors. In doing so, it supports mapping using dense visual imagery without the computational burden of standard GP regression, making it suitable for online on-platform applications. Moreover, it accommodates noisy sensors with altitude-dependent performance, as relevant for UAV-based systems. Secondly, an online, adaptive informative path planning algorithm is introduced for generating continuous trajectories to collect data in resource-constrained missions. A key feature of the method is that it uses the output from a discrete grid search as an informed prior to initialize a trajectory optimization routine and thereby improve its convergence in a large and complex objective space. This strategy also enables trading off between computational efficiency and solution accuracy for deployment on systems with limited computing power. Finally, methods are proposed to account for the robot pose uncertainty in active sensing tasks. Unlike prior work, the approach propagates this uncertainty into both the mapping and planning modules towards improving the robustness and accuracy of information gathering. A new utility function is developed that allows the robot to automatically trade off between exploiting the existing map to maintain good localization and exploring areas to acquire new data in a principled way, without relying on any manually-tuned parameters. The formulation is derived in the context of a GP-based monitoring scenario and is also applicable across different learning problems.

The developed framework is modular, and can be tailored to a wide range of active sensing problems. Extensive simulation studies were conducted to evaluate the approach, examining how it performs against existing methods both as an integrated system as well as in terms of its key components. The main findings show that the proposed approach effectively: (a) produces maps with similar certainty and accuracy in significantly less time compared to current planning strategies; (b) can focus on adaptively mapping specific areas of interest; and (c) improves upon both field map accuracy and

robot localization by accounting for the pose uncertainty in informative planning. Results using an experimental dataset demonstrate system integration and validation in a photorealistic UAV-based terrain monitoring scenario. Finally, field tests are presented to demonstrate the algorithms implemented and running in real-time on robots for various data gathering tasks, including vegetation mapping on a farm. The framework is made publicly available as an open-source package.

Zusammenfassung

Roboterplattformen erlauben eine Vielzahl neuer Möglichkeiten für die Datenerfassung im Rahmen von Überwachungs- und Analyseaufgaben, wie beispielsweise Geländeerkundung, Landwirtschaft, Observierung und Schadensanalysen. Die zentrale Herausforderung hierbei die Vorgehensweise, damit die wichtigsten Daten erhoben werden, in unbekanntem Gelände und limitierten Resourcen der Plattform. Diese Arbeit untersucht hierfür das Problem der aktiven Wahrnehmung: Wie soll sich die Einheit verhalten, um die Effizienz der Datenerhebung zu maximieren? Dafür werden neue Methoden für die Landschaftskartierung und die informative Pfadplanung, den beiden wichtigsten Elementen dieses Prozesses, vorgeschlagen. Die Strategien werden im Rahmen von zwei verschiedenen Projekten entwickelt: (a) Geländeüberwachung und (b) aktives Erfassen bei Lokalisierungsunsicherheiten, mit Schwerpunkt auf unbemannten Luftfahrzeugen (UAV).

Der Hauptbeitrag dieser Arbeit ist ein informativer Planungsrahmen, welcher für allgemeine Aufgaben der aktiven Wahrnehmung anwendbar ist. Der Ansatz integriert einzelne Beiträge auf dreierlei Ebenen, mit dem Ziel, die Herausforderungen im Zusammenhang mit der Informationsbeschaffung in dreidimensionalen, fehlerbehafteten Geländemodellen anzugehen. Zuerst wird eine neue Methode zur Geländeobservierung vorgestellt. Die Strategie basiert auf einem "Gaußschen Prozess" (GP) als Prädiktionsschritt für eine rekursive, bayes'sche Datenfusion mit Sensoren mit statistisch verteilter, variabler Auflösung. Dabei wird die Kartierung mit hochauflösenden Bildern unterstützt, ohne die rechnergestützte Belastung durch die Standard-GP-Regression, wodurch diese Strategie für Online-Plattformanwendungen geeignet ist. Außerdem ermöglicht es den Gebrauch von verrauschten Sensoren mit höhenabhängiger Leistung, welche für UAV-basierte Systeme relevant sind. Zweitens wird ein adaptierbarer Algorithmus zur Echtzeit- Pfadplanung vorgestellt. Dieser generiert kontinuierliche Trajektorien für Operationen mit eingeschränkten Resourcen. Ein wesentliches Merkmal des Verfahrens ist, dass als Prädiktion zur Pfadgenerierung eine Rastersuche verwendet wird. Dabei wird die Konvergenz in grossen, komplizierten Zielräumen verbessert. Zusätzliche wird auch eine Abwägung zwischen rechnerischer Effizienz und Genauigkeit der Lösung ermöglicht. Schließlich werden Methoden vorgeschlagen, um der Orientierungsungenauigkeit des Roboters Rechnung zu tragen. Anders, als bei bisherigen Ansätzen wird diese Abweichung sowohl in die Kartierung wie auch die Pfadplanung propagiert, um die Robustheit und Genauigkeit der Informationsbeschaffung zu verbessern. Eine neue Nutzenfunktion wird entwickelt, welche es dem Roboter ermöglicht, einen Kompromiss einzugehen: soll die existierende Karte benützt werden, um eine genaue Positionierung zu erreichen, oder sollen neue Gebiete erkundet werden, ohne dabei manuell eingestellte Parameter zu verwenden? Die Formulierung dieser Funktion wird im Rahmen eines GP-basierten Überwachungsszenarios hergeleitet und ist auf verschiedene Lernprobleme anwendbar.

Das entwickelte Framework ist modular aufgebaut und kann auf eine Vielzahl von Problemen des aktiven Erfassens zugeschnitten werden. Umfangreiche Simulationen wurden durchgeführt, um sowohl Ansatz als Ganzes, wie auch seine Bestandteile zu bewerten und mit existierenden zu vergleichen. Die Hauptergebnisse zeigen, dass der vorgeschlagene Ansatz tatsächlich: a) Karten mit ähnlicher Präzision und Genauigkeit in deutlich kürzerer Zeit erstellt; (b) sich auf die adaptive Zuordnung spe-

zifischer Interessensgebiete fokussieren kann; und (c) sowohl die Genauigkeit der Feldkarte als auch die der Roboterlokalisierung durch Berücksichtigung der Orientierungsunsicherheit bei der informativen Planung verbessert. Experimentelle Ergebnisse zeigen die Integration und Validierung in einem fotorealistischen Szenario zur Geländeüberwachung. Schließlich werden Feldtests vorgestellt, um die Algorithmen in Echtzeit, auf laufenden Plattformen mit verschiedenen Aufgaben im landwirtschaftlichen Rahmen zu demonstrieren. Das Framework wird als Open-Source-Paket öffentlich zugänglich gemacht.

Acknowledgements

The journey towards this doctoral thesis was both a marathon and a sprint. Without the support from many people, I could not have reached the finish line.

First of all, I wish to profoundly thank Prof. Dr. Roland Siegwart for the opportunity to conduct my studies at the Autonomous Systems Lab. It has been a great privilege to learn in an environment which brings to life the exciting technologies that are revolutionizing our future. Furthermore, to my co-referee Dr. Teresa Vidal-Calleja: thank you for your useful advice, creative ideas, unwavering enthusiasm, as well as an unforgettable trip to Sydney! It has been an incredible experience being part of a collaboration that spans across continents, and I hope we continue to work together in the future. Finally, I am grateful to Dr. Juan Nieto for his guidance and feedback during these years; for putting up with my questions; and for teaching me how to become a researcher.

There are a number of people that deserve special mention for their contributions towards this thesis. I would like to express my appreciation to Dr. Jen Jen Chung and Dr. Gregory Hitz for sharing great ideas and being excellent mentors. Many pages of this work have their roots in our brainstorming sessions. To the Flourish crew at the Lab, Dr. Inkyu Sa and Dr. Raghav Khanna: thank you for being an awesome project team. It has been a pleasure to share the excitement of flying UAVs with you. To Dr. Frank Liebisch: thank you for helping us give those UAVs purposeful things to do by sharing your crop science with us. Finally, I am indebted to Anurag Sai and Michel Breyer, who are not only great co-workers, but also amazing friends.

Furthermore, I would like to thank Prof. Dr. Kostas Alexis, Dr. Zetao Chen, Dr. Enric Galceran, Dr. Abel Gawel, Dr. Margarida Faria, Dr. Paul Furgale, Dr. Maani Jadidi, Dr. Mina Kamel, Dr. Nicholas Lawrance, Dr. Michela Longhi, Alex Millane, Dr. Helen Oleynikova, Dr. Mark Pfeiffer, Martin Schwendener, Dr. Ulrich Schwesinger, Dr. Joan Solà, and Dr. Zachary Taylor for their assistance and support. Thank you also to Luciana Borsatti, Cornelia Della Casa, Regina Notz, and Michael Riner for making sure our administrative systems are always running.

My gratitude extends to family and friends in places both near and far. From Belgrade, to Abu Dhabi, to Bath, to Zürich, to Sydney, thank you for everyone that has supported me during this journey! To my BodyCombat sisters and the extended family at Silhouette Fitness: thank you for fighting beside me, in more ways than one. To the 'AISA Girls': thank you for the best high school experience ever. To Raphaël: thank you for sharing words of encouragement at times when I needed to hear them the most. To my best friends, Kiana, Manon, and Mariam: thank you for staying closest to me, even though you are the furthest ones away.

Last, but never least, I would like to thank my grandparents, for helping me make the final stretch, as well as my parents, for always believing in my (informative?) path.

Without you all, I wouldn't have conquered.

Zurich, June 3, 2019 *Marija Popović*

Financial Support

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644227 and from the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0029.

Contents

Abstract i

Zusammenfassung iii

Acknowledgements v

List of Figures ix

List of Tables xi

List of Algorithms xiii

Acronyms xv

List of Symbols xvii

1 Introduction 1
- 1.1 Motivation and Objectives . . . 2
  - 1.1.1 Terrain Monitoring . . . 3
  - 1.1.2 Active Sensing Under Localization Uncertainty . . . 4
- 1.2 Approach . . . 4
  - 1.2.1 Terrain Monitoring . . . 5
  - 1.2.2 Active Sensing Under Localization Uncertainty . . . 6
- 1.3 Organization . . . 7

2 Related Work and Contributions 9
- 2.1 Literature Review . . . 9
  - 2.1.1 Environmental Mapping . . . 9
  - 2.1.2 Informative Path Planning . . . 11
  - 2.1.3 Active Sensing Under Localization Uncertainty . . . 12
- 2.2 Contributions . . . 14
  - 2.2.1 Environmental Mapping . . . 14
  - 2.2.2 Informative Path Planning . . . 15
  - 2.2.3 System Integration and Experimental Results . . . 16
  - 2.2.4 Active Sensing Under Localization Uncertainty . . . 16
- 2.3 List of Publications . . . 17
- 2.4 Supervised Student Projects . . . 19

3 Environmental Mapping 21
- 3.1 Discrete Variable Mapping . . . 21
- 3.2 Continuous Variable Mapping . . . 23

3.2.1 Gaussian Processes . . . 24
3.2.2 Sequential Data Fusion . . . 25
3.2.3 Altitude-dependent Sensor Model . . . 26
3.3 Summary and Discussion . . . 27

4 Informative Path Planning **31**
4.1 Problem Setup . . . 31
4.2 Algorithm . . . 32
4.2.1 Asynchronous Sensor . . . 33
4.2.2 Constant-frequency Sensor . . . 35
4.3 Utility Definition . . . 35
4.3.1 Discrete Variable Mapping . . . 36
4.3.2 Continuous Variable Mapping . . . 37
4.4 Grid Search . . . 39
4.5 Optimization . . . 40
4.6 Summary and Discussion . . . 42

5 System Integration and Experimental Results **43**
5.1 System Integration . . . 43
5.2 Simulation Results . . . 46
5.2.1 Discrete Variable Mapping . . . 46
5.2.2 Continuous Variable Mapping . . . 51
5.3 RIT-18 Mapping Scenario . . . 62
5.4 Experimental Results . . . 68
5.4.1 Target Mapping . . . 68
5.4.2 Color Mapping . . . 71
5.4.3 Agricultural Monitoring . . . 74
5.5 Summary and Discussion . . . 79

6 Active Sensing Under Localization Uncertainty **81**
6.1 Problem Setup . . . 81
6.2 Environmental Mapping Under Pose Uncertainty . . . 82
6.3 Informative Planning . . . 83
6.3.1 Utility Definition . . . 84
6.3.2 Uncertainty Prediction . . . 86
6.4 Experimental Results . . . 88
6.4.1 Comparison Against Planning Benchmarks . . . 89
6.4.2 Evaluation of Field Mapping Under Uncertainty . . . 93
6.4.3 Temperature Mapping . . . 95
6.5 Summary and Discussion . . . 98

7 Conclusion and Future Outlook **101**
7.1 Summary of Contributions . . . 101
7.2 Future Outlook . . . 102

Bibliography **107**

Curriculum Vitae **121**

List of Publications **123**

List of Figures

1.1 Active sensing applications 2
1.2 Informative planning system framework 5

3.1 Example sensor model for discrete variable mapping 23
3.2 Example inverse sensor model for continuous variable mapping 27
3.3 Overview of the approach for continuous variable mapping 28

4.1 Schematic explanation of the replanning routine using an asynchronous sensor 34
4.2 Visualizations of 3-D lattices for the grid search in the planning algorithm 39

5.1 Informative planning system framework 44
5.2 Example simulation results of the informative planner for continuous variable mapping . . 45
5.3 Discrete mapping: Comparison of planning strategies 48
5.4 Discrete mapping: Comparison against "lawnmower" coverage benchmark 49
5.5 Discrete mapping: Comparison of utility functions 50
5.6 Discrete mapping: Comparison of CMA-ES optimization strategies 51
5.7 Continuous mapping: Comparison of planning strategies 54
5.8 Continuous mapping: Detailed comparison against "lawnmower" coverage benchmark . . . 56
5.9 Continuous mapping: Comparison of optimization methods - map uncertainty decay rates 59
5.10 Continuous mapping: Evaluation of the adaptive replanning scheme 60
5.11 Continuous mapping: Example result of the adaptive replanning scheme 61
5.12 RIT-18 mapping: Simulation setup in RotorS 63
5.13 RIT-18 mapping: Example image classification result 64
5.14 RIT-18 mapping: Altitude-dependent sensor models 64
5.15 RIT-18 mapping: Comparison against "lawnmower" coverage benchmarks 66
5.16 RIT-18 mapping: Qualitative result of the adaptive informative planner 67
5.17 Target mapping: Experimental setup 69

5.18 Target mapping: System diagram 69
5.19 Target mapping: Quantitative experimental results - information metrics 70
5.20 Target mapping: Qualitative experimental results 71
5.21 Color mapping: Experimental setup 72
5.22 Color mapping: Experimental results 74
5.23 Agricultural monitoring: Experimental setup 75
5.24 Agricultural monitoring: Experimental results - second (main) field trial 77
5.25 Agricultural monitoring: Experimental results - first (preliminary) field trial 78

6.1 Variation of the Rényi term $\log \alpha^{\frac{1}{\alpha-1}}$ 85
6.2 Schematic explanation of the uncertainty propagation scheme 87
6.3 Overview of the approach for active sensing under uncertainty 88
6.4 Active sensing under pose uncertainty: Comparison of planning strategies 92
6.5 Active sensing under pose uncertainty: Example result of the informative planner 94
6.6 Active sensing under pose uncertainty: Evaluation of field mapping under uncertainty . . 95
6.7 Temperature mapping: Experimental setup 96
6.8 Temperature mapping: System diagram 96
6.9 Temperature mapping: Experimental results 97

List of Tables

5.1 Continuous mapping: Comparison of optimization methods - information metrics 58

List of Algorithms

1 REPLAN_PATH procedure - asynchronous sensor 33
2 REPLAN_PATH procedure - constant-frequency sensor 33
3 Time-varying utility function 37
4 CMA-ES optimization routine 41

Acronyms

AMCL	Adaptive Monte Carlo Localization
BO	Bayesian Optimization
CDF	Cumulative Distribution Function
CMA-ES	Covariance-Matrix Adaptation Evolution Strategy
DoF	Degree of Freedom
ExG	Excess Green index
FoV	Field of View
GP	Gaussian Process
GSD	Ground Sample Distance
IP	Interior Point
KF	Kalman Filter
MCL	Monte Carlo Localization
MLL	Mean Log Loss
MPC	Model Predictive Control
NDVI	Normalised Difference Vegetation Index
NLML	Negative Log of the Marginal Likelihood
NP	Non-deterministic Polynomial-time
POMDP	Partially Observable Markov Decision Process
RIG	Rapidly exploring Information Gathering
RMSE	Root Mean Squared Error
ROS	Robot Operating System
ROVIO	Robust Visual Inertial Odometry
SA	Simulated Annealing
SLAM	Simultaneous Localization And Mapping

SDK Software Development Kit

SE Squared Exponential

TOED Theory of Optimal Experimental Design

UI Uncertain Input

UAV Unmanned Aerial Vehicle

WMLL Weighted Mean Log Loss

WRMSE Weighted Root Mean Squared Error

List of Symbols

$+$	Superscript denoting prior
$-$	Superscript denoting posterior
$\bar{\sigma}^2(\cdot)$	Mean variance of a set
β	Confidence interval scaling factor
$\boldsymbol{\kappa}$	Random vector for the motion noise
$\boldsymbol{\mu}$	Mean vector of a multivariate Gaussian
$\boldsymbol{\sigma}_{\text{CMA}}$	Initial step sizes in the CMA-ES
δ_f	Threshold on free state
δ_o	Threshold on occupied state
ϵ	Random variable for the observation noise
$\Gamma[\cdot]$	Gamma function
γ_1, γ_2	Coefficients in an altitude-dependent sensor model
$\hat{\mathbf{u}}$	Control input vector
$\hat{\mathrm{H}}(\cdot)$	A-optimality criterion
λ	Eigenvalue of a covariance matrix
λ_{CMA}	Population size in the CMA-ES
$\log(\cdot)$	Natural logarithm
$\mathbb{E}[\cdot]$	Expected value of a random variable
$\mathbb{R}$	Set of real numbers
$\mathbb{V}[\cdot]$	Variance of a random variable
$\boldsymbol{\Sigma}$	Covariance matrix of a pose estimate
$\mathbf{c}$	Control waypoint
$\mathbf{c}^*$	Next-best control waypoint
$\mathbf{H}$	Measurement model

$\mathbf{I}_n$	Identity matrix of size n
$\mathbf{K}$	Kalman gain
$\mathbf{m}_0$	Initial offspring the CMA-ES
$\mathbf{P}$	Covariance matrix of a multivariate Gaussian
$\mathbf{p}$	Pose in the workspace
$\mathbf{Q}$	Covariance matrix of the motion noise
$\mathbf{R}$	Measurement noise matrix
$\mathbf{S}$	Covariance innovation
$\mathbf{u}$	Noisy control input vector
$\mathbf{v}$	Measurement innovation
$\mathbf{x}$	Point in the environment
$\mathbf{x}_*$	Query point
θ	Set of hyperparameters
$\mathcal{C}$	Set of control waypoints
$\mathcal{C}^g$	Set of global control waypoints
$\mathcal{C}^i$	Set of intermediate control waypoints
$\mathcal{E}$	Environment (input space)
$\mathcal{GP}$	Gaussian process
$\mathcal{L}$	Set of points in the workspace
$\mathcal{O}(\cdot)$	Big O notation
$\mathcal{X}$	Set of possible points in the environment
$\mathcal{X}_*$	Set of query points in the environment
$\mathcal{X}_-$	Set of uninteresting points in the environment
$\mathcal{X}_I$	Set of interesting points in the environment
$\mathcal{X}_V$	Set of visible points in the environment
$\mathcal{Z}$	General environment model
$\mathrm{H}(\cdot)$	Shannon's entropy of a random variable
$\mathrm{H}_\alpha(\cdot)$	Rényi's entropy of order α
$\mathrm{I}(\cdot)$	Utility function
$\mathrm{L}(\cdot)$	Log likelihood function

μ_s	Mean value of the sensor measurement
μ_{th}	Threshold on adaptive planning objective
Ψ	Set of possible trajectories in the workspace
ψ	Trajectory
ψ^*	Optimal trajectory
σ_f^2	Signal variance
σ_n^2	Variance of the observation noise
σ_s^2	Variance of the sensor noise
RAND()	Random number generator in the range [0,1]
$\tilde{k}(\cdot,\cdot)$	Expected/modified covariance function of Gaussian processes
$\mathrm{Tr}(\cdot)$	Trace of a matrix
$\Delta\sigma^2$	Uncertainty difference
B	Budget
b	Normalized blue color channel
C	Convergence threshold in the CMA-ES
D	Number of dimensions of the environment
$f(\cdot)$	Latent function of Gaussian processes
f_*	Predicted value of latent function
g	Normalized green color channel
h	Altitude
H_n	Hermite polynomials
i	Map cell index
k	Order of polynomial trajectory segment
$K(\cdot,\cdot)$	Covariance matrix of Gaussian processes
$k(\cdot,\cdot)$	Covariance function of Gaussian processes
$K_v(\cdot)$	Modified Bessel function of the second kind of order v
l	Characteristic length scale
m	Number of measurements
$m(\cdot)$	Mean function of Gaussian processes
N	Number of control waypoints

n	Number of training points
n_*	Number of query points
N_{ITERS}	Maximum number of iterations in the CMA-ES
$p(\cdot)$	Probability measure
r	Normalized red color channel
s_f	Image resolution scaling factor
t	Time/time index
y	Target variable
z	Measurement
$\text{COST}(\cdot)$	Cost function
$\text{MEASURE}(\cdot)$	Measurement function
$\text{OBJECTIVE}(\cdot)$	Objective (fitness) function in the CMA-ES

1

Introduction

The past several decades have witnessed incredible breakthroughs in the fields of robotic technologies and autonomous systems. The maturation of these disciplines, coupled with rapid advances in related technologies, have enabled robots to steadily advance into our everyday routines. As they become more user-friendly, sophisticated, and affordable, intelligent machines are promising to revolutionize the ways in which we work and live; growing more commonplace in our workplaces, private homes, as well as in society as a whole.

The potential applications for robotic platforms are extensive and expanding at an extraordinary rate. With their rising accessibility, these devices offer time, labor, and cost savings across all sectors of the economy. Their growing autonomous capabilities present significant benefits, especially in tasks that involve mundane or repetitive procedures, demand high levels of accuracy or precision, and transpire in risky or hazardous environments. In particular, one promising application domain is that of exploration and data acquisition (Bajcsy et al., 2018; Bajcsy, 1988 and S. B. Thrun et al., 1992). Here, robots are ripe to replace traditional campaigns based on static sensors, manual sampling, or conventional manned platforms, which can be unreliable, costly, and even dangerous (Dunbabin and Marques, 2012 and Manfreda et al., 2018). The use-cases for safe, cost-efficient solutions include agriculture (Sa, Z. Chen, et al., 2018; Colomina et al., 2014; Liebisch et al., 2017; Vivaldini, Guizilini, et al., 2016; Popović, Hitz, et al., 2017 and Albani et al., 2018), surveillance (Gao et al., 2017; Semsch et al., 2009; Colomina et al., 2014 and Nex et al., 2014), infrastructure inspection (Bircher et al., 2018; Ezequiel et al., 2014 and Papachristos, Khattak, et al., 2017), post-disaster assessment (Ezequiel et al., 2014; Hollinger, S. Singh, et al., 2009 and Anil Meera et al., 2019), and environmental monitoring (Hollinger and Sukhatme, 2014; Dunbabin and Marques, 2012; Manfreda et al., 2018; Detweiler et al., 2015 and Popović, Vidal-Calleja, Hitz, Chung, et al., 2019). In these scenarios, and many more, intelligent systems enable gathering valuable information at unprecedented levels of resolution and autonomy.

However, to fully exploit their potential, many open challenges remain to be addressed. The era of robotics-based monitoring has opened up many exciting directions for research under the theme of *active sensing* (Bajcsy, 1988 and Cao et al., 2013). The main idea is to grant robots a level of cognition similar to that of humans: an intelligent agent can reason and decide on the next sensing actions to maximize the information gathered about a target environment given its current belief state. In doing so, it can trade off between the finite quantity of sensing resources (e.g., energy consumption, mission time, travel distance) and map completeness; thereby allowing it to collect the most useful data given platform-specific constraints. To tackle this problem, this thesis aims to develop strategies for planning efficient informative paths with resource-limited systems in uncertain environments.

(a) (b) (c) (d)

Figure 1.1: Examples of various active sensing applications targeted in this thesis. (a) and (b) illustrate UAV-based data collection in agricultural and atmospheric sampling scenarios, (c) shows a ground robot performing a mine inspection task, and (d) depicts lake monitoring using an autonomous surface vehicle.

1.1 Motivation and Objectives

An *informative path planning* strategy is crucial in order to fully automate and streamline collecting the most useful data in active sensing tasks. Specifically, this work examines a class of missions in which the value of sensor information is unevenly distributed across a target environment. In this setting, the robotic agent is required to maximize the long-term expected reward, or information gain, of acquired measurements subject to a finite quantity of sensing resources. In contrast to conventional, one-shot planning methods, such as goal-based (Richter et al., 2013 and Dijkstra, 1959) or coverage-based approaches (Cabreira et al., 2019; Nex et al., 2014 and Galceran et al., 2013), active strategies permit improving the quality of information gathering in such scenarios while respecting the resource constraints. This field of research, however, is relatively young, with currently available solutions that tend to be limited to specific application domains.

The main objective of this thesis is to create a modular framework for informative path planning that can be tailored to any general active sensing problem, paving the way for robots as efficient data acquisition tools. More specifically, it focuses on *Unmanned Aerial Vehicles (UAVs)* as a target platform due to their recent popularity (Manfreda et al., 2018; Colomina et al., 2014; Carrio et al., 2017 and Nex et al., 2014), as well as the more complex research challenges they present in comparison to ground vehicles by being able to navigate and sense in 3-D environments. Nonetheless, the methods introduced within this work are application-agnostic, and can be used on any mobile device equipped with a suitable sensor suite. Figure 1.1 illustrates several examples of potential use-cases in various domains.

Within this thesis, new solutions are presented in the two core tasks that are key requirements for autonomous data collection in unknown environments: *environmental mapping* and *informative planning*. The main aims are to develop and investigate:

1. planning algorithms for data acquisition in 3-D space given limited computing power and sensing resources;
2. mapping methods for accurate environment representation that capture the model uncertainties; and
3. utility functions coupling the above to achieve predictive planning for active sensing.

The objectives of this thesis are split based on two motivating applications with unique setups. The main use-case investigated is that of *terrain monitoring*. In this scenario, a UAV equipped with a

downward-facing sensor is tasked to survey an environment from different altitudes. This problem sets the context for the majority of this work (corresponding to Chapters 3–5) as well as the building block for the general methods introduced within it. Then, in Chapter 6, the interesting problem of *active sensing under explicit localization uncertainty* is considered to render data collection more accurate and robust. A detailed overview of these applications is provided in the following subsections.

1.1.1 Terrain Monitoring

In many UAV-based monitoring scenarios, including agriculture (Sa, Z. Chen, et al., 2018; Colomina et al., 2014; Liebisch et al., 2017; Vivaldini, Guizilini, et al., 2016 and Albani et al., 2018) search and rescue (Anil Meera et al., 2019; Hollinger, S. Singh, et al., 2009; Colomina et al., 2014 and Ezequiel et al., 2014), wildfire tracking (Cabreira et al., 2019; Manfreda et al., 2018 and Detweiler et al., 2015), and forestry (Vivaldini, Guizilini, et al., 2016; Vivaldini, Martinelli, et al., 2018 and Detweiler et al., 2015), a key task is surveying specific terrain characteristics in a target area of interest. Related fields analogously refer to this application domain as aerial *remote sensing*; Nex et al. (2014) and Colomina et al. (2014) provide excellent up-to-date reviews on this topic from practical perspectives.

To plan missions for aerial data collection in such scenarios, the conventional approach is to solve the coverage path planning problem, which aims at exploring every location in the environment by flying the UAV in a predefined geometric pattern, usually at a fixed altitude (Galceran et al., 2013 and Cabreira et al., 2019). Coverage-based algorithms are considered the state-of-the-art in commercial UAV products, e.g., the DJI series[1], and enable exhaustively mapping terrain at a uniform resolution in a simple manner. However, the main drawback is that they are *passive*. In other words, the limited resources of the platform are not exploited to improve the performance and time efficiency of the data acquisition process. In such settings, an informed, *active* strategy that allows the robot to decide on its next sensing actions has powerful practical implications. For instance, using a UAV to quickly and precisely pinpoint the weeds on a farm can impart valuable data for crop management decisions to reduce pesticide usage and optimize yield (Cardina et al., 1997; Liebisch et al., 2017 and Detweiler et al., 2015).

Developing an intelligent monitoring system in this setting presents several interesting research questions. A key challenge is fusing the visual information received from different altitudes into a compact probabilistic map. Using the current map, the planning unit must search for informative trajectories in the large 3-D space above the monitored area, which poses a complex optimization problem. An important consideration is the relationship between sensor performance and flying altitude. Since the same point in the field can be observed from different heights, a suitable trade-off has to be found between spatial coverage and spatial resolution, while accounting for limited computational and sensing resources.

Within the terrain monitoring application, the objective of this thesis is to develop an informative planning framework for UAVs. This work was initially motivated by the Flourish project[2], which had the goal of developing robotic solutions for precision agriculture, as depicted in Figure 1.1(a). The main idea behind the project was to convey the aerial crop data acquired by a small UAV to a vehicle which then performs targeted autonomous intervention on the ground. However, the developments here are not restricted to this problem setup. The ultimate vision is to create a modular system that can be tailored to any given active sensing task.

[1] `dji.com`
[2] `flourish-project.eu`

1.1.2 Active Sensing Under Localization Uncertainty

Building upon the developments from the previous sub-section, the last part of this thesis generalizes and furthers the active sensing problem. Here, the aim is to map a 2-D or 3-D environmental (field) phenomenon, e.g., temperature (Popović, Vidal-Calleja, Chung, et al., 2019), wind speed (Neumann et al., 2012; Lawrance et al., 2011 and Chung, 2014), signal strength (Binney, Krause, et al., 2013; M. G. Jadidi et al., 2016 and Hollinger and Sukhatme, 2014), light intensity (A. Singh, Ramos, et al., 2010 and Dunbabin and Marques, 2012), etc., using measurements collected by an on-board sensor. In any of these settings, traditional coverage-based algorithms suffer from the same inefficiencies as described above. Their practicality in terms of mapping could be especially limited in 3-D setups based on UAVs due to the high dimensionality of the environment that needs to be explored.

Most existing strategies for active sensing (Hitz, Galceran, et al., 2017; Hollinger and Sukhatme, 2014; A. Singh, Ramos, et al., 2010 and Vivaldini, Martinelli, et al., 2018) incorrectly assume perfect pose information, which is an implicit requirement for creating robust, high-quality reconstructions of an initially unknown environmental field. Consider, for instance, a scenario where the robot odometry drifts significantly. If the robot is continuously directed to explore new areas, rather than improve its localization, this will cause new measurements to be registered at inconsistent locations and is likely to lead to inaccurate field maps.

Despite usually being readily available, estimates of the robot pose uncertainty are often not taken into account to improve upon information gathering behavior. A major issue is the fact that the uncertainty measures of the field map and robot pose represent fundamentally different quantities (Carrillo et al., 2018). As a result, mathematically coupling them in a single *uncertainty-aware* planning objective is a challenging problem that remains an open and actively pursued research question.

This work investigates methods of explicitly accounting for the robot pose uncertainty in informative path planning problems. The main motivation is to improve the accuracy and robustness of field reconstructions by allowing the robot to automatically trade off between gathering new information (*exploration*) and maintaining good localization (*exploitation*). These developments in this sub-section are designed to integrate into the general framework described in Section 1.1.1, such that they can be applied in any field mapping scenario where an estimate of the robot state is available. The following sub-section provides further details on how these two tasks are interconnected.

1.2 Approach

The thesis studies informative planning approaches for active sensing in order to bridge the gap between modern intelligent robots and practical data acquisition applications. Figure 1.2 illustrates a generic structure of the system architecture proposed for such tasks. This diagram represents the overarching approach of this thesis, and is relevant to both of its target applications. For each setting, contributions in *environmental mapping* and *informative planning* are presented that integrate into this framework. The red dashed line designates the additional data flows emerging for the problem of active sensing under explicit localization uncertainty.

A brief overview of the general approach follows. During a mission, the proposed planner uses the environmental field maps built online to generate paths for maximum gain in an information-theoretic measure. This metric represents the value of acquired data and is chosen to reflect the aim of the active sensing mission, as described below. The strategy proceeds in a finite-horizon fashion, alternating between replanning and plan execution until the travel cost reaches a prespecified budget. A key aspect of the approach is its generic formulation; the interfaces and modules can be easily tailored to a particular problem setup.

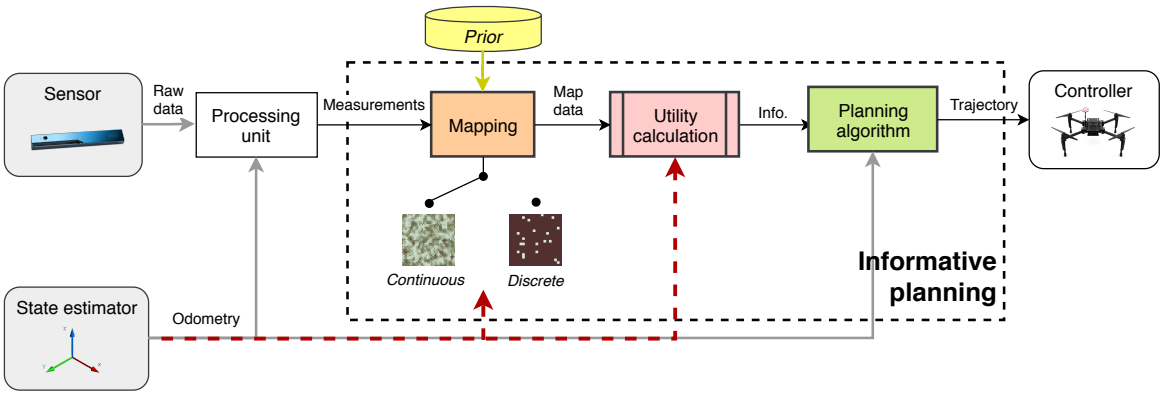

Figure 1.2: System diagram showing the key elements of the proposed informative path planning framework for active sensing and its associated data streams. The red line indicates the propagation of the robot pose uncertainty into the mapping and planning modules for the uncertainty-aware sensing approach introduced in Chapter 6.

1.2.1 Terrain Monitoring

This part of the work corresponds to the majority of the thesis and introduces the general pipeline shown in Figure 1.2. Two key research questions are investigated: (1) how to accurately represent the target environment for informative planning, and (2) how to efficiently plan information-rich paths in the 3-D space above the terrain. Within (1), mapping methods for discrete or continuous variables are addressed to accommodate a variety of practical applications. To link the problems of environmental mapping and predictive planning, the study also encapsulates the development of altitude-dependent sensor models for data acquisition. The following paragraphs summarize the approaches pursued for each of the two aspects.

Environmental Terrain Mapping

Methods of field map representation are investigated for monitoring discrete or continuous target variables on a terrain using sensors whose performance varies with altitude. To monitor a discrete variable, an active classification problem is considered in which the environment is represented by a discretized 2-D occupancy map (Elfes, 1989). To monitor a continuous variable, a more sophisticated method is proposed using Gaussian Processes (GPs) (Rasmussen et al., 2006) to encode spatial correlations common in environmental distributions. Inspired by Vidal-Calleja et al. (2014), the approach leverages a GP to initialize a recursive Bayesian filtering procedure with probabilistic sensors. A key advantage of this strategy is computational efficiency in handling dense visual data; the burden of applying GPs directly, as commonly done in related studies (Hollinger and Sukhatme, 2014; Hitz, Galceran, et al., 2017 and Binney and Sukhatme, 2012), is replaced with constant processing time in the number of measurements.

In Figure 1.2, a discrete or continuous map representation can thus be selected depending on the type of sensor data received in a given monitoring scenario. Among many others, potential use-cases include spatial occupancy (O'Callaghan et al., 2012; Charrow et al., 2015 and Carrillo et al., 2018), target occupancy (Popović, Hitz, et al., 2017; Hollinger, S. Singh, et al., 2009 and Anil Meera et al., 2019), and semantic segmentation (Carrio et al., 2017; Berrio et al., 2017 and Sa, Z. Chen, et al., 2018) for a discrete mapping scenario, and signal strength (Hollinger and Sukhatme, 2014; M. G. Jadidi et al., 2016 and Binney, Krause, et al., 2013), temperature (Ezequiel et al., 2014;

Dunbabin and Marques, 2012 and Colomina et al., 2014), and crop canopy cover (Detweiler et al., 2015; Anthony et al., 2014 and Dunbabin and Marques, 2012) for a continuous mapping scenario.

Online Informative Path Planning

The planning algorithm introduced in this work is general and designed for online, adaptive active sensing applications. To cater for practical robot dynamics, candidate plans are represented as smooth polynomial trajectories in continuous space following the optimization method presented by Richter et al. (2013). As described above, the proposed approach generates fixed-horizon plans to maximize an information-theoretic objective given the current state of the environmental field map.

The replanning strategy consists of two stages. First, an initial trajectory solution is derived through a coarse grid search in the robot workspace above the terrain. This step proceeds sequentially, in a greedy manner, so that a rough, discrete solution can quickly be obtained. Then, the trajectory is refined to maximize a chosen information objective, as discussed below. In contrast to prior work on similar spline-based optimization problems (Hitz, Galceran, et al., 2017 and Morere et al., 2017), a key insight is that the discrete search enables informed initialization to obtain faster convergence. For the optimization step, this work proposes using a generic global evolutionary routine, the Covariance-Matrix Adaptation Evolution Strategy (CMA-ES), studied by Hansen (2006). This decision is evaluated in detail in Section 5.2.2 through a comparison against alternative routines. As a general comment, it is worth mentioning that the replanning procedure is entirely agnostic to the problem scenario.

A core ingredient of the planning algorithm is the *utility*, or *objective function*, which captures the mission-specific interests for data-driven sensing. Following prior studies (Hitz, Galceran, et al., 2017; Sim et al., 2005; Krause, A. Singh, et al., 2008; Lim et al., 2015 and Bourgault et al., 2002), this work presents utility functions for both discrete and continuous variable mapping. The information gain for exploration gain in these scenarios is derived based on the definition of Shannon's entropy.

Moreover, this work considers an *adaptive planning* setup where the objective depends on the actual values of the measurements taken in addition to their locations. Rather than committing to a precomputed path, the *online* nature of the algorithm allows it to adapt plans to new data collected during a mission. Properties of adaptivity are vital in many practical applications in complex, unknown environments, such as finding function extrema (Marchant and Ramos, 2014 and A. Singh, Krause, and Kaiser, 2009), classifying level sets (Hitz, Gotovos, et al., 2014 and Gotovos et al., 2013), and identifying targets (Anil Meera et al., 2019 and Lim et al., 2015). For this problem, a utility function is proposed that exploits the principles of uncertainty-aware classification to create adaptive plans focusing on a certain value range of the monitored variable (Gotovos et al., 2013 and Srinivas et al., 2012). In this manner, the planning aims within the framework can be easily set via the utility function design.

1.2.2 Active Sensing Under Localization Uncertainty

The final contribution of this thesis generalizes and extends the informative planning pipeline described above to active sensing problems in 2-D or 3-D space. The proposed approach for these settings also accounts for the robot pose uncertainty to achieve more robust data acquisition with greater accuracy. As shown by the red dashed line in Figure 1.2, the localization uncertainty, supplied by a state estimation unit, e.g., particle filter or AR tag localization system, is propagated into both the mapping and planning modules to allow for uncertainty-aware sensing in a tightly coupled framework. Further details are introduced in the following.

Mapping Under Uncertain Inputs

The problem setup considers a continuous, spatially correlated, field, e.g., temperature, wind speed, etc. To capture the field, the proposed mapping strategy exploits a GP model given the robot pose uncertainty on its input training points. This uncertainty is folded into the map inference by extending the expected kernel technique of M. G. Jadidi et al. (2016) to planning problems in 3-D setups. The essence of the method lies in taking the expectation of the GP covariance function over the uncertain inputs using a computationally efficient approximation. Thereby, this approach enables maintaining robust maps for planning informative trajectories.

Uncertainty-aware Informative Planning

A fundamental aspect of the extension to planning under uncertainty is a new utility function which simultaneously accounts for both the field mapping and robot localization objectives. This component addresses the exploration-exploitation trade-off for problems in active sensing; the motivation is to allow the robot to gather new sensory data in unknown areas while remaining well-localized with respect to the known map.

Namely, based on the ideas of Carrillo et al. (2018), a planning objective is formulated for GP-based field mapping scenarios using the concept of Rényi's entropy applied for continuous target distributions. The main idea is to discount the potential information gained from future sensor measurements based on the localization uncertainty predicted based on the robot motion. This relationship couples the uncertainties of the field map and robot pose in a unified manner, without any heuristic manual tuning requirements or environment-dependent parameters. Moreover, it lends itself to a variety of continuous field monitoring tasks in which high-quality mapping is essential, as well as to learning problems in general (Whitehead et al., 2008).

1.3 Organization

This thesis is organized in seven chapters and structured as follows. Chapter 2 provides a literature review on the major research fields covered in this thesis and summarizes the key contributions. The remaining chapters then cover in detail the different elements introduced above. Chapter 3 and Chapter 4 describe the techniques proposed to address the problems of environmental mapping and informative planning in active sensing applications, respectively. In Chapter 5, the theoretical frameworks are integrated into a single system and the proposed approach is evaluated extensively in both simulated and physical experiments. The final contribution of this research, addressing the problem of active sensing under localization uncertainty, is presented in Chapter 6. Finally, Chapter 7 closes the thesis with an outlook towards future work and possible research directions.

2

Related Work and Contributions

2.1 Literature Review

Active sensing strategies are the subject of a large and growing body of literature in robotics, computer vision, and related fields. The origins of this paradigm applied on intelligent systems stem from the seminal work of Bajcsy (1988). Since then, advances in sensing and autonomous technologies have unlocked their potential for data acquisition in a wide variety of practical applications, including surveillance, aquatic monitoring, search and rescue, and infrastructure inspection.

The following sub-sections provide an overview of relevant prior studies based on two main research streams: (1) methods for probabilistic environmental mapping, and (2) algorithms for informative planning. Then, the survey is extended to address methods for incorporating the robot localization uncertainty into active sensing decisions.

2.1.1 Environmental Mapping

In problems of autonomous data gathering, methods for environmental field modeling are fundamental to capture the target variable of interest as a basis for informative planning. This section surveys recent work on probabilistic representations based on the two distinct applications examined in this thesis: monitoring (1) discrete-valued, and (2) continuous-valued environmental phenomena. A key assumption behind this work is that the target distribution is static, i.e., does not vary with time. The following review is therefore limited to this scope.

A discrete variable mapping scenario is one in which the sensor data collected is spatially uncorrelated or independent. Example applications include monitoring spatial occupancy (O'Callaghan et al., 2012; Charrow et al., 2015 and Carrillo et al., 2018), target occupancy (Popović, Hitz, et al., 2017; Hollinger, S. Singh, et al., 2009 and Anil Meera et al., 2019), semantic segmentation (Carrio et al., 2017; Berrio et al., 2017 and Sa, Z. Chen, et al., 2018), etc. The occupancy grid framework, introduced by Elfes (1989), is the predominantly used map representation for such types of tasks. In this approach, the environment is discretized into cells, each associated with an independent probability that the corresponding area is occupied by an obstacle. When new measurements are received, a Bayesian procedure is employed to update the occupancy status of the cells based on a probabilistic sensor noise model.

Several recent works (Berrio et al., 2017; Costante et al., 2017; Fankhauser et al., 2016 and Dang et al., 2018) extend these concepts to general segmentation problems by using multi-layer maps. Delmerico et al. (2016) consider a semantically-annotated map representation to classify different types of terrain in an active exploration problem. In this work, a similar algorithm is applied in order

to construct discrete maps in a computationally efficient manner, which can be easily integrated with different types of sensing modalities.

However, many natural and urban phenomena exhibit complex interdependencies where the above assumption of measurement independence does not hold. This problem involves modeling a scalar variable in a continuous domain. Relevant scenarios include monitoring signal strength (Hollinger and Sukhatme, 2014; M. G. Jadidi et al., 2016 and Binney, Krause, et al., 2013), temperature (Ezequiel et al., 2014; Dunbabin and Marques, 2012 and Colomina et al., 2014), crop canopy cover (Detweiler et al., 2015; Anthony et al., 2014 and Dunbabin and Marques, 2012), elevation (Vasudevan et al., 2009), etc.

A popular approach for capturing such relationships is using a Gaussian Process (GP) model, which is a non-parametric supervised learning method designed to solve regression and classification problems (Rasmussen et al., 2006). The regression property of this framework enables predicting the field distribution in an environment given a collection of observations made within it. Historically, GP models are rooted in the field of geostatistics, where this capability was originally exploited to estimate mineral resources through a process known as *kringing* (Matheron, 1973).

In the context of modern machine learning, there are many scenarios where GPs have been similarly applied to model various environmental distributions. One well-known use-case is in the problem of designing a static sensor network, where the aim is to maximize the information about a spatial phenomenon (Krause, A. Singh, et al., 2008). In mobile robotics, they are attracting widespread interest in active sensing applications (Hollinger and Sukhatme, 2014; Hitz, Galceran, et al., 2017; Binney and Sukhatme, 2012; Cao et al., 2013 and K.-C. Ma, Z. Ma, et al., 2018) to collect data accounting for the correlated structure of the field map. More recent work has also considered GP models as probabilistic classifiers for occupancy map building problems (O'Callaghan et al., 2012 and M. G. Jadidi et al., 2017). K. Yang et al. (2014) demonstrate the use of this representation to plan informative paths for an Unmanned Aerial Vehicle (UAV) mapping in continuous 3-D space.

A powerful feature of the GP framework is that it permits using different kernels to express the data dependencies within the underlying target distribution, thus overcoming the assumption of uncorrelated measurements. As an example of its flexibility, A. Singh, Ramos, et al. (2010) present an approach for creating covariance functions that can capture complex spatio-temporal dynamics. Most importantly for planning, the GP model is able to provide formal measures of the predictive uncertainty in the environment, e.g., based on an entropy criterion derived from the model variance distribution (Krause, A. Singh, et al., 2008), which can be used to guide robots to explore more uncertain (unknown) areas.

Although GPs were initially unpopular for online learning applications due to their computational complexity, scaling to large-scale/high-resolution datasets is now more feasible using various methods. One strategy is to perform regression with sparse approximations, e.g., based on inducing inputs or properties of conditional independence (Rasmussen et al., 2006 and Quiñonero-Candela et al., 2007). Recently, Sun, Vidal-Calleja, and Miró (2015) and Sun, Vidal-Calleja, and Miro (2017) introduced submapping techniques which can be applied directly to accelerate the Bayesian data fusion procedure presented in this work. The main idea is to split the original map of the environment into multiple smaller submaps for incorporating new measurements. This mechanism effectively circumvents the need to invert a large covariance matrix that has the same dimensionality as the number of input measurements within the GP. Following a different approach, Vasudevan et al. (2009) propose a multi-resolution GP-based representation with local approximations for modeling large-scale terrain.

This thesis introduces a GP-based environmental field modeling method for informative planning in terrain monitoring setups. In a similar vein as the works discussed above, the GP is used to

handle spatially correlated and uncertain data. Rather than using a standard regression approach to incorporate new sensor data in the model, the proposed strategy exploits a recursive filtering procedure based on the approach of Vidal-Calleja et al. (2014). A key difference compared to previous work is that the proposed mechanism also features a scalable measurement model to accommodate variable-resolution sensors, and is applied in altitude-dependent settings.

2.1.2 Informative Path Planning

Path planning for mobile robots has been studied extensively in the last decades. However, unlike in conventional goal-based planning, where the objective is to find the shortest route between two locations (Dijkstra, 1959), paths in informative planning problems are subject to a resource budget, which limits the path length. In active sensing scenarios, this constraint effectively limits the number of measurements that can be taken, and hence the value of acquirable information, given a platform-specific budget, e.g., time, energy, or distance. Furthermore, in such tasks, the measures used to quantify information are often submodular (Krause and Guestrin, 2011), meaning that the utility of new data is dependent on previous observations. The submodularity property, which arises if the measurement locations exhibit correlations in space (and potentially time), can have the effect of making the objective space highly non-linear and the resulting optimization problem challenging to solve, especially in large, complex environments. Unfortunately, these search problems have been shown to be Non-deterministic Polynomial-time (NP)-hard (A. Singh, Krause, Guestrin, et al., 2009) or even PSPACE-hard (Reif, 1979), depending on the form of the objective function and path representation considered.

Formally, this thesis defines the general *informative path planning problem* as described below. The aim is to find an optimal trajectory ψ^* in the space of all possible continuous trajectories Ψ for maximum gain in some information-theoretic measure:

$$\begin{aligned} \psi^* = \underset{\psi \in \Psi}{\operatorname{argmax}} \, & \mathrm{I}(\textsc{measure}(\psi)) \,, \\ \text{s.t. } & \textsc{cost}(\psi) \leq B \,. \end{aligned} \tag{2.1}$$

The function $\textsc{measure}(\cdot)$ obtains a finite set of measurements along trajectory ψ in an environment, and $\textsc{cost}(\cdot)$ provides the corresponding cost, which cannot exceed a predefined budget B. The operator $\mathrm{I}(\cdot)$ defines the informative objective quantifying the utility of the acquired measurements, and is also referred to as the *utility function*. In Chapter 4, the terms above are discussed more concretely in the context of this work.

A subtle variation of the problem in Equation 2.1 is to consider the maximization of *information gain rate*, as opposed to information gain only:

$$\begin{aligned} \psi^* = \underset{\psi \in \Psi}{\operatorname{argmax}} & \frac{\mathrm{I}(\textsc{measure}(\psi))}{\textsc{cost}(\psi)} \,, \\ \text{s.t. } & \textsc{cost}(\psi) \leq B \,, \end{aligned} \tag{2.2}$$

where the information value of a potential trajectory ψ is divided by its cost. This formulation is relevant for problems that require comparing the value of actions over different time and length scales. Several examples of concrete applications can be found in recent literature (Charrow et al., 2015; K.-C. Ma, Z. Ma, et al., 2018 and Viseras et al., 2019). This thesis addresses both types of planning problems, as elaborated in the following chapters.

In its most general form, the informative planning task corresponds to the problem of sequential decision-making under uncertainty, which can be expressed as a Partially Observable Markov Decision Process (POMDP) (Kaelbling et al., 1998). Recently, substantial progress has been made in

developing near-optimal POMDP algorithms (M. Chen et al., 2016 and Kurniawati et al., 2008) to improve computational efficiency. Unfortunately, despite these efforts, solving such models with large search/action spaces remains an open challenge, motivating simpler, more scalable methods.

The NP-hard sensor placement problem (Krause, A. Singh, et al., 2008) addresses selecting the most informative measurement sites in a static setting. Discrete planning algorithms, e.g., branch and bound (Binney and Sukhatme, 2012), build upon this task by performing combinatorial optimization over a grid. These solvers are typically limited in resolution and scale exponentially with the problem instance. Greedy methods with limited look-ahead (Chekuri et al., 2005 and Hitz, Gotovos, et al., 2014), including frontier-based approaches (Bircher et al., 2018), have been applied to address this issue. However, a major drawback is that they tend to converge to locally suboptimal solutions, since they do not guarantee optimal behavior beyond the horizon length.

Generally, planning strategies in continuous space offer improved scalability by leveraging sampling-based methods (Hollinger and Sukhatme, 2014 and Bircher et al., 2018) or splines (Vivaldini, Guizilini, et al., 2016; Hitz, Galceran, et al., 2017; Charrow et al., 2015 and Morere et al., 2017) directly in the robot workspace or configuration space. This thesis follows the latter class of approaches in defining smooth polynomial trajectories (Richter et al., 2013) that are optimized globally for an information objective. The spline optimization problem setup of this work most closely resembles the ones studied by Hitz, Galceran, et al. (2017) and Morere et al. (2017). However, a key contribution with respect to previous methods is that the proposed scheme uses an informed initialization procedure to obtain faster convergence. Recently, Heng et al. (2015) and Charrow et al. (2015) have proposed similar two-step planning strategies for autonomous exploration in cluttered environments. A distinguishing aspect is that the algorithm presented here does not rely on geometric frontiers. Instead, it finds informative trajectories by performing global optimization within a bounded volume, which makes it less coupled to the internal structure of the environment and therefore more generally applicable.

Finally, informative planning methods can be categorized as being (1) non-adaptive (offline) or (2) adaptive (online). Non-adaptive approaches, e.g., coverage-based methods (Galceran et al., 2013 and Cabreira et al., 2019), explore an environment using a sequence of pre-determined actions. In contrast, adaptive approaches (Hitz, Galceran, et al., 2017; Girdhar et al., 2015 and Lim et al., 2015) allow plans to change as data are collected, so that future sensing actions are conditioned on previously acquired data to capture mission-specific interests. The works of Lim et al. (2015) and Hollinger, Englot, et al. (2013) include an excellent discussion on the properties of adaptivity and its benefits. Recently, Sadat et al. (2015) devised an adaptive coverage-based planner for UAV-based terrain monitoring problems similar to the ones considered in this thesis. Although it exploits the ability to fly at different altitudes, their algorithm assumes discrete viewpoints and does not support probabilistic data acquisition. In contrast, the proposed approach uses uncertain sensor models for mapping and performs incremental finite-horizon replanning in the continuous robot workspace.

2.1.3 Active Sensing Under Localization Uncertainty

This part of the literature review focuses on methods for addressing localization uncertainty in active sensing problems, which are further examined in Chapter 6. The discussion is broken down in terms of the two topics introduced in the previous sub-sections.

The scope of this survey considers GP-based methods for environmental field modeling, as they are the map representation adopted in this thesis. Most works using GPs, including the terrain mapping strategy introduced in Chapter 3, assume that the training data used for prediction is inherently noise-free. Common practice is to opt for the maximum likelihood representation of the dense map.

However, if uncertainties in the robot pose are significant, this can lead to poor mapping performance because new measurements are incorporated at the wrong locations. Further, unreliable maps can mislead predictive planning algorithms and hinder the efficiency of data collection.

Unfortunately, propagating the input uncertainty with a dense map representation is a computationally challenging task. Previous solutions have addressed this issue by using analytical (Girard, 2004) and heteroscedastic approximation methods (Mchutchon et al., 2011; M. G. Jadidi et al., 2016 and Oliveira et al., 2017). To account for localization uncertainty within the proposed informative planning approach, this thesis leverages the expected kernel technique of M. G. Jadidi et al. (2017). The main idea is to integrate over uncertain training data assuming that the query points for GP regression are deterministic. Specifically, this work extends the original approach to more complex planning problems in 3-D environments, and integrates it with a new uncertainty-aware utility function. This coupling enables robust, tractable mapping under pose uncertainty for online sensing applications.

In the area of informative planning, relatively limited research has been invested in active sensing scenarios where robot localization is uncertain. Similar problem settings have been tackled in the contexts of belief-space planning (Bry et al., 2011; Papachristos, Khattak, et al., 2017 and Costante et al., 2017) and active Simultaneous Localization And Mapping (SLAM) (Valencia et al., 2012 and Bourgault et al., 2002), which enable the robot to plan based on the distribution of its own state as well as that of the environment. Active SLAM considers this formulation for occupancy mapping, where the goal is to maintain good localization as an unknown environment is explored. In contrast, and similarly to Papachristos, Khattak, et al. (2017) and Costante et al. (2017), the problems addressed in this thesis treat the processes of map building and robot localization as being decoupled, i.e., involving two separate streams of sensor data. An important distinction here is that the mapping objective is to reconstruct a continuous field that is independent of the environmental features used for localization.

An open question for current research is formulating utility functions to adaptively trade off between gathering new information (exploration) and maintaining good localization (exploitation) in a principled manner. This problem arises due to the different ways in which the environmental field and robot pose are modeled. As a result, extracted measures of the map and pose uncertainty, e.g., based on entropy (Cover et al., 2006), constitute two fundamentally different quantities with numerical scales that cannot be compared directly. This characteristic makes them difficult to combine mathematically in a single utility function that effectively balances between the objectives of field mapping and robot localization. Carrillo et al. (2018) offers an in-depth analysis of scaling issues arising in autonomous exploration tasks where a Gaussian distribution (continuous) is considered to represent the robot pose in an occupancy grid map (discrete), using uncertainty measures defined by Shannon's entropy.

A common solution to the scaling problem is to apply heuristic methods (Valencia et al., 2012), e.g., utility functions based on a weighted linear combination of the map and pose uncertainties (Bourgault et al., 2002). However, these approaches involve parameters that are difficult to tune manually, and often depend on the specifications of the mission scenario. Alternatively, in a modified problem setting, multi-layer strategies can be used to switch between objectives at separate planning stages (Costante et al., 2017; Papachristos, Khattak, et al., 2017 and Heng et al., 2015).

Unlike the methods above, this thesis follows Carrillo et al. (2018) in using Rényi's entropy to discount information gain based on predicted localization uncertainty. The proposed utility function shares the benefit of coupling the map and robot uncertainties in a mathematically sound way, without any manual tuning requirements. The core difference compared to prior work is that the utility formulation is developed for a continuous mapping scenario using a GP field model, instead

of an occupancy grid. Moreover, the planning approach is based on maps that are built under pose uncertainty for more robust data acquisition.

2.2 Contributions

This thesis makes four main contributions:

- **Environmental mapping:** An efficient, variable-resolution mapping method for informative planning using probabilistic sensors in terrain monitoring types of setups. This approach was first published in Popović, Vidal-Calleja, Hitz, Sa, et al. (2017).
- **Informative path planning:** A fixed-horizon algorithm for generating informative, dynamically feasible trajectories in continuous 2-D or 3-D space, based on optimization with an informed initialization procedure. This approach was first published in Popović, Hitz, et al. (2017).
- **System integration and experimental results:** The unification of the above methods into a single informative planning framework for active sensing, its extensive evaluation in simulation experiments, and its real-time application in field deployments. This work was published in Popović, Vidal-Calleja, Hitz, Chung, et al. (2019), with results drawn from Popović, Vidal-Calleja, Hitz, Sa, et al. (2017) and Popović, Hitz, et al. (2017). An open source implementation is provided[1].
- **Active sensing under localization uncertainty:** A generalization and extension of the above informative planning system to active sensing problems that explicitly account for the robot localization uncertainty. The work in this context was published in Popović, Vidal-Calleja, Chung, et al. (2019).

The following four sub-sections provide a compressed summary of the contributions above, which are further discussed in later chapters. Section 2.3 records all publications originating from the doctoral studies, while Section 2.4 overviews the students supervised by the author and their corresponding project names.

2.2.1 Environmental Mapping

This thesis investigates probabilistic methods for environmental map representation for online active sensing problems. The most significant contribution is a new, computationally efficient multi-resolution mapping strategy for informative planning based on GP models. The approach is developed in the context of UAV-based terrain monitoring applications. As such, it is suitable for monitoring continuous, spatially correlated environmental phenomena, e.g., temperature, wind speed, signal strength, light intensity, etc., using probabilistic camera-based sensors to acquire data at potentially different resolutions.

As discussed in Section 2.1.1, in related scenarios, most existing approaches similarly exploit GPs to model the underlying structure and uncertainty of environmental fields to guide predictive planning algorithms (Hitz, Galceran, et al., 2017; Hollinger and Sukhatme, 2014; Vasudevan et al., 2009; Rasmussen et al., 2006 and Marchant and Ramos, 2014). However, the main issue with applying these methods directly for terrain monitoring is the escalating computational load as dense imagery data accumulate over time. Moreover, limited studies have addressed data fusion for mapping using

[1] `github.com/ethz-asl/tmplanner`

input images at varying resolutions and noise levels. This is a critical requirement in setups featuring sensors whose performance depends on altitude, i.e., higher altitudes enable larger area coverage but yield less reliable measurements.

The method introduced in this thesis addresses these issues by using a GP model as a prior for recursive Bayesian data fusion with probabilistic sensors. Inspired by the work of Vidal-Calleja et al. (2014), this strategy enables online mapping in a more computationally efficient manner. By operating on a fixed map size, the cubic complexity of GP regression is replaced with constant processing time in the number of measurements. To cater for different sensors with altitude-dependent behavior, a key insight is that the measurement model for data acquisition is manipulated to permit fusing heterogeneous, variable-resolution information into a single environmental field map. As in standard GP-based mapping, this approach enables capturing the uncertainty and completeness of the field map in a statistically sound manner.

The proposed mapping approach was extensively evaluated in simulation and successfully implemented on UAV platforms. These experimental results are presented and discussed in detail in Chapter 5.

2.2.2 Informative Path Planning

This thesis develops an online adaptive planning algorithm to obtain information-rich paths subject to platform-specific resource constraints. The contributed approach has several desirable properties; namely, it:

1. generates dynamically feasible trajectories in continuous 2-D or 3-D space;
2. satisfies sensing resource limitations;
3. uses sufficiently low computing power for online and real-time application; and
4. can provide adaptive, any-time solutions, which are useful in real-world scenarios.

The presented strategy combines relevant ideas from several different streams of literature. The main concepts were introduced in Section 2.1.2. To satisfy requirements for dynamic smoothness, candidate plans are represented as continuous polynomial trajectories following the optimization procedure of Richter et al. (2013). During a mission, the proposed algorithm generates trajectories for maximum information gain in a finite-horizon manner, alternating between replanning and plan execution until a prespecified budget of sensing resources is exhausted.

The replanning procedure consists of two main steps. Following a discrete planning paradigm (Binney and Sukhatme, 2012; Lim et al., 2015 and Chekuri et al., 2005), the first step executes a sequential, greedy search over a predefined grid in the robot workspace. This enables quickly obtaining an initial solution to the informative planning problem. Then, in the second step, the Covariance-Matrix Adaptation Evolution Strategy (CMA-ES) technique (Hansen et al., 2009) is used to globally optimize the trajectory for a specified information objective. The principal advantage of this mechanism over existing spline-based optimization approaches (Hitz, Galceran, et al., 2017; Vivaldini, Martinelli, et al., 2018 and Morere et al., 2017) is that it uses informed initialization from the grid search to improve the solution convergence given limited computational power, while relying on only a few parameters that must be manually tuned. Moreover, its computational requirements can be easily adapted to the available hardware resources. For example, by increasing the discretization of the initial grid search or reducing the number of optimizer iterations, one can trade off solution accuracy for algorithmic efficiency.

The proposed informative planning algorithm and its internal components were extensively evaluated in various simulated and real-world settings. These experimental findings are presented and discussed in detail in Chapter 5.

2.2.3 System Integration and Experimental Results

This thesis then unifies the preceding theoretical contributions in mapping and planning into a single framework. The key contributions in this context are:

1. the development of a informative planning architecture for general active sensing problems;
2. its comprehensive evaluation in simulated experiments, including a photorealistic scenario using a publicly available dataset; and
3. its online application on robotic systems under real-world conditions and with real-time requirements.

An implementation of the proposed approach is released as an open-source software package for usage and further development by the community. This package also includes supporting documentation and example experimental results. By addressing the lack of accessible, modular tools for informative planning within current work, this development intends to bridge the gap between academic research and real-world practices.

Applications of the approach are demonstrated in a series of experiments focusing on terrain monitoring scenarios. First, extensive results from simulations are presented to evaluate both the key internal elements of the framework (mapping and planning), as well as how it behaves as an integrated system. Benchmarking experiments illustrate that the proposed approach outperforms existing strategies, and enables mapping specific areas of interest with greater efficiency. The studies reveal that informative framework achieves $> 50\%$ time savings for $> 500\,\mathrm{s}$ missions when compared against a traditional coverage-based strategy executing a "lawnmower" path (Section 5.2.2). Second, real-time and integration capabilities are validated on a photorealistic mapping scenario using a publicly available dataset. Finally, field trials are conducted to show the algorithms implemented on physical UAV platforms for various monitoring tasks.

2.2.4 Active Sensing Under Localization Uncertainty

Finally, this thesis extends the proposed framework by introducing methods to account for the uncertainty in the robot localization. The problem setup considers a general active sensing scenario in which the aim is to map a continuous 2-D or 3-D field, e.g., of temperature, humidity, pressure, etc., using point measurements taken by an on-board sensor. The robot pose estimate is assumed to be provided by a probabilistic localization or SLAM back end system, e.g., particle filter, Kalman Filter (KF), graph SLAM system, etc.

Most existing strategies for active sensing assume perfect information about the robot pose, which is an implicit requirement for accurately reconstructing an environmental field of interest. Despite often being readily available, state estimates are often disregarded towards improving the quality of acquired data. As discussed in Section 2.1.3, a major challenge is coupling the robot localization and field map uncertainties in a single utility function since they are based on fundamentally different representations. This discrepancy also arises in the type of scenario examined in this work. Namely, the continuous field is modeled as the realization of a GP, whereas the robot pose is treated as a multivariate Gaussian distribution. Measures of uncertainty in the two models, e.g., based on entropy or mutual information criteria (Krause, A. Singh, et al., 2008), are not directly comparable due to the fact that they are obtained from their respective variances, which have different units and scales.

A common solution to this issue is to design heuristic utility formulations that depend on adjusting manual parameter values (Bourgault et al., 2002 and Valencia et al., 2012). The motivation behind this work is to circumvent such procedures by introducing a new approach for active sensing under robot localization uncertainty. In the scenario described above, the proposed approach accounts for

the pose uncertainty in two places: (1) it is propagated into the environmental field model to account for the additional noise it induces, and (2) it is included as a shaping factor in a utility function that defines the informative planning task. For (1), GP models with Uncertain Inputs (UIs) are used to map a continuous field under pose uncertainty. To this end, this work extends the methods introduced by M. G. Jadidi et al. (2017) to planning problems in 3-D settings. For (2), a new utility function is presented that jointly considers the uncertainty in the field map and robot pose in a principled manner, without relying on any manually-tuned, task-specific parameters. This enables the robot to adaptively trade off between exploring new areas and exploiting prior information in a variety of environments, as required. The key contribution of this work is that the formulation is developed for the problem of mapping a continuous field using a GP-based model.

To demonstrate system integration, the proposed methods are consolidated with the previously introduced framework for informative planning. The entire approach is evaluated extensively in simulation scenarios featuring a 3-D graph SLAM back end. The experimental results indicate that the proposed approach achieves reductions of up to 45.1% and 6.3% in mean pose uncertainty and total map error compared against existing methods. Finally, a proof of concept deployment is presented on a ground robot performing a temperature mapping task.

2.3 List of Publications

This section provides a complete list of publications, in reverse chronological order, that were published during the doctoral studies.

Publications Included in This Thesis

The following publications are related to the main contributions of this thesis. The results of these publications are included in the document at hand.

- Popović, M., Vidal-Calleja, T., Chung, J. J., Nieto, J., and Siegwart, R. (2019). "Informative Path Planning and Mapping for Active Sensing Under Localization Uncertainty". *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Under review. arXiv: 1902.09660
- Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660
- Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1382–1388. DOI: 10.1109/IROS.2017.8202317
- Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., and Galceran, E. (2017). "Online Informative Path Planning for Active Classification Using UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5753–5758. DOI: 10.1109/ICRA.2017.7989676

Related Publications

The results of this thesis directly contributed to the following related work. However, no content is included here.

- Blum, H., Rohrbach, S., Popović, M., Bartolomei, L., and Siegwart, R. (2019). "Active Learning for UAV-based Semantic Mapping". In: *2nd Workshop on Informative Path Planning and Active Sampling, RSS'19*. MIT Press
- Anil Meera, A., Popović, M., Millane, A., and Siegwart, R. (2019). "Obstacle-aware Adaptive Informative Path Planning for UAV-based Target Search". In: *IEEE International Conference on Robotics and Automation*. Accepted. IEEE
- Causa, F., Popović, M., Fasano, G., Grassi, M., Nieto, J., and Siegwart, R. (2019). "Navigation aware planning for tandem UAV missions in GNSS challenging environments". In: *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics. DOI: 10.2514/6.2019-0916
- Papachristos, C., Kamel, M., Popović, M., Khattak, S., Bircher, A., Oleynikova, H., Dang, T., Mascarich, F., Alexis, K., and Siegwart, R. (2019). "Autonomous Exploration and Inspection Path Planning for Aerial Robots Using the Robot Operating System". In: *Robot Operating System (ROS): The Complete Reference (Volume 3)*. Ed. by A. Koubaa. Springer, pp. 67–111. DOI: 10.1007/978-3-319-91590-6_3

Other Publications

The author also contributed to the following publications, unrelated to the work presented in this thesis.

- Faria, M., Marín, R., Popović, M., Maza, I., and Viguria, A. (2019). "Efficient Lazy Theta* Path Planning over a Sparse Grid to Explore Large 3D Volumes with a Multirotor UAV". *Sensors* 19.1. DOI: 10.3390/s19010174
- Bähnemann, R., Pantic, M., Popović, M., Schindler, D., Tranzatto, M., Kamel, M., Grimm, M., Widauer, J., Siegwart, R., and Nieto, J. (2019). "The ETH-MAV Team in the MBZ International Robotics Challenge". *Journal of Field Robotics* 36.1, pp. 78–103. DOI: 10.1002/rob.21824
- Sa, I., Kamel, M., Burri, M., Bloesch, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2018). "Build Your Own Visual-Inertial Drone: A Cost-Effective and Open-Source Autonomous Drone". *IEEE Robotics & Automation Magazine* 25.1, pp. 89–103. DOI: 10.1109/MRA.2017.2771326
- Longhi, M., Taylor, Z., Popović, M., Nieto, J., Marrocco, G., and Siegwart, R. (2018). "RFID-Based Localization for Greenhouses Monitoring Using MAVs". In: *IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communications*. IEEE, pp. 905–908. DOI: 10.1109/APWC.2018.8503764
- Sa, I., Popović, M., Khanna, R., Chen, Z., Lottes, P., Liebisch, F., Nieto, J., Stachniss, C., Walter, A., and Siegwart, R. (2018). "WeedMap: A Large-Scale Semantic Weed Mapping Framework Using Aerial Multispectral Imaging and Deep Neural Network for Precision Farming". *Remote Sensing* 10.9. DOI: 10.3390/rs10091423

- Miki, T., Popović, M., Gawel, A., Hitz, G., and Siegwart, R. (2018). "Multi-Agent Time-Based Decision-Making for the Search and Action Problem". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2365–2372. DOI: 10.1109/ICRA.2018.8460996
- Sa, I., Chen, Z., Popović, M., Khanna, R., Liebisch, F., Nieto, J., and Siegwart, R. (2018). "weedNet: Dense Semantic Weed Classification Using Multispectral Images and MAV for Smart Farming". In: *IEEE Robotics and Automation Letters*. IEEE, pp. 588–595. DOI: 10.1109/LRA.2017.2774979
- Sa, I., Kamel, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2017). "Dynamic System Identification, and Control for a cost-effective and open-source Multi-rotor MAV". in: *Field and Service Robotics*. Springer, pp. 605–620. DOI: 10.1007/978-3-319-67361-5_39
- Vetrella, A. R., Sa, I., Popović, M., Khanna, R., Nieto, J., Fasano, G., Accardo, D., and Siegwart, R. (2017). "Improved Tau-Guidance and Vision-Aided Navigation for Robust Autonomous Landing of UAVs". In: *Field and Service Robotics*. Springer, pp. 115–128. DOI: 10.1007/978-3-319-67361-5_8
- Liebisch, F., Popović, M., Pfeifer, J., Khanna, R., Lottes, P., Pretto, A., Sa, I., Nieto, J., Siegwart, R., and Walter, A. (2017). "Automatic UAV-based field inspection campaigns for weeding in row crops". In: *EARSeL SIG Imaging Spectroscopy Workshop*
- Pfeifer, J., Khanna, R., Constantin, D., Popović, M., Galceran, E., Kirchgessner, N., Achim, W., Siegwart, R., and Liebisch, F. (2016). "Towards automatic UAV data interpretation for precision farming". In: *International Conference on Agricultural Engineering*

2.4 Supervised Student Projects

This section provides a list of all student projects supervised by the author as part of the doctoral studies. Citations are provided to works where the outcome resulted in a scientific publication.

Master Thesis

Six months full time research project

- Silvan Rohrbach, 2019: "Active Deep Learning with UAVs" (Blum et al., 2019)
- Salomé Schärer 2019: "Learning to Predict Wind Flows for Fixed-wing UAVs"
- Giuseppe Rizzi, 2019: "Cooperative Multi-agent Decision-making Using Distributed Databases"
- Federico Proni, 2018: "Autonomous Path Planning for 3-D Exploration"
- Ajith Anil Meera, 2018: "Informative Path Planning for Search and Rescue Using UAV" (Anil Meera et al., 2019)
- Takahiro Miki, 2017: "Planning for Exploration" (Miki et al., 2018)

Semester Thesis

Semester-long part time research project during the master studies

- Fabian Stelling, 2018: "Towards Active Object Classification with PointNet"

Bachelor Thesis

Semester-long part time research project

- Omar Elkhatib, 2017: "Control Allocation of a Tilting Rotor Hexacopter"

Course Projects

Semester-long part time research project during the master studies

- Yuwen Chen and Pascal Schoppmann (Course: *Perception and Learning for Robotics*), 2018: "Planning for Data"

3

Environmental Mapping

This chapter develops environmental representation methods for active sensing applications using probabilistic sensors for data acquisition. The proposed approach focuses on, but is not limited to, terrain monitoring setups featuring sensors whose performance depends on altitude. Material from the following first author publications is incorporated:

Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1382–1388. DOI: 10.1109/IROS.2017.8202317

Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660

Popović, M., Vidal-Calleja, T., Chung, J. J., Nieto, J., and Siegwart, R. (2019). "Informative Path Planning and Mapping for Active Sensing Under Localization Uncertainty". *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Under review. arXiv: 1902.09660

In line with the literature review in Section 2.1.1, the following sections present methods for mapping (1) discrete-valued and (2) continuous-valued target variables in the context of the terrain monitoring problem. The main intention is that, in an active sensing scenario, one of these two representations can be chosen to accommodate the requirements of a particular mapping task. To illustrate this, Chapter 5 presents the integration of the proposed methods into a unified framework for informative planning, followed by evaluation experiments.

3.1 Discrete Variable Mapping

The task of monitoring a discrete-valued variable is studied as an active classification problem with unique labels. The terrain environment $\mathcal{E} \subset \mathbb{R}^2$ within which measurements are taken is discretized and represented using a 2-D occupancy grid $\mathcal{X}$ (Elfes, 1989). The grid map consists of a tesselation of n regularly spaced cells $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, where $\mathbf{x}_i$ denotes a vector of (x, y) spatial co-ordinates corresponding to the location of a particular cell. Each cell is associated with an independent Bernoulli random variable indicating the probability $p(\mathbf{x}_i)$ of being occupied by a target, such that $p(\mathbf{x}_i) = 0$ if the cell is target-free or $p(\mathbf{x}_i) = 1$ if it is occupied. In the types of problems examined in this work, the occupancy target could be represented by, e.g., the presence of weed infestations on a farmland or disaster victims on a debris field.

For simplicity, a binary classification problem is first considered in which the grid representation consists of a single layer. Extensions to multi-label problems are discussed further below. The

occupancy mapping algorithm makes two standard assumptions: (1) the cells are independent of one another, so that the probability of a particular map is equivalent to the product of the individual marginal probabilities of its cells $p(\mathcal{X}) = \prod_i p(\mathbf{x}_i)$, and (2) the Unmanned Aerial Vehicle (UAV) is capable of estimating its own pose and the pose of its sensor in its local frame. For generality, the UAV pose $\mathbf{p}$ is determined by its complete six Degrees of Freedom (DoFs) as specified by its (x, y, z) position and roll, pitch, yaw angles in the 3-D workspace above the terrain. Note that it is possible to relax assumption (2) by using modern Simultaneous Localization And Mapping (SLAM) solutions to model the uncertainty in the pose as well as the map (Cadena et al., 2016).

In this setup, measurements are taken with a downward-looking sensor mounted on the UAV. The sensor provides inputs to a data processing unit to classify them. At mission time t, for each cell $\mathbf{x}_i$, a binary classification label $z_{i,t} \in \{1, 0\}$ is obtained and used to update the map. Note that extracting the measurement labels typically involves downsampling the raw dense sensor data, e.g., images or point clouds, in order to perform data fusion at the same resolution as defined by the grid. For the sake of brevity, the remainder of this section drops the cell index i, i.e., a cell $\mathbf{x}_i = \mathbf{x}$ is associated with a measurement label $z_{i,t} = z_t$.

The objective of the mapping algorithm is to estimate the posterior of each grid cell $\mathbf{x}$ given the data $p(\mathbf{x} \,|\, z_{1:t}, \mathbf{p}_{1:t})$, where $z_{1:t}$ is the history of the measurement labels received from time 1 to t, and $\mathbf{p}_{1:t}$ is the corresponding UAV trajectory. The posterior can be estimated according to (S. Thrun et al., 2006):

$$p(\mathbf{x} \,|\, z_{1:t}, \mathbf{p}_{1:t}) = \left(1 + \frac{1 - p(\mathbf{x} \,|\, z_t, \mathbf{p}_t)}{p(\mathbf{x} \,|\, z_t, \mathbf{p}_t)} \frac{1 - p(\mathbf{x} \,|\, z_{1:t-1}, \mathbf{p}_{1:t-1})}{p(\mathbf{x} \,|\, z_{1:t-1}, \mathbf{p}_{1:t-1})} \frac{p(\mathbf{x})}{1 - p(\mathbf{x})}\right), \tag{3.1}$$

where the term $p(\mathbf{x} \,|\, z_t, \mathbf{p}_t)$ corresponds to the probability of the cell being occupied given the latest sensor measurement label, and all other measurements $p(\mathbf{x} \,|\, z_{1:t-1}, \mathbf{p}_{1:t-1})$ from time 1 to $t-1$. The term $p(\mathbf{x})$ represents the prior probability of occupancy, which is uniformly set to 0.5 for mapping in an initially unknown environment.

Let $\mathcal{X}_V \subset \mathcal{X}$ denote the subset of grid cells within the perception field from a UAV pose $\mathbf{p}_t$ when the sensor is triggered. A binary Bayesian filter is used to update the observed cells $\mathbf{x} \in \mathcal{X}_V$ with the new measurement labels from the classifier. To perform this step efficiently, the log-odds representation is commonly used to express occupancy status in terms of probability. It is defined by:

$$\mathrm{L}(\mathbf{x}) = \log\left(\frac{p(\mathbf{x})}{1 - p(\mathbf{x})}\right). \tag{3.2}$$

Then, for each observed cell $\mathbf{x}$ receiving a corresponding measurement label z_t, a log likelihood update is performed:

$$\mathrm{L}(\mathbf{x} \,|\, z_{1:t}, \mathbf{p}_{1:t}) = \mathrm{L}(\mathbf{x} \,|\, z_{1:t-1}, \mathbf{p}_{1:t-1}) + \mathrm{L}(\mathbf{x} \,|\, z_t, \mathbf{p}_t) - \mathrm{L}(\mathbf{x}), \tag{3.3}$$

where the term $\mathrm{L}(\mathbf{x})$ is calculated from the map prior probability $p(\mathbf{x})$, and drops out in the case of an unknown environment. $\mathrm{L}(\mathbf{x}_i \,|\, z_t, \mathbf{p}_t)$ denotes the inverse sensor model capturing the measurement output for the observed cell $\mathbf{x}$. Note that the term *sensor model* in this context refers to the accuracy in the *classifier* output, which depends on the raw data it is provided as input. To develop this model in a terrain monitoring scenario, this work considers sensors whose performance is related to the UAV altitude h above the environment.

As an illustrative example, Figure 3.1 shows the sensor model for a hypothetical camera-based binary classifier labeling observed cells as '1' (occupied by target) or '0' (target-free). For each of the two classes, probability curves are used to define the confidence of making a true positive classification $p(z = 1 \,|\, \mathbf{x} = 1, \mathbf{p})$ (blue) and a false positive classification $p(z = 1 \,|\, \mathbf{x} = 0, \mathbf{p})$ (orange), with respect

to the true state of a map cell $\mathbf{x}$. Note that, whereas Figure 3.1 considers a *forward* model to characterize the performance of the sensor, the *inverse* model for the map update in Equation 3.3 can be easily derived by a straightforward application of Bayes' rule.

The way in which the curves vary with the UAV altitude h accounts for the fact that the performance of the classifier is poorer with measurement labels derived from images that have lower resolutions. In other words, at higher altitudes, the UAV is able to observe a larger area of the terrain (more cells in $\mathcal{X}_V$), but the level of certainty in the corresponding classifications is reduced (the probability $p(z \,|\, \mathbf{x}, \mathbf{p})$ approaches 0.5). In the extreme case, at $h > 30\,\text{m}$, the curves flatten out here assuming no reliable observations can be made. As a side comment, it is worth mentioning that *low* probability values for the orange curve $p(z = 1 \,|\, \mathbf{x} = 0, \mathbf{p})$ designate *better* performance, as they represent the chances of making an *incorrect* false positive classification.

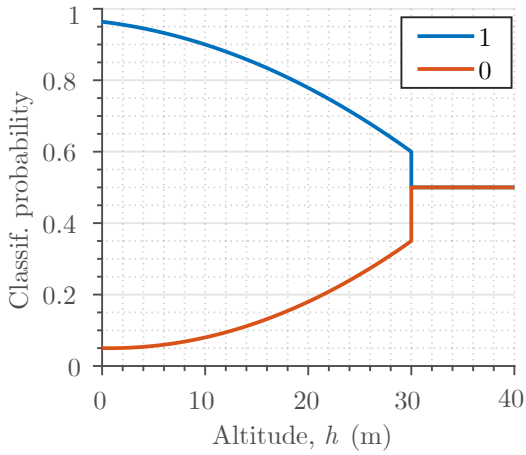

Figure 3.1: Sensor model for a typical camera-based binary classifier operating above a terrain. The blue and orange curves depict the probabilities of observing a measurement label '1' for a map cell containing '1' or '0', respectively, i.e., $p(z = 1 \,|\, \mathbf{x} = 1, \mathbf{p})$ and $p(z = 1 \,|\, \mathbf{x} = 0, \mathbf{p})$, respectively. As the altitude h increases, both curves approach a classification probability of 0.5 in order to reflect poorer performance.

From a practical point of view, the sensor model in Figure 3.1 can be developed in a data-driven manner, as commonly done in conventional occupancy grid mapping (Elfes, 1989). In the active classification scenario, the modeling strategy involves a Monte Carlo-type accuracy analysis of raw outputs obtained from the classifier. Using this approach, the number of true and false positive classifications are recorded at various altitudes, e.g., via confusion matrices, in order to develop an empirical predictive model of how the sensor unit behaves. Section 5.3 demonstrates applying these principles to design the model of a neural network-based classifier. Further, in a follow-up study to this work, Anil Meera et al. (2019) follow a similar method to capture the altitude-dependency of an on-board human detector in a UAV-based search and rescue problem.

The described approach can be easily extended to mapping multiple class labels by maintaining layers of occupancy grids for each one. In Section 5.3, this approach is used to map a publicly available dataset for semantic segmentation, and a more detailed discussion can be found there.

3.2 Continuous Variable Mapping

To monitor a target variable with a continuous domain, this sub-section introduces a more sophisticated mapping method using Gaussian Processes (GPs) to encode the spatial correlation structure common in environmental phenomena. The proposed strategy is applicable for the terrain monitoring type of scenario, where distinct sub-regions within a 2-D area of interest are observed as sensory inputs. In this setup, a GP model is used to initialize a recursive filtering procedure with probabilistic sensors at potentially different resolutions. As outlined in Section 2.2.1, this scheme effectively

replaces the computational burden of applying GPs directly with constant processing time in the number of measurements. The following sub-sections describe the method for creating prior maps before detailing the Bayesian data fusion technique.

3.2.1 Gaussian Processes

In this work, GPs are used to model continuous-valued target distributions in a probabilistic and non-parametric manner. A brief introduction to this framework is presented in Section 2.1.1 in the context of related work. This sub-section overviews the key theoretical ideas relevant to this thesis. The following discussion is formulated in a general way, in order to be applicable for both the 2-D terrain monitoring setup as well as the 3-D active sensing problems studied in Chapter 6, since they both address the problem of reconstructing a continuous environmental field. For a more comprehensive treatment of GPs, the reader is directed to the book of Rasmussen et al. (2006) as an authoritative reference.

A GP is a machine learning technique suitable for modeling spatially correlated and stochastic data in a statistically sound way. It is a non-parametric approach in the sense that an explicit functional relationship between the input and output is not specified. Instead, statistical inference is employed to learn the dependencies between different points in a dataset. The target data is assumed to be a multivariate Gaussian distribution, with one dimension for each input variable.

In this work, the target field variable for mapping is assumed to be a continuous scalar (1-D) latent function: $f : \mathcal{E} \rightarrow \mathbb{R}$, where $\mathcal{E} \subset \mathbb{R}^D$ is the D-dimensional input space, i.e., environment within which measurements are taken. Using the GP, a Gaussian correlated prior is placed over the function space, which is fully characterized by a mean function $m(\mathbf{x}) \triangleq \mathbb{E}[f(\mathbf{x})]$ and a covariance function $k(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}')]$, where $\mathbb{E}[\cdot]$ denotes the expectation operator. $\mathbf{x}$ and $\mathbf{x}'$ are both input points, which are D-dimensional vectors composed of the spatial co-ordinates corresponding to a location in the environment $\mathcal{E}$. Note that $D = 2$ in the terrain monitoring problem, as the surveyed area is on the ground, whereas $D = 3$ in the general active sensing scenario, as the field also extends in the vertical direction.

The GP is therefore written as:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}),\, k(\mathbf{x}, \mathbf{x}'))\,. \tag{3.4}$$

Let $\{(\mathbf{x}_i,\, y_i)\} \,|\, i = 1, \ldots, n\}$ be a training dataset of n observations, where $\mathbf{x}_i$ is a D-dimensional input training vector and y_i denotes the corresponding value of the scalar target (dependent variable), i.e., a measurement registered by the sensor at that location. The input points are aggregated in the $D \times n$ matrix $\mathbf{X}$, and the associated target values are used to construct the column vector $\mathbf{y}$, such that the training dataset can be also written as $(\mathbf{X},\, \mathbf{y})$. The measurements are drawn from a noisy process:

$$y_i = f(\mathbf{x}_i) + \epsilon\,, \tag{3.5}$$

where ϵ denotes the observation noise, which is assumed to be independent and identically Gaussian distributed according to $\mathcal{N}(0,\, \sigma_n^2)$.

Consider a set of n_* query (test) points $\mathbf{x}_*$ arbitrarily specified within the environment, and collected in the $D \times n_*$ query matrix $\mathbf{X}_*$. The aim is to evaluate the conditional Gaussian distribution $\mathbf{f}_* \,|\, \mathbf{X}, \mathbf{y}, \mathbf{X}_* \sim \mathcal{N}(\mathbb{E}[\mathbf{f}_*],\, \mathbb{V}[\mathbf{f}_*])$, where $\mathbb{V}[\cdot]$ is the variance operator and $\mathbf{f}_* \triangleq f(\mathbf{X}_*)$ represents the latent function values corresponding to the query points in $\mathbf{X}_*$. The conditional can be found by

applying the classic predictive equations for GP regression, which are defined as follows (Rasmussen et al., 2006):

$$\boldsymbol{\mu} = \mathbb{E}[\mathbf{f}_*] = m(\mathbf{X}_*) + K(\mathbf{X}_*, \mathbf{X}) \times [K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n]^{-1} \times (\mathbf{y} - m(\mathbf{X}))\,, \tag{3.6}$$

$$\mathbf{P} = \mathbb{V}[\mathbf{f}_*] = K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X}) \times [K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n]^{-1} \times K(\mathbf{X}_*, \mathbf{X})^\top\,, \tag{3.7}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix and σ_n^2 is a hyperparameter representing the observation noise variance. The $n_* \times n$ matrix $K(\mathbf{X}_*, \mathbf{X})$ denotes the cross-correlation terms between the query and observed points, and $K(\mathbf{X}, \mathbf{X})$ and $K(\mathbf{X}_*, \mathbf{X}_*)$ are the $n \times n$ and the $n_* \times n_*$ joint covariance matrices for the observed input and query points, respectively. Note that $K(\cdot,\cdot)$ corresponds to $k(\cdot,\cdot)$ for only one element. The variables $\boldsymbol{\mu}$ and $\mathbf{P}$ are introduced above to represent the GP mean and covariance for purposes of convenience.

The covariance, or kernel, function $k(\cdot,\cdot)$ determines the generalization properties of the GP model and is used to evaluate dependencies between measurements. It is chosen to describe the characteristics of the latent function f with respect to a particular application. To describe environmental phenomena, there exist a number of well-known covariance functions common in geostatistical analysis and spatial statistics, e.g., the Squared Exponential (SE), the rational quadratic, or the Matérn family of functions. The free parameters of the covariance function, called hyperparameters θ, control the relations between data points within the GP. The hyperparameter values are optimized to match the properties of f by training on multiple maps obtained previously at the required resolution. A common method is by minimizing the Negative Log of the Marginal Likelihood (NLML) function:

$$\log p(\mathbf{y} \,|\, \mathbf{X}, \theta) = -\frac{1}{2}\mathbf{y}^\top (K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n)^{-1}\mathbf{y} - \frac{1}{2}\log\left(K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n\right) - \frac{n}{2}\log 2\pi\,. \tag{3.8}$$

A detailed derivation and discussion is provided by Rasmussen et al. (2006).

In the active sensing setup studied in Chapter 6, Equations 3.6 and 3.7 are applied to infer the distribution of the environmental field at the query points in $\mathbf{X}_*$. The computation is performed repeatedly, for online mapping and planning, as measurements are accumulated in the training dataset $\{(\mathbf{x}_i,\, y_i) \,|\, i = 1,\, \ldots,\, n\}$ during a mission. This corresponds to the way in which GPs are applied to solve conventional regression problems.

In the terrain monitoring setup, however, the learned hyperparameters and a correlated prior field map $p(\mathbf{f} \,|\, \mathbf{X})$ are leveraged to recursively fuse new sensor data as they arrive. Using this approach, Equation 3.7 is applied only once, at the start of the procedure, to capture the spatial correlations within the distribution. Moreover, the training dataset does not grow in size, enabling constant time updates in the number of measurements. To this end, the values at the input training points $\mathbf{X}$ need to be initialized. For unknown environments, they can simply be set to a constant uniform mean. For known environments, this can be done by training the GP model from available data and performing inference at the same or different resolutions, as required. The following sub-section describes the data fusion technique.

3.2.2 Sequential Data Fusion

As mentioned, a key component of the proposed method is a map update procedure based on recursive filtering. Given a uniform mean and the spatial correlations captured by Equation 3.7, the map $p(\mathbf{f} \,|\, \mathbf{X}) \sim \mathcal{GP}(\boldsymbol{\mu}^-,\, \mathbf{P}^-)$ is used as a prior for incorporating new sensor measurements.

Let $\mathbf{z} = [z_1,\, \ldots,\, z_m]^\top$ denote m new independent measurements received at points $[\mathbf{x}_1,\, \ldots,\, \mathbf{x}_m]^\top \subset \mathbf{X}$ in the environment $\mathcal{E}$. Each measurement is modeled assuming a sensor with Gaussian noise as

$p(z_i \,|\, f_i, \mathbf{x}_i) \sim \mathcal{N}(\mu_{s,i},\, \sigma^2_{s,i})$, with $\sigma^2_{s,i}$ corresponding to its variance level, which can be constant or non-constant. To fuse the measurements $\mathbf{z}$ with the prior map $p(\mathbf{f} \,|\, \mathbf{X})$, the maximum *a posteriori* estimator is used. It is formulated as:

$$\underset{\mathbf{f}}{\operatorname{argmax}}\, p(\mathbf{f} \,|\, \mathbf{z}, \mathbf{X})\,. \tag{3.9}$$

The Kalman Filter (KF) update equations are applied directly to compute the posterior density $p(\mathbf{f} \,|\, \mathbf{z}, \mathbf{X}) \propto p(\mathbf{z} \,|\, \mathbf{f}, \mathbf{X}) \times p(\mathbf{f} \,|\, \mathbf{X}) \sim \mathcal{GP}(\boldsymbol{\mu}^+,\, \mathbf{P}^+)$ (Reece et al., 2010):

$$\boldsymbol{\mu}^+ = \boldsymbol{\mu}^- + \mathbf{K}\mathbf{v}\,, \tag{3.10}$$

$$\mathbf{P}^+ = \mathbf{P}^- - \mathbf{K}\mathbf{H}\mathbf{P}^-\,, \tag{3.11}$$

where $\mathbf{K} = \mathbf{P}^-\mathbf{H}^\top\mathbf{S}^{-1}$ is the Kalman gain, and $\mathbf{v} = \mathbf{z} - \mathbf{H}\boldsymbol{\mu}^-$ and $\mathbf{S} = \mathbf{H}\mathbf{P}^-\mathbf{H}^\top + \mathbf{R}$ are the measurement and covariance innovation terms. $\mathbf{R}$ is a diagonal $m \times m$ matrix of altitude-dependent variances $\sigma^2_{s,i}$ associated with each noisy measurement z_i, and $\mathbf{H}$ is an $m \times n$ matrix denoting a linear measurement model that intrinsically selects part of the latent state $[f_1, \ldots, f_m]^\top$ observed through $\mathbf{z}$. The information to account for measurements at different resolutions is incorporated according to the measurement model $\mathbf{H}$ in a simple manner, as detailed in the following sub-section.

The map updates in Equations 3.10 and 3.11 are repeated sequentially, every time new data are registered, and can be performed in constant time. Note that, as all models are linear in this case, the KF update above produces the optimal posterior solution. Finally, an important benefit of this approach is that is agnostic to the type of sensor used, since it permits fusing heterogeneous data into a single representation of the environmental field map.

3.2.3 Altitude-dependent Sensor Model

As an example, this sub-section details an altitude-dependent sensor model for a downward-facing camera used to take measurements of a terrain, e.g., a farmland or debris field. In contrast to the discrete case examined in Section 3.1, the environmental field model in this setup needs to express uncertainty with respect to a target distribution defined in a continuous domain. To do this, the visual data is considered to degrade with increasing altitude in two ways: (a) noise and (b) resolution. The proposed model accounts for these issues in a probabilistic manner as detailed below.

This work assumes a sensor model contaminated by altitude-dependent Gaussian noise. For each observed point $\{\mathbf{x}_i \,|\, i = 1, \ldots, m\}$, the camera provides a measurement z_i capturing the target field variable f_i as $\mathcal{N}(\mu_{s,i},\, \sigma^2_{s,i})$, where $\sigma^2_{s,i}$ is the noise variance level expressing uncertainty in z_i. Similarly to the sensor in Section 3.1, to account for lower-quality images taken with larger camera footprints on the terrain, $\sigma^2_{s,i} = \sigma^2_s$ is modeled as increasing with the UAV altitude h using:

$$\sigma^2_s = \gamma_1(1 - e^{-\gamma_2 h})\,, \tag{3.12}$$

where γ_1 and γ_2 are positive constants defining the variation of sensor noise with altitude.

Figure 3.2 illustrates an example noise model for a camera-based sensor defined by Equation 3.12 with $\gamma_1 = 0.2$ and $\gamma_2 = 0.05$. The measurements z_i denote the physical values of the continuous target variable being measured, e.g., green biomass level or temperature. As for the discrete classifier discussed in Section 3.1, this model can be obtained by analyzing how the sensor behaves at different altitudes using previously acquired datasets. An important consideration in these studies is that the ground truth data of the field distribution must be available in order to assess the sensor accuracy. Reliable values of these measurements can be difficult to obtain in practical settings, as further discussed in Section 5.2.2.

The approach defines altitude envelopes corresponding to different resolution scales with respect to the initial points $\mathbf{X}$ specified on the terrain. This is motivated by the fact that the Ground Sample Distance (GSD) ratio (in m/px) depends on the altitude of the sensor and its fixed intrinsic resolution. To handle data received from variable altitudes, adjacent $\mathbf{x}_i$ are indexed by a single sensor measurement z_i through the measurement model $\mathbf{H}$. At lower altitudes (higher GSDs, corresponding to the maximum mapping resolution in $\mathbf{X}$), $\mathbf{H}$ is simply used to select the part of the state observed with a scale of 1. However, at higher altitudes (lower GSDs), the elements of $\mathbf{H}$ used to map multiple f_i to a single z_i are scaled by the square inverse of the resolution scaling factor s_f. Note that the fusion procedure described in the previous sub-section is always performed at the maximum mapping resolution, so that the proposed model $\mathbf{H}$ considers the low-resolution measurements as a scaled average of the high-resolution map.

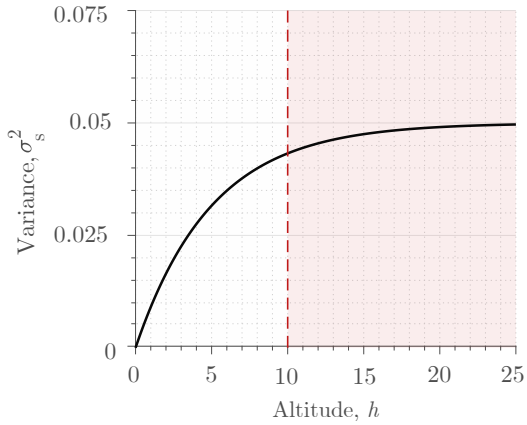

Figure 3.2: Inverse sensor noise model for a camera providing measurements as $\mathcal{N}(\mu_s, \sigma_s^2)$ with $\gamma_1 = 0.2$, $\gamma_2 = 0.05$ in Equation 3.12. The level of uncertainty σ_s^2 in the measurements increases with the UAV altitude h to represent degrading image quality. The dotted line at $h = 10\,\text{m}$ indicates the altitude above which the resolution of the camera images scales down by a factor of $s_f = 2$.

To close this section, Figure 3.3 presents an example application of the proposed mapping method for continuous field mapping to highlight its key features. The map in (a) represents the ground truth corresponding to an environmental field distribution on a terrain. The maps in (b)-(e) illustrate sequentially fusing two measurements taken at different altitudes (top row) into a single probabilistic map (bottom row), i.e., (d) and (e) visualize the results of fusing first (b) then (c), respectively, assuming that the map is initialized with a uniform prior mean of 50%. The measurement shown in (c), taken at a higher altitude, has greater coverage in terms of area visibility, but a reduced resolution scale compared to the one depicted in (b). By closely inspecting the final field map in (e), upon fusing (c), it can be seen that the off-center values are more widely diffused compared to those in the center, where the higher-quality measurement in (b) was registered. These observations demonstrate the ability of the approach to handle sensory inputs at variable resolutions and capture spatial correlations through the GP kernel.

3.3 Summary and Discussion

This chapter developed environmental mapping methods for informative planning with a focus on terrain monitoring applications. Strategies were presented for mapping both discrete and continuous target variables using probabilistic sensors for data acquisition. For the case of a discrete-valued target variable, the area of interest is represented as an occupancy grid in the context of an active classification task. The proposed approach is simple, efficient, and can be easily extended to accommodate multiple class labels by maintaining different occupancy grid layers for each one.

The main contribution of the work is a more sophisticated GP-based mapping strategy for the case of monitoring a continuous environmental field. Namely, the spatial correlation encoded in a

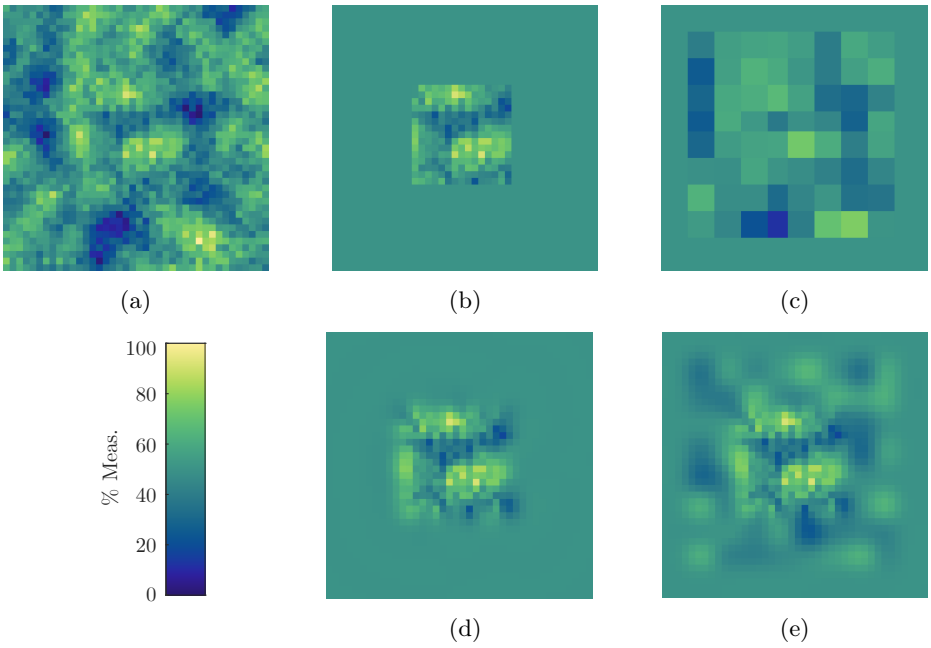

Figure 3.3: Overview of the main contribution of this chapter: a GP-based strategy for monitoring a continuous environmental field in terrain mapping applications. An example ground truth target distribution is shown in (a). (b) and (c) depict variable-resolution measurements taken using a downward-facing camera from 9 m and 20 m altitudes, respectively. (d) and (e) illustrate maps resulting from fusing the data sequentially. The variable diffusion effects show that the presented method can handle uncertain measurements at different resolutions while capturing spatial correlations.

GP model is exploited as a prior for Bayesian data fusion with probabilistic sensors. This feature enables incorporating altitude-dependent sensor models for aerial imaging with only constant-time measurement updates.

The inherent uncertainty measures in the map representations above can be used to plan informative trajectories using the algorithms introduced in Chapter 4. The integration of the proposed mapping and planning methods is demonstrated in Chapter 5 alongside experimental results.

The Bayesian fusion algorithm introduced in this chapter effectively bypasses the computational load of traditional GP-based mapping strategies. However, its advantages in terms of efficiency decrease with very large datasets, i.e., mapping large environments or mapping at high resolutions, due to the necessity to invert a large covariance matrix during the data fusion update. As a result, the approach suffers from cubic time complexity $\mathcal{O}(n^3)$ in the number of input points n in the training dataset. Sun, Vidal-Calleja, and Miró (2015) and Sun, Vidal-Calleja, and Miro (2017) introduce a submapping technique that is a possible solution to this problem. The key idea is to split the GP into multiple conditionally independent submaps for the fusion step, before backpropagating resulting information to recover the full joint map. Alternatively, the approximation methods discussed by Rasmussen et al. (2006) and Vasudevan et al. (2009) could be investigated.

In this thesis, it is assumed that the GP hyperparameters θ are pre-learned using reliable training data and fixed during the information gathering task. However, this approach is not suitable for situations where the structure of the monitored environmental phenomenon is entirely unknown *a priori* and needs to be explored. In fact, applying the wrong set of hyperparameters would lead to

a mismatch in the model, causing incorrect sensing actions to be planned. In such a scenario, one solution is to adapt the hyperparameters online based on the measurements collected. The use of this strategy to improve mapping performance has been the subject of several recent studies (Viseras et al., 2019 and K.-C. Ma, Liu, et al., 2017) and presents an interesting avenue for future work.

For data fusion, further empirical studies are necessary to investigate the implications of the Gaussian noise assumption made for the sensor in Section 3.2.3. This question is particularly interesting for sensors relying on deep learning architectures. Having only recently acquired real-time capabilities, such devices are becoming increasingly popular in UAV-based monitoring applications (Carrio et al., 2017; Sa, Z. Chen, et al., 2018 and Dang et al., 2018). Therefore, the ability to model them accurately will play an instrumental role in future active sensing systems.

4

Informative Path Planning

This chapter develops an online motion planning algorithm for generating informative trajectories to collect data in active sensing scenarios. It incorporates material from the following first author publications:

Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., and Galceran, E. (2017). "Online Informative Path Planning for Active Classification Using UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5753–5758. DOI: 10.1109/ICRA.2017.7989676

Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660

As discussed in the context of prior work in Section 2.1.2, the proposed method belongs to a class of adaptive, fixed-horizon algorithms for informative planning in resource-constrained missions. Essentially, the planner uses environmental map models built online to optimize continuous trajectories for maximum gain in an information-theoretic metric reflecting the mission aim. The following sections specify the problem setup before detailing the specifics of the planning routine. A key ingredient in the approach is the utility function, which encapsulates the planning objective and is used to evaluate the quality of information gathered along candidate paths. Several possible definitions are examined in Section 4.3 with respect to the mapping methods introduced in the previous chapter. In Chapter 5, the ideas are then unified into a single system for informative planning, which is finally evaluated against state-of-the-art methods.

4.1 Problem Setup

This chapter begins with a brief clarification about the two variations on the general informative path planning problem introduced in Section 2.1.2. Namely, a distinction is drawn between planning to maximize information gain (Equation 2.1) and information gain *rate* (Equation 2.2). Though these two formulations involve the same general terms, described further below, the rate-based variant is more applicable for tasks where the information quality of paths needs to be compared over different time and length scales. The main idea is to prevent the robot from expending its sensing resources in long detours to explore the environment, even if they lead to collecting valuable measurements. The remainder of this chapter focuses on this particular problem setup as it is relevant for the majority of scenarios addressed by the thesis, including the experiments in Chapter 5. However, the presented methods could be equally applied in the information-only case. This is demonstrated and further discussed in Chapter 6, where the proposed utility function $\mathrm{I}(\cdot)$ penalizes more costly paths, instead of the path planning problem itself.

Next, the terms in Equation 2.2 are defined in the context of robotics-based active sensing. The specific methods presented here are relevant for the experiments conducted in Chapters 5 and 6. However, note that they can be easily adapted for a wide variety of different applications. In the following exposition, it is assumed that the robot has perfect information about its pose, as provided by an external state estimation system.

To address the dynamic smoothness and agility requirements of physical platforms, such as small agile Unmanned Aerial Vehicles (UAVs), this thesis proposes parametrizing paths as continuous trajectories directly in the robot workspace. A polynomial trajectory ψ is represented by a sequence of N control waypoints to visit $\mathcal{C} = \{\mathbf{c}_1, \ldots, \mathbf{c}_N\}$ connected using $N-1$ k-order spline segments. In a 3-D workspace, the control waypoint $\mathbf{c}_i$ consists of (x, y, z) spatial co-ordinates defining position, with angles of roll, pitch, and yaw assumed to be zero for planning purposes. The first waypoint $\mathbf{c}_1$ is clamped to coincide with the initial (current) position of the robot.

Given a maximum reference velocity and acceleration, the trajectory is optimized for smooth, feasible minimum-snap dynamics following the method presented by Richter et al. (2013). Essentially, using this approach, the polynomial trajectory is expressed in terms of its end-point derivatives, which allows for efficient optimization in an numerically stable unconstrained quadratic program. Note that alternative strategies can be used to represent smooth paths, including the popular cardinal basis spline function, adopted by Hitz, Galceran, et al. (2017) in a similar optimization problem, or the Bézier curves studied by Gan et al. (2009). The planning strategy in these cases would simply proceed in the same fashion with different optimization variables.

As discussed in Section 2.1.2, the function $\textsc{measure}(\cdot)$ in Equation 2.1 provides a set of discrete measurements along a trajectory ψ, as determined by the properties of this sensor. Two possible definitions for this term are investigated; this work considers (1) an asynchronous sensor, providing measurements triggered at specified positions as the robot travels along a trajectory ψ, and (2) a constant-frequency sensor, providing measurements at equal time intervals given the robot velocity profile derived from the trajectory. These cases are distinguished as they have different implications on the planning procedure introduced in the following sub-section.

Finally, the function $\textsc{cost}(\cdot)$ obtains the cost of sensing resources incurred along ψ. This thesis examines a time-limited problem subject to a budget B, which is allocated based on practical considerations, such as finite battery life or flight time. However, the way in which the cost is defined can be easily modified to accommodate other platform-specific constraints, e.g., energy consumption or travel distance.

4.2 Algorithm

This section explains the operation of the algorithm for planning informative trajectories. First, a general description of the method is presented. The sub-sections that follow then highlight some subtle differences in the approach between setups using sensors with asynchronous and constant-frequency measurement triggers. Here, the main motivation is to improve the quality of predictive plans by catering for the ways in which various devices operate, and thus to also demonstrate how the framework can be easily adapted for different problem settings.

A fixed-horizon approach is used to plan adaptively as data are collected. During a mission, the scheme alternates between replanning and execution until the elapsed time t exceeds a prespecified budget B. The replanning scheme consists of two stages, and is shown in Algorithms 1 and 2 for the cases of the asynchronous and constant-frequency sensor, respectively. For clarity, these algorithms are formatted so that their line numbers correspond to equivalent steps of the algorithm. First, an

initial trajectory, defined by N fixed control waypoints $\mathcal{C}$, is derived through a search over a set of coarse grid points $\mathcal{L}$ (Lines 3-6) directly in the robot workspace. This step proceeds sequentially, selecting points in a greedy manner, so that a rough solution to the complex planning problem can be obtained quickly. Then, the trajectory is refined to maximize the utility function representing the mission aim. In this step, this work proposes employing an evolutionary routine, the Covariance-Matrix Adaptation Evolution Strategy (CMA-ES) (Line 7), to carry out global optimization in the continuous domain of possible trajectories in the workspace.

In Algorithms 1 and 2, note that the input variable $\mathcal{Z}$ is used as an abstraction to symbolize a general model of the environment $\mathcal{E}$ within which measurements are taken. In the context of this work, $\mathcal{Z}$ could correspond to a map representation based either on an occupancy grid or a Gaussian Process (GP) for the cases of monitoring a discrete or continuous target variable, respectively. Chapter 3 offers more extensive discussions on these models, including possible strategies to perform map updates during the mission.

Algorithm 1 REPLAN_PATH procedure - asynchronous sensor

Input: Current model of the environment $\mathcal{Z}$, initial position $\mathbf{c}_1$, number of control waypoints N, grid points $\mathcal{L}$
Output: Waypoints defining next polynomial plan $\mathcal{C}$

1: $\mathcal{Z}' \leftarrow \mathcal{Z}$ ▷ Create a local copy of the map.
2: $\mathcal{C}^g \leftarrow \mathbf{c}_1$; $\mathcal{C}^i \leftarrow \emptyset$ ▷ Initialize control points.
3: **while** $N \geq |\mathcal{C}^g \cup \mathcal{C}^i|$ **do**
4: $\quad \mathbf{c}^* \leftarrow$ Select next-best point in $\mathcal{L}$ using Equation 2.2 ▷ Perform grid search.
5: $\quad \mathcal{Z}' \leftarrow$ PREDICT_MEASUREMENT($\mathcal{Z}'$, $\mathbf{c}^*$)
6: $\quad \mathcal{C}^g \leftarrow \mathcal{C}^g \cup \mathbf{c}^*$; $\mathcal{C}^i \leftarrow \mathcal{C}^i \cup$ ADD_INTERMEDIATE_POINTS($\mathbf{c}^*$) ▷ Add point(s) to solution.
7: $\mathcal{C} \leftarrow \mathcal{C}^g \cup \mathcal{C}^i$; $\mathcal{C} \leftarrow$ CMAES($\mathcal{C}$, $\mathcal{Z}$) ▷ Optimize polynomial trajectory.

Algorithm 2 REPLAN_PATH procedure - constant-frequency sensor

Input: Current model of the environment $\mathcal{Z}$, initial position $\mathbf{c}_1$, number of control waypoints N, grid points $\mathcal{L}$
Output: Waypoints defining next polynomial plan $\mathcal{C}$

1: $\mathcal{Z}' \leftarrow \mathcal{Z}$ ▷ Create a local copy of the map.
2: $\mathcal{C} \leftarrow \mathbf{c}_1$ ▷ Initialize control points.
3: **while** $N \geq |\mathcal{C}|$ **do**
4: $\quad \mathbf{c}^* \leftarrow$ Select next-best point in $\mathcal{L}$ using Equation 2.2 ▷ Perform grid search.
5: $\quad \mathcal{Z}' \leftarrow$ PREDICT_MEASUREMENT($\mathcal{Z}'$, $\mathbf{c}^*$)
6: $\quad \mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}^*$ ▷ Add point to solution.
7: $\mathcal{C} \leftarrow$ CMAES($\mathcal{C}$, $\mathcal{Z}$) ▷ Optimize polynomial trajectory.

4.2.1 Asynchronous Sensor

For the case of a sensor providing measurements at unevenly spaced time intervals, this work distinguishes between (1) "global" control waypoints $\mathcal{C}^g$ and (2) "intermediate" control waypoints $\mathcal{C}^i$ used to define a candidate robot trajectory. As shown in Line 6 of Algorithm 1, the key difference between the two types of points is as follows: $\mathcal{C}^g$ are chosen as measurement points *directly on the coarse grid points* in $\mathcal{L}$ during the sequential greedy search (Line 4), whereas $\mathcal{C}^i$ are the points *interpolated* at a fixed frequency between each consecutive pair in $\mathcal{C}^g$ to add Degrees of Freedom (DoFs) to the

polynomial path. In doing so, the algorithm prioritizes the points $\mathcal{C}^g$ where the potential value of information is precisely evaluated, whereas the points $\mathcal{C}^i$ correspond to the data collected provided the sensor has enough time available between successive measurements. Note that, in this setup, the measurement points and control points are coincident such that $\mathcal{C} = \mathcal{C}^g \cup \mathcal{C}^i$.

Differentiating between the two points allows for a more realistic predictive planning procedure given the asynchronous nature of the sensor. Moreover, it enables investigating the effects of using different sets of points as optimization variables in the second step of the algorithm (Line 7). Specifically, this work examines: (1) globally optimizing all the control waypoints $\mathcal{C}$, similarly to the constant-frequency sensor case addressed below, and (2) optimizing the intermediate control waypoints $\mathcal{C}^i$ only for intersegment refinements along the trajectory. Hereafter, these optimization methods are referred to as "global" and "local" strategies, respectively.

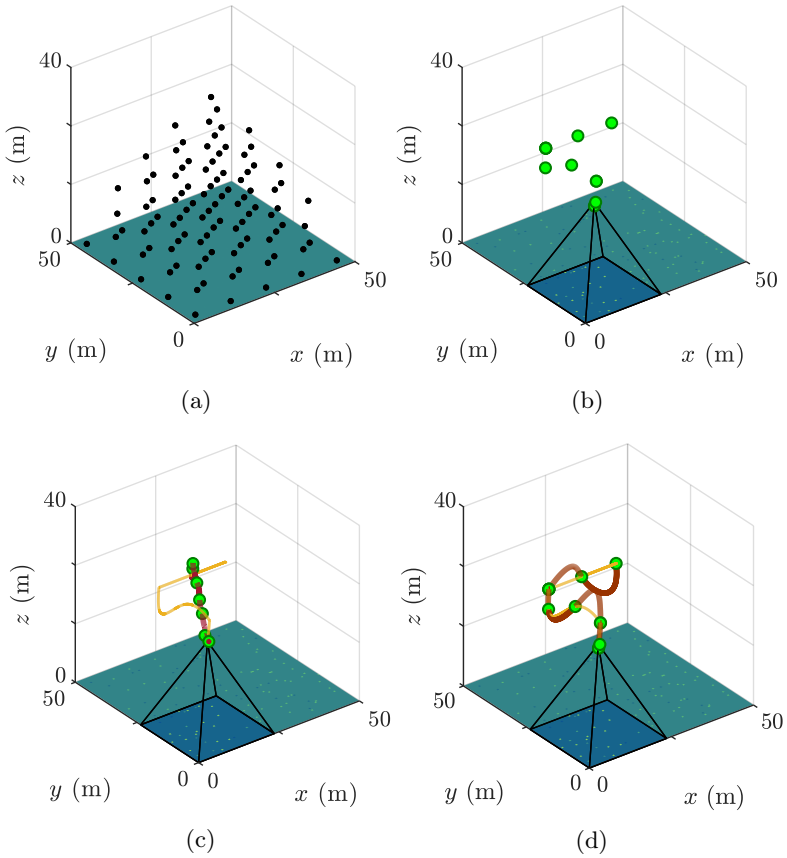

Figure 4.1: Schematic explanation of using different CMA-ES optimization strategies in Algorithm 1. (a) visualizes an example 3-D grid of points used for the discrete search in the first step of the procedure and (b) shows the control waypoints (green) in the resulting solution (Section 4.4). (c) and (d) depict subsequent "global" and "local" optimizations using the CMA-ES (Section 4.5). The orange and maroon curves correspond to trajectories before and after optimization, respectively. The pyramid shows the camera footprint corresponding to the initial robot position.

Figure 4.1 is provided as a visual aid to help understanding the key steps of the replanning routine in Algorithm 1. The plot in (a) exemplifies a set of points $\mathcal{L}$ in a 3-D robot workspace, which are used in the sequential grid search. The corresponding discrete solution resulting from this procedure

is represented by the green points in (b). The maroon curves in (c) and (d) depict the respective trajectories obtained upon optimizing (1) $\mathcal{C}$ and (2) $\mathcal{C}^i$ only, as described above, using the discrete solution (orange curve) to initialize the CMA-ES routine. Further details on these steps can be found in Sections 4.4 and 4.5.

4.2.2 Constant-frequency Sensor

This work also examines a constant sampling sensor frequency, whereby the spacing of measurement points is computed with respect to the dynamic profile of the robot trajectory. This corresponds to the behavior of simpler and more practical sensing devices, e.g., a camera or sensor suite (Hitz, Galceran, et al., 2017) operating at a fixed frequency. As shown in Algorithm 2, in this case, the measurement points (for data acquisition) and the control waypoints (for trajectory optimization) are not necessarily coincident. Instead, the proposed approach directly subsamples parameterized trajectories to calculate the accumulated information gain.

Having described the main mechanism underlying the proposed approach, the following sub-sections discuss possible objectives for informative planning and further detail its two stages.

4.3 Utility Definition

The utility, or information gain, function $\mathrm{I}(\cdot)$ in Equation 2.2 is an integral component of the informative planning algorithm as it determines the reward achieved by the robot for performing particular sensing actions. In other words, it specifies the objective of the planning task in terms of the aims of the data acquisition mission. This section presents possible definitions for the utility function that quantify the information value associated with potential sensor measurements. The proposed formulations are based on the discrete and continuous map representations introduced in Chapter 3. Note that, for consistency with the mapping methods, the planning objectives here assume that the observed environmental phenomenon does not vary with time.

In the spirit of previous work (Krause, A. Singh, et al., 2008; Carrillo et al., 2018; Charrow et al., 2015 and Bourgault et al., 2002), this thesis examines utility functions for the objective of *autonomous robotic exploration*. The basic idea is to encourage the robot to acquire new data in unknown areas in order to quickly obtain a high-fidelity map of the environment. To achieve this, one needs to evaluate the utility, or "usefulness", of a sensor measurement that would be obtained from a specific robot pose $\mathbf{p}$. This step corresponds to Line 4 of Algorithms 1 and 2. Note that, in the algorithms, $\mathbf{c}$ denotes a control waypoint (a position) parameterizing a *planned* polynomial trajectory, whereas here $\mathbf{p}$ is a specific pose along it from where the measurement was *actually* registered. As $\mathbf{p}$ is defined in a general manner, it enables capturing the true noisy motion of the robot for map building, e.g., the orientation angles of the robot are allowed to be non-zero due to its physical dynamics, although the planning procedure assumes otherwise. A further minor remark is that the term *measurement* in the context here refers to a single sensor trigger, whereas in Chapter 3 it designates a piece of data used to update the map, i.e., one measurement per location in the map.

In information theory, Shannon's entropy is often used to measure of the amount of uncertainty associated with a random variable (Cover et al., 2006). The proposed approach considers this metric as a natural way of describing the uncertainty in the field map. Following the conventional approach to informative path planning (Bourgault et al., 2002), the robot can then decide on the next sensing actions to maximize the map fidelity.

Generally, the objective of the pure exploration problem is thus to maximize the reduction of entropy $\mathrm{H}(\mathcal{Z})$, i.e., maximize the information gain, in the environmental field model $\mathcal{Z}$:

$$\mathrm{I}(\mathbf{p}) = \mathrm{H}(\mathcal{Z}^-) - \mathrm{H}(\mathcal{Z}^+)\,, \tag{4.1}$$

where the minus and plus superscripts denote the prior and posterior distributions on the model $\mathcal{Z}$, respectively, given a new measurement taken from a pose $\mathbf{p}$. Note that $\mathcal{Z}$ is used as an abstraction to symbolize a general map representation, as described above. The following two sub-sections introduce more concrete formulations for the two cases of discrete and continuous variable monitoring.

4.3.1 Discrete Variable Mapping

Let X be a discrete random variable with possible values $\{x_1, \ldots, x_n\}$ and a probability distribution function $p(x_i) = p(X = x_i)$. Shannon's entropy $\mathrm{H}(X)$ of X is defined as (Cover et al., 2006):

$$\mathrm{H}(X) = -\sum_{i=1}^{n} p(x_i) \log p(x_i)\,. \tag{4.2}$$

In the scenario of mapping a discrete-valued variable, the map of the environment is represented as an occupancy grid $\mathcal{X}$ (Section 3.1). Each single grid cell $\mathbf{x}_i \in \mathcal{X}$ is a random variable associated with two possible occupancy states (occupied and free) and a corresponding probability of being occupied $p(\mathbf{x}_i)$. Hence, the value of Shannon's entropy for the cell can be written by applying Equation 4.2 considering the two complementary states (Bourgault et al., 2002):

$$\mathrm{H}(\mathbf{x}_i) = -p(\mathbf{x}_i) \log p(\mathbf{x}_i) - (1 - p(\mathbf{x}_i)) \log (1 - p(\mathbf{x}_i))\,. \tag{4.3}$$

As the cells in the grid map are assumed to be spatially independent, the total entropy value can be obtained through a simple summation:

$$\mathrm{H}(\mathcal{X}) = -\sum_{\mathbf{x}_i \in \mathcal{X}} p(\mathbf{x}_i) \log p(\mathbf{x}_i) + (1 - p(\mathbf{x}_i)) \log (1 - p(\mathbf{x}_i))\,. \tag{4.4}$$

The definition in Equation 4.5 can be substituted into Equation 4.1 in order to evaluate the expected information gain associated with a particular measurement taken from $\mathbf{p}$:

$$\mathrm{I}(\mathbf{p}) = \mathrm{H}(\mathcal{X}^-) - \mathrm{H}(\mathcal{X}^+)\,, \tag{4.5}$$

where the superscripts are used to indicate the prior and posterior states of the occupancy grid map, as above.

Alongside the pure exploration case, this work also examines an adaptive planning setup where the information objective depends on the actual values of the measurements taken in addition to their location. Planning in this scenario is achieved by the means of a threshold-based strategy. Namely, Equation 4.1 is modified so that the elements mapping to the occupancy probability $p(\mathbf{x}_i)$ of each cell $\mathbf{x}_i \in \mathcal{X}$ are excluded from the objective computation, provided they do not satisfy a requirement that defines interest-based planning.

For monitoring a discrete target variable, this task is cast as an active classification problem. To encourage the classification of a binary variable, i.e., occupancy, the cells in the grid map $\mathcal{X}$ are partitioned into two sets: (1) occupied by target $p(\mathbf{x}_i) = 1$ and (2) target-free $p(\mathbf{x}_i) = 0$ based on user-defined confidence thresholds δ_o and δ_f on occupancy for the occupied and free states,

respectively. This is similar to the manner in which known and unknown space is distinguished in conventional occupancy mapping (Hornung et al., 2013). Applying the thresholds leaves a subset of grid cells in the map $\mathcal{X}_U$ which are considered unclassified:

$$\mathcal{X}_U = \{\mathbf{x}_i \in \mathcal{X} \,|\, \delta_f < p(\mathbf{x}_i) < \delta_o\}\,. \tag{4.6}$$

The proposed strategy then seeks to maximize the reduction of the subset $\mathcal{X}_U$ for a measurement taken from $\mathbf{p}$:

$$\mathrm{I}(\mathbf{p}) = |\mathcal{X}_U^-| - |\mathcal{X}_U^+|\,. \tag{4.7}$$

The definition in Equation 4.7 can be used as the utility function instead of that in Equation 4.5 for informative planning problems where the aim is to classify the grid cells within specific thresholds as quickly as possible.

Algorithm 3 Discrete variable mapping using a time-varying objective

1: **if** $t/B <$ RAND() **then** ▷ Select information objective based on time.
2: $\quad\mathbf{c}^* \leftarrow$ Select next-best point in $\mathcal{L}$ using Equation 2.2 with Equation 4.5 ▷ Exploration gain.
3: **else**
4: $\quad\mathbf{c}^* \leftarrow$ Select next-best point in $\mathcal{L}$ using Equation 2.2 with Equation 4.7 ▷ Classification gain.

Building upon these concepts, a time-varying parameter can be introduced in the informative planning algorithm to dynamically bias the information gathering objective from Equation 4.5 towards Equation 4.7. Intuitively, the idea is to first guide the robot to explore an initially unknown environment, before increasingly focusing on improving the confidence in its classification. The relevant parameter is defined in terms of the remaining proportion of the time budget as t/B. Using this planning strategy, Line 4 in Algorithms 1 and 2 is replaced by Algorithm 3, which simply selects one of the two objectives based on the elapsed mission time t. Note that the function RAND() is used to seed a random number in the range $[0, 1]$ such that the change between objectives occurs as a gradual transition, instead of based on a hard, manually-defined switch.

4.3.2 Continuous Variable Mapping

For a continuous random variable X with a probability distribution $p(x)$, Shannon's entropy is equivalently defined as (Cover et al., 2006):

$$\mathrm{H}(X) = -\int_{\mathbb{R}} p(x) \log p(x) dx\,, \tag{4.8}$$

provided the integral above exists.

To monitor a continuous-valued environmental phenomenon (Section 3.2), the field map is represented as the realization of a GP. In this case, if X follows a D-variate Gaussian distribution $X \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{P})$, the entropy in Equation 4.8 is (Rasmussen et al., 2006):

$$\mathrm{H}(X) = \frac{1}{2} \log |\mathbf{P}| + \frac{D}{2} \log(2\pi e)\,. \tag{4.9}$$

Unfortunately, the calculation in Equation 4.9 involves deriving the determinant of the covariance matrix $\mathbf{P}$ of the multivariate distribution, which is has a cubic time complexity $\mathcal{O}(n^3)$ in the number of input data points n in the GP model. To circumvent this computationally expensive step, principles

from the Theory of Optimal Experimental Design (TOED) (Fedorov, 1972) are applied to measure the scalar uncertainty associated with $\mathbf{P}$ more efficiently. Following prior studies (Sim et al., 2005 and Carrillo et al., 2018), this work considers the A-optimal information criterion $\hat{\mathrm{H}}(X)$, which accounts for the sum of the eigenvalues λ_i of $\mathbf{P}$, or equivalently, the total variance of its states, as given by the matrix trace:

$$\hat{\mathrm{H}}(X) = \hat{\mathrm{H}}(\mathbf{P}) = \mathrm{Tr}(\mathbf{P}) = \sum_{i=1}^{D} \lambda_i \,. \tag{4.10}$$

In particular, the A-optimal information measure is chosen as it incorporates the uncertainty in all the dimensions of the map state space (compared to D-optimality) without requiring all the eigenvalues λ_i to be computed directly (E-optimality) (Costante et al., 2017). The interested reader is referred to the work of Sim et al. (2005) and experimental results therein for a more extensive analysis of the advantages of using this approach.

Now, with the definition in Equation 4.10, the utility function for pure exploration in Equation 4.1 can be formulated as:

$$\mathrm{I}(\mathbf{p}) = \hat{\mathrm{H}}(\mathbf{P}^{-}) - \hat{\mathrm{H}}(\mathbf{P}^{+}) \,, \tag{4.11}$$

which corresponds to minimizing the mean squared error of the GP model of the field map.

Similarly to the previous sub-section, this work also examines problems of mapping continuous variables that require adaptive sensing behavior. This property is very valuable for practical applications in environmental monitoring, such as finding function extrema (Marchant and Ramos, 2014 and A. Singh, Krause, Guestrin, et al., 2009), classifying level sets (Hitz, Gotovos, et al., 2014 and Gotovos et al., 2013), or focusing on specific value ranges.

To demonstrate adaptive planning for such tasks, consider a mission where the aim is to focus specifically on regions of interest which have higher values of the latent target parameter, e.g., areas of high vegetation cover in an agricultural field. A threshold μ_{th} is specified to separate the (a) "interesting" value range (above) and the (b) "uninteresting" value range (below) in the environment into complementary subsets (a) $\mathcal{X}_I$ and (b) $\mathcal{X}_-$ of all points contained in the map. The partitioning strategy is described below. The main idea is then to only include the interesting subset of points $\mathcal{X}_I$ in calculating the information value of potential measurements in Equation 4.11, i.e., the utility associated with the points $\mathcal{X}_-$ yielding low values of the target parameter is ignored.

Let $\mathcal{X}_*$ denote the set of points in the GP model at which the values of the underlying distribution are to be inferred. To define the points in $\mathcal{X}_*$, two subtly different cases are distinguished: (a) in a standard GP regression approach, $\mathcal{X}_*$ corresponds to the query points in the matrix $\mathbf{X}_*$, whereas (b) in mapping using Bayesian data fusion, $\mathcal{X}_*$ corresponds to the input training points in $\mathbf{X}$, provided the map updates are performed at a fixed resolution and no regression step takes place. For generality, the asterisk subscript is added on the set. The upper subset of interesting cells $\mathcal{X}_I \subset \mathcal{X}_*$ is therefore:

$$\mathcal{X}_I = \{\mathbf{x}_{*i} \,|\, \mathbf{x}_{*i} \in \mathcal{X}_* \,\wedge\, f(\mathbf{x}_{*i}) \geq \mu_{th}\} \,, \tag{4.12}$$

where $f(\mathbf{x}_{*i})$ represents the value of the latent target parameter at the location corresponding to $\mathbf{x}_{*i}$. Note that $\mathcal{X}_I = \mathcal{X}_* \setminus \mathcal{X}_-$.

Since the latent distribution $f(\mathbf{x}_{*i})$ in Equation 4.12 is unknown, one needs to estimate it accounting for the inherent uncertainty of the GP model. To this end, the proposed approach leverages the principles of bounded uncertainty-aware classification presented by Gotovos et al. (2013) and Srinivas et al. (2012).

The subset of interesting locations $\mathcal{X}_I$ is now defined based on the mean $\mu(\mathbf{x}_{*i})$ and variance $\sigma^2(\mathbf{x}_{*i})$ of each cell $\mathbf{x}_{*i} \in \mathcal{X}_*$ as:

$$\mathcal{X}_I = \{\mathbf{x}_{*i} \,|\, \mathbf{x}_{*i} \in \mathcal{X}_* \wedge \mu(\mathbf{x}_{*i}) + \beta\sigma^2(\mathbf{x}_{*i}) \geq \mu_{th}\}\,, \tag{4.13}$$

where β is a design parameter tuned to scale the confidence interval for classification, i.e., it specifies with what certainty a point has to lie below the threshold before it is classified.

As a comment to close this section, note that the general utility function formulation in Equation 4.1 defines $\mathrm{I}(\mathbf{p})$ for a single robot pose $\mathbf{p}$. To determine the information value of a complete trajectory $\mathrm{I}(\psi)$, the same principles can be applied by fusing a sequence of measurements and computing the total change in the map.

Having laid out the fundamental elements in the informative planning problem, the following sections detail the main steps of the proposed replanning procedure.

4.4 Grid Search

The first step of the replanning scheme supplies an initial solution for the subsequent optimization routine. To achieve this, the planner performs a discrete grid search based on a coarsely defined set of grid points $\mathcal{L}$ in the robot workspace. As an example, Figure 4.2 illustrates two 3-D lattices applicable for a terrain monitoring scenario featuring a UAV equipped with a downward-facing camera. Note that the pyramidal geometry of the grid structures arises due to the perception field of the camera, which increases with altitude.

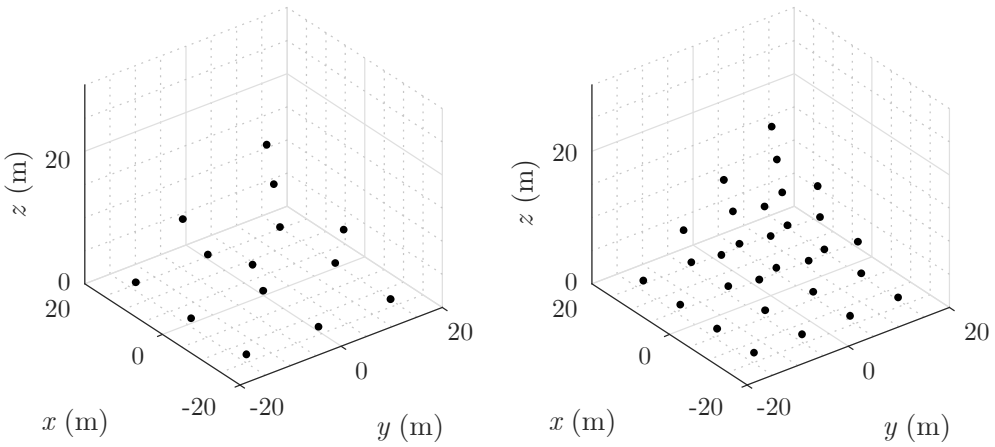

Figure 4.2: Visualizations of (a) 14-point and (b) 30-point 3-D lattice grids $\mathcal{L}$ for obtaining an initial trajectory solution in a $40\,\mathrm{m} \times 40\,\mathrm{m} \times 30\,\mathrm{m}$ robot workspace above a monitored terrain. The density of the points in $\mathcal{L}$ can be chosen to trade-off between solution accuracy and computational efficiency in the grid search. Note that the points are sparser at higher altitudes to ensure that measurements remain within bounds as the perception field of the downward-facing camera increases.

In the planning algorithm, a discrete low-accuracy solution neglecting the sensor dynamics is obtained efficiently by using the points in $\mathcal{L}$ to represent candidate measurement sites and assuming a constant robot velocity. In the related task of exploring unknown cluttered environments, a common strategy

is to drive the robot to take measurements on the map frontiers, which are detected in an occupancy grid as the boundaries of explored and unexplored areas (Yamauchi, 1997 and Charrow et al., 2015). However, this method is not applicable for the illustrated scenario, where the Field of View (FoV) of the sensor is always unobstructed by obstacles.

Instead, the proposed approach conducts a sequential greedy search for N waypoints (Line 3), where the next-best control waypoint $\mathbf{c}^*$ (Line 4) is found by evaluating the informative planning problem in Equation 2.2 with the chosen utility function $\mathrm{I}(\cdot)$ over the points in $\mathcal{L}$. Importantly, for prediction, this step assumes that no prior knowledge about the value of future measurements is available. The search is thus conditioned on the *most likely estimate* of the current field map, i.e., the prediction considers that the maximum probability distribution of the map state will be re-observed. In the discrete mapping scenario, this involves classifying the occupancy grid cells with $p(\mathbf{x}_i) \geq 0.5$ as being occupied and those with $p(\mathbf{x}_i) < 0.5$ as being free. Such a partition is required in order to identify which of the two classification curves in Figure 3.1 determines the inverse sensor model in the map update procedure (Equation 3.3). In the continuous mapping scenario, the most likely estimate of the map simply corresponds to the current mean $\boldsymbol{\mu}$ of the field distribution. Note that Equation 3.11 updates the map using the Bayesian data fusion procedure, whereas the regression equations in Equations 3.6 and 3.7 are applied for a standard GP-based mapping approach, as demonstrated in Chapter 6.

Once the next waypoint $\mathbf{c}^*$ is selected, a new measurement is simulated in the map $\mathcal{Z}$ using the same strategy (Line 5). This point is then added to the initial trajectory solution (Line 6).

As depicted in Figure 4.2, the length scales of $\mathcal{L}$ can be defined based on the computational resources available and the level of accuracy desired in the discrete grid search. The denser grid in (b), in comparison to (a), procures better initial solutions at the expense of longer evaluation times. The ability to manage this trade-off is a key feature of the proposed approach: by tuning the algorithm parameters, the system can be adapted to accommodate the computational requirements of different hardware platforms.

4.5 Optimization

The second step of the replanning scheme optimizes the coarse grid search solution for $\mathcal{C}$ by computing $\mathrm{I}(\cdot)$ for a sequence of measurements taken along the corresponding trajectory, as defined in Section 4.3. By exploiting the grid-based result, the optimizer is initialized with an informed prior to speed up its convergence. Note that this step is solver-agnostic, and can be performed with any optimization routine. Specifically, the approach in this work applies the CMA-ES, as discussed in more detail below.

The CMA-ES is a generic global, derivative-free optimization routine based on the concepts of evolutionary algorithms. It has been successfully applied to solve high-dimensional, non-linear, non-convex problems in a continuous domain, including for trajectory generation (Hitz, Galceran, et al., 2017), well placement (Bouzarkouna et al., 2012), and feedback control (Hansen et al., 2009). In this work, the choice of optimization method is motivated by the non-linearity of the objective space in Equation 2.2 as well as by the above mentioned results. To evaluate this decision, Section 5.2.2 presents experimental results comparing the CMA-ES against alternative optimizers.

As an evolutionary strategy, the CMA-ES operates through the repeated interplay of variation and selection. The main idea is to evolve iterations (*generations*) of candidate solutions (*offspring*) by stochastic sampling during the search for a function optimum. Importantly, during the optimization routine, the algorithm is able to adaptively increase or decrease the search space, by modifying

the covariance matrix of the Gaussian distribution from which offspring in the next generation are sampled.

Algorithm 4 presents a brief overview of the algorithm, adapted from the work of Hitz, Galceran, et al. (2017). As inputs, the CMA-ES expects a fitness function OBJECTIVE$(\cdot)$ to evaluate new offspring and an initial guess $\mathbf{m}_0$ representing the sampling distribution mean. In the informative planning problem, the offspring correspond to candidate trajectories traveled by the robot, and OBJECTIVE$(\cdot)$ is the utility function I$(\cdot)$ that defines their information quality. The key insight of this work is to initialize $\mathbf{m}_0$ with the greedy grid-based output $\mathcal{C}$, thereby providing the optimizer with an informed prior for where to start its search. As described in Section 4.1, a constraint is placed on $\mathbf{c}_1$ to coincide with the initial robot position. Note, also, that a a coordinate-wise boundary handling algorithm (Hansen et al., 2009) is applied to constrain the measurement points along trajectories to lie within the feasible robot workspace, e.g., the volume above a surveyed terrain.

The main parameters required by the method are the population size λ_{CMA} (number of offspring), and the initial step size σ_{CMA} (standard deviation), the latter of which controls the steps taken during the optimization search. Note that σ_{CMA} is a vector such that different step sizes can be specified for each co-ordinate of the search space. To terminate the search, a convergence criterion C and/or a maximum number of optimizer iterations N_{ITERS} can be defined.

In the algorithm, at each iteration n, a new generation of offspring $\mathbf{x}^{(1,\dots,\lambda_{\text{CMA}})}$ is sampled from a multivariate Gaussian distribution according to a mean offspring $\mathbf{m}$ and a covariance matrix $\mathbf{C}$ (Line 6). The new offspring are then evaluated using the fitness function (Line 7) and sorted in order of best performance (Line 8). The result is then used to update the sampling distribution for the following generation (Line 9). This procedure also modifies the step size σ_{CMA} using an *evolution path* which maximizes the likelihood of previously successful search steps.

Algorithm 4 CMA-ES optimization routine

Input: Objective function OBJECTIVE$(\cdot)$, initial offspring $\mathbf{m}_0$, population size λ_{CMA}, initial step size $\boldsymbol{\sigma}_{\text{CMA}}$, convergence threshold C, maximum number of iterations N_{ITERS}

Output: Optimized solution $\mathbf{x}^*$

1: $\mathbf{C} \leftarrow \mathbf{I}$ ▷ Initialize the covariance matrix to identity.
2: $\mathbf{m} \leftarrow \mathbf{m}_0$ ▷ Initialize the mean offspring of the current generation.
3: $n \leftarrow 0$ ▷ Initialize the iteration counter.
4: **while** $(\text{Tr}(\boldsymbol{\sigma}^2_{\text{CMA}}\mathbf{C}) > C) \wedge (n < N_{\text{ITERS}})$ **do** ▷ Iterate until convergence or max. number of iterations reached.
5: **for** $g \in \{1, , \dots , \lambda_{\text{CMA}}\}$ **do** ▷ Sample a new generation.
6: $\mathbf{x}^{(g)} \leftarrow \mathcal{N}(\mathbf{m}, \boldsymbol{\sigma}^2_{\text{CMA}}\mathbf{C})$ ▷ Sample a new offspring.
7: $f^{(g)} \leftarrow \text{OBJECTIVE}(\mathbf{x}^{(g)})$ ▷ Evaluate the fitness of the offspring.
8: $\mathbf{x}^{(1,\dots,\lambda_{\text{CMA}})} \leftarrow \text{SORT}(\mathbf{x}^{(1,\dots,\lambda_{\text{CMA}})}, f^{(1,\dots,\lambda_{\text{CMA}})})$ ▷ Sort solutions in order of fitness.
9: $\mathbf{m}, \mathbf{C}, \boldsymbol{\sigma}_{\text{CMA}} \leftarrow \text{UPDATE}(\mathbf{x}^{(1,\dots,\lambda_{\text{CMA}})}, \mathbf{m}, \mathbf{C}, \boldsymbol{\sigma}_{\text{CMA}})$ ▷ Update the mean, covariance matrix, and step size.
10: $n \leftarrow n + 1$ ▷ Increment the iteration counter.
11: $\mathbf{x}^* \leftarrow \mathbf{m}$ ▷ Return the best solution candidate.

Further details on this routine go beyond the scope of this thesis and are provided in the in depth review by Hansen (2006), which derives the necessary *internal* parameters for optimal updates in Line 9 and presents a corresponding convergence analysis. The implementation in this work considers these optimal parameters, considering the inputs to Algorithm 4 as *external* parameters that need to be defined by the user. These variables can be set manually based on application-specific requirements, e.g., decreasing the number of iterations for the optimization process lowers its computational intensity, but sacrifices the quality of the output solution.

4.6 Summary and Discussion

This chapter introduced an online algorithm for informative path planning in resource-constrained missions. The proposed approach operates in an adaptive, fixed-horizon manner by optimizing smooth trajectories in continuous 2-D or 3-D space to maximize an informative objective. To do this efficiently, the replanning procedure exploits a coarse grid-based solution in the robot workspace to initialize an evolutionary optimization routine. A valuable practical feature is that one can trade off between the solution quality and computational complexity of the algorithm, allowing it to run on a variety of platforms. The integration of the planning scheme with different environmental mapping methods is demonstrated in Chapter 5 alongside experimental results.

Although it enables adaptive data acquisition, one limitation of the algorithm is that its output plans are limited to a fixed horizon. As a result, the path quality strongly depends on the horizon lookahead, which must be defined manually by the user. This could be especially problematic as candidate paths are defined by a specified number of control waypoints, which makes them comparable in complexity, but not necessarily in terms of time duration, length, or number of measurements acquired. It is thus relevant to investigate more carefully how the complexity of path parameterization influences information gathering in different environments. One potential improvement is to adopt a receding-horizon approach in which only the next-best step of a new plan is executed before replanning, as illustrated by Bircher et al. (2018).

Related to this point, a further issue is that the strategy does not exploit any knowledge about the remaining budget when generating new plans. This could be problematic if the robot requires some resources to finalize the mission, e.g., to perform a safe landing within a time limit or to return to the starting point. Section 4.3.1 made a brief mention of utility functions which vary with time that are a possible solution. Similar ideas are introduced by Palazzolo et al. (2018). Essentially, the time-dependency could serve as a signal to gradually direct the robot towards a particular area or action before the sensing task is complete.

A utility function with these properties is also useful from an application point of view. In the context of continuous variable monitoring, it could be coupled with the principles of uncertainty-aware classification in Section 4.3.2 to define very precise objectives over the duration of the mission. Such capabilities are very relevant for practical sensing scenarios and worth exploring further.

Given its modularity, a natural direction for future studies is to adapt the algorithm for various informative planning setups, e.g., using alternative methods to represent candidate paths or environmental field maps. In particular, it would be interesting to explore applications with inherently similar objectives but different practical requirements, e.g., soaring with an unpowered glider in wind (Chung, 2014 and Lawrance et al., 2011) or monitoring bacteria with an autonomous surface vessel in a lake Hitz, Gotovos, et al., 2014; Hitz, Galceran, et al., 2017, or entirely different tasks, e.g., robotic manipulation (Bajcsy et al., 2018).

Finally, with the increasing applicability of multi-agent systems, one could extend the planning approach for environmental mapping with several robots. Many new intriguing research questions emerge in this problem setup. Within the algorithm, key challenges would include reformulating the objective function (Section 4.3) to achieve coordinated behavior, as well as improving the computational efficiency of the optimization procedure (Section 4.5). Chapter 7 provides a further outlook in this direction.

5

System Integration and Experimental Results

In this chapter, the theoretical contributions in mapping and planning from Chapters 3 and 4 are unified into a single informative planning framework applicable for general active sensing tasks. Material from the following first author publications is incorporated:

Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., and Galceran, E. (2017). "Online Informative Path Planning for Active Classification Using UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5753–5758. DOI: 10.1109/ICRA.2017.7989676

Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1382–1388. DOI: 10.1109/IROS.2017.8202317

Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660

This chapter begins with a brief overview of the proposed infrastructure for informative planning. The methods are then extensively evaluated in a wide range of simulated environments for Unmanned Aerial Vehicle (UAV)-based terrain monitoring tasks. The studies investigate how the approach behaves both as an integrated framework and in terms of its key components. System validation is performed in a photorealistic mapping problem using RIT-18, a publicly available dataset of aerial multispectral imagery. Finally, the algorithms are demonstrated running in real-time on UAV platforms in various practical scenarios.

An open-source implementation of the framework is released for usage and further development by the community. Supporting documentation and example experimental results are also included. This package is available at: `github.com/ethz-asl/tmplanner`.

5.1 System Integration

As a recapitulation, Figure 5.1 depicts the generic architecture of the proposed informative planning framework for active sensing applications. This system consolidates the previously introduced mapping and planning methods, and represents the overarching approach of this thesis. The dashed red lines in the diagrams correspond to additional data connections required for the problem of uncertainty-aware sensing, which is studied in Chapter 6. As shown, the framework is capable of

mapping either discrete or continuous target variables based on measurements extracted from a sensing unit, e.g., a depth or a multispectral camera. A suitable map representation between the two can be chosen based on the nature of the observed environmental phenomenon. During a mission, the informative path planning algorithm uses the environmental maps built online to optimize trajectories for maximum gain in an information-theoretic metric reflecting the mission aim. As discussed in depth in Chapter 1, a key feature of the approach is its modular construction, which enables it to be adapted for any active sensing scenario. The various experiments presented in this chapter demonstrate its flexibility.

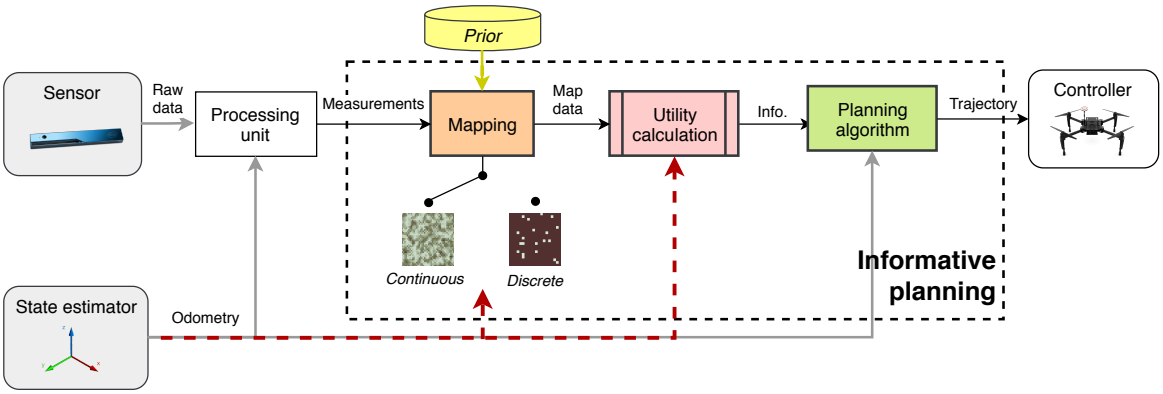

Figure 5.1: System diagram showing the key elements of the proposed informative path planning framework for active sensing and its associated data streams. The system is modular and can be easily implemented on a variety of systems. The red line indicates the dataflow emerging in uncertainty-aware sensing problems (Chapter 6).

To begin with, Figure 5.2 presents an example illustrating the progression of the proposed framework. The setup considers mapping a continuous field distribution on an initially unknown terrain using a UAV equipped with a downward-facing camera. The UAV workspace is the volume above the terrain and the camera sensor model assumes degrading image quality at higher altitudes. The field map is represented as a Gaussian Process (GP) model (Section 3.2). For adaptive planning, as described in Section 4.3.2, a base threshold $\mu_{th} = 40\%$ is set to focus on the more interesting, higher-valued range of the target parameter. This value also allows for including unobserved cells in the information objective, which are initialized uniformly with an uninformed mean prior of 50%.

The first and second rows in Figure 5.2 visualize the planned UAV trajectories and field maps, respectively, as images are registered at different times during the mission. The top-left plot depicts the first planned trajectory before (orange) and after (colored gradient) applying the Covariance-Matrix Adaptation Evolution Strategy (CMA-ES) routine. As shown, the optimization step shifts the initial measurement points (squares) to high altitudes, allowing for low-resolution, high-uncertainty data to be collected quickly before the map is refined (second and third columns). A qualitative comparison with ground truth on the bottom-left confirms that the method performs well, producing a fairly complete map in a short period of time with most uninteresting regions (hatched areas) identified. In the following sections, similar simulation setups are considered to evaluate the system both qualitatively and quantitatively.

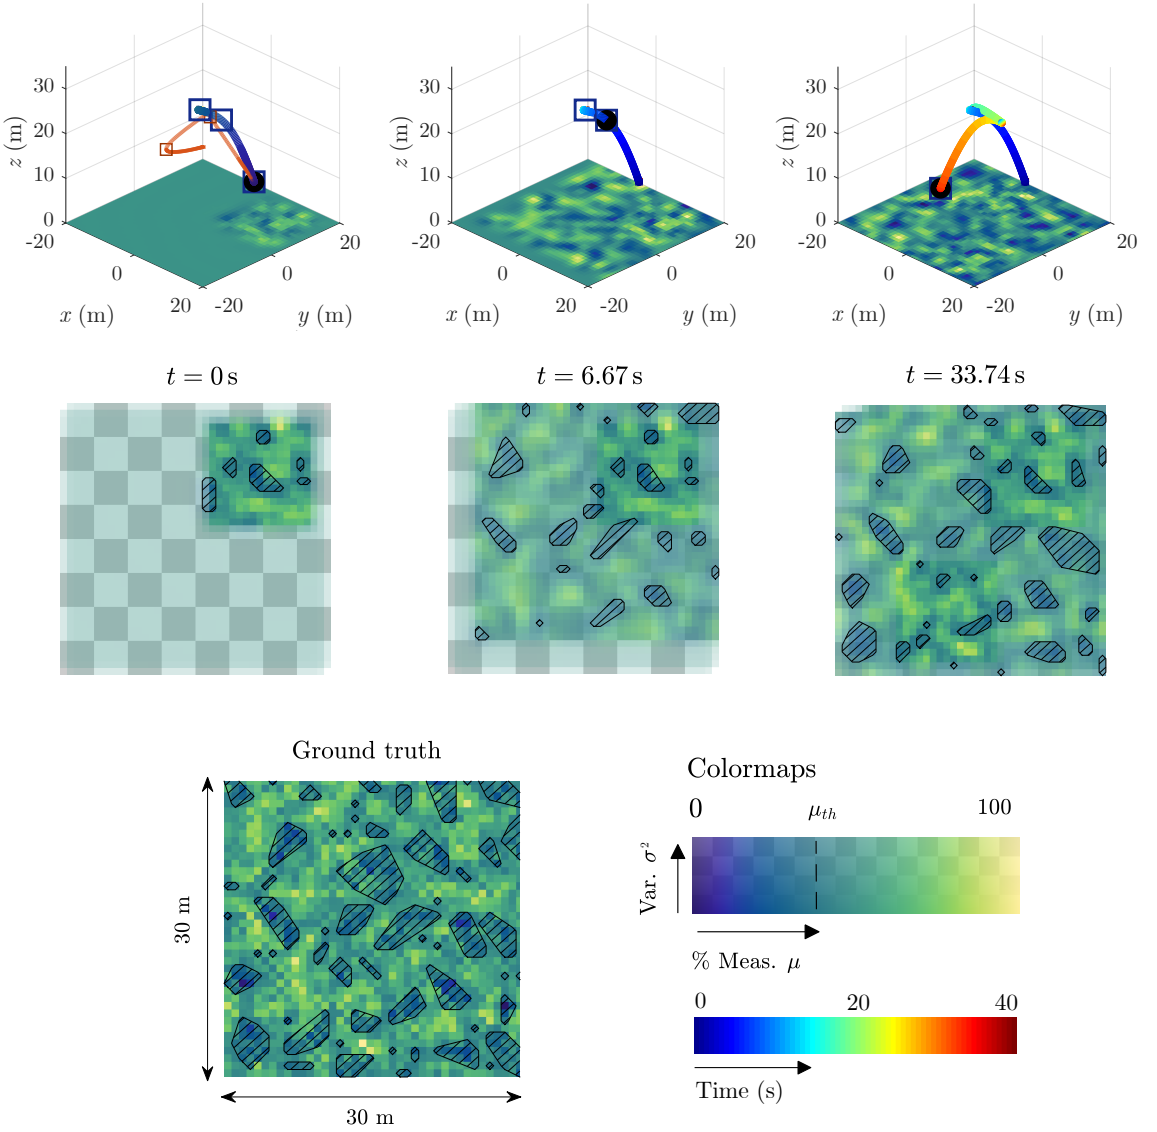

Figure 5.2: Example simulation results of the proposed informative planning framework (Figure 5.1) for UAV-based terrain monitoring. The colormaps are shown on the bottom-right. Bluer and yellower shades represent lower and higher values of the target parameter, respectively. In the maps, opacity indicates the model uncertainty (variance, σ^2), with the checkerboard added for visual clarity and the hatched sections denoting uninteresting areas with $< \mu = 40\%$. The ground truth is shown on the bottom-left. The three columns depict the trajectories (first row) and maps (second row) at different snapshots of the mission at times $t = 0\,\text{s}$, $6.67\,\text{s}$, and $33.74\,\text{s}$. In the top plots, the black dot indicates the current UAV position while the squares show the measurement points. The top-left plot illustrates an example trajectory before (orange) and after (colored gradient) optimization using the CMA-ES. Note that the GP means are rendered in the top trajectory plots.

5.2 Simulation Results

This section presents a series of simulation experiments used to evaluate the proposed approach and its key components. The experimental setup models a synthetic information gathering problem in a terrain monitoring application using a UAV equipped with a camera-based sensor. A wide range of scenarios are studied, targeting the reconstruction of both discrete and continuous environmental phenomena. The aim and outcome of each set of experiments is described in the following subsections.

It is important to note that, to the best of the author's knowledge, there is no informative planning method that procures provably optimal results when operating in the space of continuous trajectories. Therefore, in the following experiments, the tested algorithms are evaluated by assessing their relative performance using information-theoretic metrics.

5.2.1 Discrete Variable Mapping

These experiments investigate monitoring a binary variable on a terrain using the occupancy grid-based mapping method (presented in Section 3.1) and an asynchronous sensor for data acquisition (Section 4.2.1). The integrated framework is first benchmarked against existing approaches. Then, a comparison is performed between the different variants of the planning algorithm, which were introduced and explained in Section 4.2.1.

Comparison Against Planning Benchmarks

The aim of these experiments is to evaluate the proposed system against state-of-the-art planning algorithms. The simulation setup features a $50\,\mathrm{m} \times 50\,\mathrm{m}$ terrain (field) environment with randomly scattered targets. To analyze how the algorithms behave with different target densities, Poisson distributions with 50 to 250 targets are randomly generated in each environment during the trials. The field map is represented by an occupancy grid with a uniform resolution of $0.5\,\mathrm{m}$ and probability thresholds of $\delta_o = 0.75$ and $\delta_f = 0.25$ defining grid cells as being occupied and target-free, respectively, for adaptive planning purposes.

To simulate incorrect classifications, uniform noise is added on the binary classifier output based on the model shown in Figure 3.1. The number of false positive cells is limited to 800 to avoid excessive noise in target-free regions, which constitute the majority of the environment. The simulated camera providing image inputs to the classifier features a square footprint with a 60° Field of View (FoV) and maximum measurement frequency of $0.2\,\mathrm{Hz}$. As explained in Section 4.2.1, the latter value is important as it defines the spacing between the interpolated intermediate control waypoints $\mathcal{C}^i$ used to represent candidate trajectories. Finally, the experimental setup assumes perfect control and state estimation, such that the UAV can track the planned trajectories without any actuation noise and has ideal knowledge of its pose. The implications associated with relaxing the latter requirement are the subject of discussions in Chapter 6.

In the following experiments, the benchmark methods considered for comparison are: (a) a traditional coverage-based strategy using a "lawnmower" pattern and (b) the Rapidly exploring Information Gathering (RIG)-tree, introduced by Hollinger and Sukhatme (2014). The latter is a state-of-the-art sampling-based planning algorithm that incrementally extends a geometric tree to find the best information gathering trajectory. Each method was tested over 100 trials with a budget of $B = 300\,\mathrm{s}$ on the duration of the mission.

The algorithms are evaluated using different performance criteria commonly studied in active classification problems; the total Shannon's entropy of the map, classification rate, and mean F2-score.

Following a similar approach to Pomerleau et al. (2013), the Cumulative Distribution Function (CDF) of entropy in the map is computed over a time histogram to summarize the variability among different planned trajectories. Intuitively, steeper (faster-rising) curves using this metric represent quicker reductions in map uncertainty and therefore better performances. The mean F2-score of the cells in the grid map is used to compare classification accuracy. This statistic was chosen as it weighs the recall of the classification output higher than its precision to place more emphasis on false negative errors. Since the ground truth maps are generated with a sparse distribution of targets, it is desirable to punish the classifier more harshly for failing to identify them.

The initial UAV position for the two informative planning schemes is specified as the center of the mapped environment with a high 40 m altitude, from which the entire area of interest can be seen. The maximum reference velocity and acceleration for minimum-snap trajectory generation are 3 m/s and 1.5 m/s^2 using polynomials of order $k = 12$. A replanning horizon of $N = 5$ control waypoints is set to limit the number of optimization variables in the CMA-ES routine, and therefore its computational complexity. As explained in detail in Section 4.2.1, both "global" and "local" strategies for the CMA-ES are considered to analyze the effects of applying different trajectory optimization strategies in this setup.

For RIG-tree, the cost and information of a vertex (measurement point) are associated with the accumulated travel time and total map entropy given the newly collected data. To compute cost, the trajectory optimization procedure is performed for each edge, assuming that measurements at each vertex are taken with the UAV at rest. Considering that the cells in the occupancy grid are assumed to be spatially independent of one another, the information objectives in this scenario are modular, i.e., the total value of a set of measurements does not depend on the order in which they are registered. Therefore, as the tree is expanded, a conservative modular pruning strategy is applied to eliminate nodes that cannot lead to the optimal solution. This effectively reduces the running time of the algorithm without sacrificing its performance. Further details are discussed by Hollinger and Sukhatme (2014) and Binney and Sukhatme (2012).

Since RIG-tree is an offline algorithm, i.e., the entire informative trajectory is computed before the start of the mission, it is not directly comparable to the methods presented in this thesis, which have online replanning capabilities. A modified version of RIG-tree is therefore designed that allows for incremental replanning and adaptivity. First, the strategy is provided with prior knowledge of the environment from a high-altitude scan. This allows for generating an initial plan. Then, it alternates between tree construction (replanning) and plan execution, updating the map with each newly acquired set of measurements. In this approach, each tree construction stage is terminated after $\sim$ 20 s to allow for approximately the same amount of time as required by the CMA-ES to optimize a single trajectory.

For the coverage-based planner, a constant altitude (14.43 m) and maximum velocity (0.844 m/s) are defined for complete coverage given the specified mission budget of 300 s. To provide a fair comparison with the informative planning methods, the benchmark design method examined several possible "lawnmower" patterns above the square terrain with altitudes determined by the camera FoV. For each pattern, the UAV velocity was modified to match the time budget, and the best-performing one among them selected based on the evaluation criteria.

Figure 5.3 summarizes how the algorithms score against the three evaluation metrics over the 100 trials. For the new planning methods introduced in this work, results using the time-varying utility function (Algorithm 3) are presented, as this approach was found to perform better than optimization objectives targeting either pure information or pure classification gain. As expected, the naïve coverage-based strategy reduces the uncertainty of the map at a uniform rate (left), since the entire environment is scanned at a fixed altitude and constant velocity. The informative planning methods

perform better in comparison as they allow the UAV to fly at variable altitudes and exploit wider perception fields.

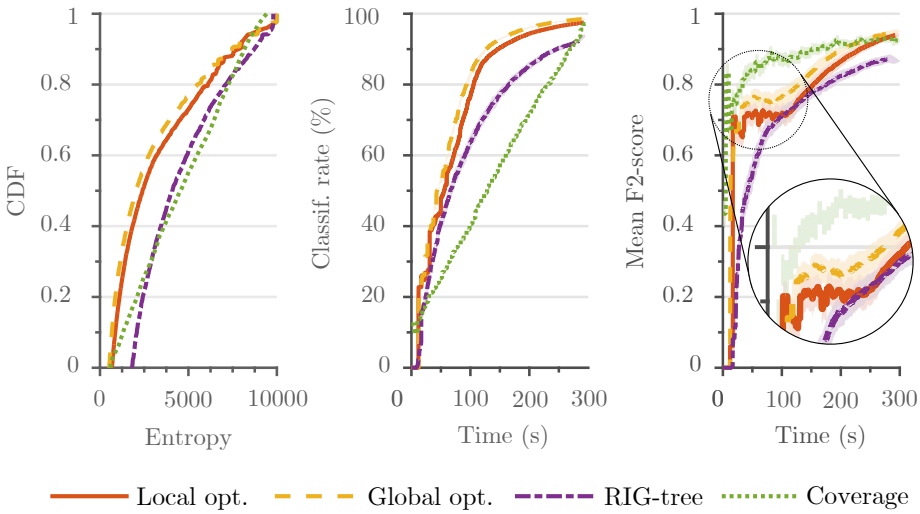

Figure 5.3: Comparison of the proposed approach using the CMA-ES against benchmarks for a fixed mission time budget of 300 s. For each algorithm, average results over 100 trials are shown with 90% confidence bounds. Using informative planning strategies, the map uncertainty (left) and classification (middle, right) improve quickly as the UAV can fly at variable altitudes to trade off between area coverage and sensor noise while acquiring new data.

In the plot on the right, the variations in mean F2-score suggest that the coverage planner yields the most accurate classification in the areas of the environment it has observed. Most likely, this occurs due to its low fixed flight altitude of 14.43 m (set based on the allocated time budget), in combination with the sensor model considered in the experiments. From the curve in Figure 3.1, it can be seen that the classification errors at this altitude are relatively low, leading to higher overall accuracy when compared with the wider altitude ranges of the informative planners that receive more noise. However, note that the mean F2-score does not account for the fact that, at the beginning of the "lawnmower" pattern, a large proportion of the environment is still completely unknown.

To reinforce these ideas, Figure 5.4 offers a direct comparison between the proposed approach and the coverage strategy in a typical mission. The trajectory produced by the informative planning algorithm (top-left) illustrates a path covering variable altitudes in resemblance to a spiral, starting with descent to the initially unknown center of the field. This type of behavior can be expected, as it enables the UAV to first collect low-quality, high-altitude data and obtain a wide overview of the environment. With this coarse information, it can then focus specifically on mapping the corners of the area and any detected regions of interest. In contrast, by being constrained to a fixed altitude, the "lawnmower" cannot exploit the probabilistic nature of the classifier and produces the result described above.

Generally, the results in Figure 5.3 show that the proposed approach (red, yellow) produces more informative paths than RIG-tree (purple) when the same amount of time is allowed for replanning. This indicates that its two-step scheme is more efficient at finding promising viewpoints in 3-D space than the incremental sampling-based technique. Moreover, unlike the tree-growing method, it does not require prior knowledge for initialization and can generate inherently smooth trajectories. Note that, in Section 5.2.2, considering a constant-frequency sensor model, an improved adaptive version

of RIG-tree is developed that removes the latter limitation by allowing for continuous transitions between successive vertices.

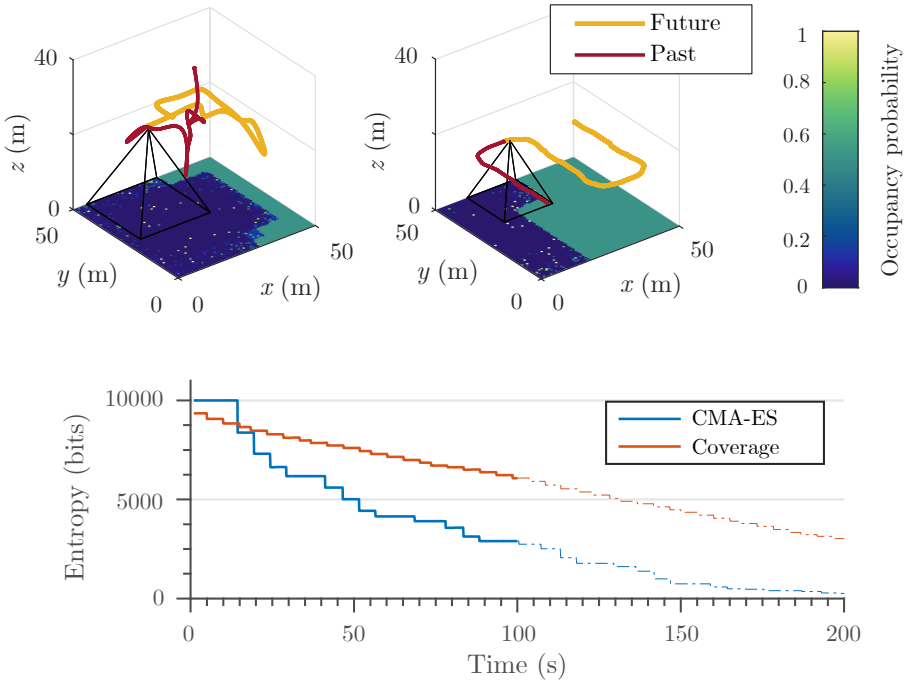

Figure 5.4: Comparison between the proposed informative planning approach and a coverage-based "lawnmower" strategy. *Top*: Visualizations of the trajectories traveled by the UAV using each method (left: informative planner, right: coverage path). The environment is assumed to be initially unknown, and the yellow spots on the terrain maps represent detected targets. The pyramid shows the camera perception field at a particular mission time and the maroon and orange lines correspond to trajectories before and after this instant, respectively. *Bottom*: Evolution of total map entropy during the mission. The solid and dotted lines are associated with the past and future trajectories. By replanning online, the proposed approach produces a map with 45% lower entropy of the coverage path in the same amount of time (100 s).

Comparing the two variants of the CMA-ES, one can perceive that the "global" CMA-ES strategy (yellow) produces both faster reductions in total map uncertainty, as well as higher gains in classification rate, compared with the "local" optimization of only intermediate points (red). This suggests that the algorithm performs better when all control waypoints defining a trajectory are passed as variables to the optimization routine. The following sub-section performs a more detailed analysis to explore these ideas.

Evaluation of Planning Strategies

The following experiments examine changing planning strategies within the proposed approach to assess their effects. The aim is to evaluate how different variations of the planning algorithm perform in a realistic active classification scenario. The experimental setup considers a UAV monitoring a binary variable with the same simulation parameters as described in the previous sub-section.

Specifically, these experiments investigate varying:

(a) *Grid search objectives*: information only (Equation 4.5), classification only (Equation 4.7), using the time-varying parameter (Algorithm 3)

(b) *Optimization methods*: without applying the CMA-ES (i.e., grid search only; omitting Line 7 in Algorithms 1 and 2), "local" CMA-ES (Figure 4.1(d)), "global" CMA-ES (Figure 4.1(c))

As in the previous experiments, simulations for each algorithm were repeated over 100 trials in environments with randomly varying distributions of targets. The same classification metrics are examined for a quantitative evaluation of the methods: the time variation of the CDF of total Shannon's entropy (map uncertainty), classification rate, and mean F2-score.

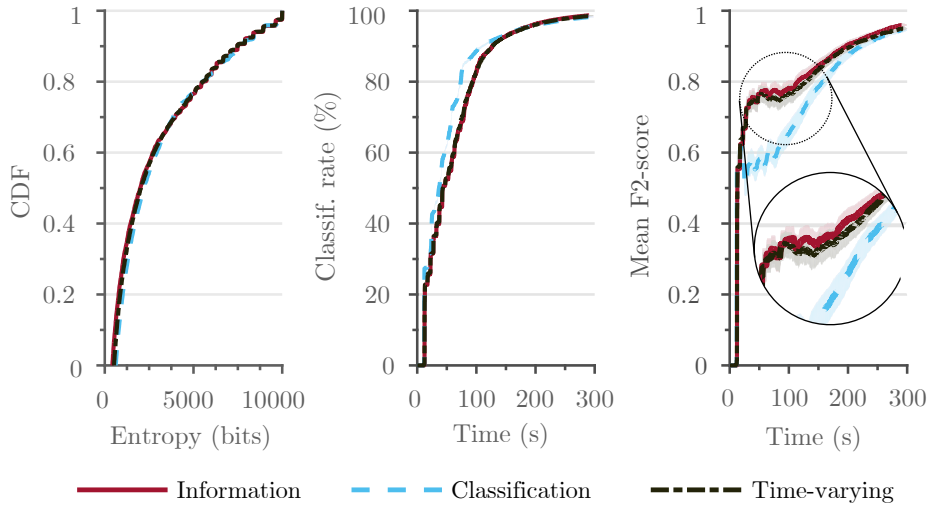

Figure 5.5: Comparison between using different utility functions to define the planning objective in the proposed approach for a fixed mission time budget of 300 s. All variants of the method optimize trajectories with the "global" CMA-ES strategy. For each algorithm, average results over 100 trials are shown with 90% confidence bounds. Using a time-varying objective enables trading off between information and classification gain during the mission to obtain maps with both low uncertainty (left) and high classification accuracy (middle, right).

The first set of experiments concerns the utility function $I(\cdot)$ in the planning procedure, which is used to evaluate the information value of candidate paths (Section 4.3). A comparison is performed between using different utility functions to obtain the initial trajectory solution in the discrete grid search. Subsequently, the results are optimized using the "global" CMA-ES approach based on the same objective. Figure 5.5 reports on the results. The curves illustrate the intuitive trade-off between area coverage and sensor noise: using the classification objective (light blue), flying at lower altitudes quickly produces accurate maps with the grid cells classified within occupancy thresholds (middle), as higher-quality images can be obtained from the camera. However, the reduction of map uncertainty (left) is limited. Moreover, initial accuracy (right) is poor since distinguishing between the two binary states (occupied and free) is not considered. In contrast, by accounting for the mission time when selecting the planning objective (black), one balances between exploiting the known model of the environment and exploring its unknown areas to quickly obtain maps that are both highly certain and correctly classified. This finding underpins the benefit of using a utility function that can vary dynamically during the mission, motivating further work in this direction.

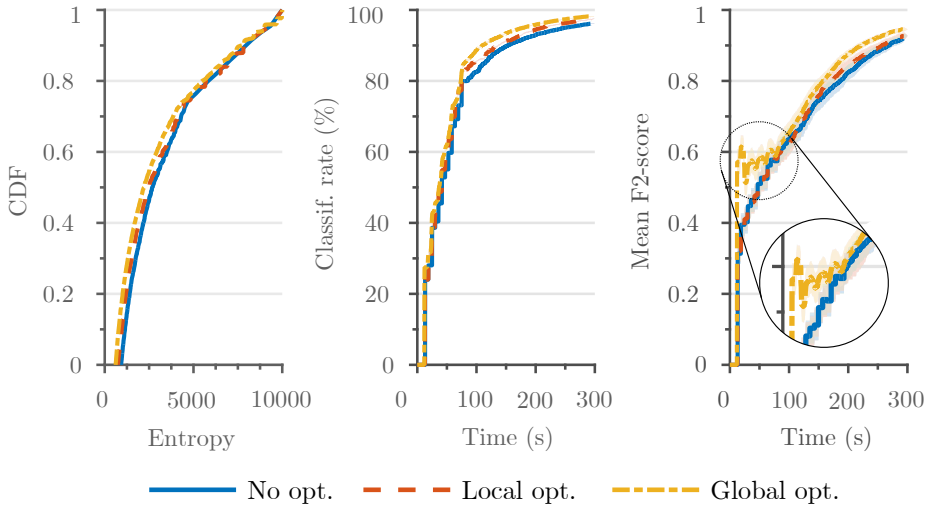

Figure 5.6: Comparison between different optimization methods using the CMA-ES in the proposed approach for a fixed mission time budget of 300 s. All variants of the method consider classification gain as the utility function for informative planning. For each algorithm, average results over 100 trials are shown with 90% confidence bounds. The "global" strategy performs significantly better as it allows for modifying the entire trajectory during the optimization procedure.

The second set of experiments concerns the second step of the replanning procedure (Section 4.5). A comparison is performed between using different approaches for optimizing informative trajectories with the CMA-ES, given the asynchronous sensor model. Figure 5.6 summarizes the experimental findings using the classification-based utility function (Equation 4.7). Although all three utility functions produced similar results, those with the classification objective were found to yield the most obvious differences between optimization methods, and are therefore included for purposes of discussion. As expected, the trajectory obtained directly from the grid search (blue) scores worst against all three evaluation criteria. This baseline result essentially corresponds to a coarse discrete solution with no steps taken to refine it further in continuous space. In contrast, using the "global" CMA-ES approach (yellow) to optimize for all control waypoints leads to the best performance. A likely reason for this is that, because the entire trajectory is passed as an optimization variable, this strategy offers the highest number of Degrees of Freedom (DoFs) for possible improvement. Subsequently, it yields highest performing results when given enough time to find a solution. Following this intuition, optimizing "locally" (orange) only for points between trajectory segments produces plans that are slightly less efficient at information gathering, but still perform better than those output from the raw grid search.

5.2.2 Continuous Variable Mapping

These experiments investigate monitoring a continuous variable on a terrain using the GP-based mapping method (introduced in Section 3.2) and a constant-frequency sensor for data acquisition (Section 4.2.2). As in the previous sub-section, benchmarking experiments are first presented. Then, detailed studies are performed to investigate the effects of using different optimization routines (Section 4.5) and evaluate the proposed adaptive scheme for adaptive replanning (Section 4.3.2).

Comparison Against Planning Benchmarks

The aim of these experiments is to evaluate the integrated system by comparing it against existing strategies. The experimental methods used in this sub-section proceed similarly to those in Section 5.2.1.

The simulations consider information gathering in a $30\,\text{m} \times 30\,\text{m}$ area. The continuous target distributions are generated as 2-D Gaussian random fields, with the mapped scalar parameter in the field ranging from 0 % to 100 %. To investigate a wide range of environments, the cluster radii of the field distributions are randomly set between 1 m and 3 m. For environmental mapping with the proposed GP-based representation, a resolution of 0.75 m is set to define uniform grids for both the training $\mathbf{X}$ and (predictive) query $\mathbf{X}_*$ data points. Assuming that the area is *a priori* unknown, the map is initialized uniformly with an uninformed mean prior of 50 %.

The Matérn family of functions is employed to capture spatial correlations within the GP model. This kernel is popularly applied in spatial and geostatistical analyses (M. Jadidi, 2017; O'Callaghan et al., 2012 and Vidal-Calleja et al., 2014) and was chosen for its ability to estimate smooth transitions between the clusters in the target distributions. It has the general form (Rasmussen et al., 2006):

$$k_{Mat}(\mathbf{x}, \mathbf{x}') \triangleq \sigma_f^2 \frac{2^{1-v}}{\Gamma(v)} \left(\frac{\sqrt{2v}d}{l} \right)^v K_v \left(\frac{\sqrt{2v}d}{l} \right), \tag{5.1}$$

where d is the Euclidean distance between input points $\mathbf{x}$ and $\mathbf{x}'$, and l and σ_f^2 are the hyperparameters of the covariance function representing the characteristic length scale and signal variance, respectively. $\Gamma(\cdot)$ is the Gamma function, $K_v(\cdot)$ is the modified Bessel function of the second kind of order v, and v is a positive parameter that determines the smoothness of the covariance.

In particular, this work applies the widely-used Matérn covariance function with $v = 3/2$:

$$k_{Mat3}(\mathbf{x},\, \mathbf{x}') \triangleq \sigma_f^2 \left(1 + \frac{\sqrt{3}d}{l} \right) \exp\left(-\frac{\sqrt{3}d}{l} \right). \tag{5.2}$$

The hyperparameter set $\theta \triangleq \{\sigma_n^2,\, \sigma_f^2,\, l\} = \{1.42,\, 1.82,\, 3.67\}$ was learned by minimizing the Negative Log of the Marginal Likelihood (NLML) in Equation 3.8 using the Polack-Ribiére conjugate gradient method. The training procedure in this work considered 4 independently generated field maps with variances modified to cover the entire range of the target parameter during inference.

To incorporate new visual imagery for online map building, Gaussian noise is simulated based on the camera-based sensor model shown in Figure 3.2, including a 10 m altitude beyond which the images scale by a factor of $s_f = 0.5$. This places a realistic limit on the quality of data that can be obtained from higher altitudes, and demonstrates the ability of the mapping method to accommodate inputs at different resolutions. The downward-facing sensor considers a square camera footprint with a 60° FoV and a constant measurement frequency of 0.15 Hz. As in the previous sub-section, the experimental setup assumes perfect state estimation and trajectory following free from any actuation and localization noise.

The proposed approach is compared against three different planning strategies: (a) a traditional "lawnmower" coverage pattern; (b) the RIG-tree algorithm, introduced by Hollinger and Sukhatme (2014) and described in Section 5.2.1; and (c) naïve random waypoint selection in the UAV workspace. Each method was tested over 30 trials with a fixed mission time budget of $B = 200\,\text{s}$.

The algorithms are evaluated using different information metrics relevant for the problem of mapping a continuous field distribution. Following the ideas in Section 4.3.2, the uncertainty in the field map

is quantified using the A-optimal information measure (Fedorov, 1972) as the trace of the covariance matrix of the GP model $\mathrm{Tr}(\mathbf{P})$. To compare performance in terms of accuracy, the total Root Mean Squared Error (RMSE) and Mean Log Loss (MLL) at the query points $\mathbf{X}_*$ in the GP are calculated with respect to the simulated ground truth field maps. As described by Marchant and Ramos (2014), the MLL is a probabilistic confidence measure that incorporates not only the prediction error, but also its associated uncertainty. It is defined by the expression:

$$MLL = \frac{\sum_{i=1}^{n_*} \left(\frac{1}{2} \log(2\pi\sigma^2(\mathbf{x}_{*i})) + \frac{(\mu_{gt}(\mathbf{x}_{*i}) - \mu(\mathbf{x}_{*i}))^2}{2\sigma^2(\mathbf{x}_{*i})} \right)}{n_*}, \tag{5.3}$$

where $\mu(\mathbf{x}_{*i})$ and $\sigma^2(\mathbf{x}_{*i})$ are the estimated mean and variance at the query point $\mathbf{x}_{*i}$ in $\mathbf{X}_*$ in the GP, respectively, and $\mu_{gt}(\mathbf{x}_{*i})$ indicates the corresponding ground truth value at that location. Note that, in this setup, $\mathbf{X}_* = \mathbf{X}$ such that Equation 5.3 is computed directly from $\mathbf{X}$ upon fusing new data, and without the need of an additional regression step.

Intuitively, all three metrics above are expected to decrease as measurements are acquired during a mission, with more rapid reductions marking greater efficiency, i.e., faster map completion, and thus better performance.

For all methods, the UAV starting position is specified as $(7.5\,\mathrm{m}, 7.5\,\mathrm{m})$ within the environment with an altitude of $8.66\,\mathrm{m}$. This particular point was chosen to assert the same initial conditions as those required to fully execute the pattern for complete coverage. For trajectory generation, the maximum reference velocity and acceleration are $5\,\mathrm{m/s}$ and $2\,\mathrm{m/s^2}$ using polynomials of order $k = 12$. The maximum number of measurements along a path is set to 10 in order to limit the computational complexity of the field map update predictions during replanning. In the proposed approach, polynomials paths are defined by $N = 5$ control waypoints and the 30-point lattice depicted in Figure 4.2(b) is employed for the initial 3-D grid search. For informative trajectory planning in continuous space, the CMA-ES optimization routine is applied with initial step sizes $\boldsymbol{\sigma}_{\mathrm{CMA}}$ of $(3\,\mathrm{m}, 3\,\mathrm{m}, 4\,\mathrm{m})$ in the (x, y, z) workspace co-ordinates, where the z-axis defines the UAV altitude, a population size of $\lambda_{\mathrm{CMA}} = 12$ according to the recommendation of Hansen (2006), and a maximum number of iterations of $N_{\mathrm{CMA}} = 45$ set based on the available computational resources. The selection of $\boldsymbol{\sigma}_{\mathrm{CMA}}$ is evaluated and discussed further in the following sub-section.

In the RIG-tree method, the control waypoints are associated with vertices, and trajectory polynomials are formed by tracing the parents of leaf vertices to the root of the current tree. For both informative planning strategies, the utility function $\mathrm{I}(\cdot)$ is set to maximize the reduction of total map uncertainty, as measured by $\mathrm{Tr}(\mathbf{P})$ (Equation 4.11). The objective also features a requirement for adaptive replanning, by considering areas with higher values of the target parameter as being interesting for mapping. A base threshold of $\mu_{th} = 40\%$ in Equation 4.13 with a scaling factor of $\beta = 3$ specifies the lower limit of the interesting value range. As an aside, note that the way in which paths are defined here permits the algorithms to be evaluated more fairly with respect to the approach in Section 5.2.1, since both informative planners are able to generate smooth trajectories.

A online variant of RIG-tree is developed to enable a direct comparison against the proposed framework. As explained in Section 5.2.1, this algorithm alternates between tree construction and plan execution in a finite-horizon fashion. The branch expansion step size in the replanning procedure is set to $10\,\mathrm{m}$. This parameter essentially controls the diffusion of the tree as it grows in the UAV workspace. Given that there is no standardized approach to tuning its value, it was chosen for best performance based on multiple empirical trials.

For the coverage-based planner, altitude ($8.66\,\mathrm{m}$) and maximum velocity ($0.78\,\mathrm{m/s}$) are defined for complete coverage given the specified mission time budget and constant measurement frequency

of the sensor. This benchmark was designed using the same method as for the experiments in Section 5.2.1: the altitude of the "lawnmower" was selected based on an evaluation of different discrete patterns covering the entire square area. Finally, using the random planning strategy, a destination is randomly sampled in the bounded workspace above the terrain. This point is then simply connected to the current UAV position to generate a trajectory, with the measurements along it taken at the specified sensor frequency.

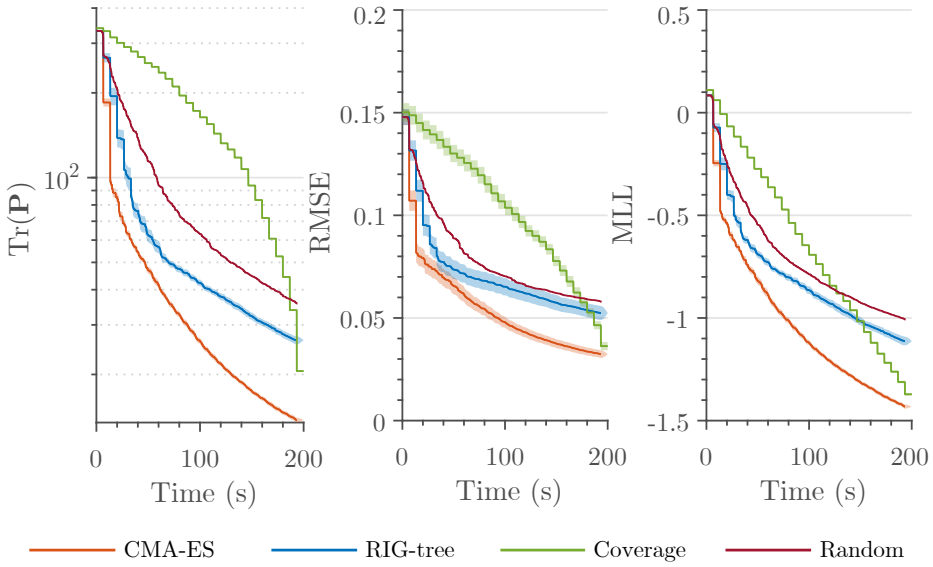

Figure 5.7: Comparison of the proposed informative planning framework using the CMA-ES against benchmarks for a fixed mission time budget of 200 s. For each algorithm, average results over 30 trials are shown with 95% confidence bounds. Using informative strategies, map uncertainty (left) and error (middle, right) reduce quickly by allowing the UAV to acquire high-altitude, low-uncertainty images before descending to improve the field map reconstruction. Note the logarithmic scale of the Tr($\mathbf{P}$) axis.

Figure 5.7 shows how the informative metrics evolve for each tested planner over the 30 trials. First, the results validate the performance of the coverage-based planner (green) as being consistent with the findings presented in Section 5.2.1. As expected, total map uncertainty (left) reduces at a uniform rate with the constant altitude and velocity of the "lawnmower" pattern. This motivates approaches based on informative planning, which are not limited to a fixed altitude and can instead compromise between area coverage and sensor noise to achieve greater efficiency. Figure 5.2 and Figure 5.4 visualize this behavior in example missions. Moreover, both the proposed algorithm (light orange) and RIG-tree (blue) perform better than the random benchmark (dark red), as the latter does not exploit a utility function to guide the selection of waypoint destinations.

As in Section 5.2.1, the proposed approach produces field maps lower in both uncertainty (left) and error (middle, right) than those of RIG-tree given the same budget on mission time. These results also mirror the studies in Section 5.2.1. Taken together, the evidence suggests that the new two-stage planner is more effective than the tree-growing strategy for terrain mapping applications targeting both discrete and continuous variables. In the experiments, the non-adaptive step size of RIG-tree was identified as a major drawback of this algorithm. Generally, it was found that higher values of this parameter allow the UAV to ascend initially, but limit incremental navigation when later refining the environmental map. Conversely, lower values enable very precise mapping, e.g., of the

corners of the environment or any detected areas of interest, but restrict the range of exploration at the beginning of the mission.

Using the simulation setup described above, further experiments were conducted to compare the proposed approach against the traditional coverage-based planning strategy. The aim is to examine the benefits of informative planning for missions of different durations more precisely. To this end, the experiments considered six different budgets $B = (100\,\text{s}, 200\,\text{s}, \ldots, 600\,\text{s})$ on mission time. For each budget, the CMA-ES-based algorithm was tested over 10 trials, giving a total of 60 simulations, and the coverage planner was run once with its deterministic path. As detailed above, for a fair evaluation, the fixed coverage altitude for each mission time was chosen for best performance among different complete "lawnmower" patterns.

The evaluation results are presented in Figure 5.8. As an example, the plots in the top row visualize the trajectories executed by both planning strategies in 200 s missions. Note that this time budget is the same as the one used in the evaluation experiments above. In general, the comparison further testifies that an informative strategy can exploit variable altitudes for data collection. The box plot on the bottom-left provides a quantitative analysis of the final uncertainties achieved in the field map, measured using $\text{Tr}(\mathbf{P})$. For comparison purposes, the results obtained using the proposed approach are normalized with the corresponding values for the coverage planner. Hence, for a given budget, percentages below 100% (orange line) mark a better performance of the new informative method. The graph on the bottom-right offers a comparison between the mission times required by the two methods to produce field maps with the same final uncertainties. In this study, the CMA-ES is given a fixed time budget, such that values above the orange line indicate a longer required time to construct a map of the same quality (worse relative performance). The motivation is to determine the amount of time savings achievable by using an intelligent planning strategy for data acquisition instead of following a traditional approach.

Figure 5.8 illustrates two important advantages of using the informative planning strategy to collect data. First, field maps are obtained with lower final uncertainty (bottom-left) as the flight altitude of the UAV is not fixed. This finding is reflected in the trajectory plots and is also in line with the discussions above. Second, significant flight time is saved (bottom-right) as lower-quality data can be collected in the early stages of the mission. This finding is consistent with the results shown in Figures 5.3 and 5.7. As a result, with increasing mission time, the marginal discrepancy between the final map reconstructions increases.

Considering the line graph in Figure 5.8, it can be observed that the proposed algorithm requires substantially less time to produce field maps with the same final uncertainty when compared against the coverage-based method. The evaluation reveals remarkable $> 50\%$ time savings for missions lasting longer than 500 s. This rapid jump occurs because the "zig-zag" pattern for the "lawnmower" in these missions must be fixed at a lower altitude to obtain reduced sensor noise for a higher achievable map quality. Specifically, the altitude required for longer missions is 6.5 m (four transversals across the terrain) compared with 8.66 m (three transversals) for shorter ones. The distance traveled by the UAV during the extra transversal has a significant cost on flight time and subsequent resources, e.g., it requires longer battery life.

Interestingly, the coverage-based planner performs better only in a 100 s mission. This occurs due to the fact that, with this budget, the informative planner lacks the time to allow the UAV to descend and refine the map with low-altitude, high-quality data. Although it is worth mentioning, that this path budget is too short to consider for a realistic mission from a practical point of view.

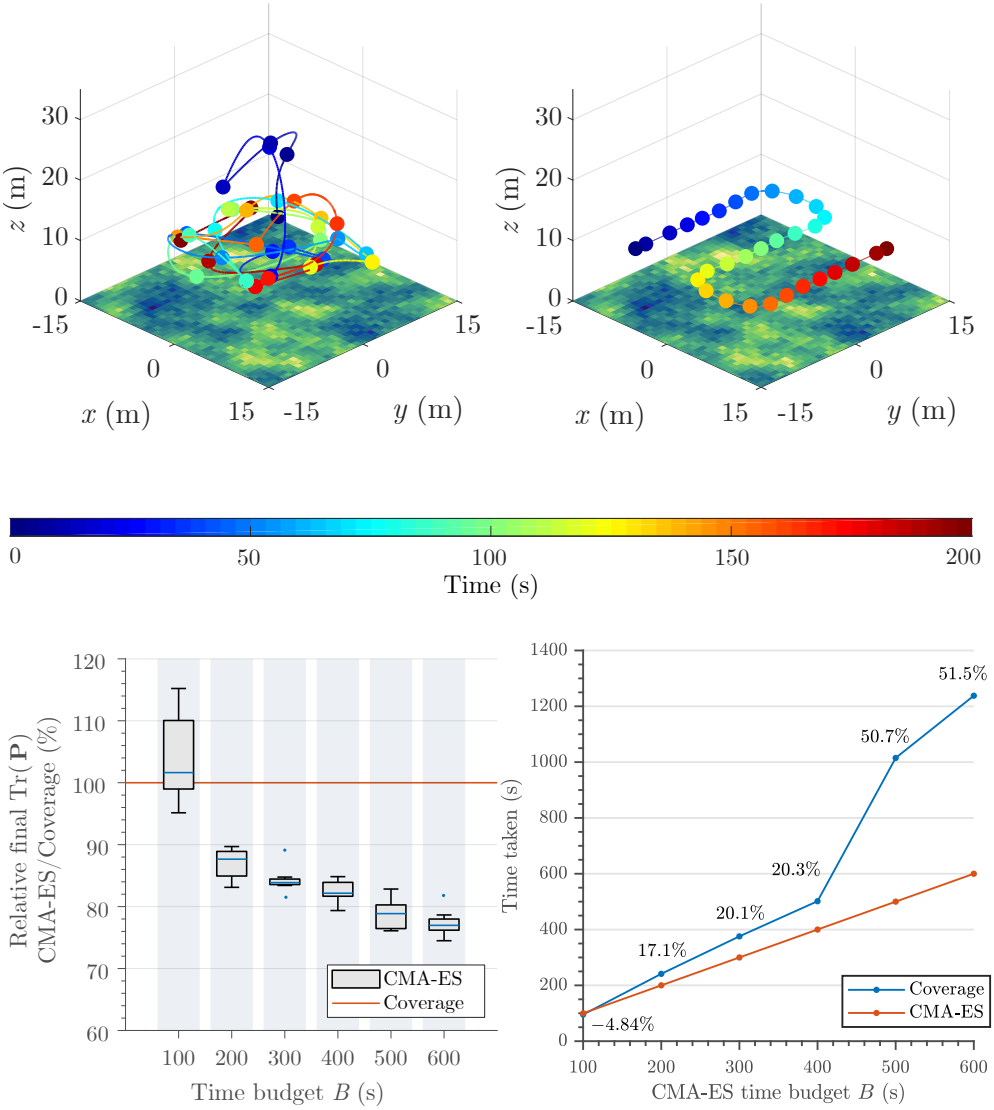

Figure 5.8: Comparison between the proposed informative planning approach using the CMA-ES and a coverage-based "lawnmower" strategy. *Top*: Visualizations of the trajectories traveled by the UAV (colored lines) using each method (left: informative planner, right: coverage path) in a 200 s mission. The spheres indicate measurement sites and the ground truth maps are rendered. *Bottom-left*: Comparison of the final map uncertainties $\mathrm{Tr}(\mathbf{P})$ for various time budgets B. Ten CMA-ES trials were run for each budget. *Bottom-right*: Comparison of times taken to achieve the same final map uncertainty, given a fixed time budget for the CMA-ES. The orange line corresponds to average values over 10 CMA-ES trials, and relative time savings using the new method are shown. By allowing for altitude variations, the proposed approach trades off between area coverage and sensor noise to quickly obtain high-confidence maps with finer quality in the same time period.

Comparison of Optimization Methods

The following experiments examine the effects of using different optimization routines on the output of the discrete grid search. The main question is: how does the choice of optimization method influence the performance of the informative planning strategy? Through this investigation, the aim is to evaluate how the CMA-ES routine, presented in Section 4.5, performs against alternative approaches. The experimental setup considers a UAV monitoring a continuous variable with the same simulation parameters as described in the previous sub-section.

These experiments assess several optimization routines commonly applied for solving problems with complex, nonlinear objective spaces. Specifically, the evaluation considers:

(a) *Lattice*: without applying an optimization strategy (i.e., grid search only; omitting Line 7 in Algorithms 1 and 2);

(b) *CMA-ES*: global gradient-free optimization based on the concepts of evolutionary algorithms (Hansen, 2006 and Hansen et al., 2009), described in Section 4.5 and studied by Hitz, Galceran, et al. (2017);

(c) *Interior point*: approximate local gradient-based optimization based on the interior point technique (Byrd et al., 2006);

(d) *Simulated annealing*: global gradient-free optimization based on the physical cooling process in metallurgy (Ingber et al., 1992), studied by Vivaldini, Martinelli, et al. (2018); and

(e) *Bayesian Optimization (BO)*: global gradient-based optimization using a GP-based process model (Gelbart et al., 2014), studied by Marchant and Ramos (2012); Marchant and Ramos (2014) and Morere et al. (2017).

Some example applications of the algorithms in the context of active sensing are provided above. For further details, the interested reader is pointed to the works cited and references therein.

Simulations for each algorithm were repeated over 30 trials in environments with randomly varying field map distributions. The experiments were performed in MATLAB on Ubuntu Linux 16.04 LTS running on a single desktop computer with a 1.8 GHz Intel i7-8550U processor and 16 GB of RAM. Standard baseline implementations were considered for each approach, with approximately the same amount of $\sim$ 20 s optimization time allocated for each. Note that the benchmarks were applied without significant effort invested into adjusting their parameters in order to make them comparable with the CMA-ES, which does not require significant tuning procedures.

The implementation details for each strategy are as follows. The inbuilt `fmincon` function is used for the local IP optimizer. Using the iterative step-wise IP algorithm, described by Byrd et al. (2006), Hessian matrices are approximated by a dense quasi-Newton strategy. For the SA method, the `simulannealbnd` function is used with an exponential cooling schedule and an initial temperature of 100. The `bayesopt` function is applied for the BO strategy. In this framework, a crucial design aspect is selecting a suitable acquisition function to manage the exploration-exploitation trade-off during the optimization procedure. Essentially, this function guides the selection of where to evaluate the objective function next based on a probabilistic belief. The BO variant in these experiments considers the commonly used time-weighted Expected Improvement criterion with an exploration ratio of 0.5 (Močkus, 1975 and Gelbart et al., 2014).

For the proposed approach, the implementation is based on the `libcmaes` library[1] for MATLAB. Additionally, two variations of the CMA-ES are investigated with initial step sizes $\boldsymbol{\sigma}_{\text{CMA}}$ of $(3\,\text{m}, 3\,\text{m}, 4\,\text{m})$

[1] `cma.gforge.inria.fr/cmaes_sourcecode_page.html`

and $(10\,\text{m}, 10\,\text{m}, 12\,\text{m})$ in the (x, y, z) workspace co-ordinates, where the z-axis defines the UAV altitude. The main idea behind this is to compare different global search behaviors, as the step size parameter effectively captures the distribution from which new solutions are sampled, and thus how well the problem domain is covered by the optimization routine.

Method	Tr($\mathbf{P}$)	RMSE	WRMSE	MLL	WMLL
Lattice	56.193	0.0624	0.0622	-0.880	-0.881
CMA-ES $(3, 4)$	**46.780**	**0.0541**	**0.0536**	**−0.976**	**−0.981**
CMA-ES $(10, 12)$	50.991	0.0599	0.0596	-0.897	-0.900
IP	51.628	0.0575	0.0574	-0.918	-0.919
SA	55.868	0.0599	0.0595	-0.866	-0.867
BO	62.121	0.0646	0.0642	-0.805	-0.808
RIG-tree	68.581	0.0696	0.0696	-0.755	-0.757
Random	92.681	0.0773	0.0767	-0.668	-0.668
Coverage	165.121	0.0972	0.0972	-0.685	-0.688

Table 5.1: Mean information metrics for all optimization methods, averaged over 30 continuous field mapping trials. The lowest uncertainties and errors obtained with the CMA-ES justify using this strategy for global optimization in the informative path planning algorithm.

Table 5.1 displays the mean results for each method averaged over the 30 simulation trials. The benchmarks from the evaluation in Section 5.2.2 are included as references for comparison. In the table, the suffixes '$(3, 4)$' and '$(10, 12)$' denote smaller and larger values of step sizes for the CMA-ES routine, respectively. Following Marchant and Ramos (2012), the quantitative study also examines weighted versions of the performance criteria: the total Weighted Root Mean Squared Error (WRMSE) and Weighted Mean Log Loss (WMLL) of the final field map. These variations multiply the original indicators by a factor that depends on the mean of the predicted value, in order to give more importance to errors in higher-valued regions. Note that, as the same information objective is used for all optimization routines, consistent trends are observed in both non-weighted and weighted versions of the metrics. Finally, Figure 5.9 depicts the mean times required by each method to reduce Tr($\mathbf{P}$) to 75% of its initial value, which is associated with a completely unknown environment. This evaluation compares the different decay rates of total map uncertainty obtained using the tested algorithms, and indicates relative time savings in terms of achievable map quality.

In Table 5.1, comparing the raw output of the lattice grid search with the CMA-ES and IP confirms the expected result: using an optimization routine reduces both map uncertainty and error. The evaluation highlights that, with the lowest values (shown in bold in Table 5.1), the proposed CMA-ES scores best against all performance criteria. This suggests that its global optimization strategy is most effective at solving the informative path planning problem examined in this work. Using larger step sizes in the 'CMA-ES $(10, 12)$' variant (corresponding to $> 33\%$ of the extent of the UAV workspace, co-ordinate-wise) yields poorer optimization performance. Intuitively, taking larger steps during the evolutionary search can lead to large random fluctuations in the solution, which reduces the rate at which the algorithm converges. Too small step sizes, on the other hand, are expected to increase the susceptibility to local optima. This observation reflects an important consideration: for the CMA-ES, special attention is needed to select step sizes that appropriately cover the application domain without

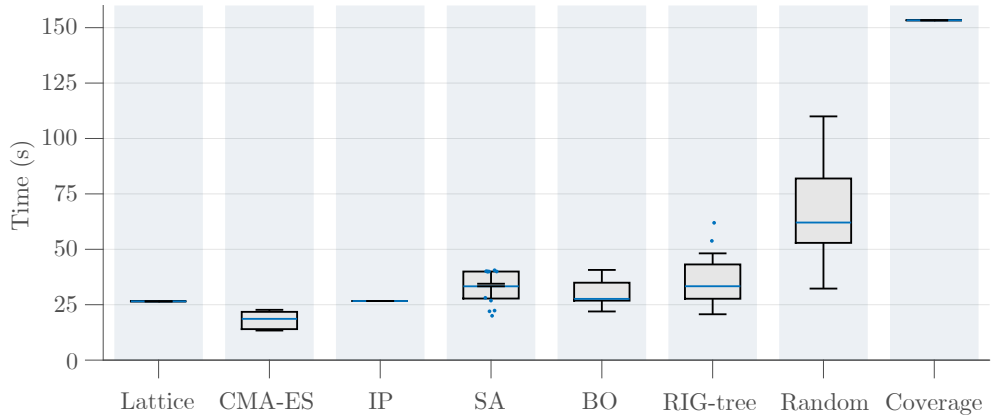

Figure 5.9: Mean times required by each optimization method to reduce field map uncertainty $\mathrm{Tr}(\mathbf{P})$ to 75% of its initial value, averaged over the 30 simulation trials. The depicted result using the CMA-ES corresponds to a search with smaller step sizes ('CMA-ES $(3, 4)$' in Table 5.1). With the lowest mean decay time of 18.1 s, this approach performs best.

overexploititative or overexploratory behavior. Similar conclusions based on empirical studies are drawn by Hitz, Galceran, et al. (2017) and Hansen (2006).

Surprisingly, applying BO produces a result poorer than one obtained by taking the solution from the lattice search directly, without applying an optimization routine at all. An explanation for this could be that the BO implementation used in the experiments was biased towards more aggressive global exploration of the objective space. In this case, one would expect candidate solutions during the search to vary erratically, leading to inefficiency. Using the CMA-ES with large initial step sizes ('CMA-ES $(10, 12)$') is expected to exhibit similar problems. In general, the performance of the BO strategy was found to be highly sensitive to the acquisition function chosen. It was therefore considered the most difficult algorithm to tune without any prior knowledge about the complex, nonlinear objective landscape.

Adaptive Replanning Evaluation

The final experiments assess the effects of using the proposed scheme for adaptive replanning in the informative planning framework. The theory relevant for this problem is introduced in Section 4.3.2.

The aim of the experiments is to evaluate the ability of the adaptive strategy to focus on specific regions of interest in different observed environments. The simulation setup and parameters are the same as in the previous two sub-sections. The experiments consider two types of monitoring scenario with different $30\,\mathrm{m} \times 30\,\mathrm{m}$ field distributions: (a) 'Split', handcrafted maps where the interesting area is well-defined by design, and (b) 'Gaussian', the uniformly distributed maps employed for the evaluation in Section 5.2.2. 'Split' fields are partitioned spatially such that half of the cells in the map are classified as being interesting based on the uncertainty-bounded criterion in Equation 4.13 with a base threshold of $\mu_{th} = 40\%$ and scaling factor of $\beta = 3$.

To evaluate the gain of replanning online, the proposed adaptive approach for informative planning is compared against itself without an interest-based objective. This benchmark considers the utility function for exploration defined in Equation 4.11, i.e., the value of information obtained from all locations in the map is treated equally, without discriminating the upper value range. As in the previous experiments, the two algorithms are run over 30 trials in each type of environment, giving a total of 120 trials.

The performance criteria considered for quantitative evaluation are: the variation of WRMSE, which places more importance on interesting areas with higher parameter values, and the uncertainty difference $\Delta\sigma^2$ between the area of interest and the rest of the total area. This indicator measures the exploiting behavior of the algorithms, and is defined by Hitz, Galceran, et al. (2017) as:

$$\Delta\sigma^2 = \frac{\bar{\sigma}^2(\mathcal{X}_-) - \bar{\sigma}^2(\mathcal{X}_I)}{\bar{\sigma}^2(\mathcal{X}_-)}, \tag{5.4}$$

where $\bar{\sigma}^2(\cdot)$ evaluates the mean variance and $\mathcal{X}_-$ and $\mathcal{X}_I$ denote the sets of uninteresting and interesting locations in the environmental field map, respectively, as defined in Section 4.3.2.

Moreover, the rate of total map uncertainty reduction in $\mathcal{X}_I$, as measured by $\mathrm{Tr}(\mathbf{P})$, is considered to investigate the relative ability of the planners to focus on the interesting regions in the field.

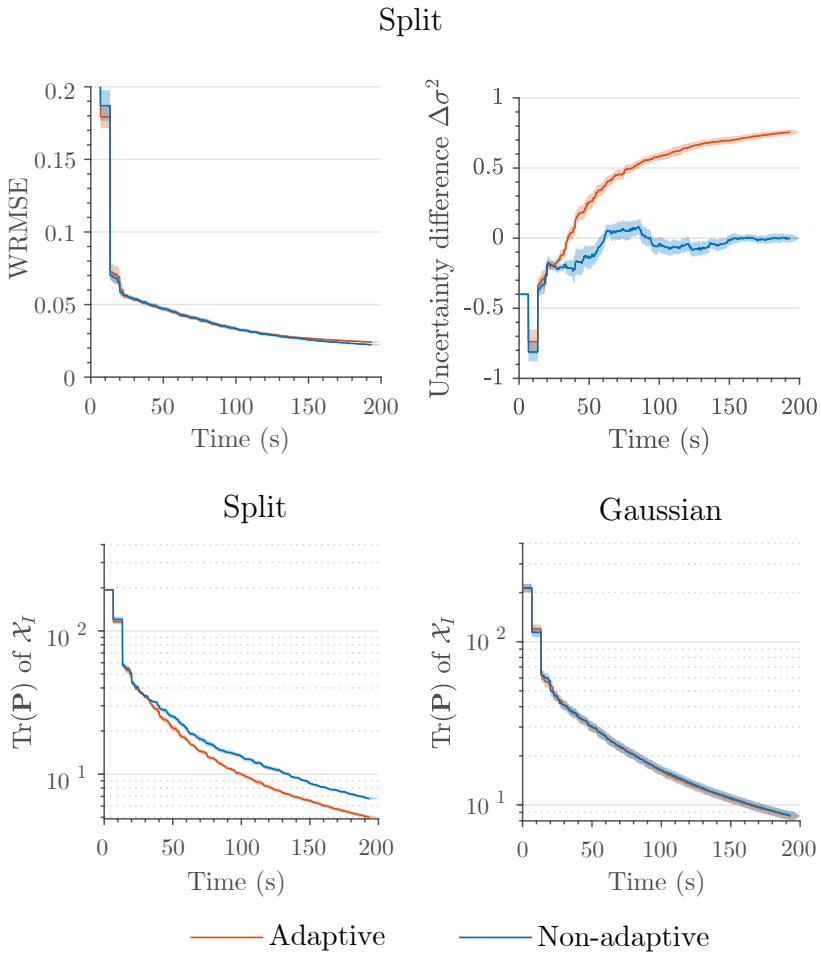

Figure 5.10: Comparison between adaptive and non-adaptive informative planning strategies in different types of continuous field mapping scenario for a fixed mission time budget of 200 s. The average results over 30 trials are shown with 95% confidence bounds. *Top*: In a 'Split' scenario, adaptivity achieves low error (left) with higher uncertainty differences (right) in interesting areas. *Bottom*: In a 'Split' scenario (left), adaptivity reduces uncertainty faster in interesting areas, while performing comparably to a standard non-adaptive approach in a 'Gaussian' scenario (right). Note the logarithmic scale of the $\mathrm{Tr}(\mathbf{P})$ axis.

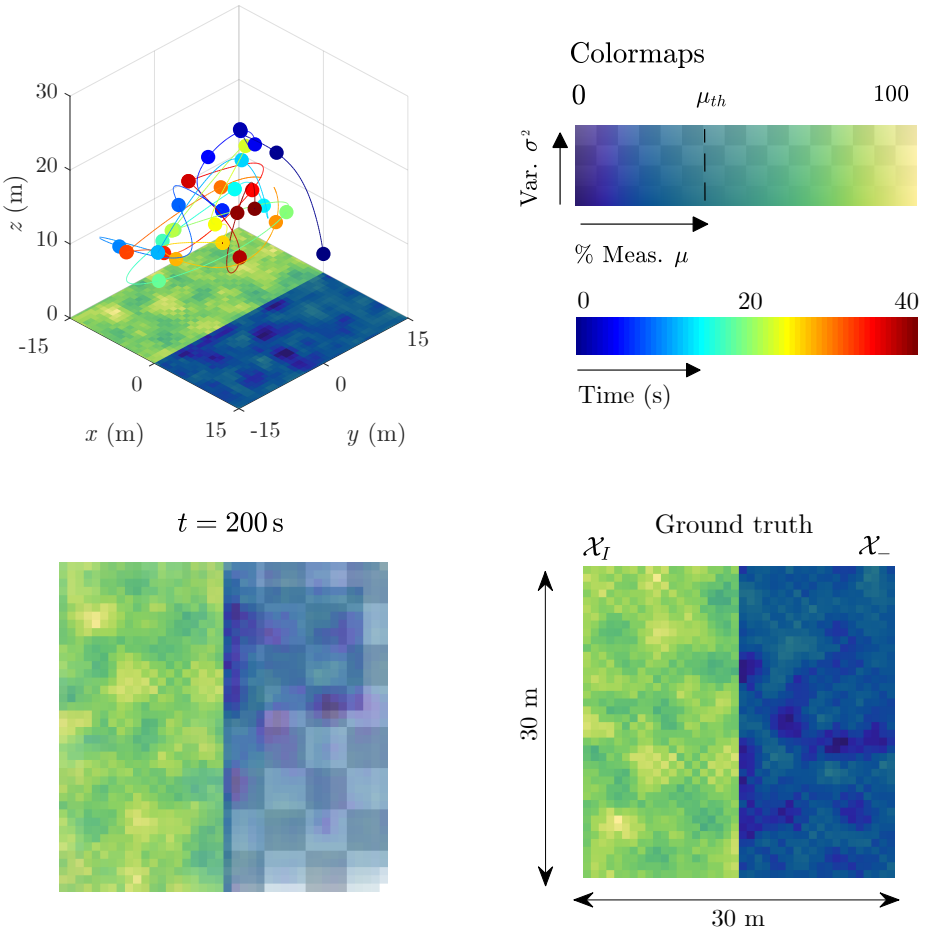

Figure 5.11: Example result of the adaptive replanning scheme to map a 'Split' field distribution in a 200 s mission. *Top*: Visualization of the trajectory traveled by the UAV (colored line). The spheres indicate measurement points and the ground truth map is rendered. The dashed line shows the threshold $\mu_{th} = 40\%$ above which map regions are considered interesting (yellower). *Bottom*: Final map reconstruction (GP mean) output by the planner (left) compared with ground truth (right). The opacity indicates model uncertainty with the checkerboard added for visual clarity. Lower opacity confirms higher certainty in the interesting area $\mathcal{X}_I$.

The results of the quantitative evaluation are reported in Figure 5.10. The plots in the top row illustrate how the algorithms perform in a 'Split' scenario using the information metrics, while those in the bottom row compare the evolution of uncertainty in interesting areas considering the two extreme types of field distribution. The trends for the 'Split' scenario underline that the relative uncertainty difference $\Delta\sigma^2$ in interesting areas increases more quickly using the adaptive replanning strategy (right), while the map WRMSE remains low (left). This confirms that, by virtue of designing its utility function, the informative planning algorithm can be carefully tailored to target certain parts of the environment, as desired. Note that both the adaptive and non-adaptive approaches perform similarly during the early phase of the mission (< 30 s) as the UAV is exploring the initially unknown area.

The bottom plots offer further insight. The observed trends suggest that the benefit of adaptivity, in terms of reducing uncertainty in areas of interest, is greater in the 'Split' type of scenario (left) when compared with the 'Gaussian' type (right). This conclusion is intuitive given the nature of the problem: since the region $\mathcal{X}_I$ is clearly distinguishable in 'Split' maps, it enables collecting data so that measurements considered to be purely informative lie within the camera FoV, given the specified planning objective. In contrast, the advantages of adaptive replanning are more limited when the field distribution is uniformly dispersed, as in 'Gaussian' maps, as there is less room to exploit the detected areas of interest.

Finally, Figure 5.11 offers a useful qualitative illustration of the proposed adaptive replanning scheme in a 'Split' scenario. Once the uninteresting (bluer) side of the map $\mathcal{X}_-$ is classified within confidence bounds, the value of future information in this area is discarded. The UAV can instead focus on taking measurements of the interesting (yellower) region $\mathcal{X}_I$ to improve the fidelity of the estimated map. This result is intuitive and consolidates the analysis above.

5.3 RIT-18 Mapping Scenario

This section demonstrates the application of the complete informative planning approach in a photorealistic terrain monitoring scenario. The aim is to validate the system for mapping using experimental data and with real-time requirements in a realistic simulation environment. Importantly, this setup also enables evaluating the system in a controlled setting with reliable ground truth data as a step towards real-world deployment.

The experiments are performed in RotorS[2], a simulation framework for UAVs based in Gazebo (Furrer et al., 2016). The ground truth terrain map for the monitoring problem is based on data from the recently published RIT-18 dataset[3] (Kemker et al., 2018), available open-source. This dataset contains high-resolution six-band visual near-infrared imagery for semantic segmentation, collected using a UAV along the coast of Lake Ontario in Hamlin, NY. For the purposes of these experiments, the surveyed region of interest is a $200\,\mathrm{m} \times 290\,\mathrm{m}$ area of the orthomosaic image taken from the RIT-18 validation fold. To reflect the nature of the dataset, which is semantically annotated, a discrete map representation is considered to represent the field environment. Figure 5.12 illustrates screenshots of the simulation setup, which runs on a single desktop computer with a 1.8 GHz Intel i7-8550U processor and 16 GB of RAM. The mapping and planning algorithms are implemented in MATLAB on Ubuntu Linux 16.04 LTS and interfaced to Robot Operating System (ROS) via the Robotics System Toolbox[4] for real-time computation.

The simulations include a AscTec Firefly hexacopter UAV[5] equipped with a downward-facing multispectral camera used to take images of the terrain. The UAV model is provided by RotorS and features its complete dynamic behavior. Based on the full state estimate feedback obtained directly from the simulation environment, the system setup uses the linear Model Predictive Control (MPC) framework of Kamel et al. (2017)[6] to track reference trajectories generated by the informative planning algorithm. The simulated camera for data acquisition has an image resolution of $360\,\mathrm{px} \times 480\,\mathrm{px}$ and a FoV of (35.4°, 47.2°) along the x- and y-axis, respectively, with reference to the vertical and horizontal directions of the terrain orthomosaic.

[2] github.com/ethz-asl/rotors_simulator
[3] github.com/rmkemker/RIT-18
[4] mathworks.com/products/robotics.html
[5] asctec.de
[6] github.com/ethz-asl/mav_control_rw

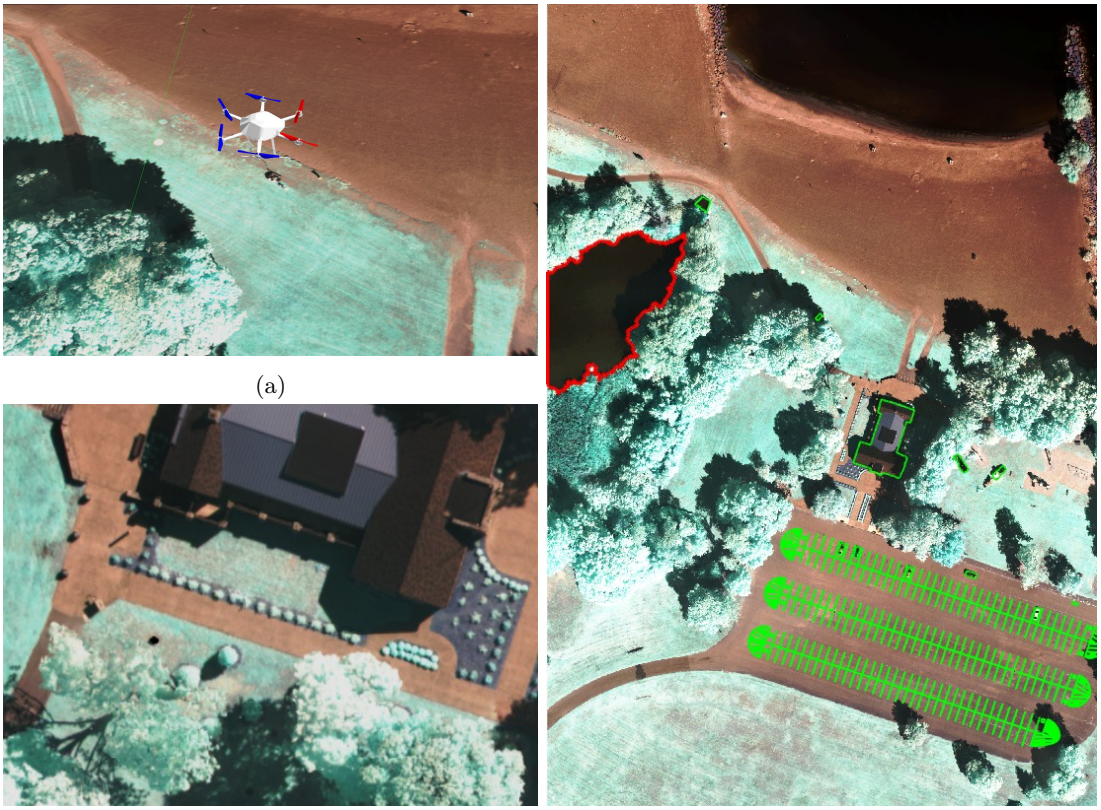

(a) (b) (c)

Figure 5.12: The simulation setup in RotorS used to validate the planning strategy in a photorealistic scenario. (a) and (b) depict the AscTec Firefly UAV and the view from its on-board camera. (c) shows an aerial view of the 200 m× 290 m surveyed area (RIT-18 validation orthomosaic). The red and green lines annotate the two classes for targeted for mapping: 'Lake' and 'BRV', respectively.

The task of mapping RIT-18 is studied as an active classification problem considering the discrete nature of its labels. To extract sensor measurements for online mapping, image semantic segmentation is performing using a deep learning framework, as is common practice in remote sensing applications (Sa, Z. Chen, et al., 2018; Kemker et al., 2018 and Carrio et al., 2017). The system setup leverages a version of the popular SegNet convolutional neural network architecture (Badrinarayanan et al., 2017), which has been modified by Sa, Z. Chen, et al. (2018) to accept multispectral as well as RGB image channels as inputs. In the experiments, the orthomosaic imagery registered from a given UAV pose is passed to the network to produce a pixel-wise segmented output.

An example of the classification procedure is shown in Figure 5.13. This figure confirms qualitatively that the network can discern most details in the dataset, i.e., the parking lot features in the depicted image. For the purposes of these experiments, the classification problem is simplified by only considering 3 class labels as targets for mapping. Derived from the 18 total labels in RIT-18, these are: (a) 'Lake'; (b) a combination of 'Building', 'Road Markings', and 'Vehicle' ('BRV'); and (c) 'Background' ('Bg'), i.e., all other labels excluding (a) and (b). These particular labels were chosen based on their distributions to obtain classification performance with a strong altitude-dependency, which is relevant for the sensor model in the informative terrain monitoring problem.

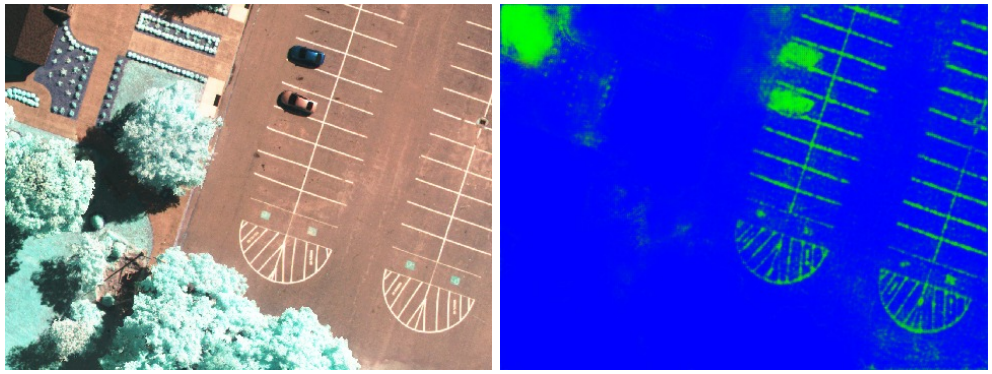

(a) (b)

Figure 5.13: Example classification result from a UAV altitude of 70 m using the modified SegNet architecture with RIT-18. (a) shows the RGB image channel input, and (b) visualizes the dense segmentation output. In (b), the probabilistic output for each class ['Lake', 'BRV', 'Bg'] is mapped to the corresponding pixel intensity on the [R, G, B] channels.

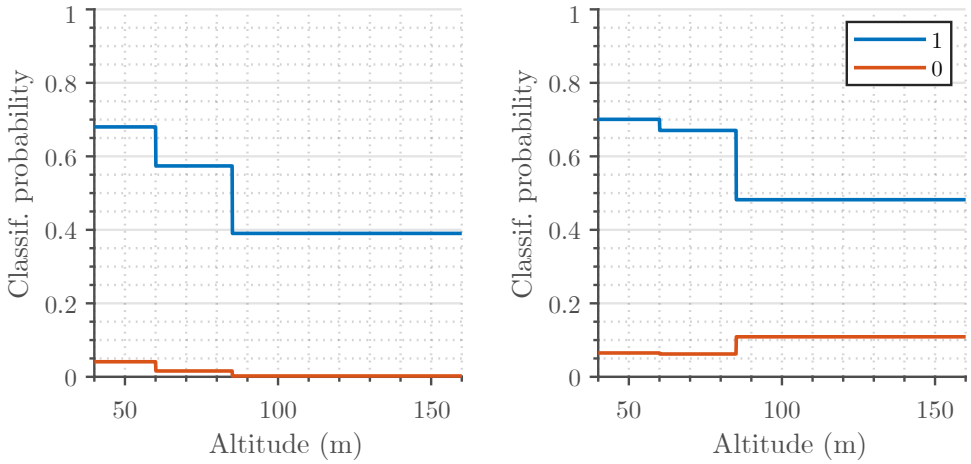

(a) (b)

Figure 5.14: Sensor models for the trained Segnet-based classifier for the (a) 'Lake' and (b) 'BRV' classes in the RIT-18 dataset. The blue and orange curves depict the probability of observing class label '1', given that the map contains '1' or '0', i.e., true and false positives, respectively. Note that the false positive probability can decrease with altitude as the classifier becomes more conservative with true output results.

A key prerequisite for informative planning is developing a reliable sensor model for online mapping and prediction. To model the modified SegNet architecture, the network was first trained on all class labels using imagery from the RIT-18 training fold. The network training procedure used (323, 72, 16) six-band images, which were simulated at 3 different altitudes, (50 m, 70 m, 100 m), in the orthomosaic according to the specifications of the multispectral camera. At altitudes of 70 m and 100 m, the training images were additionally downscaled to exaggerate the effects of pixel mixing at lower resolutions. The Caffe deep learning framework[7] (Jia et al., 2014) is employed for the network

[7] `caffe.berkeleyvision.org`

implementation, with computations running on a desktop computer with an Nvidia Titan X Pascal GPU module.

Using the trained network, classification accuracy was assessed by using RIT-18 validation fold data to compute confusion matrices at each altitude for the 3 classes of interest (30% train and 70% test split, with a higher proportion of training data considered at lower altitudes). Figure 5.14 illustrates the altitude-dependent sensor models that were derived from this analysis. The functions shown were formed piece-wise, by associating intermediate altitudes with the closest discrete performance statistics available. In general, note that the altitude range in these experiments is wider compared with those in the previous simulation studies as the area of interest for mapping is much greater.

A discrete method is used to map the target region, as described in Section 3.1. The online estimated map maintains one independent occupancy grid layer for each of the 3 classes, with the cells in each layer corresponding to the probability of being occupied by an object of that class. Each layer of the map has a uniform resolution of 5 m, and all cells are initialized with an uninformed probability of 0.5, assuming that the area is initially unknown.

The sensor models in Figure 5.14 are applied to predict new measurements during replanning, conditioned on the most likely states given the current map estimates. For fusing new data, the semantic segmentation output, e.g., as shown in Figure 5.13(b), is projected on the occupancy grids for each class, and likelihood updates are performed with the maximum pixel probabilities mapping to each cell. Note that, unlike the pixel-wise classifier output, the proposed discrete mapping strategy does not enforce the probabilities of a cell across the layers to sum to 1, as a cell may contain objects from multiple classes. This approach is contrasted with the methods of Berrio et al. (2017), who examine a semantic reconstruction problem in which only one true label is given to a particular cell.

In this experimental setup, the utility function for informative planning is to efficiently map the 'BRV' class, which would be useful, e.g., for identifying man-made features in search and rescue scenarios. The proposed informative planning approach using the CMA-ES is evaluated against the "lawnmower" coverage-based method, which is perceived as the naïve choice of algorithm in such types of application. To investigate how mapping performance varies with altitude, simulations are performed with two coverage patterns at fixed altitudes of 157 m and 104 m, denoted 'Cvge. 1' and 'Cvge. 2', respectively. In addition, and similarly to the evaluation in Section 5.2.2, both non-adaptive and adaptive versions of the proposed approach are studied in order to assess the benefits of using an adaptive strategy to map regions of interest. The performance criteria extracted to evaluate the methods quantitatively are: the variations of total map entropy (uncertainty) and RMSE with respect to the ground truth labels in RIT-18.

All methods are given an equal 400 s budget B on mission time. The computational requirements of the online network classifier are limited by assigning a constant measurement frequency of 0.1 Hz for the camera. The UAV is also allowed to stop while processing images, before new data is fused into the map. As in the previous experiments, the trajectory generation routine is performed using polynomials of order $k = 12$ with a maximum reference velocity and acceleration of 15 m/s and 20 m/s^2. For planning, candidate paths are defined by $N = 5$ control waypoints. The UAV starting position in the proposed approach is specified as (33 m, 46 m) within the bottom-left corner of the orthomosaic map with 104 m altitude, in order to achieve consistency with the lower-altitude "lawnmower" pattern ('Cvge. 2').

In the proposed approach, the initial grid search is carried out over a scaled version of the 30-point lattice shown in Figure 4.2(b), stretched to cover the UAV workspace volume above the rectangular orthomosaic image. The CMA-ES optimization routine runs with initial step sizes $\boldsymbol{\sigma}_{\text{CMA}}$ of (50 m, 60 m, 40 m). These parameters were selected based on the empirical conclusions drawn in

Section 5.2.2. To express the requirement for adaptive replanning, a low threshold of $\delta_f = 0.4$ (Equations 4.6 and 4.7) is applied on the occupancy grid layer of the target class ('BRV'). This reflects the mission aim of focusing on regions that contain objects of this class. Finally, the "lawnmower" pattern for the coverage benchmarks was designed for a fair time-based evaluation using the principles detailed in the evaluation experiments in Sections 5.2.1 and 5.2.2.

Figure 5.15 compares the performance of the planning algorithms using the information metrics. As found in Sections 5.2.1 and 5.2.2, the total uncertainty of the map reduces uniformly using the coverage-based method. The plot in the middle shows that, due to the layout of the RIT-18 orthomosaic, the "lawnmower" pattern observes the interesting areas featuring the 'BRV' class only towards the end of the mission. This reflects the fact that a simple coverage path is a naïve strategy for exploration, since it neglects the internal structure of the mapped environment. In the interesting regions, 'Cvge. 2' (dark red) achieves higher-quality mapping than 'Cvge. 1' (green) as its lower altitude allows for more accurate classifications based on the sensor models in Figure 5.14. As previously discussed, this highlights the altitude-dependent nature of the problems and motivates using informative strategies to cover all three dimensions of the UAV workspace.

Using an adaptive replanning scheme (orange), both total map uncertainty (middle) and error (right) decay rapidly in the areas of interest. Such behavior further validates the finding in Section 5.2.2 that the proposed approach is effective for mapping in adaptive scenarios. However, the curves on the left suggest that a non-adaptive strategy (blue) performs better in terms of reducing overall map uncertainty. Intuitively, without a specific objective, the approach achieves pure exploration and is therefore quickest at covering the map as a whole. Taken together, the results imply a key advantage of the framework: it can easily be tailored to balance between exploration (uniform uncertainty reduction) and exploitation (mapping a target class) in a particular scenario, as desired.

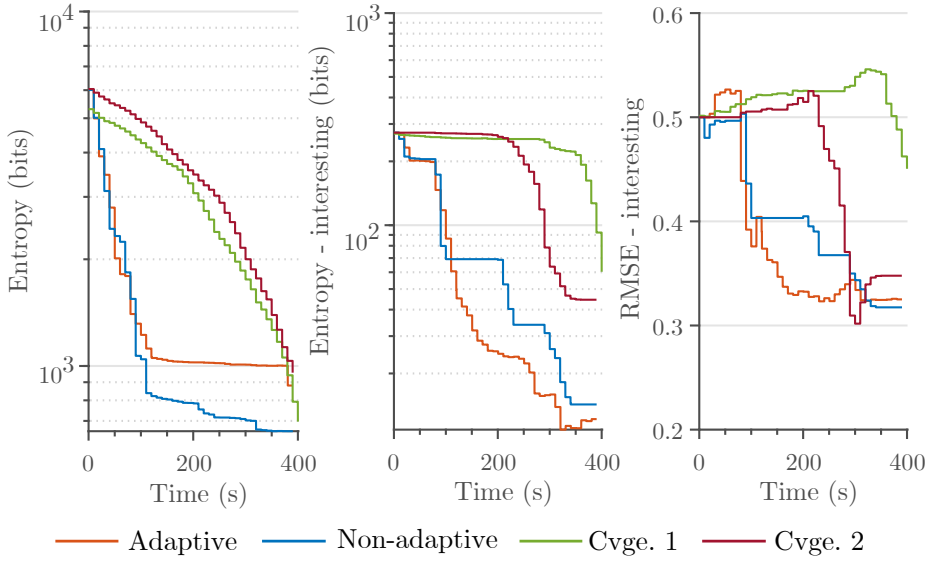

Figure 5.15: Comparison of the proposed informative approach using the CMA-ES against fixed-altitude coverage benchmarks ('Cvge. 1' = 157 m, 'Cvge. 2' = 104 m) in a 400 s photorealistic mapping scenario. By planning adaptively, map uncertainty (middle) and error (right) in interesting areas ('BRV' class of RIT-18) reduce most rapidly, while yielding lowest overall map uncertainty (left). Note the logarithmic scale of the Entropy axis.

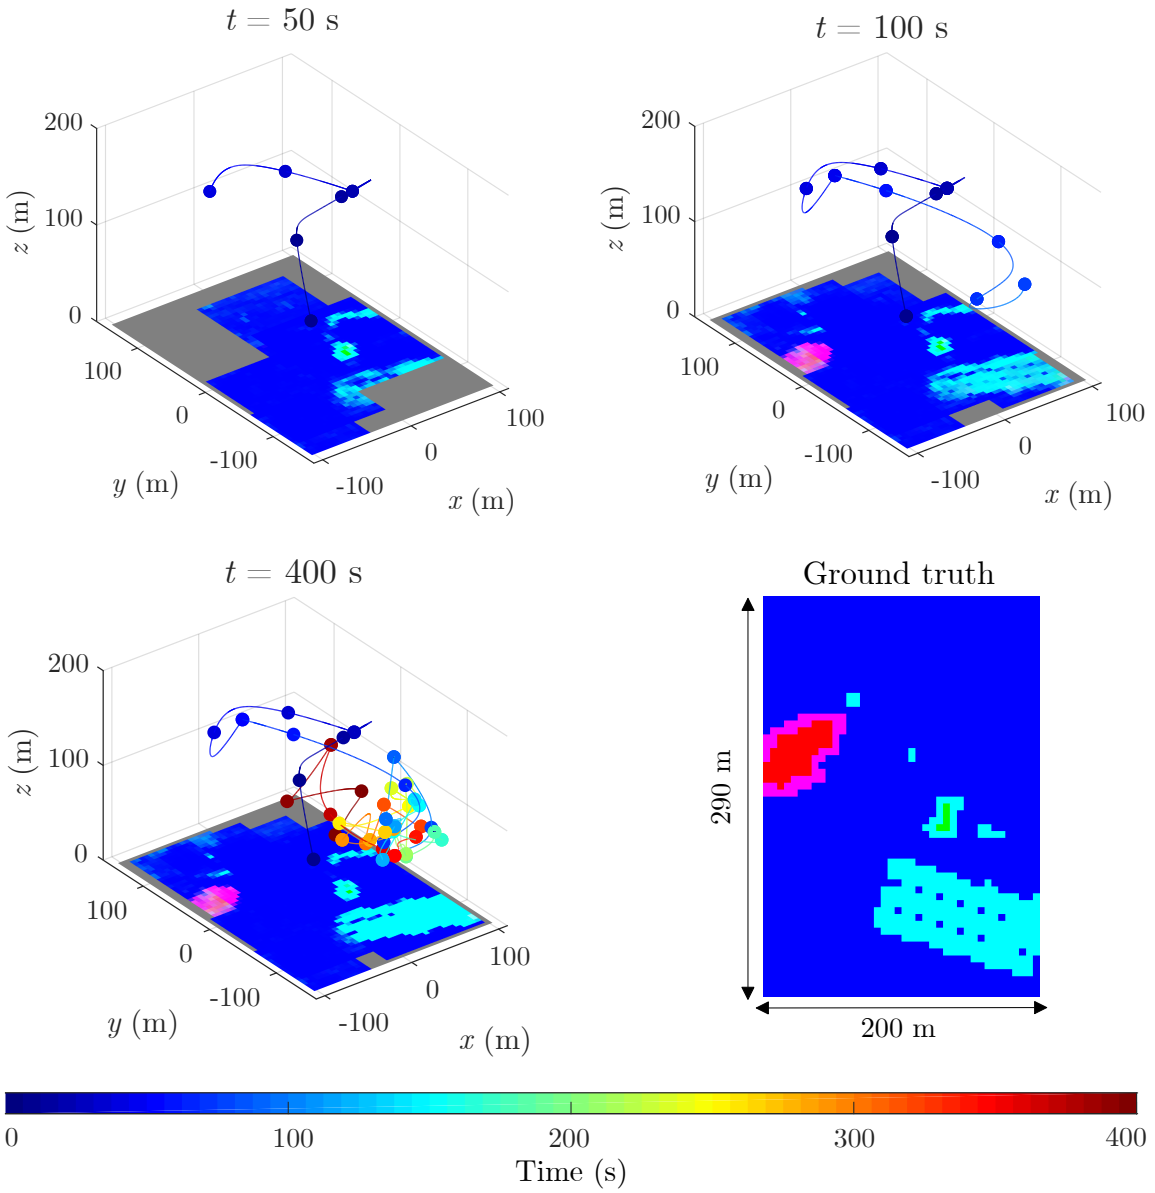

Figure 5.16: Visualization of the trajectory traveled by the UAV (colored line) using an adaptive planning strategy in the 400 s terrain monitoring mission. The three plots depict different snapshots of the mission at times $t = 50$ s, 100 s, and 400 s. The RIT-18 ground truth is shown on the bottom-right. In the trajectory plots, the spheres indicate measurement sites and the current occupancy maps are rendered. The colors portray composites of the three-layer grid-based map representation, with the cell probabilities for each class ['Lake', 'BRV', 'Bg'] mapped to the corresponding intensities on the [R, G, B] channels. Gray indicates unobserved space. The sequence shows that the planner quickly explores the area to later focus on more closely mapping interesting regions ('BRV' class). Note that the magenta and cyan cells indicate the presence of two classes, ['Bg', 'Lake'] and ['Bg', 'BRV'], respectively.

For a qualitative study, Figure 5.16 illustrates the trajectory traveled by the proposed adaptive planner in the 400 s mission. It can be observed that, at the start of the mission (< 100 s), the UAV uniformly explores the unknown terrain. Then, once they have been discovered, it concentrates on gathering information in regions with a high probability of containing objects from the 'BRV' class. In Figure 5.16, these correspond to green cells, or cyan for cells that contain both the 'BRV' and 'Bg' classes. The final visualization at time $t = 400$ s shows that the fidelity of the predicted map in these areas is improved as low-altitude measurements are accumulated. Note that the two small cars to the right of the building and above the parking lot (visible in Figure 5.12(c)) are mapped incorrectly as the SegNet model is limited in segmenting out very fine details given the data it was trained on. Considering the richness of the RIT-18 dataset in terms of both resolution and labeled object classes, an interesting direction for future work is to explore using different segmentation approaches and target classes for mapping using informative planning methods.

In summary of this section, the experimental findings presented demonstrated the successful application of the proposed informative planning approach in a realistic terrain monitoring scenario and its practical ability to identify and focus on specific areas of interest.

5.4 Experimental Results

This section presents a series of experiments showing the proposed informative planning framework deployed on physical UAV platforms. The aim is to validate the real-world efficacy of the complete system in different terrain monitoring applications. The first two experiments demonstrate proofs of concepts in discrete and continuous mapping scenarios. Then, the algorithms are implemented in an agricultural field deployment using a UAV to capture the spread of vegetation on a farm. All tests presented are fully autonomously executed and running online with real-time requirements.

5.4.1 Target Mapping

This set of experiments demonstrates the application of the proposed framework on a UAV performing a target mapping task in an indoor environment. The aim is to validate the aspects of the approach relevant for a discrete monitoring scenario (Section 5.2.1) in real-world conditions.

The goal of the experiments is to reconstruct a distribution of AR tags, which are representative of practical targets, e.g., weeds on a farmland or victims on a disaster site. Various tag arrangements are studied to investigate how the proposed method behaves in different types of scenario. The experiments are performed using an AscTec Pelican quadrotor UAV platform in an empty indoor environment with a volume of $4\,\text{m} \times 4\,\text{m} \times 3\,\text{m}$ defining its workspace. As depicted in Figure 5.17(a), $6.4\,\text{cm} \times 6.4\,\text{cm}$ AR tags are securely placed flat on the ground as targets for mapping. State estimation is provided by an external motion capture system (Vicon) and planned trajectories are tracked using MPC (Kamel et al., 2017). Mapping, high-level control and informative planning modules run on an on-board computer with a 3.2 GHz Intel NUC i7-5557U processor, 16 GB of RAM, and running Ubuntu Linux 14.04 LTS with ROS as middleware.

The sensor for data acquisition is a downward-facing FMVU-03MTM-CS 0.3MP Point Grey camera with an image resolution of $752\,\text{px} \times 480\,\text{px}$ and a FoV of $(54.9°, 36.7°)$. To extract measurements for online mapping, the AprilTag library[8] (Olson, 2011) is used to detect the tag positions on the ground. Figure 5.17(b) shows an example tag detection obtained from the algorithm. A maximum sensor frequency of 0.4 Hz is considered to allow for enough space between successive measurements as the UAV travels. The sensor model used for planning has the form shown in Figure 3.1 with scaling factors applied to accommodate the dimensions of the indoor environment.

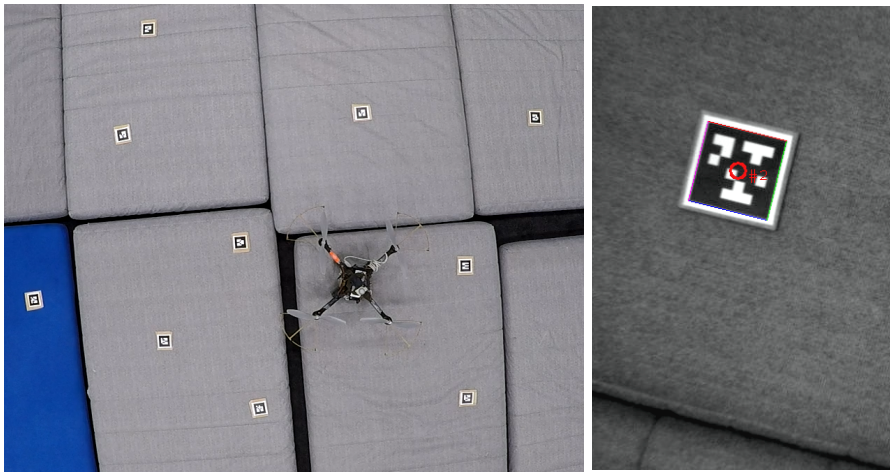

(a) (b)

Figure 5.17: (a) shows an aerial view of the experimental setup with AR tags as ground targets. The red circle in (b) exemplifies an AprilTag detection output overlaid on part of a camera image. When a measurement is taken, the detected tag pose is projected onto an occupancy grid map, which is then used for informative planning online.

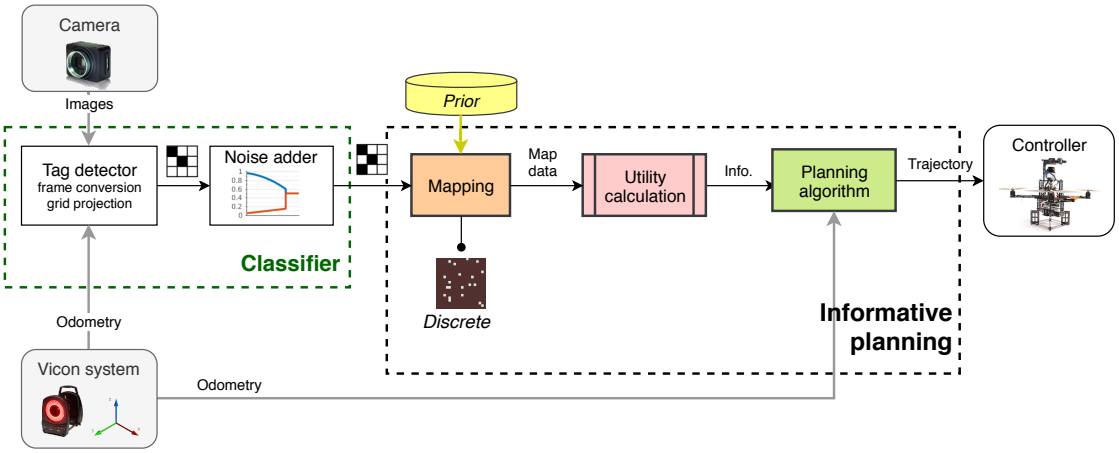

Figure 5.18: System diagram for the target (AR tag) mapping experiments. Localization is provided by the Vicon system. An image-based tag detection algorithm with optionally added noise represents the classifier unit providing measurements for discrete map building online. Based on this, the planning unit outputs informative trajectories, tracked using model predictive control (MPC).

The binary mapping method introduced in Section 3.1 is used to represent the environment given the discrete nature of the tags. A uniform resolution of 0.15 m is considered for the occupancy grid. When an image is registered, any tags detected in it are projected onto the grid to update its cells according to the sensor model. As in the simulation trials in Section 5.2.1, it is possible to generate uniform noise on the output of the tag detection algorithm to emulate sensor performance that depends more strongly on altitude. In practice, however, the tag positions computed in the raw

[8] `april.eecs.umich.edu/software/apriltag.html`

detections were found uncertain enough so that additional noise was unnecessary. Figure 5.18 shows a system diagram illustrating how the various modules used in the experimental system interact. Note that this setup corresponds to an adaptation of the general architecture shown in Figure 5.1 for the specific target mapping application.

The experiments are further designed to demonstrate the applicability of approach for mapping different target distributions. To this end, three distinct experimental scenarios are considered, featuring: (a) 10 randomly distributed tags; (b) 14 tags grouped together in a cluster; and (c) no tags, i.e., the area is completely free from targets. In each scenario, the mission time budget is $B = 150\,\mathrm{s}$ and the environment is assumed to be initially unknown.

The position for the first measurement is specified as 2 m above the map center, from where the UAV initially takes off. For trajectory optimization, the maximum reference velocity and acceleration are $0.5\,\mathrm{m/s}$ and $1.5\,\mathrm{m/s^2}$. Finally, for informative planning with an asychnrononously triggered sensor, the utility function $\mathrm{I}(\cdot)$ is chosen to maximize map uncertainty reduction (Equation 4.5) and the "global" CMA-ES variant is applied to optimize all control waypoints defining a trajectory. Note that this setup demonstrates the integration of all key elements of the proposed strategy.

The quantitative results of the trials are reported in Figure 5.19. These plots indicate how the approach performs in the various scenarios against the information metrics relevant for active classification problems. Generally, the trends observed confirm that total map uncertainty (left) is successfully reduced, while classification accuracy (right) improves, independently of the target layout in the experiments. Note that accuracy was evaluated with respect to the ground truth map of tags derived at the same resolution as the occupancy grid.

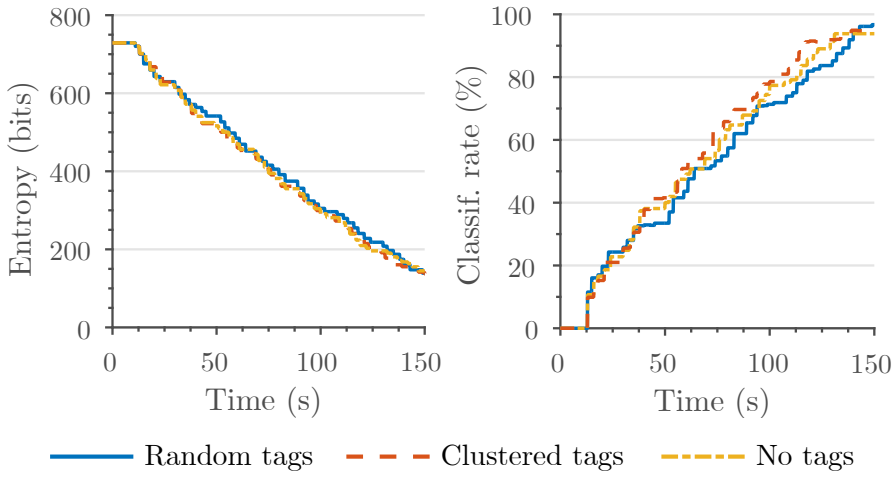

Figure 5.19: Information metrics over time for the proposed informative planning approach in scenarios with different target tag distributions. The planning algorithm aims to minimize the uncertainty in the map by exploring the environment as quickly as possible. The expected performance against the metrics (left: reducing uncertainty, right: improving classification) validates the real-time implementation of the method. Note that replanning is time taken into account.

As an example, the trajectory traveled in the scenario featuring a random tag distribution is visualized in Figure 5.20. Qualitatively, this result is very similar to those obtained in the simulations in Section 5.2. Starting from the center, the UAV descends a little to carefully map the corners of the observed area. The lighter cells in the rendered occupancy grid indicate the locations of the successfully detected tags (9 out of 10).

A closer inspection of the final map reveals an interesting subtlety that explains why the algorithm performs most poorly in this type of scenario (blue curve in Figure 5.19). Namely, during the trial, minor inaccuracies in the tag detection algorithm often caused the occupancy grid cells adjacent to those containing the targets to be updated incorrectly, given its limited resolution. The misclassified areas could then remain highly uncertain until multiple reliable measurements were received, which effectively limits the rate at which the area is mapped. To improve performance, it would be relevant to develop sensor models that capture this behavior more faithfully for predictive planning.

A video of the experiment with the random tag distribution is available at: `youtu.be/1cKg1fjT54c`.

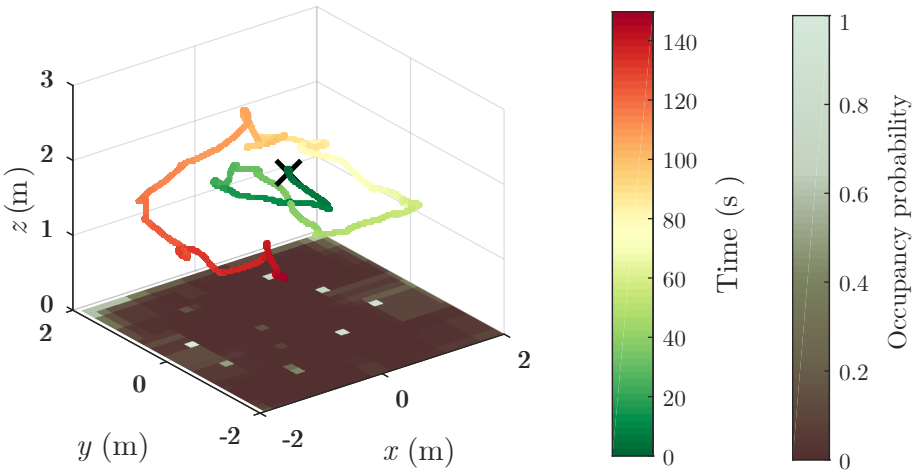

Figure 5.20: Visualization of the trajectory traveled by the proposed planner (colored line) in a 200 s mission to map a random distribution of tags on the ground. The final occupancy grid map is rendered. The black 'x' indicates the initial position of the UAV. From here, an outward path is followed (green-red) to explore the edges and corners of the environment. The lighter cells in the map correspond to 9 out of 10 successfully detected tags.

Overall, these results demonstrate the first successful application of the proposed framework in a real-world monitoring scenario. The experiments on the UAV platform also show its ease of deployment on a physical system.

5.4.2 Color Mapping

In these experiments, the performance of the framework is demonstrated on a UAV mapping the color saturation level in an indoor environment. The purpose is to validate the aspects of the approach relevant for monitoring a continuous field distribution (Section 5.2.2) in real-world conditions. Note that the experimental procedure in this section proceeds similarly to that of Section 5.4.1 for a discrete mapping scenario.

In the experiments, the objective is to reconstruct the distribution of color saturation over an area of interest, which would be useful, e.g., for monitoring vegetation indices to pinpoint weeds in precision agriculture (Section 5.4.3). As in the previous sub-section, trials are performed with various target distributions for experimental validation in different types of environment. The UAV platform is a DJI Matrice M100 quadrotor navigating in an empty indoor environment with a cubical volume of $2\,\text{m} \times 2\,\text{m} \times 2\,\text{m}$. For mapping, a downward-facing Intel RealSense ZR300 depth camera is used to

observe a white sheet with blobs of green paint. The sheet is placed flatly on the ground to represent terrain, as shown in Figure 5.21(a). The Vicon external motion capture system and MPC (Kamel et al., 2017) are used for state estimation and trajectory tracking.

Low-level functionality runs on an on-board computer with a 3.2 GHz Intel NUC i7-5557U processor and 16 GB of RAM. This implementation is based on a modified version of the DJI ROS Software Development Kit (SDK)[9] (Sa, Kamel, Burri, et al., 2018). High-level functionality, including mapping and informative planning, runs in MATLAB on a a single desktop computer with a 1.8 GHz Intel i7-8550U processor and 16 GB of RAM, with the Robotics System Toolbox as the real-time communication interface.

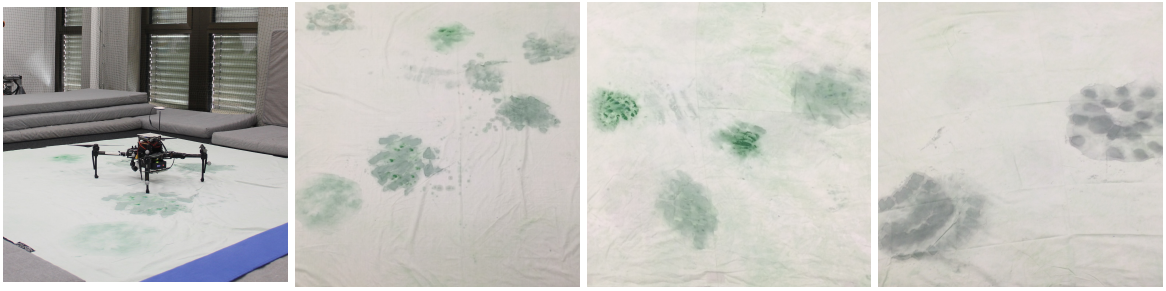

(a) (b) (c) (d)

Figure 5.21: (a) shows a side view of the experimental setup. The painted sheets in (b)-(d) represent different distributions of color saturation used as targets for mapping in the mission trials. For online mapping using the GP-based method, measurements are extracted from projected color pointclouds provided by the on-board depth camera. Note that the green painted regions correspond to areas more highly saturated in color.

Using the GP-based mapping strategy, a uniform resolution of 0.1 m is set for both the training and predictive data grids. As in the simulation experiments in Section 5.2.2, the map is initialized uniformly with a mean prior of 50 % saturation level, assuming that the area is initially unknown. The isotropic Matérn 3/2 kernel in Equation 5.2 is applied to encode the structure of the data within the GP model. The set of hyperparameters for this kernel $\theta \triangleq \{\sigma_n^2, \sigma_f^2, l\} = \{0.50, 0.22, 2.08\}$ was trained by maximizing the NLML in Equation 3.8 using conjugate gradients, as before. Data for the training procedure was obtained by flying the UAV manually over a part of one of the sheets before the experimental trials.

The Intel ZR300 depth camera has a FoV of (19.0°, 23.0°) and provides colored 3-D pointcloud data at a constant frequency of 0.2 Hz as sensory input for online mapping. The performance of the sensor is based on the model depicted in Figure 3.2. To determine its variation with altitude, the coefficients in Equation 3.12 are specified as $\gamma_1 = 0.1$ and $\gamma_2 = 0.2$, with a 1.5 m altitude above which images scale by $s_f = 0.5$. When a pointcloud is registered, the camera FoV is projected on the grid map, and the saturation levels of groups of points falling within each of its cells are simply averaged for the Bayesian map update procedure. A depth camera was opted for in order to provide a suitable variation in noise level within the altitude range. Moreover, in comparison with dense RGB images, the sparsity of the point cloud implies that less information is lost in downsampling with respect to the resolution of the field map. Note that the system diagram for these trials is similar to the one shown in Figure 5.18, with the interfaces modified to accommodate the color mapping scenario. This

[9] `github.com/ethz-asl/mav_dji_ros_interface`

provides further evidence that the software framework can be easily implemented and reused in a broad range of missions.

Three experimental trials with unique target distributions were performed to show that the proposed approach can map different continuous and realistic fields. The painted sheets in the trials are depicted in Figures 5.21(b)–5.21(d). Note that green paint is used to distinguish regions that are more highly saturated in color, and is deliberately spread unevenly in order to diversify the experimental scenarios. A mission time budget of $B = 130\,\text{s}$ is specified for each test.

The position for the first measurement is 0.8 m above the field center, and the environment is assumed to be *a priori* unknown. Polynomial paths of order $k = 12$ are defined by $N = 4$ control waypoints with a maximum reference velocity and acceleration of $1.5\,\text{m/s}$ and $2\,\text{m/s}^2$ for trajectory generation. For informative planning, the utility function is set to reduce map uncertainty based on Equation 4.11. A base threshold $\mu_{th} = 20\%$ and scaling parameter of $\beta = 3$ in Equation 4.13 are considered to define an adaptive requirement in this scenario, i.e., the aim is to accurately map areas with saturation level above this confidence bound. To obtain an initial trajectory solution, the grid search is performed over a 30-point lattice that has a similar structure to the one in Figure 4.2(b), but features a wider distribution of points at higher altitudes to account for the narrow FoV of the depth camera. For subsequent optimization using the CMA-ES, an initial step size $\boldsymbol{\sigma}_{\text{CMA}}$ of 0.2 m in each co-ordinate is set to obtain symmetrical search properties in the cubical workspace of the UAV.

Figure 5.22 presents the results obtained in the three experimental trials. The proposed approach is quantitatively validated by studying the evolution of total map uncertainty during the missions, as measured by $\text{Tr}(\mathbf{P})$. Unfortunately, an analysis of map error, as in the simulations, is not applicable due to the absence of ground truth data in the physical setup. As an example, Figure 5.22(a) provides a visualization of the trajectory traveled to reconstruct the distribution in Figure 5.21(c), which corresponds to the second experiment. The lighter grid cells in the final estimated map, shown rendered, correspond to successfully detected (painted) areas with higher levels of color saturation.

Considering Figure 5.22, a comparison between the mapping output and the physical sheet confirms qualitatively that the mapping method captures the spread of green paint, as expected. The trajectory in (a) is also visually consistent with the findings presented earlier in this chapter, e.g., the example shown in Figure 5.2. Upon ascending to a high altitude ($\sim 20\,\text{m}$), the UAV proceeds to explore the environment and refine the map with higher-quality data. In these experiments, however, the descent to lower altitudes towards the end of the mission is less evident due to the narrow FoV of the depth camera, which limits the information gain achievable with the sensor model. In (b), the expected reductions in map uncertainty over time validate the real-time applicability for monitoring in different scenarios.

As a final remark, it is worth mentioning that the replanning procedure takes $\sim 17\,\text{s}$ in this proof of concept setup. Computation time reduces substantially with an on-board implementation completely based in ROS, as demonstrated in the following sub-section. Moreover, the ideas presented in Section 3.3 could be applied to improve the efficiency of the map update procedure, which corresponds to the main computational bottleneck within the system. Further ideas are discussed in Section 5.5.

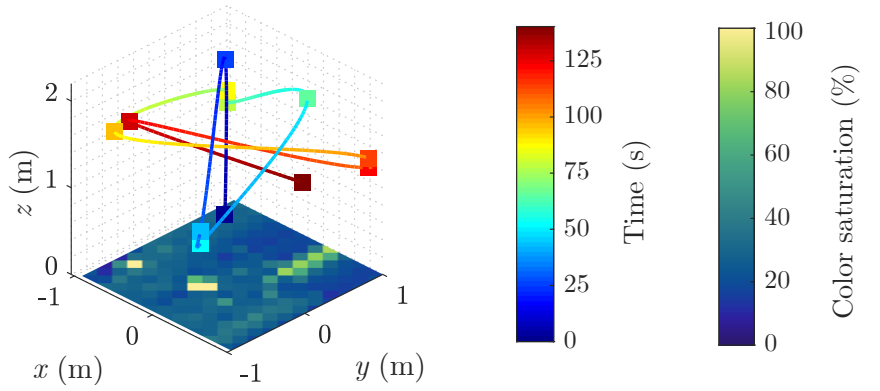

(a)

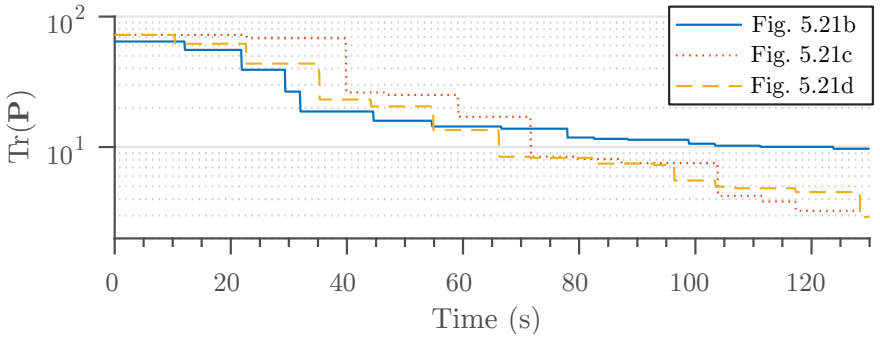

(b)

Figure 5.22: (a) visualizes the trajectory traveled by the proposed planner (colored line) in a 130 s mission to map the saturation level of the sheet shown in Figure 5.21(c). The squares indicate measurement sites. The mean of the final GP field map is rendered. The lighter cells in the map correspond to successfully detected areas of green paint that are more highly saturated in color. The trends in (b) validate that map uncertainty reduces in scenarios featuring different sheets (Figure 5.21). Note that planning time is taken into account.

5.4.3 Agricultural Monitoring

Finally, this section presents experimental results from field deployments implementing the informative planning framework on a UAV to monitor the distribution of vegetation in a field. The aim is validate the system for performing a practical sensing task in a challenging outdoor environment, with all algorithms running in on-board and in real-time.

The experiments were conducted on an agricultural field at the Research Station for Plant Sciences Lindau of ETH Zurich in Switzerland (Lat. 47.450040°, Lon. 8.681056°) in September 2018. Figure 5.23 shows various images of the on-field setup. The UAV platform in (a) is the DJI Matrice M100 quadrotor monitoring a 20 m × 20 m area within the field with maximum and minimum altitudes of 21 m and 8 m, respectively. As shown in (b), the controlled field features an central area of crops planted in row arrangements, surrounded by dense distributions of weed on its edges. The width of the crop row area measures ∼ 18 m. Vegetation mapping is performed using RGB imagery from a downward-facing Intel RealSenseZR300 camera with a resolution of 1920 px × 1080 px and a FoV of (68.0°, 47.2°). Example images taken from different altitudes are shown in (c) and (d). In (c), the boundary between the weed (left) and crop (right) areas of the field can be clearly distinguished.

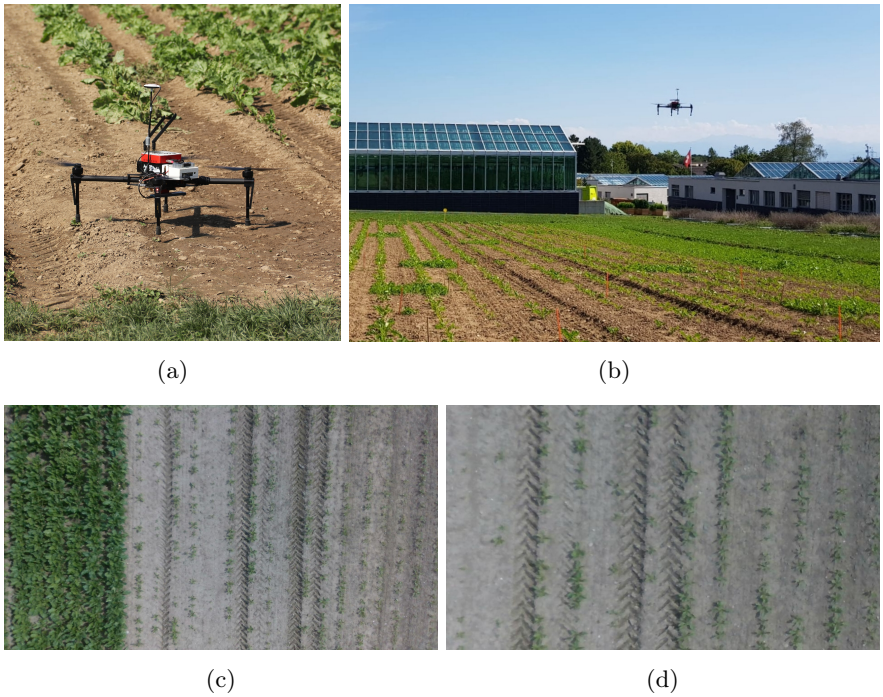

(a) (b)

(c) (d)

Figure 5.23: (a) shows a close-up of the DJI Matrice M100 used in the outdoor deployments. (b) illustrates the UAV executing a mission, flying to monitor the vegetation distribution in a field of interest. (c) and (d) exemplify RGB images taken by the on-board camera at different altitudes, which are used for mapping online. Note that, in (b), the UAV is positioned approximately above the boundary between the controlled crop (left) and weed (right) areas, visible in (c).

In this setup, state estimation is based on the Robust Visual Inertial Odometry (ROVIO) framework[10] (Bloesch et al., 2015) with MPC (Kamel et al., 2017) used for tracking trajectories output by the planning algorithm. All computations, including modules for environmental mapping and informative planning, are run on an on-board computer with a 3.2 GHz Intel NUC i7-5557U processor, 16 GB of RAM, and running Ubuntu Linux 16.04 LTS with ROS as middleware. Further details about the platform, including the modified DJI SDK package, are discussed by Sa, Z. Chen, et al. (2018).

The proposed GP-based mapping method is used to reconstruct the continuous distribution of vegetation in the field. Specifically, the goal of the field tests is to map the level of Excess Green index (ExG) in the area of interest based on the RGB images. This vegetation index is commonly used in agriculture to identify greenness levels for segmenting plants and monitoring crop health (Sa, Z. Chen, et al., 2018). It is defined by (W. Yang et al., 2015).

$$ExG = 2g - r - b\,, \tag{5.5}$$

where r, g, and b are the normalized red, green, and blue color channels in the RGB color space. For experimental purposes, note that the range of the ExG is scaled between 0 and 1 based on the maximum and minimum magnitudes measured on the experimental field in previously acquired datasets. Hereafter, this value is referred to as the "normalized ExG".

[10] `github.com/ethz-asl/rovio`

For mapping, a uniform resolution of 0.5 m is set for both the training and predictive data grids in the GP model. With the assumption that the field contains only soil and no plants, all cells in the map are initialized with a mean prior of 0 normalized ExG. Internally, the GP leverages the isotropic Matérn 3/2 kernel to approximate the spread of vegetation in the field. The hyperparameters of the kernel $\theta \triangleq \{\sigma_n^2, \sigma_f^2, l\} = \{0.50, 0.50, 1.76\}$ were learned by minimizing the NLML using conjugate gradients based on training images from a manually acquired dataset. The dataset was recorded as the UAV was flown at a fixed altitude of 8 m over the area, which corresponds to the minimum allowable level considered in the experiments, and the maximum Ground Sample Distance (GSD) (in m/px) obtainable in the images, i.e., highest possible resolution for mapping.

Measurements are extracted from RGB images taken by the Intel ZR300 camera at a constant frequency of 0.20 Hz. The sensor noise model considered is defined by Equation 3.12 with coefficients $\gamma_1 = 0.05$ and $\gamma_2 = 0.2$ determining its altitude dependency. Note that the effects of measurement downscaling at higher altitudes are not applicable in the field test setup due to the inherently high camera resolution, i.e., in Figure 3.2, the scaling factor is $s_f = 1$ across the altitude range.

To update the map with a newly acquired image, the pixels are first projected on the ground plane by applying a homography transformation given the camera parameters and the UAV pose estimated from visual inertial odometry. The projected image pixels are then averaged per cell (downsampled), and the normalized ExG values for each cell are computed based on the individual color channels by applying Equation 5.5. The step of fusing the data with the current estimated map then proceeds according to the procedure described in Section 3.2.2.

Two field trials using this setup were performed on different days. First, this section presents the main results obtained from the later experiment. At this time, the system pipeline was fully integrated and well-tested. Preliminary findings from the first experiment are then introduced in order to convey some interesting insights for further discussion. For simplicity, a mission budget B was not specified in these trials. Instead, the algorithm was allowed to generate finite-horizon plans until the estimated map output was perceived visually as being complete.

In the planning strategy, polynomial trajectories of order $k = 12$ are defined by $N = 3$ control waypoints and optimized for a maximum reference velocity and acceleration of 5 m/s and 3 m/s^2. The utility function for exploration targets map uncertainty reduction (Equation 4.11) with no threshold assigned for adaptive replanning, i.e., the aim is to reconstruct the fields in a uniform manner, as quickly as possible. The initial grid search considers a coarse 14-point lattice and the CMA-ES is set initial step sizes $\boldsymbol{\sigma}_{\text{CMA}}$ of 4.5 m in each co-ordinate of the UAV workspace with a population size of $\lambda_{\text{CMA}} = 10$ and a maximum number of iterations of $N_{\text{CMA}} = 20$.

The results from the first field trial are reported in Figure 5.24. Similarly to the indoor mapping experiment in Section 5.4.2, the ground truth data of normalized ExG levels in the field are not available. Instead, the approach is validated by assessing the progression of total map uncertainty. This analysis is shown on the right, which confirms that uncertainty is reduced over time. Note that the curve in the plot is offset by ~ 100 s as data recording was triggered before the UAV took off to the first measurement point.

The sequence of plots on the left provides qualitative insight: the estimated map does become more complete during the course of the mission, as images of the field are accumulated. The yellower parts of the maps, corresponding to areas with high values of the normalized ExG, indicate successfully identified weeds on the edges of the physical field (visible in Figure 5.23(c)). Towards the central area, a close look at the bluer parts with lower ExG reveals that even the details of the crop rows (visible in Figures 5.23(c) and 5.23(d)) are mapped correctly. These findings validate the mapping strategy, and, from a practical point of view, represent valuable data that could be used by a farmer to guide decisions on how best to manage the crops in this field.

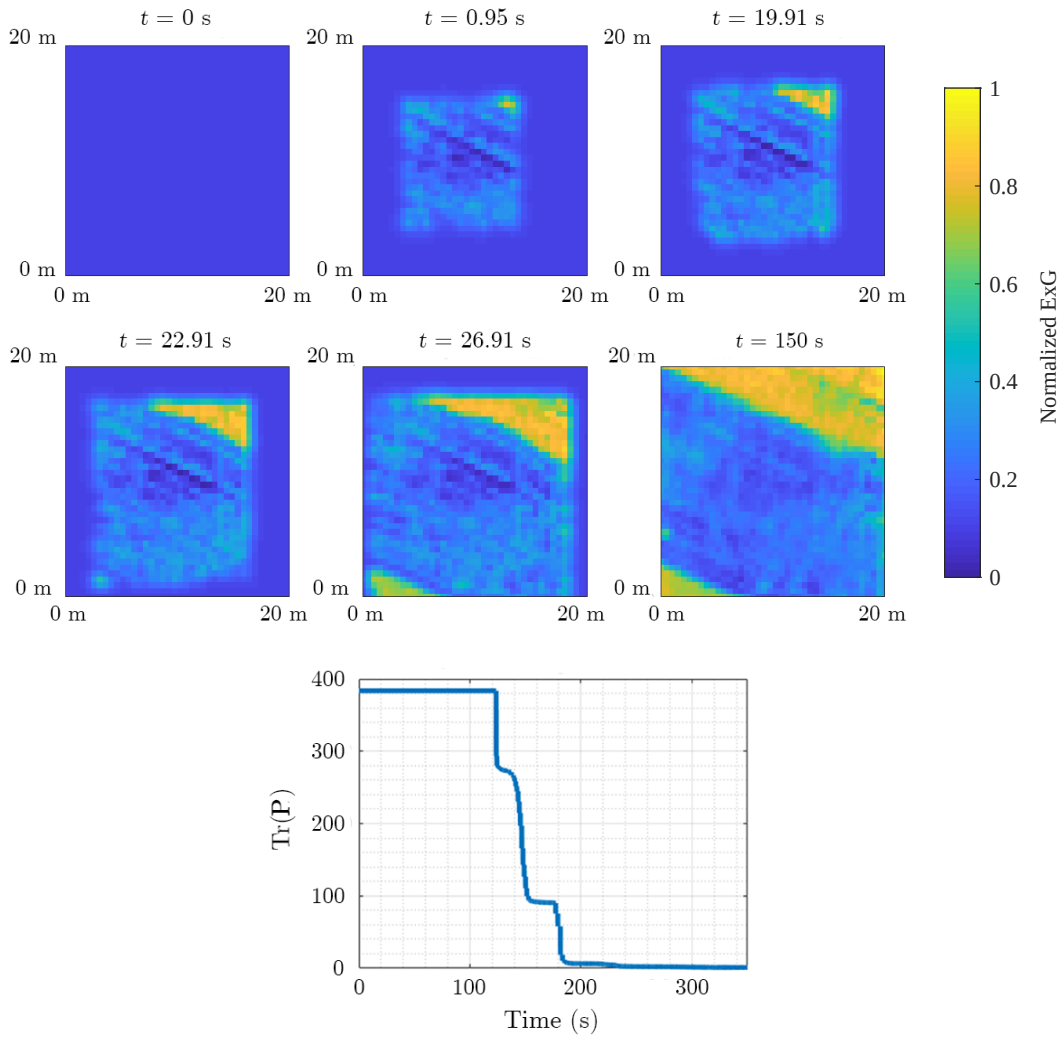

Figure 5.24: Experimental results from the second (final) field trial using the informative planning framework for UAV-based vegetation monitoring. *Top-right*: Colorbar, indicating the normalized Excess Green Index (ExG) level (plant greenness). *Top-left*: Estimated map reconstructions (GP means) of the normalized ExG on the field at different times t during the mission. The bluer and yellower regions correspond to successfully identified crop row and weed regions, respectively. *Bottom*: Evolution of total map uncertainty over time. The results confirm improving map completeness as the number of measurements increases.

A video of this experiment is available at: `youtu.be/5dK8LcQH85o`.

Though the on-field findings demonstrate the proposed approach successfully applied in a real-world scenario and under challenging conditions, they also point to several areas where the experiment could be improved. First, with the on-board system implementation deployed on a realistic, but relatively small, field, the replanning step took $\sim$ 13 s. This is substantially faster than the time reported

for the experiments in Section 5.4.2 using the ROS-MATLAB bridge. However, considering that the algorithm parameters here were set to minimize computational effort, i.e., by coarsely discretizing the initial grid search, reducing the complexity of informative plans, and minimizing optimization time, further work is necessary for the approach to scale to larger environments or applications requiring higher mapping resolution. As mentioned in Section 5.4.2, more efficient mapping approaches would yield significant speedups. Moreover, since plants are typically spread widely in big fields, the potential accuracy loss using approximation-based methods is expected to be minimal.

Second, in Figure 5.24, the final map reconstruction at a mission time of $t = 150\,\mathrm{s}$ shows some blurring artifacts, especially among the crop rows in the center. This was found to occur when multiple images of the same area were incorporated into the map, with small inconsistencies between them due to, e.g., variable lighting conditions during the mission or inaccuracies in the UAV state estimation/control. A more realistic sensor model designed based on empirical data, as in Section 5.3, is a partial remedy for this issue. Ultimately, however, mapping accuracy will be limited to that achievable with respect to the realistic mission conditions and hardware devices. Finally, a natural next step is to explore monitoring different vegetation indicators using the same setup, e.g., the Normalised Difference Vegetation Index (NDVI) based on multispectral imagery. This would extend the robustness and applicability of acquired data.

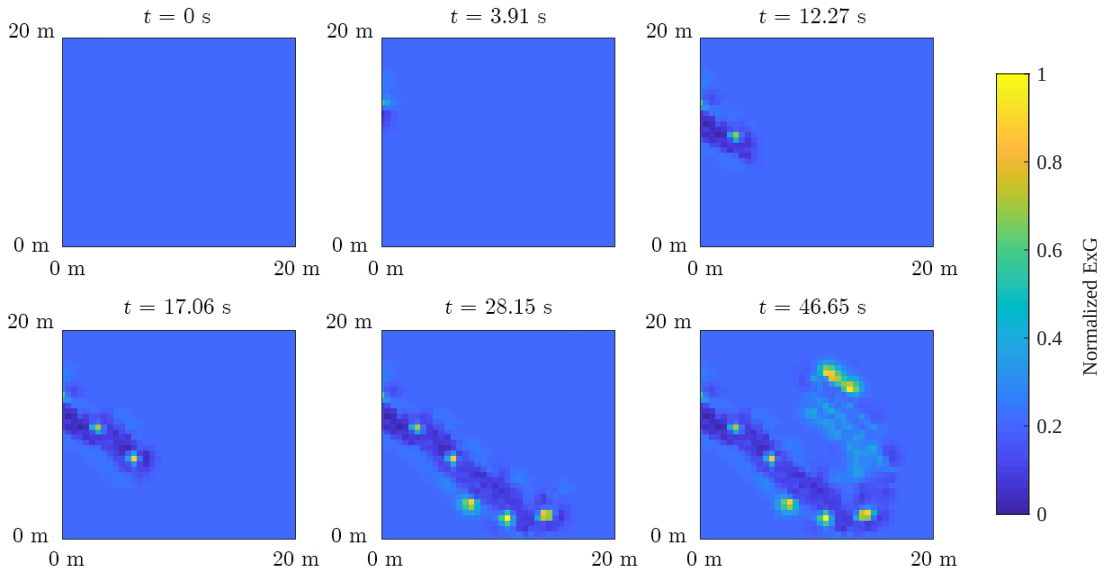

Figure 5.25: Experimental results from the first (preliminary) field trial using the proposed informative planning framework for UAV-based vegetation monitoring. The sequence of plots depicts the estimated map reconstructions (GP means) at different times t during the mission. The yellow spots correspond to successfully pinpointed "hotspot" patches of weeds growing in the field, thus validating the mapping approach. Unfortunately, the experiment was terminated early due to technical issues.

In closing, this chapter presents results from a preliminary experiment as an additional interesting point of discussion. This first test was performed on the same field as the second one shown in Figure 5.24. The main difference is that, as the first trial took place several days in advance, the field featured several controlled "hotspot"-like patches of weed growth throughout the central crop area. During the second trial, the crops on the field had matured enough so that the differences between them and the weed patches were not as evident. Note that the GP hyperparameters mentioned above

were retrained to match the differences in structure for this distribution. Also, for testing purposes, the UAV starting point was specified as 8 m above the left edge of the field (near the ground base station), and the maximum altitude for planning was limited.

The first experiment is interesting because pointpointing the weed patches is a very relevant practical task for precision agriculture. Figure 5.25 shows the qualitative mapping results obtained from this trial. Unfortunately, the mission was terminated early due to drifts in the UAV state estimate, and more testing is necessary to draw further conclusions. However, the yellower features in the preliminary maps clearly demonstrate the ability of the approach to locate the points of high vegetation (normalized ExG) on the field. Such data could contribute greatly towards developing site-specific weed treatment strategies and advancing sustainable agriculture practices (Cardina et al., 1997 and Peña et al., 2013)

5.5 Summary and Discussion

This chapter opened with a brief overview of the proposed approach for informative planning as an integrated system. It then presented a series of experiments demonstrating its performance for a diverse range of terrain monitoring tasks. First, the framework was evaluated extensively through simulations of both discrete and continuous mapping problems. Experimental results were presented comparing the approach against existing methods, as well as assessing its various internal elements. The findings exposed the benefits of using informative planning for sensing missions with finite resources, in terms of both the speed at which data is collected and its quality. Then, system validation was conducted in a photorealistic mapping scenario using experimental data from a publicly available dataset. Finally, this chapter provided field deployments implementing the algorithms on UAVs in various scenarios, including for monitoring distributions of vegetation in agricultural settings.

Some interesting discussion topics about the experiments were offered in the previous sub-sections. A takeaway message is that the efficiency of the continuous mapping method (Sections 5.2.2, 5.4.2 and 5.4.3) needs to be improved if the approach is to scale to larger environments. For reference, in Section 5.4.3, the replanning step took $\sim 13\,\text{s}$ for a grid size of $40\,\text{cells} \times 40\,\text{cells}$ (terrain size of $20\,\text{m} \times 20\,\text{m}$) with carefully optimized parameters. The map update step is the main computational bottleneck since it has cubic time complexity and is often repeated. Namely, this procedure is executed each time a candidate point is evaluated during the initial grid search, as well as multiple times for each trajectory sampled by the CMA-ES optimizer. Therefore, approximation-based techniques for mapping, as discussed in Section 3.3, are necessary to extend the applicability of the approach, for instance, to larger farms.

Developing models that accurately capture sensor properties for mapping and predictive planning was a recurring problem. Without solid empirical data, the experimental sensor models were mostly tuned heuristically to obtain altitude-dependent performance. However, the results, e.g., in Sections 5.4.2 and 5.4.3, clearly illustrate situations where mapping fidelity suffered due to a mismatch between the forward-simulation of the sensor and its real-world behavior. The neural network-based sensor model in Section 5.3 is an exception and good example of a more realistic design. In general, the field of sensor modeling is very related to the subjects of this thesis, and it requires further research to improve the reliability of informative plans.

The field tests successfully validated the proposed approach applied in real-world conditions. Evaluative studies, however, were not performed. In future work, it would be relevant to compare the method against benchmark strategies, as in the simulated experiments, including commercial on-field tools, e.g., the DJIFlightPlanner software[11], to examine its practical benefit. Also, more quantitative

[11] `djiflightplanner.com`

experimental results are generally required. As shown by M. G. Jadidi et al. (2016) and Hollinger and Sukhatme (2014), on field sites, one can construct a proxy for ground truth using a coverage-type survey of the area of interest. This approach enables calculating the error and assessing the accuracy of the reconstructions.

On that note, it is worth emphasizing the general lack of available software tools for informative planning in the state-of-the-art. Considering that this field is expanding rapidly, it would benefit researchers to have standardized benchmark datasets and algorithms against which to evaluate new approaches. One barrier towards this is that planning strategies are usually application-specific, e.g., by assuming a particular map representation, making them not directly comparable. The open-source package introduced in this thesis, and its dissemination by Papachristos, Kamel, et al. (2019), is an important step in this direction.

Finally, a key assumption behind the framework is an obstacle-free robot workspace. Building upon this work, Anil Meera et al. (2019) extend the methods in this thesis for informative planning in cluttered environments, e.g., a disaster site with buildings in a search and rescue scenario. These methods could be directly integrated into the presented system and tested as an additional feature.

6

Active Sensing Under Localization Uncertainty

This chapter introduces methods of accounting for the uncertainty in the robot localization in general active sensing problems. It incorporates material from the following first author publication:

Popović, M., Vidal-Calleja, T., Chung, J. J., Nieto, J., and Siegwart, R. (2019). "Informative Path Planning and Mapping for Active Sensing Under Localization Uncertainty". *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Under review. arXiv: 1902.09660

This part of the thesis builds upon the informative planning framework presented in Chapter 5. The motivation is to improve upon the quality of environmental field reconstructions by factoring the robot pose uncertainty into the decision-making process. After introducing the problem setup, this chapter describes a technique for propagating the pose uncertainty into the Gaussian Process (GP) mapping framework. Then, as a key contribution, it presents a new tuning-free utility function that incorporates both the robot localization and field mapping objectives. Finally, the approach is evaluated in simulation experiments and validated with results obtained from a temperature monitoring scenario.

6.1 Problem Setup

The active sensing problem studied in this chapter focuses on a robot mapping a scalar environmental field, e.g., of temperature, humidity, pressure, etc., using point measurements (samples) taken by an on-board sensor. In contrast to the terrain monitoring type of scenario, it encompasses tasks of mapping both 2-D or 3-D field distributions in which the same dimensions represent the workspace of the robot. For informative planning under pose uncertainty, an estimate of the robot position and orientation is assumed to be provided by a probabilistic localization or Simultaneous Localization And Mapping (SLAM) back end system, e.g., particle filter, Kalman Filter (KF), graph SLAM system, etc. As is commonly done, the pose is represented as a Gaussian distribution in the workspace $\mathcal{N}(\mathbf{p}, \mathbf{\Sigma})$.

In this setup, Equation 2.1 formulates the informative path planning objective. Note that the problem of maximizing information gain, rather than information gain *rate* (Equation 2.2), as in the previous chapters, is considered here. This is because the proposed utility function $\mathrm{I}(\cdot)$ implicitly penalizes the value of measurements obtained over longer paths, which cause the robot pose uncertainty to grow as it explores unknown areas.

The terms in this general equation are defined in the same way as in Section 4.1. In summary: (a) candidate paths are represented as smooth polynomial trajectories; (b) the sensor for data acquisition

is assumed to take measurements along trajectories at a constant frequency; and (c) a time-based budget is considered for the resource-constrained mission. For further details on these definitions, the reader is referred to the sub-section mentioned above.

Figure 5.1 shows a system diagram summarizing the proposed approach. The dashed red line indicates the propagation of the robot pose uncertainty into the informative planning framework, corresponding to the developments in this chapter. As shown by the connections, a key feature of the method is that this information is jointly accounted for in both the mapping (Section 6.2) and planning (Section 6.3) modules to improve the robustness of data collection.

6.2 Environmental Mapping Under Pose Uncertainty

This section presents a method for mapping under pose uncertainty as the basis of the informative planning approach. The environmental field of interest is represented as a realization of a GP in order to capture its spatial correlation structure and predictive uncertainty. This probabilistic framework is introduced in Section 3.2.1 and the reader is referred to there for the theoretical background necessary to follow the discussions below. Further details are available in the book of Rasmussen et al. (2006)

For online map building in this chapter, the standard GP regression equations (Equations 3.6 and 3.7) are applied as measurement samples are collected during a mission. For convenience, they are repeated here:

$$\boldsymbol{\mu} = \mathbb{E}[\mathbf{f}_*] = m(\mathbf{X}_*) + K(\mathbf{X}_*, \mathbf{X})[K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n]^{-1} \times (\mathbf{y} - m(\mathbf{X}))\,,$$
$$\mathbf{P} = \mathbb{V}[\mathbf{f}_*] = K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X})[K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n]^{-1} \times K(\mathbf{X}_*, \mathbf{X})^\top\,.$$

The strategy for propagating uncertainty into the GP model is based on the expected kernel technique of M. G. Jadidi et al. (2017). The main contribution of this thesis is that the concepts are extended for active sensing problems from 2-D to 3-D scenarios, which poses a more complex optimization problem in terms of finding informative paths. Moreover, in Section 6.3.1, they are integrated with a new utility function to perform not only mapping, but also planning, under pose uncertainty.

In particular, the method is concerned with the covariance function $k(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}')]$ of the GP model, which captures its internal correlated structure and is used to build the kernel matrices $K(\cdot,\cdot)$ in the equations above. The key idea is to redefine the covariance function k considering that the training points aggregated in the matrix $\mathbf{X}$ are Uncertain Inputs (UIs) corrupted by the measurement noise of the sensor. The query points $\mathbf{X}_*$ used for regression are assumed to be deterministic.

Let $X \in \mathcal{X}$ represent a random variable associated with an UI and distributed according to a probability distribution $p(x)$. The modified covariance function $\tilde{k}$ can be computed by taking the expectation over $p(x)$ using:

$$\tilde{k} = \mathbb{E}[k] = \int_X k p(x) dx\,. \tag{6.1}$$

Once the expected kernel matrices are calculated using the expression above, Equations 3.6 and 3.7 can be applied in the standard way to evaluate the predictive conditional distribution for the deterministic query points $\mathbf{X}_*$ in the environment.

Given that the integration in Equation 6.1, in general, analytically intractable, it is possible to solve it using numerical approximation methods. To this end, this work employs the Gauss-Hermite quadrature rule (Davis et al., 2007). Assuming a Gaussian distribution for the robot pose $\mathcal{N}(\mathbf{p}, \boldsymbol{\Sigma})$,

this strategy is chosen as it provides a better accuracy and efficiency trade-off compared to a Monte-Carlo technique, which is studied as an alternative in the original work of M. G. Jadidi et al. (2017).

The Gauss-Hermite quadrature approximates the value of integrals of the kind $\int_{-\infty}^{\infty} \exp(-x^2) f(x) dx$ by using the formula:

$$\int_{-\infty}^{\infty} \exp(-x^2) f(x) dx \approx \sum_{i=1}^{n} w_i f(x_i) \,, \tag{6.2}$$

where n is the number of sample points used in the approximation, x_i are the roots of the Hermite polynomial H_n, and the associated weights w_i are given by:

$$w_i = \frac{2^{n-1} n! \sqrt{\pi}}{n^2 [H_{n-1}(x_i)]^2} \,. \tag{6.3}$$

More extensive details are discussed by Davis et al. (2007).

Consider the multivariate Gaussian distribution of a D-dimensional noisy input to be given by $\mathcal{N}(\tilde{\mathbf{x}}, \boldsymbol{\Sigma}_x)$. M. G. Jadidi et al. (2017) show that, through a change of variable such that $\mathbf{L}\mathbf{L}^\top = 2\boldsymbol{\Sigma}_x$ and $\mathbf{u} = \mathbf{L}^{-1}(\mathbf{x} - \tilde{\mathbf{x}})$, where $\mathbf{L}$ is a lower triangular matrix computed via the Cholesky decomposition, Equation 6.1 can be approximated as:

$$\tilde{k} = (2\pi)^{-\frac{D}{2}} \sum_{i_1=1}^{n} \dots \sum_{i_D=1}^{n} \bar{w} k_{i_{1:D}} \,, \tag{6.4}$$

where $\bar{w} \triangleq \prod_{j=1}^{D} w_{i_j}$, u_{i_j} correspond to the roots of H_n, $\mathbf{u}_{i_{1:D}} \triangleq [u_{i_1}, \dots, u_{i_D}]^\top$, and $k_{i_{1:D}}$ is the covariance function evaluated at $\mathbf{x}_{i_{1:D}} = \mathbf{L}\mathbf{u}_{i_{1:D}} + \tilde{\mathbf{x}}$.

When $D = 2$, i.e., the environment is planar, Equation 6.4 simplifies to:

$$\tilde{k} = \frac{1}{2\pi} \sum_{i_1=1}^{n} \sum_{i_2=1}^{n} \bar{w} k_{i_{1:2}} \,. \tag{6.5}$$

When $D = 3$, i.e., the environment is volumetric, Equation 6.4 simplifies to:

$$\tilde{k} = \frac{1}{(2\pi)^{\frac{3}{2}}} \sum_{i_1=1}^{n} \sum_{i_2=1}^{n} \sum_{i_3=1}^{n} \bar{w} k_{i_{1:3}} \,. \tag{6.6}$$

Note that Equation 6.6 corresponds to the additional dimension required for the active sensing problems studied in this thesis.

6.3 Informative Planning

This section describes the proposed method for informative planning under pose uncertainty. The general path planning scheme was introduced and detailed in Chapter 4 as a core component of the overall approach of this thesis. The new uncertainty-aware utility function (Section 6.3.1) is emphasized as the major contribution within the algorithm. Methods of uncertainty propagation for predictive planning using this approach are discussed in Section 6.3.2.

6.3.1 Utility Definition

This sub-section introduces a new utility, or information gain, function $\mathrm{I}(\cdot)$ in Equation 2.1 for active sensing problems under pose uncertainty. The proposed formulation is suitable for environmental field mapping using a GP model to represent the underlying scalar distribution. Its main benefit over existing methods is that the uncertainty of both the field map and robot pose are accounted for in a unified manner, without relying on any manual, heuristic tuning procedures or problem-specific parameters. The main motivation is to enable the robot to adaptively trade off between gathering new information about the field of interest (exploration) and maintaining good localization (exploitation) in a variety of environments. Intuitively, such a scheme is expected to achieve more conservative mapping behavior compared to planning for pure information gain, since it encourages the robot to incorporate new measurements into the map at the correct locations.

The utility function is developed based on a reasoning approach similar to that of Carrillo et al. (2018). From this perspective, in fact, this thesis presents a parallel formulation of their utility for monitoring a continuous environmental phenomenon, instead of building an occupancy map in an autonomous exploration task.

The key idea is to discount the expected information gain for sensing actions by the uncertainty in the robot localization. To achieve this, the value of information in the GP field model is quantified by applying the concept of Rényi's entropy. Rényi's entropy was introduced by Rényi (1961) as a generalization of the well-known Shannon's entropy. It is a parametric family of entropies representing the most general class of information measure that preserves the additivity of independent events and complies with the axioms of probability (Xu et al., 2010). In particular, Rényi's entropy has the powerful property of being able to define several measurements of uncertainty (or dissimilarity) within a given distribution by setting a free parameter α, as shown below. This property has been exploited for statistical analyses in many fields, e.g., thermodynamics, quantum mechanics, biomedical engineering, etc. (Xu et al., 2010 and Golshani et al., 2010).

The subject of this work is concerned with Rényi's entropy applied to Gaussian distributed variables. Let X be a random variable having an absolutely continuous distribution with density function $f(x)$. Golshani et al. (2010) define Rényi's entropy of order α as:

$$\mathrm{H}_\alpha(X) = \frac{1}{1-\alpha} \log \int_{\mathbb{R}} f^\alpha(x) dx \qquad \alpha > 0,\, \alpha \neq 1 \tag{6.7}$$

and $\mathrm{H}_1(X) = \lim_{\alpha \to 1} \mathrm{H}_\alpha(X) = -\int_{\mathbb{R}} f(x) \log f(x) dx$ converges to Shannon's entropy (Equation 4.8) in the limiting case, if both integrals exist. In this expression, $\alpha \in [0, 1) \cup (1, \infty)$ is a free parameter, restricted in this work to the range $(1, \infty)$.

Now assuming that X follows a Gaussian distribution $X \sim \mathcal{N}(\mu,\, \sigma^2)$, Rényi's entropy is derived as:

$$\begin{aligned} \mathrm{H}_\alpha(X) &= \frac{1}{2} \log\left(2\pi\sigma^2\right) - \frac{1}{2(1-\alpha)} \log \alpha \\ &= \frac{1}{2} \log\left(2\pi\sigma^2 \alpha^{\frac{1}{\alpha-1}}\right). \end{aligned} \tag{6.8}$$

As in Equation 6.7, the convergence to Shannon's entropy $\mathrm{H}_{\alpha \to 1}(X) = \mathrm{H}(X)$ in the limit can be verified.

The next step is to apply this definition to quantify the uncertainty of the GP environmental field model in terms of Rényi's entropy. Following the strategy in Section 4.3, the criterion for A-optimal design (Fedorov, 1972) is used to measure information in the planning objective based on the sum

of the variances in the model. For a multivariate Gaussian distribution $X \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{P})$, considering the expression in Equation 6.8, the A-optimality metric with Rényi's entropy is computed as:

$$\begin{aligned} \hat{\mathrm{H}}_\alpha(X) = \hat{\mathrm{H}}_\alpha(\mathbf{P}) &= \frac{1}{2} \log \left(2\pi \ \mathrm{Tr}\left(\mathbf{P}\right) \alpha^{\frac{1}{\alpha-1}}\right) \\ &\propto \ \log \left(\mathrm{Tr}\left(\mathbf{P}\right) \alpha^{\frac{1}{\alpha-1}}\right), \end{aligned} \tag{6.9}$$

where $\mathbf{P}$ is the covariance matrix of the GP model. As discussed in Section 6.2, $\mathbf{P}$ is obtained by applying the GP regression equation in Equation 3.7. To this end, a specified covariance function k or its modified variant $\tilde{k}$ in Equation 6.1 can be used for mapping problems without or with UIs, respectively.

The question then arises as to how adjusting the free parameter α affects the uncertainty measure. Consider the behavior of α in the range $\alpha \in (1, \infty)$. By using l'Hopital's rule, it can be verified that, as $\alpha \to 1$, the Rényi term $\lim_{\alpha \to 1} \log \alpha^{\frac{1}{\alpha-1}} = \log e$, such that $\hat{\mathrm{H}}_{\alpha \to 1}(X) = \hat{\mathrm{H}}(X)$ (higher uncertainty) and Rényi's entropy reduces to Shannon's entropy. Conversely, as $\alpha \to \infty$, the result is $\lim_{\alpha \to \infty} \log \alpha^{\frac{1}{\alpha-1}} = \log 1 = 0$, such that $\hat{\mathrm{H}}_{\alpha \to \infty}(X) = 0$ (lower uncertainty).

Figure 6.1 provides a visual aid to illustrate this relationship. Interestingly, Jumarie (1990) interprets the decay as measuring the efficiency of a subjective observer considering the distribution X, where α is a gain coefficient that increases the value of information for $\alpha > 1$. Thus, α naturally enables discounting the Shannon's uncertainty of the model. This property is exploited by the proposed utility function for planning under pose uncertainty, as described below.

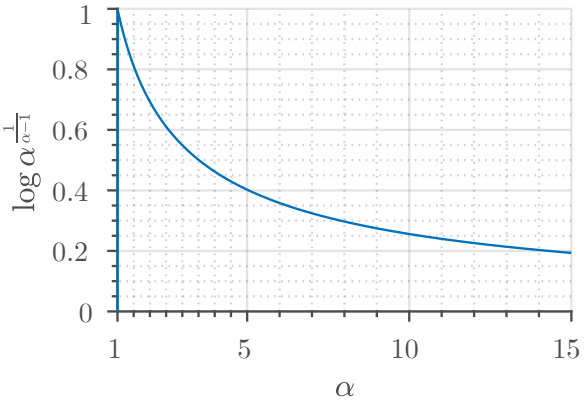

Figure 6.1: Variation of the Rényi term $\log \alpha^{\frac{1}{\alpha-1}}$ in the range $\alpha \in (1, \infty)$. The decay of the term with increasing α is exploited in the proposed utility function to discount the value of information acquired from future measurements based on the robot localization uncertainty.

The ideas above can now be combined in the definition of mutual information (transinformation), presented by Jumarie (1990):

$$\mathrm{I}_c(X) = \mathrm{H}(X) - \mathrm{H}_{\alpha=c}(X), \tag{6.10}$$

which is the difference between Shannon's entropy of the the probability distribution X and Rényi's entropy of the same distribution with $\alpha = c$. For $\alpha > 1$, Rényi's entropy effectively increases the value of information as $\alpha \to 1$, $\mathrm{I}_c(X) \to 0$ (less information in X); as $\alpha \to \infty$, $\mathrm{I}_c(X) \to \mathrm{H}(X)$ (more information).

Now, this interpretation is contextualized for active sensing tasks in order to develop the uncertainty-aware utility function. For predictive planning, it is necessary to measure the information gain following a sensor measurement taken at a robot pose $\mathbf{p}$ with an associated covariance $\boldsymbol{\Sigma}$. To achieve this, the utility formulation computes the mutual information between the training data within the current GP field model and the new observation based on the definition in Equation 6.10.

Hence, it is defined using:

$$\mathrm{I}_{\alpha(\boldsymbol{\Sigma})}(\mathbf{p}) = \hat{\mathrm{H}}(\mathbf{P}^{-}) - \hat{\mathrm{H}}_{\alpha(\boldsymbol{\Sigma})}(\mathbf{P}^{+}) \,, \tag{6.11}$$

where the minus and plus superscripts denote the prior and posterior covariance matrix $\mathbf{P}$ of the GP model, respectively.

In Equation 6.11, the main insight is to scale Rényi's entropy of the posterior GP model $\hat{\mathrm{H}}_{\alpha(\boldsymbol{\Sigma})}(\mathbf{P}^{+})$ by the robot pose uncertainty $\boldsymbol{\Sigma}$ through the Rényi parameter $\alpha(\boldsymbol{\Sigma})$. Using this strategy, the value of expected information gain is effectively coupled to the future actions of the robot, which enables managing the trade-off between robot localization and field mapping in a meaningful way, without the need for any additional parameters.

The final open issue is how to relate α to $\boldsymbol{\Sigma}$ in this formulation. Intuitively, the expected information gain should decrease when the robot has either a high uncertainty in the field map or a high uncertainty in its pose. The former exploratory behavior is captured by Shannon's definition. The latter exploititative behavior is encoded in the posterior Rényi's entropy via α. Once again, it is worth considering the two extreme scenarios. When the robot is lost (high values of $\boldsymbol{\Sigma}$), $\alpha \to 1$ is desired to maximize the field map uncertainty; when localization is perfect (low values of $\boldsymbol{\Sigma}$), $\alpha \to \infty$ is desired to minimize it.

To meet these requirements, this study considers a simple relationship between α and the A-optimality criterion, presented by Carrillo et al., 2018:

$$\alpha(\boldsymbol{\Sigma}) = 1 + \frac{1}{\mathrm{Tr}\,(\boldsymbol{\Sigma})} \,. \tag{6.12}$$

In this work, the A-optimality measure is chosen over alternative information criteria (D-optimality and E-optimality) based on the findings of Sim et al. (2005). Moreover, in the experiments in Section 6.4, it was found to have the highest variability in empirical values, making it the most sensitive to changes.

As a concluding remark, it is important to note here the integration of the Rényi-based information objective in Equation 6.9 with the modified covariance function in Equation 6.1 for environmental field modeling. By combining the two, the robot pose uncertainty is propagated into both the mapping and planning modules.

6.3.2 Uncertainty Prediction

A key requirement for predictive planning is being able to propagate the robot localization uncertainty for a potential future action. Various strategies, based on different front and back end methods, have been put forward in recent literature to address this issue (Kaess et al., 2009; Carrillo et al., 2018; Cadena et al., 2016 and Papachristos, Khattak, et al., 2017). This thesis considers methods of uncertainty propagation considering two different state estimation setups: (1) a graph-based SLAM back end system (for unknown environments, as in Sections 6.4.1 and 6.4.2), and (2) a Monte Carlo Localization (MCL) approach (for known environments, as in Section 6.4.3).

The graph-based SLAM solution considered is most closely related to the approach of Carrillo et al. (2018). Given a candidate trajectory, the main idea is to extend the current graph-based SLAM representation using a series of heuristic rules and then use optimization to extract the final predicted value of $\boldsymbol{\Sigma}$ corresponding to this plan. In contrast to their pose graph SLAM formulation with dense mapping, however, this work also includes the states of landmarks in a full optimization problem, which are received as Euclidean point measurements in 2-D or 3-D space, e.g., AR tags detected in images (Solà, 2017).

Figure 6.2 shows a schematic 2-D example of the uncertainty propagation scheme. To extend the graph from the current robot pose $\mathcal{N}(\mathbf{p}_1, \mathbf{\Sigma}_1)$, a candidate trajectory (dashed line) is interpolated at a specified fixed frequency to add $K-1$ new nodes (hollow circles) with probabilistic odometry constraints as edges, giving rise to the sequence $\{\mathcal{N}(\mathbf{p}_1, \mathbf{\Sigma}_1), \ldots, \mathcal{N}(\mathbf{p}_K, \mathbf{\Sigma}_K)\}$.

For each pair of consecutive new nodes with mean poses $\mathbf{p}_k$ and $\mathbf{p}_{k+1}$, the simulated control input vector $\hat{\mathbf{u}}$ is considered as being corrupted by Gaussian noise, given by:

$$\hat{\mathbf{u}} = \mathbf{u} + |\mathbf{u}| \cdot \boldsymbol{\kappa}\,, \tag{6.13}$$

where $\mathbf{u}$ is the true control vector corresponding to the transition between $\mathbf{p}_k$ and $\mathbf{p}_{k+1}$ as given by the odometry model and $\boldsymbol{\kappa}$ is a vector of sampled variances drawn from the Gaussian noise distribution $\mathcal{N}(0, \mathbf{Q})$ defined by the user. In other words, the control noise variance level is proportional to the magnitude of the input, in order to cater for the fact that longer motion steps are likely to have higher actuation errors associated with them. By setting the elements of the covariance matrix $\mathbf{Q}$, the uncertainty in the control system can be defined in each dimension of the robot workspace.

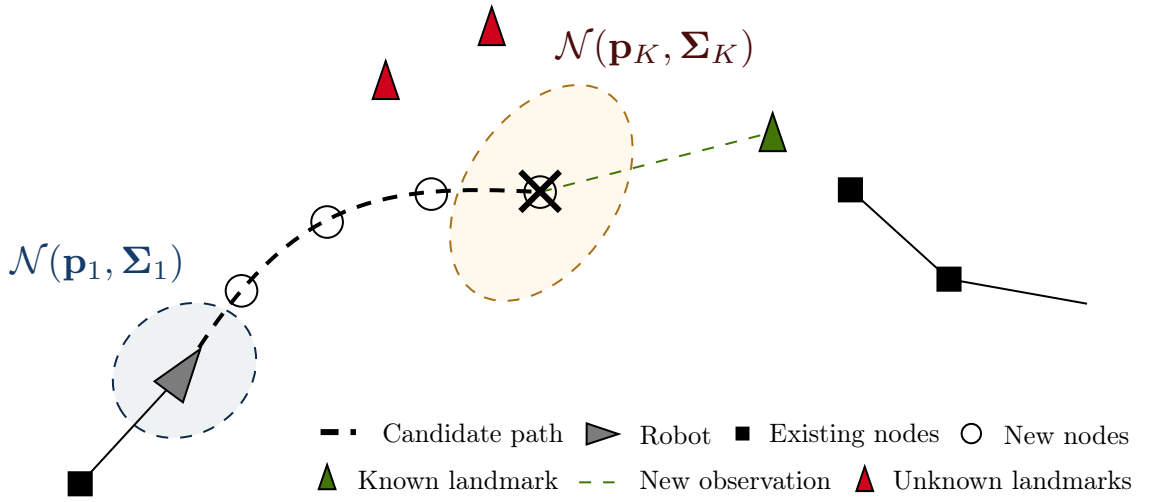

Figure 6.2: Schematic explanation of the uncertainty propagation scheme using a graph-based SLAM back end based on point landmarks. The robot is on the bottom-left and the dashed line is a candidate path to a target location 'X'. Example uncertainty ellipses at the initial and final positions are shown. To predict the evolution of the robot pose $\mathcal{N}(\mathbf{p}, \mathbf{\Sigma})$, new nodes (hollow circles) are interpolated along the path to extend the existing graph. The green dotted line corresponds to the potential re-observation of an existing landmark (green triangle), while new landmarks (red triangles) in unknown space are ignored. The extended graph is optimized to obtain the final pose estimate $\mathcal{N}(\mathbf{p}_K, \mathbf{\Sigma}_K)$ for uncertainty-aware planning.

Unlike in dense mapping applications, predicting the future uncertainty in loop closure events is simpler with point landmarks since the measurements taken to localize the robot are binary, i.e., a landmark is either visible from a given pose or it is not. To address potential loop closures, the proposed approach simulates re-observations to the known landmarks (green triangle) already maintained in the state vector of the current graph, given the sensor model. In unknown space, it is simply assumed that no new landmarks will be detected, causing the localization uncertainty to grow. Once the extended graph has been created, it is then optimized using the standard sparse QR factorization method (Solà, 2017). Finally, the predicted covariance matrix $\mathbf{\Sigma}_K$ is recovered from the most recent node in the optimized graph solution, following the technique of Kaess et al. (2009).

Furthermore, this work also proposes a simple uncertainty propagation method for the case of a localization-only problem in a known environment. A particle filter-based MCL algorithm is assumed (S. Thrun et al., 2006). Given a known map in this scenario, the predicted robot uncertainty can be obtained by forward simulating the motion model and expected sensor measurements. In Section 6.4.3, this approach is demonstrated for localizing a ground robot equipped with a laser scanner in a cluttered indoor area.

6.4 Experimental Results

This section presents a series of experiments using the proposed approach for active sensing under pose uncertainty. In Sections 6.4.1 and 6.4.2, comprehensive evaluation studies are conducted in various simulated scenarios. Considering that the planning and mapping tasks in the framework are tightly intertwined, the methods are benchmarked against existing algorithms in terms of these two aspects separately. Finally, Section 6.4.3 presents a proof of concept deployment using a ground robot to map temperature in an indoor environment.

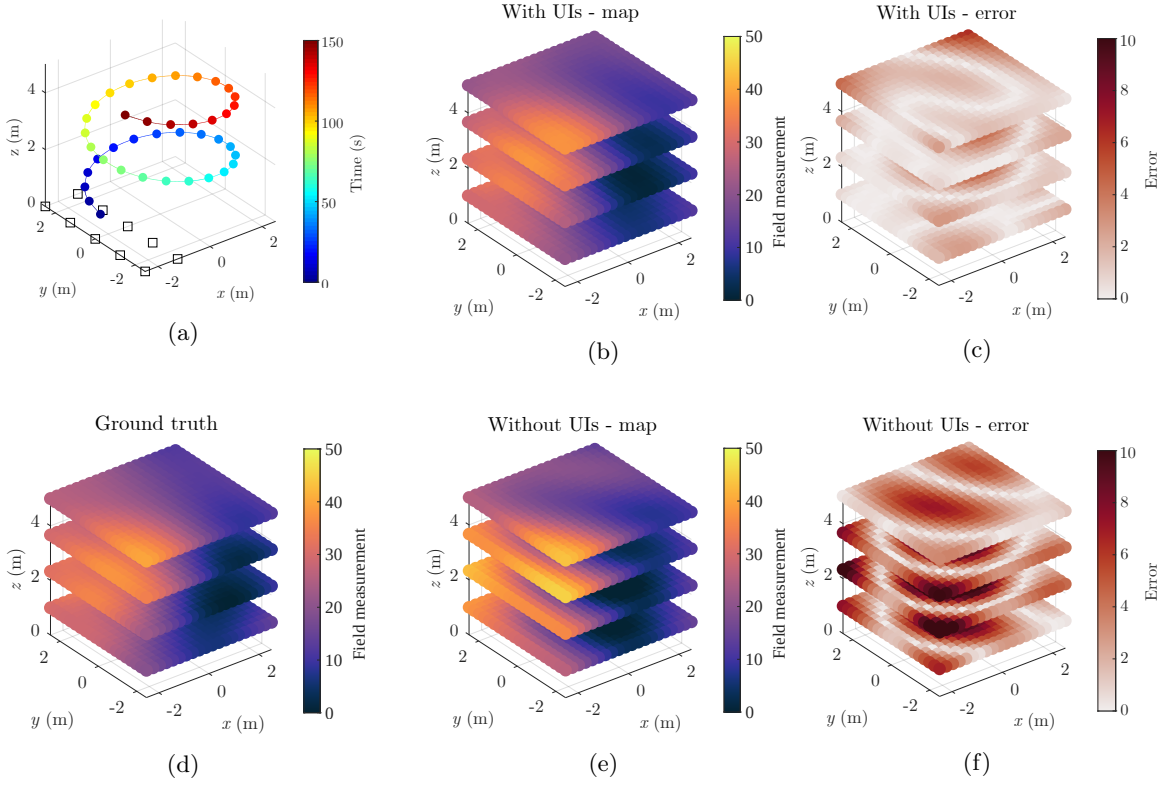

Figure 6.3: Overview of the approach for active sensing under pose uncertainty. (a) represents a deterministic spiral trajectory traveled by a robot. The squares indicate point landmarks on the ground used for localization. The spheres represent sites where measurements of the ground truth field map in (d) are taken. (b) and (e) show the final maps built using the proposed mapping method (with uncertain inputs (UIs)) and standard GP-based mapping (without UIs), respectively. (c) and (f) visualize the corresponding error distributions. Accounting for the pose estimation uncertainty, the new strategy yields more consistent, higher-quality reconstructions.

To start with, Figure 6.3 presents an simulated example showing the key features of the proposed approach. The setup considers mapping a continuous field in a 3-D environment using a Unmanned Aerial Vehicle (UAV) equipped with a sensor to take point measurements. A downward-facing camera is used to detect the point landmarks on the ground for localization. For illustrative purposes, the example considers the sensor continuously sampling the field along a deterministic spiral trajectory as shown in (a). The plots in (b) and (e) visualize the final field reconstructions build using the proposed method for mapping under pose uncertainty (with UIs) and a standard GP-based strategy assuming perfect localization (without UIs), respectively. The corresponding error distributions with respect to the ground truth in (d) are shown in (c) and (f). The lighter shades in (c) confirm that the new approach achieves significantly lower error by being able to produce consistent map estimates. Quantitatively, the approach with UI produces a final map with 2.47 times lower total Root Mean Squared Error (RMSE). In the following two sections, similar simulation setups are used to evaluate the algorithms.

6.4.1 Comparison Against Planning Benchmarks

The aim of these experiments to evaluate the new uncertainty-aware utility function (Section 6.3.1) by comparing it against existing strategies. To focus on assessing performance in terms of planning, all methods in this sub-section use the GP-based field representation with UIs (Section 6.2) as the method for online mapping under pose uncertainty.

The simulations are designed to validate the mapping and planning algorithms working in tandem in a complex 3-D environment. To this end, the setup considers information gathering in a $5\,\mathrm{m} \times 5\,\mathrm{m} \times 4\,\mathrm{m}$ volume, representing both the extent of the the environmental field as well as the robot workspace. Five different continuous field distributions are generated as 3-D Gaussian random fields with the target scalar parameter ranging from 0 to 50. A visualization of an example ground truth distribution is presented in Figure 6.3(d).

For GP-based mapping, a lattice grid of query points $\mathbf{X}_*$ with a resolution of $0.25\,\mathrm{m} \times 0.25\,\mathrm{m} \times 1\,\mathrm{m}$ (2880 points total) is specified in the environment. Since the ground truth map is also defined based on this grid spacing, a spline-based interpolation method is used to sample point measurements as the robot moves in continuous space. Within the GP model, the environmental phenomenon is characterized using a constant mean function $m(\mathbf{x}) = m_{gp}$ and the widely-used isotropic Squared Exponential (SE) covariance function. The latter has the form (Rasmussen et al., 2006):

$$k_{SE}(\mathbf{x}, \mathbf{x}') \triangleq \sigma_f^2 \exp\left(-\frac{d^2}{2l^2}\right), \tag{6.14}$$

where d is the Euclidean distance between inputs $\mathbf{x}$ and $\mathbf{x}'$, and l and σ_f^2 are hyperparameters representing the characteristic length scale and signal variance, respectively. Note that this is equivalent to the Matérn covariance function in Equation 5.1 in the limit as $v \to \infty$. The hyperparameter set of the GP $\theta \triangleq \{m_{gp}, \sigma_n^2, \sigma_f^2, l\} = \{17.61, 2.13 \times 10^{-11}, 1141.38, 2.18\}$ was trained by minimizing the Negative Log of the Marginal Likelihood (NLML) in Equation 3.8 using the Polack-Ribiére conjugate gradient method in each environment.

The modified version of the SE covariance function, defined in Equation 6.1 is used to propagate the pose uncertainty into the map. It is approximated using the Gauss-Hermite quadrature rule in Equation 6.6 with 5 sample points. For data acquisition, measurements of the field are taken at a constant frequency of 0.25 Hz using a point-based sensor centered on the robot. During the mission, the samples observed are added as input training points $\mathbf{X}$ to the GP model. The regression procedure in Equations 3.6 and 3.7 is performed in an online fashion during replanning, with the

uncertainty of the map expected to decrease as the robot explores the environment, provided the localization uncertainty stays bounded.

The simulation environment is conceived to emulate a UAV-based monitoring scenario. For localization, 10 visual 3-D point landmarks are placed on the ground in a uniform grid arrangement on one side of, and 1 m below, the target field, as shown in Figure 6.3(a). To detect the landmarks, the robot is equipped with a downward-facing depth camera with a Field of View (FoV) of (47.9°, 36.9°) and standard deviations of 1.0 px and 0.1 m to model measurement errors in pixel and depth based on a pinhole projection camera model. The landmarks are distributed in such a way in order to make state estimation more difficult and highlight the benefits of informative planning under pose uncertainty. Moreover, this layout is representative of realistic environments where access to the entire environment is not possible, e.g., a UAV localizing based on AR tags in a greenhouse to precisely monitor various plant parameters.

The landmark locations are assumed initially unknown such that a SLAM approach is necessary. The framework is interfaced with a graph-based SLAM back end (Solà, 2017) using odometry and landmark measurements to constrain the graph states. For predictive planning, the robot pose uncertainty is propagated using the method described in Section 6.3.2. To extend the structure of the current graph for a candidate action, a constant frequency of 0.5 Hz is used to interpolate new nodes along the future trajectory using an odometry motion model. The covariance matrix $\mathbf{Q}$ specifying the control noise distribution is considered as a diagonal matrix with values of 0.01 in all three co-ordinate dimensions. As mentioned, graph optimization is performed using the sparse QR factorization method, and the predicted covariance matrix of the robot state is extracted from the final node of a given trajectory through the computations developed by Kaess et al. (2009)

The experimental procedure follows in a similar way as for the simulations in Chapter 5, whereby existing strategies are considered as evaluation benchmarks. To investigate the benefits of accounting for the pose uncertainty during planning, the Rényi-based utility function is also compared against variations of the proposed informative planning strategy using different information objectives. Therefore, the main question under investigation is: how does the choice of the utility function affect performance, provided that the mapping method and planning algorithms are the same? The following strategies are considered:

(a) Proposed approach using the new Rényi-based utility function
Equation 2.1 with Equation 6.11 ('Rényi')

(b) Proposed approach using a utility function that only minimizes Tr($\mathbf{P}$)
Equation 2.1 with Equation 4.11 ('Unc.')

(c) Proposed approach using a rate-based utility function that only minimizes Tr($\mathbf{P}$)
Equation 2.2 with Equation 4.11 ('Unc. rate')

(d) Proposed approach using a linear composite utility function
Equal weighting of the map and pose uncertainties, tuned according to the method of Bourgault et al. (2002) ('Weighted')

(e) The Rapidly exploring Information Gathering (RIG)-tree algorithm using a utility function that only minimizes Tr($\mathbf{P}$)
Introduced by Hollinger and Sukhatme (2014) and described in Sections 5.2.1 and 5.2.2 ('RIG-tree - Unc.')

(f) The RIG-tree algorithm using the new Rényi-based utility function
('RIG-tree - Rényi')

(g) Random waypoint selection ('Random')

Each method above was run over 50 mission trials in each of the 5 environments with a fixed mission time budget of $B = 150\,\text{s}$, giving a total of 250 trials.

As in Section 5.2.2, without an optimal solution available for this problem, the tested algorithms are assessed based on the variation of different information metrics during a mission. To evaluate *mapping* performance, the criteria are: the total map uncertainty, quantified by the trace of the covariance matrix of the GP field model $\text{Tr}(\mathbf{P})$ and the total map RMSE with respect to the ground truth map. Similarly, to evaluate *planning* performance, the measures are: the trace of the robot covariance matrix $\text{Tr}(\mathbf{\Sigma})$ and the robot pose RMSE with respect to the ground truth trajectory, obtained from the simulation environment. Intuitively, all metrics are expected to decrease over time, with lower values indicating better performance.

The initial robot position is specified as $(2\,\text{m}, 2\,\text{m}, 1\,\text{m})$ in the environment considering the ground with landmarks as the reference plane in the z co-ordinate. No prior information about the landmarks or target field is assumed. For minimum-snap trajectory generation, the maximum reference velocity and acceleration are $1.5\,\text{m/s}$ and $3\,\text{m/s}^2$ using polynomials of order $k = 12$. For the variants of the proposed informative planning approach, polynomial paths are defined by $N = 4$ control waypoints. The 3-D search for an initial trajectory is performed using a $(5.75\,\text{m}, 5.75\,\text{m}, 2\,\text{m})$ resolution grid (27 points total) centered in the robot workspace, with scales chosen to trade-off between solution accuracy and computational feasibility. The CMA-ES runs with initial step sizes $\boldsymbol{\sigma}_{\text{CMA}}$ of $(1\,\text{m},\ 1\,\text{m},\ 0.5\,\text{m})$ in the (x, y, z) co-ordinates set empirically based on the results of prior trials, a population size of λ_{CMA} of 10, and a maximum number of iterations N_{CMA} of 20.

For a fair evaluation, an online, finite-horizon variant of the RIG-tree algorithm is developed using the principles described in Sections 5.2.1 and 5.2.2. As in Section 5.2.2, the planning strategy operates in the space of continuous trajectories using the constant-frequency sensor model. Essentially, the method samples a tree of possible plans in the robot workspace, in which control waypoints are associated with vertices, and polynomials are formed by tracing the parents of leaf vertices to the root. For brevity, further details concerning the algorithm are not repeated here and the reader is referred to Sections 5.2.1 and 5.2.2.

Importantly, for this method, two utility functions are considered: in 'RIG-tree - Unc.', the objective of only maximizing the reduction of map uncertainty, as measured by $\text{Tr}(\mathbf{P})$ (Equation 2.1 with Equation 4.11), in order to correspond to original implementation of the algorithm, presented by (Hollinger and Sukhatme, 2014); and, in 'RIG-tree - Rényi', the new Rényi-based objective (Equation 2.1 with Equation 6.11), in order to demonstrate its applicability with a different planner. For replanning, the branch step size parameter used to expand the tree is tuned as $5\,\text{m}$ to achieve best performance in the environment, based on a manual search over a discrete range of values.

Finally, using the naïve random planning strategy, a specified number of waypoint destinations are randomly sampled in the workspace, and a trajectory is generated by connecting them sequentially the the current position of the robot. In these experiments, 4 waypoints per plan are considered to ensure that the lengths of these trajectories are roughly comparable to those produced by the proposed algorithm with $N = 4$. Though it has similar effects, note that this approach is different from the random planner designed in Section 5.2.2, which performs point-to-point navigation.

Figure 6.4 shows how the informative metrics evolve for the tested planning algorithms over the 250 trials. As expected, in Figure 6.4(b), the informative strategies perform better than the random benchmark (blue) using the mapping criteria (top row) as they are guided by planning objectives.

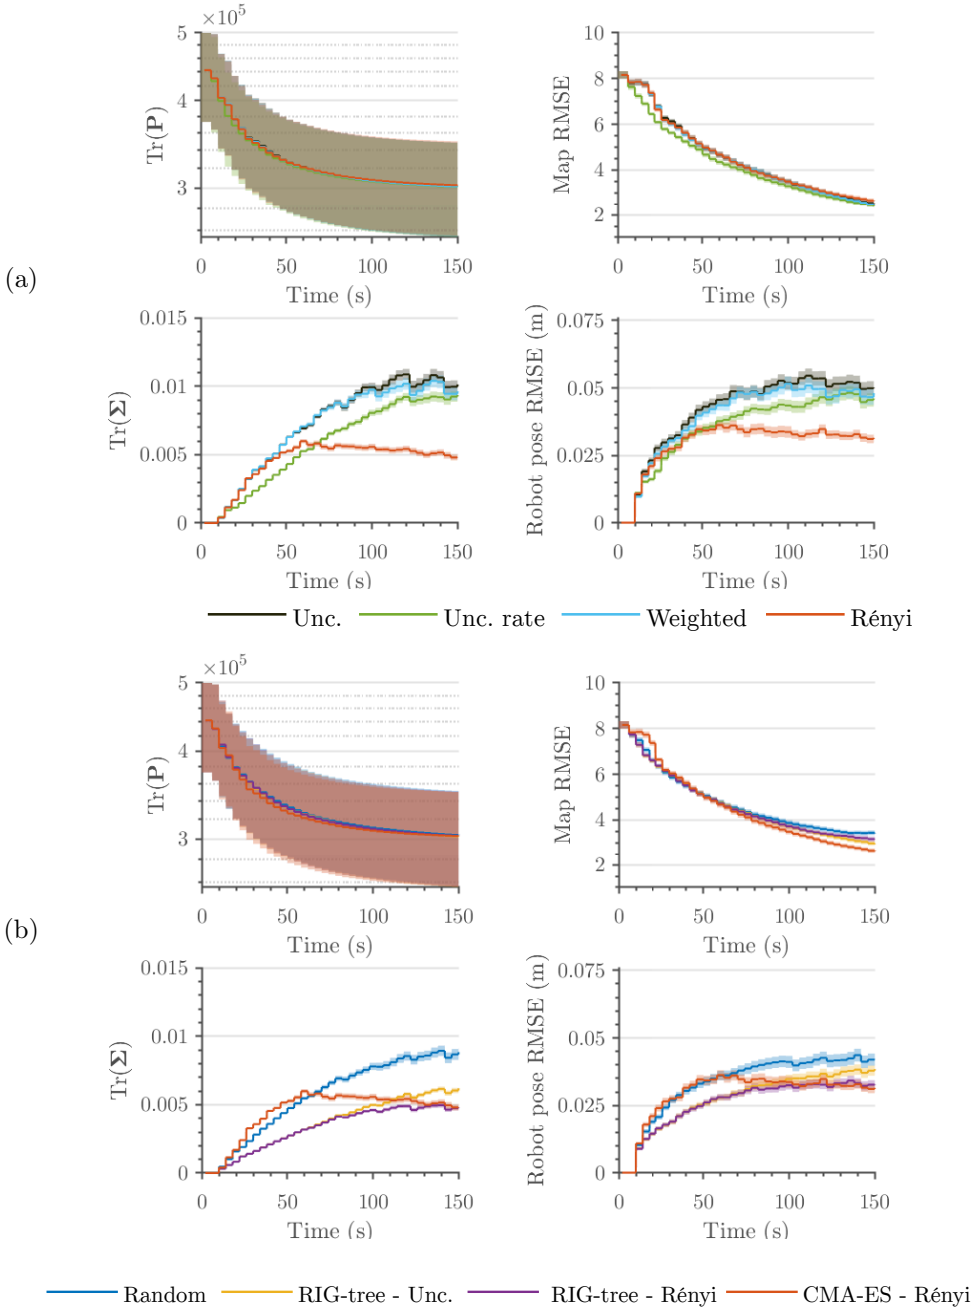

Figure 6.4: Comparison of the proposed two-step CMA-ES-based replanning framework using the Rényi-based utility function against (a) the same framework using different information objectives and (b) different planning benchmarks. All methods are given a fixed budget of 150 s and use GP-based field mapping with UIs. For each algorithm, average results over 250 trials are shown with 95% confidence bounds. By considering both the robot pose and field map uncertainties, the new utility function more quickly achieves higher-quality mapping (top row) with improved localization (bottom row). Note that $\mathrm{Tr}(\mathbf{P})$ is on a logarithmic scale.

Looking at Figure 6.4(a), a surprising result to emerge is that the random planner produces slightly improved localization (bottom row) compared to using the proposed approach with utility functions that target pure uncertainty reduction (black, green). Possibly, this is because these strategies are biased towards exploring the environment very aggressively, causing the robot pose uncertainty to escalate when measurements are taken above the area from which no landmarks can be seen. In contrast, random navigation is likely to enable some regular re-observations as landmarks haphazardly enter the camera FoV.

In Figure 6.4(a), the poor localization performance obtained with the pure exploration strategies (black, green) and the weighted utility function (cyan) motivates an approach which incorporates the pose uncertainty of the robot into the planning objective. This solution is provided by the Rényi-based utility function (red). The results show that the new objective achieves comparable rates of reduction in total map uncertainty and error, while significantly improving upon pose estimation in the same scenario. Therefore, it effectively trades off between gathering new sensor information and maintaining good localization with respect to the landmarks, as desired.

Another interesting observation is that, among the three variants of the proposed informative planning approach, the uncertainty-only utility function (black) performs the worst against all evaluation criteria. As opposed to the rate-based and Rényi-based objectives, it does not exploit any knowledge about the trajectory dynamics to create new plans; allowing the robot to travel long, high-cost paths to gather information in the least conservative way. This finding justifies the planning strategy adopted in Chapter 4.

As expected, in Figure 6.4(b), the RIG-tree variant using the Rényi-based (purple) objective provides better localization when compared against planning for map uncertainty reduction only (yellow). This further validates the expected behaviour of the proposed utility function, as described above, using a different replanning routine. An unexpected result is that the RIG-tree algorithm appears to be relatively successful at keeping the robot well-localized, whereas its capabilities in terms of field mapping are much poorer. Such behavior could be related to the branch step size parameter, which limits the range of navigation achievable during the mission, i.e., a lower value for the step size imposes shorter distances between successive waypoints along a trajectory, which constrains the robot motion. By traveling shorter trajectories, it is possible that the robot re-observes the landmarks more frequently at the cost of restricted exploration. Further insights along these lines can be found in Section 5.2.2. In general, one would expect this parameter to depend highly on the environment, especially in the problem examined here in which its internal structure, i.e., landmark arrangement, is important as well as its size. Although future studies are required to draw more complete conclusions, the apparent sensitivity of the algorithm to this parameter emphasizes the benefits of using the proposed tuning-free approach.

As an example, Figure 6.5 illustrates the result of a simulation trial performed using the Rényi-based utility function for planning in the CMA-ES based framework. The trajectory visualized in (a) confirms that the robot successfully explores the entire environment, while also re-visiting the space above the known landmarks to maintain good localization. The final field map produced, shown in (b), has a total RMSE value of 1.11, which is 1.86 times lower than the one produced by the naïve spiral path in Figure 6.3 considering the same problem scenario and mapping approach. This improvement justifies the use of the proposed informative planning strategy, as opposed to a simpler approach based on deterministic sampling.

6.4.2 Evaluation of Field Mapping Under Uncertainty

The next set of experiments evaluates the effects of incorporating UIs in the GP environmental field model (Section 6.2). The aim is to assess the benefits of mapping under the robot pose uncertainty

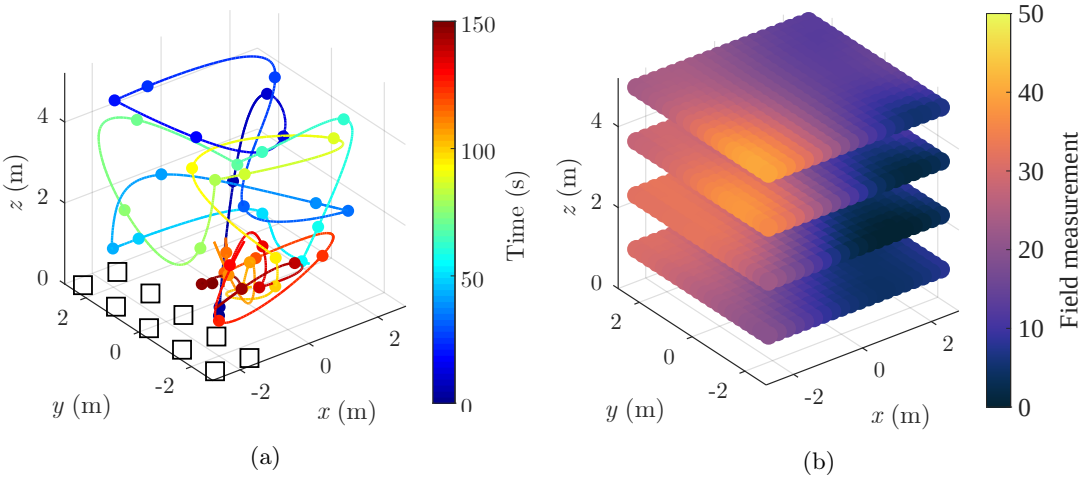

Figure 6.5: Example simulation result of informative planning using the proposed Rényi-based utility function in a 150 s mission. (a) shows the trajectory traveled by the robot during the mission. The squares indicate point landmarks used for localization. The spheres represent sites where measurements are taken to produce the final field map in (b). By balancing between gathering information and keeping the landmarks in view, the new planning objective achieves a field map with 1.86 lower RMSE compared to the spiral path shown in Figure 6.3.

in active sensing problems. The setup used for these simulations is the same as in the previous subsection, i.e., a robot monitoring a continuous variable in a 5 different $5\,\mathrm{m} \times 5\,\mathrm{m} \times 4\,\mathrm{m}$ environment using a point-based sensor to take measurements, and 10 point landmarks used for localization based on a depth camera and a graph SLAM back end.

To investigate mapping performance in isolation, the proposed Rényi-based utility function is used to define the planning objective in both of the two compared methods. The experiments consider the following strategies: (a) with and (b) without applying the modified SE covariance function in Equation 6.6 for building field maps online. These methods are referred to as 'With UIs' and 'Without UIs', respectively. Note that (b) corresponds to a benchmark mapping approach using standard GP regression, as shown in Figure 6.3. As in Section 6.4.1, the simulations were repeated over 50 trials each of the 5 different environments. The same criteria are computed to evaluate the methods in terms of field mapping and robot localization performance. However, for these tests, the total map uncertainty metric $\mathrm{Tr}(\mathbf{P})$ is omitted. As discussed by M. G. Jadidi et al. (2017), this is because the variance scales in the GPs built using (a) and (b) are not directly comparable, since the propagation of UIs in (a) adds inherent uncertainty to the model.

The results of the experiments are shown in Figure 6.6. The observed trends in the localization metrics (middle, right) confirm that the proposed mapping approach with UIs (red) achieves more conservative exploratory behavior than the benchmark (blue), since the uncertainty of the robot pose remains low. Simultaneously, it permits more accurate reconstructions of the target field (left). This supports the hypothesis that the modified covariance function can handle localization errors to build maps with better consistency and quality for more reliable planning, as desired. A visual validation of this result is provided by the example in Figure 6.3.

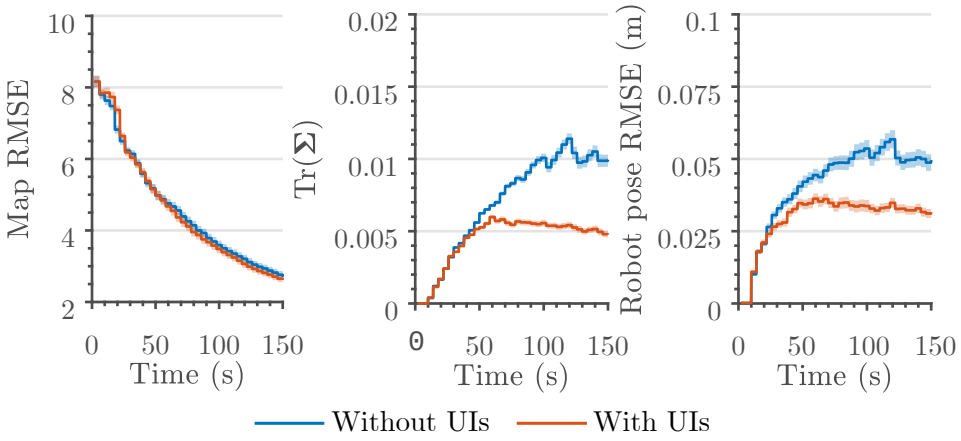

Figure 6.6: Comparison of GP-based field mapping with and without uncertain inputs (UIs) using the CMA-ES-based replanning framework and the Rényi-based planning objective for a fixed mission time budget of 150 s. For each algorithm, average results over 250 trials are shown with 95% confidence bounds. By accounting for the robot pose uncertainty, the new approach with UIs achieves significantly more conservative mapping behavior (middle, right) with improved accuracy in the field reconstructions (left).

6.4.3 Temperature Mapping

The final experiment demonstrates the complete active sensing system deployed on a ground robot for temperature mapping in an indoor environment. The aim is to validate the integrated framework and its interfaces in real-world conditions.

Figure 6.7 shows the setup. The robot platform, shown in (a), is a TurtleBot3 Waffle navigating in an empty 2.8 m × 2.8 m area. For mapping, a varying temperature gradient is generated by using a 2400 W convector radiator placed at one corner of the environment as a heat source, as shown in the thermal images in (b) and (c). Point measurements of the field are taken using a LM35 linear temperature sensor mounted on top of the robot frame, which has a sensitivity of 10 mV/°C. For state estimation, the Adaptive Monte Carlo Localization (AMCL) algorithm is used considering a known environment based on a prior map and data from an on-board 360° LDS-01 laser distance scanner. To ensure that enough features are available for reliable dense mapping, the area outside the workspace bounds is cluttered with objects.

Low-level functionality runs on an on-board Intel Joule 570x module with a 1.7 GHz processor and 4 GB of RAM running Ubuntu 16.04 LTS. The implementation is based on the TurtleBot3 packages for Robot Operating System (ROS)[1]. High-level functionality, including mapping and informative planning, runs in MATLAB on a a single desktop computer with a 1.8 GHz Intel i7-8550U processor and 16 GB of RAM, running Ubuntu Linux 16.04 LTS and the Robotics System Toolbox as the real-time communication interface. Note that the uncertainty propagation step is implemented in MATLAB using the AMCL implementation from the toolbox and a pre-loaded occupancy map of the environment.

The proposed GP-based technique with UIs is used to map the temperature field distribution under pose uncertainty. For regression, the query points are spaced on a grid in the environment with a

[1] `github.com/ROBOTIS-GIT/turtlebot3`

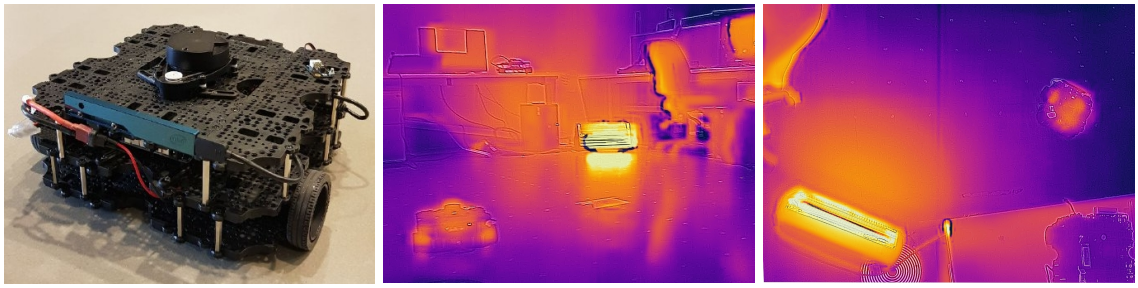

(a) (b) (c)

Figure 6.7: (a) shows a close-up of the TurtleBot3 deployed for temperature monitoring. The LM35 temperature sensor is attached on the top-right. (b) and (c) visualize thermal imagery of the experimental setup from side and aerial viewpoints, respectively. The robot and radiator are visible. The yellower shades in the images (close to the radiator) correspond to the more heated areas of the environment that are mapped in the experiment.

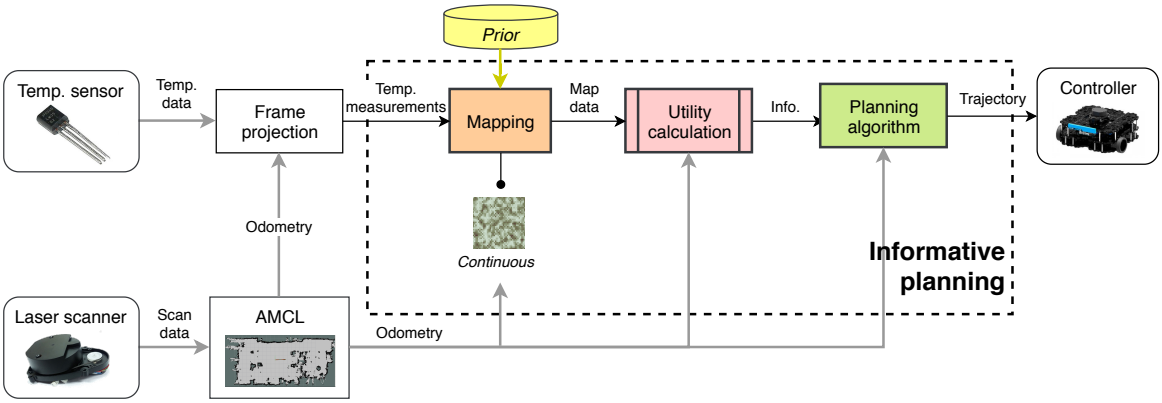

Figure 6.8: System diagram for the temperature mapping experiments. The Adaptive Monte Carlo Localization (AMCL) algorithm is used to localize the robot in a known map using laser scan data. The pose estimate is propagated into the mapping and planning modules to direct the acquisition of measurements for building online maps of a target temperature distribution.

uniform resolution of 1.4 m (9 points total), set for computational feasibility. As in the preceding subsections, the isotropic SE kernel in Equation 6.14 is applied to capture the correlation structure of the field within the GP. The hyperparameter set $\theta \triangleq \{m_{gp}, \sigma_n^2, \sigma_f^2, l\} = \{23.64, 3.6 \times 10^{-3}, 1.21, 5.23\}$ was obtained by minimizing the NLML function using manually acquired data within the target heated area. The Gauss Hermite quadrature approximation in Equation 6.5 with 5 sample points is used to compute the modified kernel in the 2-D mapping application.

To demonstrate all aspects of the proposed approach, the planning objective is defined using the Rényi-based utility function in Equation 6.11. A mission time budget of $B = 600\,\text{s}$ is specified for the test. The position for the initial measurement is $(0.2\,\text{m}, 0.2\,\text{m})$ within the corner of the environment opposite the radiator. Considering that the robot is controlled with a differential drive, each candidate plan is piecewise linear as defined by 3 control waypoints with a constant velocity of 0.26 m/s. Measurements along the path are sampled at 0.25 Hz to map the temperature distribution. Due to the latency of the LM35 sensor, note that the robot was allowed to wait between waypoints for the temperature readings to stabilize before performing the map update.

For planning under pose uncertainty, the method for predicting the robot covariance matrix is based on localization in a known environment using AMCL, as described in Section 6.3.2. Each candidate plan is interpolated at a frequency of 2 Hz. At each sampled point, the robot pose is estimated by forward simulating a differential drive odometry model and laser scans projected in the known occupancy map. The odometry model assumes a Gaussian noise distribution $\mathcal{N}(0, 0.2)$ applied to both rotational and translational motion.

Putting together the elements above, Figure 6.8 presents a system diagram illustrating the infrastructure of the integrated system. Using the proposed strategy, the robot odometry data is propagated into both the mapping and planning modules. Note that this setup can be easily recognized as an

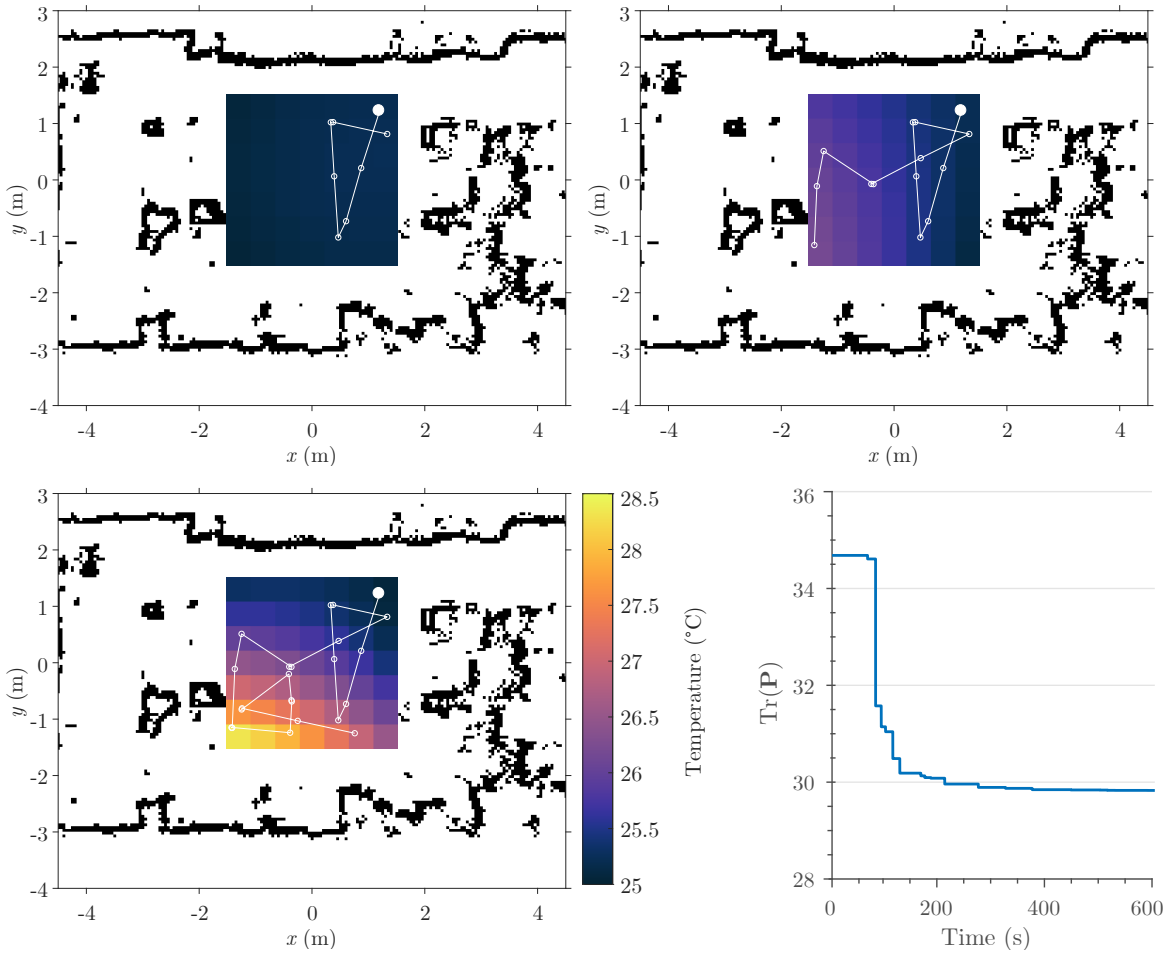

Figure 6.9: Experimental results from using the proposed active sensing framework to map the indoor temperature distribution in Figure 6.7 in a 600 s mission. The sequence of three plots depicts the trajectories (white lines) and temperature field maps (colored gradients) at different snapshots of the mission at times $t = 100$ s, 350 s, and 600 s. The white circles represent measurement sites, with the large solid one indicating the initial robot position. Yellower shades correspond to hotter regions. The sequence shows that the new strategy quickly explores the area, successfully detecting the heated corner (bottom-left) where the radiator is located. The curve on the bottom-right validates the reduction in the field map uncertainty over time. Note that planning time is taken into account.

adaptation of the general architecture shown in Figure 5.1 for the problem of active sensing under pose uncertainty.

The results of the field test are summarized in Figure 6.9. As in the continuous field mapping experiments in Sections 5.4.2 and 5.4.3, without reliable ground truth data, the approach is validated by assessing the variation of the total map uncertainty over time. The three figures visualize the state of the temperature field (GP means) at various stages of the mission within the occupancy map of the indoor space. The white lines represent the histories of the trajectories traveled by the robot at those instances. As expected, the field map becomes qualitatively more complete over time as the heated (yellower) region in front of the radiator is discovered. Moreover, the plot on the bottom-left validates that the total uncertainty of the map reduces as data are acquired. These findings provide a proof of concept implementation of the proposed approach. Towards further validation, the next step is to assess how the method performs for this task when compared against the benchmarks introduced in the previous sub-sections.

In the field maps rendered in Figure 6.9, it can be seen that the mean temperature value of the trained GP model (23.64°) is lower than the range of measurements obtained from the sensor during the mission ($25 - 28.5°$). This is because, as the robot was mapping the field, the heat from the radiator diffused throughout the environment, gradually increasing its temperature. However, at the time during which the training data was collected, the room was relatively cool. These effects made it very difficult to obtain a static temperature distribution to satisfy the assumptions underlying the GP-based mapping strategy. One could address this by using an alternative experimental setup with a target phenomenon that is easier to control, e.g., light. More interestingly, future work could tackle the problem of including an additional state within the map representation to capture the temporal dynamics of the field distribution.

As a final remark, note that the replanning step in this setup took $\sim 10\,\text{s}$. Though this is acceptable for a proof of concept in a small environment, methods of improving computational efficiency are necessary for the approach to scale to more practical applications. In the experiment, two major bottlenecks were identified: the GP regression step for mapping with UIs and the pose uncertainty propagation procedure using AMCL for each candidate trajectory. As demonstrated in Section 5.4.3, an implementation completely based in ROS would speed up the algorithms significantly. Ultimately, however, further studies addressing these issues directly will need to be undertaken.

6.5 Summary and Discussion

This chapter introduced methods to address the active sensing problem in scenarios where robot localization is uncertain. Building upon the framework developed in Chapter 5, the proposed approach accounts for the pose uncertainty in both the mapping and planning algorithms. The mapping strategy leverages a GP model which accommodates UIs by computing the expectation of the covariance function in 2-D or 3-D space. For planning, a new tuning-free utility function was presented that tightly couples the uncertainties in the robot pose and field map. This is achieved by applying the concept of Rényi's entropy considering that the target environmental phenomenon is represented as a continuous field.

System integration was demonstrated by interfacing these methods with the planning algorithm from Chapter 4. Extensive evaluation results from simulations show that the approach is able to reconstruct 3-D target fields more accurately, and in a more conservative manner, in comparison to existing methods. Finally, a proof of concept implementation is demonstrated using a ground robot to monitor a temperature distribution in an indoor environment.

A major factor governing the performance of the approach is the way in which the robot localization uncertainty is predicted to evaluate candidate plans. This aspect is particularly important since a reliable future pose estimate is needed by both the mapping and planning modules to determine the next-best sensing actions. As described in Section 6.3.2, to propagate the pose uncertainty, the SLAM solution proposed in this thesis optimizes an extended graph representation with additional nodes interpolated along a potential trajectory. Though this strategy was accurate enough for experimental purposes, its main disadvantage is that it requires manually tuning various parameters, i.e., the frequency at which the trajectory is sub-sampled, the covariance matrix $\mathbf{Q}$ of the noise distribution associated with the robot motion, the sensor noise distribution, and the predicted visibility of future landmarks. Moreover, the method is restricted to a graph-based SLAM back end with point landmarks, and therefore not suitable for dense occupancy mapping scenarios, e.g., as considered in the setup in Section 6.4.3.

Generally speaking, the task of pose uncertainty propagation raises many open research questions that must be addressed to improve upon information gathering behavior. As noted by Cadena et al. (2016), a major challenge is predicting the effects of uncertain loop closure events in an efficient, principled manner. To this end, Carrillo et al. (2018) propose an approach which is similarly limited by relying on heuristic rules to extrapolate a graph. Considering a different setup, Papachristos, Khattak, et al. (2017) introduce a belief propagation technique relying on a modified version of the Robust Visual Inertial Odometry (ROVIO) framework. Their pipeline would be very interesting to integrate with the informative planning strategies here, especially in the context of UAV-based applications. Potentially, it could be used as a means of improving the map fidelity in the agricultural monitoring experiments presented in Section 5.4.3.

Section 6.4.3 briefly discussed the computational efficiency of the approach. During the experiment, the steps of GP inference and pose uncertainty propagation were identified to be the major bottlenecks hindering performance. The GP-based mapping method suffers from the usual cubic complexity in the number of input training points, which increases as measurements are collected during a mission. This problem is further exacerbated as the Gauss Hermite quadrature approximation is needed to account for UIs.

One solution is to consider a Bayesian technique for fusing new data into a prior map, as illustrated in Section 3.2 for the terrain monitoring problem. Such an approach would not only improve efficiency considerably, but also present an alternative way of accounting for measurement noise by using a probabilistic sensor model. In future work, it would be very interesting to investigate how this strategy performs for planning under pose uncertainty. Regarding the second bottleneck, any possible strategies of accelerating the uncertainty prediction procedure are likely to depend on the problem setup and the prediction method itself. For instance, using graph-based SLAM, one could study sparsification methods for more scalable implementations (Vallvé et al., 2018).

From the planning perspective, an open question for future work concerns methods of mapping the α parameter in the Rényi-based utility function to the uncertainty of the robot. This work employs the simple relationship with the A-optimality criterion indicated in Equation 6.12. However, as discussed by Carrillo et al. (2018), many more sophisticated formulations, e.g., using exponential or logarithmic functions, or techniques, e.g., Bayesian optimization, could be examined.

A further idea is to assess how the utility function behaves in an adaptive planning scenario, as described in Section 4.3.2. By introducing an additional consideration in the planning objective, this requirement opens abundant avenues for future research.

7

Conclusion and Future Outlook

A core challenge in active sensing is deciding how an agent should act to maximize the information it acquires about an uncertain environment, subject to a limited amount of sensing resources. To address this question, this thesis developed informative planning methods for general active sensing tasks. Specifically, it examined two distinct applications, with a focus on Unmanned Aerial Vehicle (UAV) systems: *terrain monitoring* and *active sensing under explicit localization uncertainty*. In each application, contributions in *environmental mapping* and *informative path planning* were presented as key components enabling autonomous data collection in unknown environments. The developments were unified into a single modular framework that can be easily deployed on a wide variety of platforms for different tasks, and further advances the role of robots as efficient monitoring tools.

Section 7.1 summarizes the contributions of this thesis and Section 7.2 provides an outlook on future research directions.

7.1 Summary of Contributions

Environmental Mapping

A new mapping method was proposed in the context of UAV-based terrain monitoring applications. The strategy uses a Gaussian Process (GP) model to capture the underlying structure and uncertainty of a continuous-valued environmental phenomenon to guide predictive planning strategies. Unlike conventional mapping approaches, which accumulate measurements in the training dataset over time, a GP is used as a prior for recursive Bayesian data fusion with probabilistic sensors. In this manner, the cubic complexity of standard GP regression is replaced with constant processing time in the number of measurements, which enables more efficient mapping in practical applications. Furthermore, a measurement model is presented that enables fusing information at different resolutions and from heterogeneous sources into a single environmental map. A key advantage of this approach is that it supports mapping with altitude-dependent sensors providing dense visual imagery, as relevant for aerial monitoring setups. The experiments in this thesis validate these aspects by leveraging the mapping method for online informative planning in various terrain monitoring scenarios.

Informative Path Planning

An online motion planning algorithm was presented for generating continuous informative trajectories in resource-constrained sensing missions. The approach operates in a fixed-horizon manner and is suitable for monitoring problems with adaptivity requirements. Essentially, the environmental map models built online are used to find trajectories that maximize an information-theoretic metric reflecting the planning objective while remaining within a budget constraint. A key feature of the

algorithm is that it uses the output from a discrete grid search as an informed prior to initialize an evolutionary optimization routine, which improves its convergence and therefore the quality of continuous-space solutions obtained on systems with limited computing power. Furthermore, it lends itself for deployment on a wide variety of platforms, as it can be easily adapted to match available hardware resources by trading off computational efficiency against optimization accuracy. The experimental results in this thesis show that the proposed approach outperforms existing planning methods in both discrete and continuous monitoring scenarios.

System Integration and Experimental Results

A informative planning framework was presented to integrate the theoretical contributions in mapping and planning into a single system. The proposed infrastructure is formulated in a general manner, making it easily adapted for any active sensing problem. An open-source implementation is provided with the aims of stimulating further development and bridging the gap between academic research and practical applications.

Various experiments were presented to validate the overall framework and its components in both simulated and real-world environments focusing on UAV-based terrain monitoring scenarios. Results from comprehensive simulation studies reveal its effectiveness in resource-constrained active sensing missions, in terms of both the speed at which data is collected and its quality. The investigations also underline the ability of the approach to focus on specific areas of interest. Finally, field tests were presented to demonstrate the algorithms implemented and running in real-time on UAV platforms for various data gathering tasks, highlighting its ease of deployment for future use-cases.

Active Sensing Under Localization Uncertainty

The informative planning framework was extended by introducing methods to account for the uncertainty in the robot localization in general active sensing problems. In contrast to previous work, the proposed approach propagates the pose uncertainty into *both* the mapping and planning modules towards improving the robustness and accuracy of reconstructions of environmental phenomena. A new utility function is presented that jointly considers the uncertainty of the environmental field map and robot pose without relying on manually-tuned, problem-specific parameters, enabling the agent to manage the exploration-exploitation trade-off in a wide variety of environments. The key contribution of this thesis is that the formulation is developed for mapping a continuous field using a GP-based model. Moreover, its applicability extends to other areas of robotics and artificial intelligence, such as reinforcement learning.

The approach was evaluated extensively in 3-D simulation scenarios with a graph-based Simultaneous Localization And Mapping (SLAM)-back end, and was shown to outperform existing methods in terms of both mapping and planning. Experimental validation of the system was performed through a proof of concept deployment, revealing its promise in future applications.

7.2 Future Outlook

The developments in this thesis have opened up many avenues for future work, described below.

Large-scale Mapping

The major computational bottleneck in the proposed approach is the map update procedure, which hinders its scalability to large-scale/high-resolution applications. Though the new mapping method based on recursive Bayesian fusion effectively handles dense visual data, it still runs in cubic time

due to the matrix inversion operation in the filtering procedure. This impairs the computational efficiency of the planning algorithm, as map updates are performed multiple times for each candidate path evaluated by the CMA-ES optimization routine, as well as for map building online. To reiterate Section 3.3, the submapping techniques of Sun, Vidal-Calleja, and Miró (2015) and Sun, Vidal-Calleja, and Miro (2017) are relevant solutions that could be directly applied to accelerate the current approach while maintaining reconstruction accuracy.

A similar bottleneck is manifested in the problem of active sensing under localization uncertainty in Chapter 6. As in standard GP regression, the mapping strategy here suffers from cubic time complexity in the number of accumulated measurements, coupled with the additional expense of integrating the covariance function over Uncertain Input (UI). Sparsification methods (Rasmussen et al., 2006 and Quiñonero-Candela et al., 2007) could be investigated towards enabling large-scale online applications.

Time-varying Utility Functions

The results in Section 5.2.1 demonstrate the effectiveness of using a time-varying objective function to direct data acquisition at different stages of a mission. This concept is worth exploring further to cater for practical considerations and enable more complex information gathering behavior. Since the general planning strategy in this thesis does not incorporate an awareness of the remaining budget, performing any maneuvers to finalize the mission, e.g., returning the robot to the initial point of deployment or landing the UAV safely, is not possible. Palazzolo et al. (2018) and Hollinger and Sukhatme (2014) show how these aspects could be easily addressed by planning with objectives that are a function of time. Moreover, as demonstrated in Section 5.2.1, such utility functions would also allow for defining precise time-varying aims in terms of the value of information acquired, e.g., by first encouraging the robot to explore an environment uniformly before focusing on specific areas of interest.

Defining α in Rényi's Entropy

In the context of active sensing under pose uncertainty, one promising direction for future work is investigating different ways of mapping the α parameter in Rényi's entropy to the predicted localization uncertainty. Following the approach of Carrillo et al. (2018), this work exampled a simple relationship using the A-optimality criterion. Experiences acquired during the experiments, however, indicate that certain uncertainty measures are more beneficial to use than others, i.e., D-optimality was found to have a much lower variability than A-optimality, making it less sensitive to uncertainty changes. An open research question is therefore incorporating more sophisticated scalar measures of uncertainty in the utility function in order to improve upon active sensing performance. As suggested in Chapter 6, this could be done by considering alternative mappings to optimal design criteria, e.g., using exponential or logarithmic functions, or different approaches altogether, e.g., Bayesian optimization. Furthermore, it would be particularly interesting to explore how these concepts extend to other areas of robotics, such as reinforcement learning (Whitehead et al., 2008), and whether these findings can be generalized.

Multi-robot Planning

A natural and exciting direction for future research is to extend the ideas in this thesis to active sensing problems with multi-robot teams. This field of study has grown considerably in the past several years, and will continue to do as as the capabilities of autonomous platforms expand. In many monitoring missions, replacing a single-robot with a multi-robot system can offer improved spatial coverage, versatility, and robustness at a potentially lower overall cost Yan et al. (2013).

One could, for example, deploy a swarm of UAVs to monitor a very large agricultural field (Albani et al., 2018). Alternatively, mobile robots could be used as *data mules* to relay information between different nodes in a network of static sensors (Dunbabin, Corke, et al., 2006). However, such scenarios introduce a host of additional considerations in the informative path planning procedure, including requirements for coordination, communication, resource sharing, etc. Key steps in extending the existing approach is reformulating the objective function so that it accounts for the actions of other agents as additional optimization variables in the CMA-ES routine, and tackling associated scalability issues.

Multi-modal Monitoring

Rapid advances in autonomous systems and sensing technologies are enabling robots to carry increasingly sophisticated sensors with higher payload capacities (Dunbabin and Marques, 2012). In future studies, it would be relevant to extend the approach to multi-modal monitoring applications that consider data streams from several different sensors. This requires specifying multi-objective goals for the path planning algorithm, which is an open research question. In the context of environmental monitoring, another interesting problem is incorporating scientific knowledge to interpret the interactions between different environmental processes, e.g., identifying geological features based on the consistency between spectrometer readings and visual imagery (Arora et al., 2019).

Persistent Monitoring

Another interesting extension is the problem of *persistent monitoring*, in which the observed environmental phenomenon is a dynamically changing distribution. In this setup, the robot is required to move perpetually between sensing locations to maintain an accurate estimate of the field. Such a scenario poses critical research challenges in terms of both environmental mapping and informative planning. For mapping, new representations are necessary that can capture temporal, as well as spatial, variations. In the proposed approach, this could be done by introducing a periodic component in the GP covariance function (Marchant, Ramos, and Sanner, 2014 and A. Singh, Ramos, et al., 2010), which necessitates reliable training datasets. Another possibility is exploiting online hyperparameter learning K.-C. Ma, Liu, et al., 2017; Viseras et al., 2019 to improve the model over time. For planning, it must be ensured that paths are planned over time horizons short enough with respect to the environment dynamics in order to adapt to the acquired data K.-C. Ma, Z. Ma, et al., 2018. In this regard, the online replanning capabilities of the proposed algorithm are of major benefit.

Sensor Modeling

A major practical issue emerging from the experiments in this thesis is the development of accurate sensor models as a basis for predictive decision-making. Since the informative planning problem evaluates possible paths, the sensor model is critical as it determines how future measurements will influence the evolution of the map, and subsequently the value of information acquired. Well-known methods exist for probabilistic sensor modeling in occupancy mapping tasks (O'Callaghan et al., 2012 and Mullane et al., 2009). However, relatively few works have analyzed the properties of sensors used to monitor environmental phenomena. This problem is particularly relevant for devices relying on deep learning architectures, which are gaining rapid momentum in many robotics applications (Carrio et al., 2017 and Sa, Z. Chen, et al., 2018).

Section 5.3 of this work and Anil Meera et al. (2019) demonstrate how this can be done by means of a data-driven accuracy analysis. However, towards fully exploiting such sensors in an informative planning framework, a true measure of the predictive *uncertainty* is required. This topic is a popular

subject of ongoing research (Pimentel et al., 2014 and Blum et al., 2019) and is anticipated to shape the future of active sensing systems.

Software Infrastructure

The lack of available software infrastructure for informative planning is identified as a major obstacle hindering further research in this area. The framework developed in this thesis is released open-source and disseminated by Papachristos, Kamel, et al. (2019) as a step towards addressing this issue. Although it suffices as a self-contained package, the active sensing community would profit significantly from benchmark problems and algorithms for researchers to compare new methods against in a standardized manner. As demonstrated in Section 5.3, in terrain monitoring applications, experimental datasets from the field of photogrammetry could be exploited as references for evaluation. This would not only be extremely beneficial for future developments, but also underline the interdisciplinary nature of the active sensing problem.

Bibliography

Albani, D., Manoni, T., Nardi, D., and Trianni, V. (2018). "Dynamic UAV Swarm Deployment for Non-Uniform Coverage". In: *International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, pp. 523–531
(Cited on pages 1, 3, 104)

Anil Meera, A., Popović, M., Millane, A., and Siegwart, R. (2019). "Obstacle-aware Adaptive Informative Path Planning for UAV-based Target Search". In: *IEEE International Conference on Robotics and Automation*. Accepted. IEEE
(Cited on pages 1, 3, 5, 6, 9, 18, 19, 23, 80, 104)

Anthony, D., Elbaum, S., Lorenz, A., and Detweiler, C. (2014). "On Crop Height Estimation with UAVs". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 4805–4812. DOI: 10.1109/IROS.2014.6943245
(Cited on pages 6, 10)

Arora, A., Furlong, P., Fitch, R., Sukkarieh, S., and Fong, T. (2019). "Multi-modal active perception for information gathering in science missions". *Autonomous Robots*. DOI: 10.1007/s10514-019-09836-5
(Cited on page 104)

Badrinarayanan, V., Kendall, A., and Cipolla, R. (2017). "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39.12, pp. 2481–2495. DOI: 10.1109/TPAMI.2016.2644615
(Cited on page 63)

Bähnemann, R., Pantic, M., Popović, M., Schindler, D., Tranzatto, M., Kamel, M., Grimm, M., Widauer, J., Siegwart, R., and Nieto, J. (2019). "The ETH-MAV Team in the MBZ International Robotics Challenge". *Journal of Field Robotics* 36.1, pp. 78–103. DOI: 10.1002/rob.21824
(Cited on page 18)

Bajcsy, R. (1988). "Active perception". *Proceedings of the IEEE* 76.8, pp. 996–1005. DOI: 10.1109/5.5968
(Cited on pages 1, 9)

Bajcsy, R., Aloimonos, Y., and Tsotsos, J. (2018). "Revisiting active perception". *Autonomous Robots* 42.2, pp. 177–196. DOI: 10.1007/s10514-017-9615-3
(Cited on pages 1, 42)

Berrio, J. S., Ward, J., Worrall, S., Zhou, W., and Nebot, E. (2017). "Fusing Lidar and Semantic Image Information in Octree Maps". In: *Australasian Conference on Robotics and Automation*.

ACRA
(Cited on pages 5, 9, 65)

Binney, J., Krause, A., and Sukhatme, G. S. (2013). "Optimizing waypoints for monitoring spatiotemporal phenomena". *The International Journal of Robotics Research* 32.8, pp. 873–888. DOI: 10.1177/0278364913488427
(Cited on pages 4, 5, 10)

Binney, J. and Sukhatme, G. S. (2012). "Branch and bound for informative path planning". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2147–2154. DOI: 10.1109/ICRA.2012.6224902
(Cited on pages 5, 10, 12, 15, 47)

Bircher, A., Siegwart, R., Kamel, M., Alexis, K., and Oleynikova, H. (2018). "Receding horizon path planning for 3D exploration and surface inspection". *Autonomous Robots* 42.2, pp. 291–306. DOI: 10.1007/s10514-016-9610-0
(Cited on pages 1, 12, 42)

Bloesch, M., Omari, S., Hutter, M., and Siegwart, R. (2015). "Robust visual inertial odometry using a direct EKF-based approach". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 298–304. DOI: 10.1109/IROS.2015.7353389
(Cited on page 75)

Blum, H., Rohrbach, S., Popović, M., Bartolomei, L., and Siegwart, R. (2019). "Active Learning for UAV-based Semantic Mapping". In: *2nd Workshop on Informative Path Planning and Active Sampling, RSS'19*. MIT Press
(Cited on pages 18, 19, 105)

Bourgault, F., Makarenko, A. A., Williams, S. B., Grocholsky, B., and Durrant-Whyte, H. F. (2002). "Information Based Adaptive Robotic Exploration". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 540–545. DOI: 10.1109/IRDS.2002.1041446
(Cited on pages 6, 13, 16, 35, 36, 90)

Bouzarkouna, Z., Ding, D., and Auger, A. (2012). "Well placement optimization with the covariance matrix adaptation evolution strategy and meta-models". *Computational Geosciences* 16.1, pp. 75–92. DOI: 10.1007/s10596-011-9254-2
(Cited on page 40)

Bry, A. and Roy, N. (2011). "Rapidly-exploring random belief trees for motion planning under uncertainty". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 723–730. DOI: 10.1109/ICRA.2011.5980508
(Cited on page 13)

Byrd, R. H., Gilbert, J. C., and Nocedal, J. (2006). "A Trust Region Method Based on Interior Point Techniques for Nonlinear Programming". *Mathematical Programming* 89.1, pp. 149–185. DOI: 10.1007/PL00011391
(Cited on page 57)

Cabreira, T., Brisolara, L., and Ferreira Jr., P. R. (2019). *Survey on Coverage Path Planning with Unmanned Aerial Vehicles*. Vol. 3. 1. MDPI, p. 4. DOI: 10.3390/drones3010004
(Cited on pages 2, 3, 12)

Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., and Leonard, J. J. (2016). "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age". *IEEE Transactions on Robotics* 32.6, pp. 1309–1332. DOI: 10.1109/TRO.2016.2624754
(Cited on pages 22, 86, 99)

Cao, N., Low, K. H., and Dolan, J. M. (2013). "Multi-Robot Informative Path Planning for Active Sensing of Environmental Phenomena : A Tale of Two Algorithms". In: *International Conference on Autonomous Agents and Multiagent Systems*, pp. 6–10
(Cited on pages 1, 10)

Cardina, J., Johnson, G. A., and Sparrow, D. H. (1997). "The Nature and Consequence of Weed Spatial Distribution". *Weed Science* 45.3, pp. 364–373
(Cited on pages 3, 79)

Carrillo, H., Dames, P., Kumar, V., and Castellanos, J. A. (2018). "Autonomous robotic exploration using a utility function based on Rényi's general theory of entropy". *Autonomous Robots* 42.2, pp. 235–256. DOI: 10.1007/s10514-017-9662-9
(Cited on pages 4, 5, 7, 9, 13, 35, 38, 84, 86, 99, 103)

Carrio, A., Sampedro, C., Rodriguez-Ramos, A., and Campoy, P. (2017). "A Review of Deep Learning Methods and Applications for Unmanned Aerial Vehicles". *Journal of Sensors*. DOI: 10.1155/2017/3296874
(Cited on pages 2, 5, 9, 29, 63, 104)

Causa, F., Popović, M., Fasano, G., Grassi, M., Nieto, J., and Siegwart, R. (2019). "Navigation aware planning for tandem UAV missions in GNSS challenging environments". In: *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics. DOI: 10.2514/6.2019-0916
(Cited on page 18)

Charrow, B., Liu, S., Kumar, V., and Michael, N. (2015). "Information-Theoretic Mapping Using Cauchy-Schwarz Quadratic Mutual Information". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 4791–4798. DOI: 10.1109/ICRA.2015.7139865
(Cited on pages 5, 9, 11, 12, 35, 40)

Chekuri, C. and Pál, M. (2005). "A Recursive Greedy Algorithm for Walks in Directed Graphs". In: *IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 245–253. DOI: 10.1109/SFCS.2005.9
(Cited on pages 12, 15)

Chen, M., Frazzoli, E., Hsu, D., and Lee, W. S. (2016). "POMDP-lite for Robust Path Planning under Uncertainty". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5427–5433. DOI: 10.1109/ICRA.2016.7487754
(Cited on page 12)

Chung, J. J. (2014). "Learning to Soar: Exploration Strategies in Reinforcement Learning for Resource-Constrained Missions". PhD thesis. University of Sydney
(Cited on pages 4, 42)

Colomina, I. and Molina, P. (2014). "Unmanned aerial systems for photogrammetry and remote sensing: A review". *ISPRS Journal of Photogrammetry and Remote Sensing* 92, pp. 79–97. DOI:

10.1016/j.isprsjprs.2014.02.013
(Cited on pages 1–3, 6, 10)

Costante, G., Delmerico, J., Werlberger, M., and Valigi, P. (2017). "Exploiting Photometric Information for Planning under Uncertainty". In: *Robotics Research*. Springer, pp. 107–124. DOI: 10.1007/978-3-319-51532-8_7
(Cited on pages 9, 13, 38)

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. 2nd. John Wiley & Sons, Inc.
(Cited on pages 13, 35–37)

Dang, T., Papachristos, C., and Alexis, K. (2018). "Autonomous Exploration and Simultaneous Object Search Using Aerial Robots". In: *IEEE Aerospace Conference*. IEEE, pp. 1–7. DOI: 10.1109/AERO.2018.8396632
(Cited on pages 9, 29)

Davis, P. J. and Rabinowitz, P. (2007). *Methods of Numerical Integration*. 2nd ed. Dover Publications
(Cited on pages 82, 83)

Delmerico, J., Mueggler, E., Nitsch, J., and Scaramuzza, D. (2016). "Active Autonomous Aerial Exploration for Ground Robot Path Planning". *IEEE Robotics and Automation Letters*. DOI: 10.1109/LRA.2017.2651163
(Cited on page 9)

Detweiler, C., Ore, J.-P., Anthony, D., Elbaum, S., Burgin, A., and Lorenz, A. (2015). "Bringing Unmanned Aerial Systems Closer to the Environment". *Environmental Practice* 17.3, pp. 188–200. DOI: 10.1017/S1466046615000174
(Cited on pages 1, 3, 6, 10)

Dijkstra, E. W. (1959). "A note on two problems in connexion with graphs". *Numerische Mathematik* 1.1, pp. 269–271. DOI: 10.1007/BF01386390
(Cited on pages 2, 11)

Dunbabin, M., Corke, P., Vasilescu, I., and Rus, D. (2006). "Data muling over underwater wireless sensor networks using an autonomous underwater vehicle". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2091–2098. DOI: 10.1109/ROBOT.2006.1642013
(Cited on page 104)

Dunbabin, M. and Marques, L. (2012). "Robots for environmental monitoring: Significant advancements and applications". *IEEE Robotics & Automation Magazine* 19.1, pp. 24–39. DOI: 10.1109/MRA.2011.2181683
(Cited on pages 1, 4, 6, 10, 104)

Elfes, A. (1989). "Using occupancy grids for mobile robot perception and navigation". *Computer* 22.6, pp. 46–57. DOI: 10.1109/2.30720
(Cited on pages 5, 9, 21, 23)

Ezequiel, C. A. F., Cua, M., Libatique, N. C., Tangonan, G. L., Alampay, R., Labuguen, R. T., Favila, C. M., Honrado, J. L. E., Canos, V., Devaney, C., Loreto, A. B., Bacusmo, J., and Palma, B. (2014). "UAV Aerial Imaging Applications for Post-Disaster Assessment, Environmental Management and Infrastructure Development". In: *International Conference on Unmanned Aircraft*

Systems, pp. 274–283. DOI: 10.1109/ICUAS.2014.6842266
(Cited on pages 1, 3, 5, 10)

Fankhauser, P. and Hutter, M. (2016). "A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation". In: *Robot Operating System (ROS): The Complete Reference (Volume 1)*. Springer, pp. 99–120. DOI: 10.1007/978-3-319-26054-9_5
(Cited on page 9)

Faria, M., Marín, R., Popović, M., Maza, I., and Viguria, A. (2019). "Efficient Lazy Theta* Path Planning over a Sparse Grid to Explore Large 3D Volumes with a Multirotor UAV". *Sensors* 19.1. DOI: 10.3390/s19010174
(Cited on page 18)

Fedorov, V. (1972). *Theory Of Optimal Experiments*. Academic Press
(Cited on pages 38, 53, 84)

Furrer, F., Burri, M., Achtelik, M., and Siegwart, R. (2016). "RotorS - A Modular Gazebo MAV Simulator Framework". In: *Robot Operating System (ROS): The Complete Reference (Volume 1)*. Springer International Publishing, pp. 595–625. DOI: DOI:10.1007/978-3-319-26054-9_23
(Cited on page 62)

Galceran, E. and Carreras, M. (2013). "A survey on coverage path planning for robotics". *Robotics and Autonomous Systems* 61.12, pp. 1258–1276. DOI: 10.1016/j.robot.2013.09.004
(Cited on pages 2, 3, 12)

Gan, S. K., Yang, K. J., and Sukkarieh, S. (2009). "3D Online Path Planning in a Continuous Gaussian Process Occupancy Map". In: *Australian Conference on Robotics and Automation*. Australian Robotics & Automation Association
(Cited on page 32)

Gao, M., Xu, X., Klinger, Y., Van Der Woerd, J., and Tapponnier, P. (2017). "High-resolution mapping based on an Unmanned Aerial Vehicle (UAV) to capture paleoseismic offsets along the Altyn-Tagh fault, China". *Scientific Reports* 7.1. DOI: 10.1038/s41598-017-08119-2
(Cited on page 1)

Gelbart, M. A., Snoek, J., and Adams, R. P. (2014). "Bayesian Optimization with Unknown Constraints". In: *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, pp. 250–259
(Cited on page 57)

Girard, A. (2004). "Approximate methods for propagation of uncertainty with Gaussian process models". PhD thesis. University of Glasgow
(Cited on page 13)

Girdhar, Y. and Dudek, G. (2015). "Modeling Curiosity in a Mobile Robot for Long-Term Autonomous Exploration and Monitoring". *Autonomous Robots* 33.4, pp. 645–657. DOI: 10.1007/s10514-015-9500-x
(Cited on page 12)

Golshani, L. and Pasha, E. (2010). "Rényi entropy rate for Gaussian processes". *Information Sciences* 180.8, pp. 1486–1491. DOI: 10.1016/j.ins.2009.12.012
(Cited on page 84)

Gotovos, A., Casati, N., Hitz, G., and Krause, A. (2013). "Active Learning for Level Set Estimation". In: *International Joint Conference on Artificial Intelligence*. AAAI Press, pp. 1344–1350
(Cited on pages 6, 38)

Hansen, N. (2006). "The CMA evolution strategy: A comparing review". *Studies in Fuzziness and Soft Computing* 192.2006, pp. 75–102. DOI: 10.1007/11007937_4
(Cited on pages 6, 41, 53, 57, 59)

Hansen, N., Niederberger, A. S. P., Guzzella, L., and Koumoutsakos, P. (2009). "A Method for Handling Uncertainty in Evolutionary Optimization With an Application to Feedback Control of Combustion". *IEEE Transactions on Evolutionary Computation* 13.2, pp. 180–197. DOI: 10.1109/TEVC.2008.924423
(Cited on pages 15, 40, 41, 57)

Heng, L., Gotovos, A., Krause, A., and Pollefeys, M. (2015). "Efficient Visual Exploration and Coverage with a Micro Aerial Vehicle in Unknown Environments". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 1071–1078. DOI: 10.1109/ICRA.2015.7139309
(Cited on pages 12, 13)

Hitz, G., Galceran, E., Garneau, M.-È., Pomerleau, F., and Siegwart, R. (2017). "Adaptive Continuous-Space Informative Path Planning for Online Environmental Monitoring". *Journal of Field Robotics* 34.8, pp. 1427–1449. DOI: 10.1002/rob.21722
(Cited on pages 4–6, 10, 12, 14, 15, 32, 35, 40–42, 57, 59, 60)

Hitz, G., Gotovos, A., Pomerleau, F., Garneau, M.-È., Pradalier, C., Krause, A., and Siegwart, R. (2014). "Fully Autonomous Focused Exploration for Robotic Environmental Monitoring". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2658–2664. DOI: 10.1109/ICRA.2014.6907240
(Cited on pages 6, 12, 38, 42)

Hollinger, G. A., Englot, B., Hover, F. S., Mitra, U., and Sukhatme, G. S. (2013). "Active planning for underwater inspection and the benefit of adaptivity". *The International Journal of Robotics Research* 32.1, pp. 3–18. DOI: 10.1177/0278364912467485
(Cited on page 12)

Hollinger, G. A., Singh, S., Djugash, J., and Kehagias, A. (2009). "Efficient Multi-Robot Search for a Moving Target". *The International Journal of Robotics Research* 28.2, pp. 201–219. DOI: 10.1177/0278364908099853
(Cited on pages 1, 3, 5, 9)

Hollinger, G. A. and Sukhatme, G. S. (2014). "Sampling-based robotic information gathering algorithms". *The International Journal of Robotics Research* 33.9, pp. 1271–1287. DOI: 10.1177/0278364914533443
(Cited on pages 1, 4, 5, 10, 12, 14, 46, 47, 52, 80, 90, 91, 103)

Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013). "OctoMap: An efficient probabilistic 3D mapping framework based on octrees". *Autonomous Robots* 34.3, pp. 189–206. DOI: 10.1007/s10514-012-9321-0
(Cited on page 37)

Ingber, L. and Rosen, B. (1992). "Genetic Algorithms and Very Fast Simulated Reannealing: A comparison". *Mathematical and Computer Modelling* 16.11, pp. 87–100. DOI: 10.1016/0895-7177

(92)90108-W
(Cited on page 57)

Jadidi, M. (2017). "Gaussian Processes for Information-theoretic Robotic Mapping and Exploration". PhD thesis. University of Technology, Sydney
(Cited on page 52)

Jadidi, M. G., Miró, J. V., and Dissanayake, G. (2016). "Sampling-based Incremental Information Gathering with Applications to Robotic Exploration and Environmental Monitoring". arXiv: 1607.01883
(Cited on pages 4, 5, 7, 10, 13, 80)

Jadidi, M. G., Miró, J. V., and Dissanayake, G. (2017). "Warped Gaussian Processes Occupancy Mapping with Uncertain Inputs". *IEEE Robotics and Automation Letters* 2.2, pp. 680–687. DOI: 10.1109/LRA.2017.2651154
(Cited on pages 10, 13, 17, 82, 83, 94)

Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. (2014). *Caffe: Convolutional Architecture for Fast Feature Embedding*. arXiv: 1408.5093v1
(Cited on page 64)

Jumarie, G. (1990). *Relative Information*. Springer. DOI: 10.1007/978-3-642-84017-3
(Cited on page 85)

Kaelbling, L., Littman, M., and Cassandra, A. (1998). "Planning and Acting in Partially Observable Stochastic Domains". *Artificial Intelligence* 101.1-2, pp. 99–134. DOI: 10.1016/S0004-3702(98)00023-X
(Cited on page 11)

Kaess, M. and Dellaert, F. (2009). "Covariance recovery from a square root information matrix for data association". *Robotics and Autonomous Systems* 57.12, pp. 1198–1210. DOI: 10.1016/j.robot.2009.06.008
(Cited on pages 86, 87, 90)

Kamel, M., Stastny, T., Alexis, K., and Siegwart, R. (2017). "Model Predictive Control for Trajectory Tracking of Unmanned Aerial Vehicles Using Robot Operating System". In: *Robot Operating System (ROS): The Complete Reference (Volume 2)*. Springer, pp. 3–39. DOI: 10.1007/978-3-319-54927-9_1
(Cited on pages 62, 68, 72, 75)

Kemker, R., Salvaggio, C., and Kanan, C. (2018). "Algorithms for Semantic Segmentation of Multispectral Remote Sensing Imagery using Deep Learning". *ISPRS Journal of Photogrammetry and Remote Sensing* 145, pp. 60–77. DOI: 10.1016/j.isprsjprs.2018.04.014
(Cited on pages 62, 63)

Krause, A. and Guestrin, C. (2011). "Submodularity and its applications in optimized information gathering". *ACM Transactions on Intelligent Systems and Technology* 2.4, pp. 1–20. DOI: 10.1145/1989734.1989736
(Cited on page 11)

Krause, A., Singh, A., and Guestrin, C. (2008). "Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies". *Journal of Machine Learning Research*

9, pp. 235–284. DOI: 10.1145/1102351.1102385
(Cited on pages 6, 10, 12, 16, 35)

Kurniawati, H., Hsu, D., and Lee, W. S. (2008). "SARSOP : Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces". In: *Robotics: Science and Systems*. MIT Press. DOI: 10.15607/RSS.2008.IV.009
(Cited on page 12)

Lawrance, N. and Sukkarieh, S. (2011). "Autonomous Exploration of a Wind Field with a Gliding Aircraft". *Journal of Guidance, Control, and Dynamics* 34.3, pp. 719–733. DOI: 10.2514/1.52236
(Cited on pages 4, 42)

Liebisch, F., Popović, M., Pfeifer, J., Khanna, R., Lottes, P., Pretto, A., Sa, I., Nieto, J., Siegwart, R., and Walter, A. (2017). "Automatic UAV-based field inspection campaigns for weeding in row crops". In: *EARSeL SIG Imaging Spectroscopy Workshop*
(Cited on pages 1, 3, 19)

Lim, Z. W., Hsu, D., and Lee, W. S. (2015). "Adaptive informative path planning in metric spaces". *The International Journal of Robotics Research* 35.5, pp. 585–598. DOI: 10.1177/0278364915596378
(Cited on pages 6, 12, 15)

Longhi, M., Taylor, Z., Popović, M., Nieto, J., Marrocco, G., and Siegwart, R. (2018). "RFID-Based Localization for Greenhouses Monitoring Using MAVs". In: *IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communications*. IEEE, pp. 905–908. DOI: 10.1109/APWC.2018.8503764
(Cited on page 18)

Ma, K.-C., Liu, L., and Sukhatme, G. S. (2017). "Informative Planning and Online Learning with Sparse Gaussian Processes". In: *IEEE International Conference on Robotics and Automation*. IEEE. DOI: 10.1109/ICRA.2017.7989494
(Cited on pages 29, 104)

Ma, K.-C., Ma, Z., Liu, L., and Sukhatme, G. S. (2018). "Multi-robot Informative and Adaptive Planning for Persistent Environmental Monitoring". *Distributed Autonomous Robotic Systems*, pp. 285–298. DOI: 10.1007/978-3-319-73008-0_20
(Cited on pages 10, 11, 104)

Manfreda, S., McCabe, M. F., Miller, P. E., Lucas, R., Madrigal, V. P., Mallinis, G., Dor, E. B., Helman, D., Estes, L., Ciraolo, G., Müllerová, J., Tauro, F., de Lima, M. I., de Lima, J. L. M. P., Maltese, A., Frances, F., Caylor, K., Kohv, M., Perks, M., Ruiz-Pérez, G., Su, Z., Vico, G., and Toth, B. (2018). "On the Use of Unmanned Aerial Systems for Environmental Monitoring". *Remote Sensing* 10.4. DOI: 10.3390/rs10040641
(Cited on pages 1–3)

Marchant, R. and Ramos, F. (2012). "Bayesian Optimisation for Intelligent Environmental Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. DOI: 10.1109/IROS.2012.6385653
(Cited on pages 57, 58)

Marchant, R. and Ramos, F. (2014). "Bayesian Optimisation for Informative Continuous Path Planning". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 6136–6143. DOI:

10.1109/ICRA.2014.6907763
(Cited on pages 6, 14, 38, 53, 57)

Marchant, R., Ramos, F., and Sanner, S. (2014). "Sequential Bayesian Optimisation for Spatial-Temporal Monitoring". In: *Conference on Uncertainty in Artificial Intelligence*
(Cited on page 104)

Matheron, G. (1973). "The Intrinsic Random Functions and Their Applications". *Advances in Applied Probability* 5.3, pp. 439–468. DOI: 10.2307/1425829
(Cited on page 10)

Mchutchon, A. and Rasmussen, C. (2011). "Gaussian Process Training with Input Noise". *Advances in Neural Information Processing Systems*, pp. 1341–1349
(Cited on page 13)

Miki, T., Popović, M., Gawel, A., Hitz, G., and Siegwart, R. (2018). "Multi-Agent Time-Based Decision-Making for the Search and Action Problem". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2365–2372. DOI: 10.1109/ICRA.2018.8460996
(Cited on page 19)

Močkus, J. (1975). "On Bayesian Methods for Seeking the Extremum". In: *Optimization Techniques IFIP Technical Conference*. Springer, pp. 400–404. DOI: 10.1007/978-3-662-38527-2_55
(Cited on page 57)

Morere, P., Marchant, R., and Ramos, F. (2017). "Sequential Bayesian Optimisation as a POMDP for Environment Monitoring with UAVs". arXiv: 1703.04211
(Cited on pages 6, 12, 15, 57)

Mullane, J., Adams, M. D., and Wijesoma, W. (2009). "Robotic Mapping Using Measurement Likelihood Filtering". *The International Journal of Robotics Research* 28.2, pp. 172–190. DOI: 10.1177/0278364908097214
(Cited on page 104)

Neumann, P. P., Asadi, S., Lilienthal, A. J., Bartholmai, M., and Schiller, J. H. (2012). "Autonomous Gas-Sensitive Microdrone: Wind Vector Estimation and Gas Distribution Mapping". *IEEE Robotics & Automation Magazine*, pp. 50–61. DOI: 10.1109/MRA.2012.2184671
(Cited on page 4)

Nex, F. and Remondino, F. (2014). "UAV for 3D mapping applications: A review". *Applied Geomatics* 6.1, pp. 1–15. DOI: 10.1007/s12518-013-0120-x
(Cited on pages 1–3)

O'Callaghan, S. T. and Ramos, F. (2012). "Gaussian process occupancy maps". *The International Journal of Robotics Research* 31.1, pp. 42–62. DOI: 10.1177/0278364911421039
(Cited on pages 5, 9, 10, 52, 104)

Oliveira, R., Ott, L., Guizilini, V., and Ramos, F. (2017). "Bayesian Optimisation for Safe Navigation under Localisation Uncertainty". In: *International Symposium of Robotics Research*. Springer. arXiv: 1709.02169v1
(Cited on page 13)

Olson, E. (2011). "AprilTag: A robust and flexible visual fiducial system". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 3400–3407. DOI: 10.1109/ICRA.2011.5979561
(Cited on page 68)

Palazzolo, E. and Stachniss, C. (2018). "Effective Exploration for MAVs Based on the Expected Information Gain". *Drones* 2.1, p. 9. DOI: 10.3390/drones2010009
(Cited on pages 42, 103)

Papachristos, C., Kamel, M., Popović, M., Khattak, S., Bircher, A., Oleynikova, H., Dang, T., Mascarich, F., Alexis, K., and Siegwart, R. (2019). "Autonomous Exploration and Inspection Path Planning for Aerial Robots Using the Robot Operating System". In: *Robot Operating System (ROS): The Complete Reference (Volume 3)*. Ed. by A. Koubaa. Springer, pp. 67–111. DOI: 10.1007/978-3-319-91590-6_3
(Cited on pages 18, 80, 105)

Papachristos, C., Khattak, S., and Alexis, K. (2017). "Uncertainty-aware Receding Horizon Exploration and Mapping using Aerial Robots". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 4568–4575. DOI: 10.1109/ICRA.2017.7989531
(Cited on pages 1, 13, 86, 99)

Peña, J. M., Torres-Sánchez, J., de Castro, A. I., Kelly, M., and López-Granados, F. (2013). "Weed Mapping in Early-Season Maize Fields Using Object-Based Analysis of Unmanned Aerial Vehicle (UAV) Images". *PLoS ONE* 8.10, pp. 1–11. DOI: 10.1371/journal.pone.0077151
(Cited on page 79)

Pfeifer, J., Khanna, R., Constantin, D., Popović, M., Galceran, E., Kirchgessner, N., Achim, W., Siegwart, R., and Liebisch, F. (2016). "Towards automatic UAV data interpretation for precision farming". In: *International Conference on Agricultural Engineering*
(Cited on page 19)

Pimentel, M. A. F., Clifton, D. A., Clifton, L., and Tarassenko, L. (2014). "Review: A Review of Novelty Detection". *Signal Processing* 99, pp. 215–249. DOI: 10.1016/j.sigpro.2013.12.026
(Cited on page 105)

Pomerleau, F., Colas, F., Siegwart, R., and Magnenat, S. (2013). "Comparing ICP variants on real-world data sets". *Autonomous Robots* 34.3, pp. 133–148. DOI: 10.1007/s10514-013-9327-2
(Cited on page 47)

Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., and Galceran, E. (2017). "Online Informative Path Planning for Active Classification Using UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5753–5758. DOI: 10.1109/ICRA.2017.7989676
(Cited on pages 1, 5, 9, 14, 17, 31, 43)

Popović, M., Vidal-Calleja, T., Chung, J. J., Nieto, J., and Siegwart, R. (2019). "Informative Path Planning and Mapping for Active Sensing Under Localization Uncertainty". *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Under review. arXiv: 1902.09660
(Cited on pages 4, 14, 17, 21, 81)

Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660
(Cited on pages 1, 14, 17, 21, 31, 43)

Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1382–1388. DOI: 10.1109/IROS.2017.8202317
(Cited on pages 14, 17, 21, 43)

Quiñonero-Candela, J., Ramussen, C. E., and Williams, C. K. I. (2007). *Approximation Methods for Gaussian Process Regression*
(Cited on pages 10, 103)

Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press. DOI: 10.1142/S0129065704001899
(Cited on pages 5, 10, 14, 24, 25, 28, 37, 52, 82, 89, 103)

Reece, S. and Roberts, S. (2010). "An Introduction to Gaussian Processes for the Kalman Filter Expert". In: *International Conference on Information Fusion*. IEEE. DOI: 10.1109/ICIF.2010.5711863
(Cited on page 26)

Reif, J. H. (1979). "Complexity of the mover's problem and generalizations". In: *Annual Symposium on Foundations of Computer Science*. IEEE, pp. 421–427. DOI: 10.1109/SFCS.1979.10
(Cited on page 11)

Rényi, A. (1961). "On Measures of Entropy and Information". *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics* 1, pp. 547–561
(Cited on page 84)

Richter, C., Bry, A., and Roy, N. (2013). "Polynomial Trajectory Planning for Aggressive Quadrotor Flight in Dense Indoor Environments". In: *International Symposium on Robotics Research*. Springer. DOI: 10.1007/978-3-319-28872-7_37
(Cited on pages 2, 6, 12, 15, 32)

Sa, I., Chen, Z., Popović, M., Khanna, R., Liebisch, F., Nieto, J., and Siegwart, R. (2018). "weedNet: Dense Semantic Weed Classification Using Multispectral Images and MAV for Smart Farming". In: *IEEE Robotics and Automation Letters*. IEEE, pp. 588–595. DOI: 10.1109/LRA.2017.2774979
(Cited on pages 1, 3, 5, 9, 19, 29, 63, 75, 104)

Sa, I., Kamel, M., Burri, M., Bloesch, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2018). "Build Your Own Visual-Inertial Drone: A Cost-Effective and Open-Source Autonomous Drone". *IEEE Robotics & Automation Magazine* 25.1, pp. 89–103. DOI: 10.1109/MRA.2017.2771326
(Cited on pages 18, 72)

Sa, I., Kamel, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2017). "Dynamic System Identification, and Control for a cost-effective and open-source Multi-rotor MAV". In: *Field and Service Robotics*. Springer, pp. 605–620. DOI: 10.1007/978-3-319-67361-5_39
(Cited on page 19)

Sa, I., Popović, M., Khanna, R., Chen, Z., Lottes, P., Liebisch, F., Nieto, J., Stachniss, C., Walter, A., and Siegwart, R. (2018). "WeedMap: A Large-Scale Semantic Weed Mapping Framework Using Aerial Multispectral Imaging and Deep Neural Network for Precision Farming". *Remote Sensing*

10.9. DOI: 10.3390/rs10091423
(Cited on page 18)

Sadat, S. A., Wawerla, J., and Vaughan, R. (2015). "Fractal Trajectories for Online Non-Uniform Aerial Coverage". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2971–2976. DOI: 10.1109/ICRA.2015.7139606
(Cited on page 12)

Semsch, E., Jakob, M., Pavlíček, D., and Pěchouček, M. (2009). "Autonomous UAV Surveillance in Complex Urban Environments". In: *International Conference on Intelligent Agent Technology*. Vol. 2. IEEE, pp. 82–85. DOI: 10.1109/WI-IAT.2009.132
(Cited on page 1)

Sim, R. and Roy, N. (2005). "Global A-optimal robot exploration in SLAM". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 661–666. DOI: 10.1109/ROBOT.2005.1570193
(Cited on pages 6, 38, 86)

Singh, A., Krause, A., Guestrin, C., and Kaiser, W. J. (2009). "Efficient Informative Sensing using Multiple Robots". *Journal of Artificial Intelligence Research* 34, pp. 707–755. DOI: 10.1613/jair.2674
(Cited on pages 11, 38)

Singh, A., Krause, A., and Kaiser, W. J. (2009). "Nonmyopic Adaptive Informative Path Planning for Multiple Robots". In: *International Joint Conference on Artificial Intelligence*, pp. 1843–1850
(Cited on page 6)

Singh, A., Ramos, F., Durrant Whyte, H., and Kaiser, W. J. (2010). "Modeling and Decision Making in Spatio-Temporal processes for environmental surveillance". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5490–5497. DOI: 10.1109/ROBOT.2010.5509934
(Cited on pages 4, 10, 104)

Solà, J. (2017). *Course on SLAM*
(Cited on pages 86, 87, 90)

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. W. (2012). "Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting". In: *IEEE Transactions on Information Theory*. Vol. 58. 5. IEEE, pp. 3250–3265. DOI: 10.1109/TIT.2011.2182033
(Cited on pages 6, 38)

Sun, L., Vidal-Calleja, T., and Miro, J. V. (2017). "Coupling conditionally independent submaps for large-scale 2.5D mapping with Gaussian Markov Random Fields". *IEEE International Conference on Robotics and Automation*, pp. 3131–3137. DOI: 10.1109/ICRA.2017.7989358
(Cited on pages 10, 28, 103)

Sun, L., Vidal-Calleja, T., and Miró, J. V. (2015). "Bayesian fusion using conditionally independent submaps for high resolution 2.5D mapping". *IEEE International Conference on Robotics and Automation*, pp. 3394–3400. DOI: 10.1109/ICRA.2015.7139668
(Cited on pages 10, 28, 103)

Thrun, S. B. and Möller, K. (1992). "Active Exploration in Dynamic Environments". *Advances in Neural Information Processing Systems* 4, pp. 531–538
(Cited on page 1)

Thrun, S. and Fox, D. (2006). *Probabilistic Robotics*. 3rd ed. MIT Press
(Cited on pages 22, 88)

Valencia, R., Miró, J. V., Dissanayake, G., and Andrade-Cetto, J. (2012). "Active Pose SLAM". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1885–1891. DOI: 10.1109/IROS.2012.6385637
(Cited on pages 13, 16)

Vallvé, J., Solà, J., and Andrade-Cetto, J. (2018). "Graph SLAM Sparsification With Populated Topologies Using Factor Descent Optimization". *IEEE Robotics and Automation Letters* 3.2, pp. 1322–1329. DOI: 10.1109/ECMR.2017.8098675
(Cited on page 99)

Vasudevan, S., Ramos, F., Nettleton, E., Durrant-Whyte, H., and Blair, A. (2009). "Gaussian Process Modeling of Large Scale Terrain". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 1047–1053. DOI: 10.1109/ROBOT.2009.5152677
(Cited on pages 10, 14, 28)

Vetrella, A. R., Sa, I., Popović, M., Khanna, R., Nieto, J., Fasano, G., Accardo, D., and Siegwart, R. (2017). "Improved Tau-Guidance and Vision-Aided Navigation for Robust Autonomous Landing of UAVs". In: *Field and Service Robotics*. Springer, pp. 115–128. DOI: 10.1007/978-3-319-67361-5_8
(Cited on page 19)

Vidal-Calleja, T., Su, D., Bruijn, F. D., and Miró, J. V. (2014). "Learning Spatial Correlations for Bayesian Fusion in Pipe Thickness Mapping". In: *IEEE International Conference on Robotics and Automation*. IEEE. DOI: 10.1109/ICRA.2014.6906928
(Cited on pages 5, 11, 15, 52)

Viseras, A., Shutin, D., and Merino, L. (2019). "Robotic Active Information Gathering for Spatial Field Reconstruction with Rapidly-Exploring Random Trees and Online Learning of Gaussian Processes". *Sensors* 19.5, p. 1016. DOI: 10.3390/s19051016
(Cited on pages 11, 29, 104)

Vivaldini, K. C. T., Guizilini, V., Oliveira, M. D. C., Martinelli, T. H., Wolf, D. F., and Ramos, F. (2016). "Route Planning for Active Classification with UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE. DOI: 10.1109/ICRA.2016.7487412
(Cited on pages 1, 3, 12)

Vivaldini, K. C. T., Martinelli, T. H., Guizilini, V. C., Souza, J. R., Oliviera, M. D., Ramos, F. T., and Wolf, D. F. (2018). "UAV route planning for active disease classification". *Autonomous Robots*, pp. 1–17. DOI: 10.1007/s10514-018-9790-x
(Cited on pages 3, 4, 15, 57)

Whitehead, S. D. and Ballard, D. H. (2008). "Active Perception and Reinforcement Learning". *Neural Computation* 2.4, pp. 409–419. DOI: 10.1162/neco.1990.2.4.409
(Cited on pages 7, 103)

Xu, D. and Erdogmuns, D. (2010). "Information Theoretic Learning: Renyi's Entropy and Kernel Perspectives". In: *Information Theoretic Learning: Renyi's Entropy and Kernel Perspectives*. Springer
(Cited on page 84)

Yamauchi, B. (1997). "A Frontier-Based Approach for Autonomous Exploration". In: *IEEE International Symposium on Computational Intelligence in Robotics and Automation*. IEEE, pp. 146–151
(Cited on page 40)

Yan, Z., Jouandeau, N., and Cherif, A. (2013). "A Survey and Analysis of Multi-Robot Coordination". *International Journal of Advanced Robotic Systems* 10.12, p. 399. DOI: 10.5772/57313
(Cited on page 103)

Yang, K., Gan, S. K., and Sukkerieh, S. (2014). *Handbook of Unmanned Aerial Vehicles*. Springer, pp. 1655–1675. DOI: 10.1007/978-90-481-9707-1
(Cited on page 10)

Yang, W., Wang, S., Zhao, X., Zhang, J., and Feng, J. (2015). "Greenness identification based on HSV decision tree". *Information Processing in Agriculture* 2.3-4, pp. 149–160. DOI: 10.1016/j.inpa.2015.07.003
(Cited on page 75)

Curriculum Vitae

Marija Popović

born May 14, 1993 in Serbia

09/2015 – 07/2019	*ETH Zurich, Switzerland* Doctoral studies at the Autonomous Systems Lab of Prof. Roland Siegwart at the Institute of Robotics and Intelligent Systems, Department of Mechanical and Process Engineering; graduated with Dr. sc. ETH Zurich.
06/2014 – 08/2014	*University of Bremen, Germany* Internship in the area of electronic chip design.
09/2011 – 07/2015	*University of Bath, United Kingdom* Undergraduate and graduate studies. Graduated with diploma in Integrated Mechanical & Electrical Engineering.
09/2008 – 07/2011	*American International School in Abu Dhabi, United Arab Emirates* High school education. Graduated with the International Baccalaureate diploma.

List of Publications

Peer-Reviewed Journals

Bähnemann, R., Pantic, M., Popović, M., Schindler, D., Tranzatto, M., Kamel, M., Grimm, M., Widauer, J., Siegwart, R., and Nieto, J. (2019). "The ETH-MAV Team in the MBZ International Robotics Challenge". *Journal of Field Robotics* 36.1, pp. 78–103. DOI: 10.1002/rob.21824

Faria, M., Marín, R., Popović, M., Maza, I., and Viguria, A. (2019). "Efficient Lazy Theta* Path Planning over a Sparse Grid to Explore Large 3D Volumes with a Multirotor UAV". *Sensors* 19.1. DOI: 10.3390/s19010174

Popović, M., Vidal-Calleja, T., Hitz, G., Chung, J. J., Sa, I., Siegwart, R., and Nieto, J. (2019). "An informative path planning framework for UAV-based terrain monitoring". *Autonomous Robots*. Under review. arXiv: 1902.09660

Sa, I., Kamel, M., Burri, M., Bloesch, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2018). "Build Your Own Visual-Inertial Drone: A Cost-Effective and Open-Source Autonomous Drone". *IEEE Robotics & Automation Magazine* 25.1, pp. 89–103. DOI: 10.1109/MRA.2017.2771326

Sa, I., Popović, M., Khanna, R., Chen, Z., Lottes, P., Liebisch, F., Nieto, J., Stachniss, C., Walter, A., and Siegwart, R. (2018). "WeedMap: A Large-Scale Semantic Weed Mapping Framework Using Aerial Multispectral Imaging and Deep Neural Network for Precision Farming". *Remote Sensing* 10.9. DOI: 10.3390/rs10091423

Peer-Reviewed Conference Proceedings

Anil Meera, A., Popović, M., Millane, A., and Siegwart, R. (2019). "Obstacle-aware Adaptive Informative Path Planning for UAV-based Target Search". In: *IEEE International Conference on Robotics and Automation*. Accepted. IEEE

Causa, F., Popović, M., Fasano, G., Grassi, M., Nieto, J., and Siegwart, R. (2019). "Navigation aware planning for tandem UAV missions in GNSS challenging environments". In: *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics. DOI: 10.2514/6.2019-0916

Popović, M., Vidal-Calleja, T., Chung, J. J., Nieto, J., and Siegwart, R. (2019). "Informative Path Planning and Mapping for Active Sensing Under Localization Uncertainty". *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Under review. arXiv: 1902.09660

Longhi, M., Taylor, Z., Popović, M., Nieto, J., Marrocco, G., and Siegwart, R. (2018). "RFID-Based Localization for Greenhouses Monitoring Using MAVs". In: *IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communications*. IEEE, pp. 905–908. DOI: 10.1109/APWC.2018.8503764

Miki, T., Popović, M., Gawel, A., Hitz, G., and Siegwart, R. (2018). "Multi-Agent Time-Based Decision-Making for the Search and Action Problem". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2365–2372. DOI: 10.1109/ICRA.2018.8460996

Sa, I., Chen, Z., Popović, M., Khanna, R., Liebisch, F., Nieto, J., and Siegwart, R. (2018). "weedNet: Dense Semantic Weed Classification Using Multispectral Images and MAV for Smart Farming". In: *IEEE Robotics and Automation Letters*. IEEE, pp. 588–595. DOI: 10.1109/LRA.2017.2774979

Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., and Galceran, E. (2017). "Online Informative Path Planning for Active Classification Using UAVs". In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 5753–5758. DOI: 10.1109/ICRA.2017.7989676

Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 1382–1388. DOI: 10.1109/IROS.2017.8202317

Sa, I., Kamel, M., Khanna, R., Popović, M., Nieto, J., and Siegwart, R. (2017). "Dynamic System Identification, and Control for a cost-effective and open-source Multi-rotor MAV". In: *Field and Service Robotics*. Springer, pp. 605–620. DOI: 10.1007/978-3-319-67361-5_39

Vetrella, A. R., Sa, I., Popović, M., Khanna, R., Nieto, J., Fasano, G., Accardo, D., and Siegwart, R. (2017). "Improved Tau-Guidance and Vision-Aided Navigation for Robust Autonomous Landing of UAVs". In: *Field and Service Robotics*. Springer, pp. 115–128. DOI: 10.1007/978-3-319-67361-5_8

Pfeifer, J., Khanna, R., Constantin, D., Popović, M., Galceran, E., Kirchgessner, N., Achim, W., Siegwart, R., and Liebisch, F. (2016). "Towards automatic UAV data interpretation for precision farming". In: *International Conference on Agricultural Engineering*

Other Publications

Blum, H., Rohrbach, S., Popović, M., Bartolomei, L., and Siegwart, R. (2019). "Active Learning for UAV-based Semantic Mapping". In: *2nd Workshop on Informative Path Planning and Active Sampling, RSS'19*. MIT Press

Papachristos, C., Kamel, M., Popović, M., Khattak, S., Bircher, A., Oleynikova, H., Dang, T., Mascarich, F., Alexis, K., and Siegwart, R. (2019). "Autonomous Exploration and Inspection Path Planning for Aerial Robots Using the Robot Operating System". In: *Robot Operating System (ROS): The Complete Reference (Volume 3)*. Ed. by A. Koubaa. Springer, pp. 67–111. DOI: 10.1007/978-3-319-91590-6_3

Liebisch, F., Popović, M., Pfeifer, J., Khanna, R., Lottes, P., Pretto, A., Sa, I., Nieto, J., Siegwart, R., and Walter, A. (2017). "Automatic UAV-based field inspection campaigns for weeding in row crops". In: *EARSeL SIG Imaging Spectroscopy Workshop*