TÉCNICO LISBOA

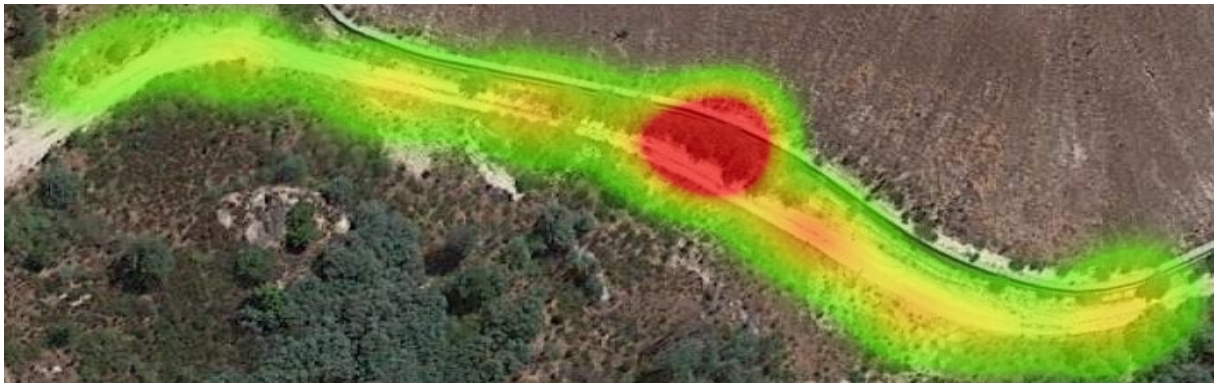

Radiological Monitor for Mobile Robots Operating in Scenarios with Nuclear Threats

Yoeri Joseph Brouwer

Thesis to obtain the Master of Science Degree in

Electrical and Computer Engineering

Supervisors: Dr. Alberto Manuel Martinho Vale
Prof. Rodrigo Martins de Matos Ventura

Examination Committee

Chairperson: Prof. João Fernando Cardoso Silva Sequeira
Supervisor: Prof. Rodrigo Martins de Matos Ventura
Member of the Committee: Prof. António Manuel dos Santos Pascoal

November 2019

Declaração

Declaro que o presente documento é um trabalho original da minha autoria e que cumpre todos os requisitos do Código de Conduta e Boas Práticas da Universidade de Lisboa.

Declaration

I declare that this document is an original work of my own authorship and that it fulfills all the requirements of the Code of Conduct and Good Practices of the Universidade de Lisboa.

Acknowledgments

This work would not have been possible without the guidance, counsel, and patience of my thesis supervisors, Doutor Alberto Manuel Martinho Vale and Professor Rodrigo Martins de Matos Ventura. I would like to express my deepest gratitude towards them for introducing me to the world of scientific research and for sharing their extensive knowledge of the area with me. They provided me with the tools I needed to successfully complete my dissertation. I am also extremely grateful to the Instituto de Plasmas e Fusão Nuclear for having granted me support and a space where I could work and grow.

I would like to extend my sincere appreciation to Doutora Tânia Rodrigues Pereira Ramos for giving me the opportunity to explore new areas and hone additional skills.

Special thanks to Professor Pedro Urbano Lima, who taught the Autonomous Systems course in the most engaging manner and showed me the elegance of probabilistic robotics. He was always available to answer my questions and inspired me to pursue this area in more depth. Thanks also to the MEEC Coordination Team for putting together an excellent curriculum and reminding me that engineering can be enjoyable and stimulating.

I cannot begin to express my thanks to my parents for showing me unrelenting support all these years, and to my brother for inspiring me to be better and for being an example of perseverance. I don't know where I would be without you.

The completion of my dissertation would not have been possible without the support, confidence, friendship, and encouragement of my girlfriend. You gave me valuable feedback, you kept me going when times got hard, and you were always there with a helping hand. Thank you Inês.

In addition, I wish to thank Duncan Schuurman for being a great friend and always showing interest in my work. Thanks should also go to my cat Artémis for keeping me company in the long nights writing this dissertation. I cannot express enough gratitude to all my friends for keeping the ground on my feet.

Finally, I want to recognize the Python community for creating excellent tools, for helping me solve technical problems, and for making software development a fun and interesting challenge.

Resumo

O desenvolvimento das capacidades de sistemas autónomos torna possível a automação de missões de inspeção e de monitorização, que são especialmente relevantes em situações envolvendo ameaças químicas, biológicas, radiológicas e nucleares, denominadas CBRN. A capacidade para utilizar agentes autónomos nestas situações pode manter os humanos fora de perigo. Esta dissertação propõe um algoritmo de planeamento de caminhos informativos como solução para gerar eficientemente mapas de grandezas escalares, motivado por aplicações em missões de inspeção radiológica. A solução proposta é baseada em conceitos de teoria de informação, processos gaussianos e parametrização de caminhos com *B-splines*, o que permite executar otimização e planeamento em espaço contínuo. O algoritmo equilibra a exploração de regiões do espaço de trabalho que não foram observadas com a exploração mais detalhada de *hotspots*, que são considerados interessantes no contexto deste trabalho. Este equilíbrio é alcançado treinando um processo gaussiano no conjunto de observações adquiridas e maximizando a utilidade de observações futuras. O desempenho do algoritmo desenvolvido é avaliado num ambiente simulado que contém várias fontes radioativas e é comparado com os resultados obtidos com um caminho *boustrophedon* e com um caminho resultante de movimento Browniano. Adicionalmente, foi implementado um algoritmo para estimar parâmetros de fontes radioativas, demonstrando que a solução proposta aumenta a precisão das posições estimadas das fontes. O algoritmo de planeamento de caminhos pode ainda ser adaptado para lidar com ameaças CBRN, tornando possível a inspeção de campos escalares de outros tipos, como por exemplo a concentração de poluentes atmosféricos.

Palavras Chave

Planeamento de caminhos; Planeamento adaptativo; Cobertura de área; Monitorização radiológica; Sensores; Veículos autónomos.

Abstract

The ever increasing capabilities of autonomous systems have opened the path for automated surveying and monitoring missions, which are especially relevant in situations involving chemical, biological, radiological, and nuclear threats, known collectively as CBRN. The ability to deploy autonomous agents in these situations can keep humans out of harm's way. This work proposes an informative path planner as a means to efficiently generate maps of scalar quantities, motivated by applications in radiological survey missions. The proposed solution is based on concepts of information theory, Gaussian processes, and path parameterization with B-splines, which enables optimization and planning in continuous space. The informative path planner makes a trade-off between exploring unobserved areas of the workspace and exploring hotspots in greater detail by training a Gaussian process on the current set of observations and maximizing the utility of the future observations. The path planner's performance is assessed in a simulated environment containing several radioactive sources and compared to results obtained with a boustrophedon path—the standard path for area coverage—and with a path resulting from Brownian motion. Additionally, a source parameter estimation algorithm was implemented and demonstrated that the informative path planner increases the source position estimate accuracy. Given that the proposed path planner is suitable for problems dealing with CBRN threats, this work opens up avenues for surveys of other scalar fields (e.g. locating the source of lethal gas attacks). This is an emerging field in the security sector which requires novel techniques, and intelligent autonomous solutions can safeguard human lives.

Keywords

Path planning; Adaptive planning; Area coverage; Radiological monitoring; Sensors; Autonomous vehicles.

Contents

1 Introduction **1**

1.1 Motivation . . . 2

1.2 Previous Work . . . 9

1.3 Objectives . . . 12

1.4 Outline of the Approach . . . 12

1.5 Contributions . . . 13

1.6 Dissertation Structure . . . 14

2 Background **15**

2.1 Concepts of Information Theory . . . 16

2.2 Offline Path Planning . . . 18

2.3 Online Path Planning . . . 20

2.4 Comparison of Approaches . . . 24

2.5 Radiological Units . . . 27

3 Proposed Solution **29**

3.1 Problem Formulation . . . 30

3.2 Assumptions . . . 35

3.3 Selected Approach . . . 37

4 Results **45**

4.1 Radiological Heatmap Simulator . . . 46

4.2 Quality Metrics . . . 48

4.3 Simulation Results . . . 50

4.3.1 Changing the Number of Sources . . . 57

4.3.2 Changing the Budget . . . 66

5 Conclusions **71**

5.1 Work Accomplished . . . 72

5.2 Future work . . . 73

References **75**

A MARIA Description **81**

B Absorbed Radiation Dose Examples **83**

List of Figures

1.1 CBRN warning signs. . . . 2
1.2 Aftermath of the two largest nuclear accidents in human history. . . . 3
1.3 The Three Mile Island nuclear power plant. . . . 4
1.4 Location of the uranium mining site near the town of Retortillo. . . . 5
1.5 Comparison between the trefoil and the new ionizing radiation warning symbols. . . . 7
1.6 Location where one of the cesium sources was found at the Lilo Training Center. . . . 8
1.7 A heatmap representation of a survey made in the field with MARIA. . . . 11
1.8 Results obtained at the IST Alameda campus with a quadcopter. . . . 11
1.9 Results of building a 3D radiological intensity map using SLAM techniques. . . . 13

2.1 Visualization of the relationships between entropy and mutual information. . . . 17
2.2 Two examples of a boustrophedon path for complete area coverage. . . . 18
2.3 Decomposition of the coverage region into subregions with optimized sweep directions. . . 19
2.4 Complete coverage path planner which covers every point in the vicinity of obstacles. . . 19
2.5 Solution to the complete coverage path planning problem applied to cooperative UAVs. . 20
2.6 The different map areas used by the path planner proposed by Zhang et al. . . . 21
2.7 Result of the path planner presented by Zhang et al. . . . 21
2.8 Hexagonal tiling over the coverage area used to generate a directed acyclic graph. . . . 22
2.9 Paths resulting from the solution proposed by Paull et al. (2014). . . . 23
2.10 Spiral path generated by the solution proposed by Paull et al. (2013). . . . 23
2.11 Simulation results of the continuous-space informative path planner proposed by Hitz et al. 24

3.1 Schematic representation of the problem variables. . . . 31
3.2 Graphs of the reward function r used in the utility function. . . . 33
3.3 Graphs of the smooth step function used in the penalty factors. . . . 34
3.4 Illustration of the signed distance function. . . . 35
3.5 Illustration of the terms used in the radiation absorption model. . . . 37
3.6 Level curves and minimizers of Himmelblau's function. . . . 38
3.7 Evolution of CMA-ES when finding the minimum of the function $f(x_1, x_2) = x_1^2 + x_2^2$. . . . 39
3.8 Evolution of CMA-ES when finding the minimum of the Himmelblau function. . . . 40
3.9 Evolution of CMA-ES when finding the minimum of the Himmelblau function. . . . 41

3.10 Evolution of CMA-ES when finding the minimum of the Himmelblau function. 42

4.1 Simple toy scenarios generated with the radiological heatmap simulator. 47
4.2 Toy scenario containing three sources where thirty observations were simulated. 47
4.3 A more complex scenario generated with the radiological heatmap simulator. 48
4.4 A simple scenario containing a single radioactive source at the origin. 50
4.5 First iteration of the informative path planner. 51
4.6 Second iteration of the informative path planner. 52
4.7 Several intermediate iterations of the informative path planner. 53
4.8 Final state of the informative path, the boustrophedon path, and the Brownian motion path. 54
4.9 Representation of the ground truth and the estimated heatmaps. 56
4.10 The ground truth for the first scenario, which contains a single source at the origin. 57
4.11 The ground truth for the second scenario, which contains three sources. 58
4.12 Graphs of $Q_{\mathbf{p}}$ and Q_A for the informative path planner and the boustrophedon paths. . . . 59
4.13 The ground truth for the third scenario, which contains four sources. 60
4.14 Graphs of $Q_{\mathbf{p}}$ and Q_A for the informative path planner and the boustrophedon paths. . . . 61
4.15 The ground truth for the fourth scenario, which contains five sources. 62
4.16 Graphs of $Q_{\mathbf{p}}$ and Q_A for the informative path planner and the boustrophedon paths. . . . 63
4.17 Graphical representation of Q_K versus the number of sources. 64
4.18 Graphical representation of $Q_{\mathbf{p}}$ versus the number of sources. 65
4.19 Graphical representation of Q_A versus the number of sources. 65
4.20 The ground truth for comparing the standard deviation distributions. 66
4.21 Histograms of the distribution of the standard deviation with $B = 120$ m. 67
4.22 Histograms of the distribution of the standard deviation with $B = 160$ m. 67
4.23 Histograms of the distribution of the standard deviation with $B = 200$ m. 67
4.24 Histograms of the distribution of the standard deviation with $B = 240$ m. 68
4.25 Histograms of the distribution of the standard deviation with $B = 280$ m. 68
4.26 Histograms of the distribution of the standard deviation with $B = 320$ m. 68
4.27 Histograms of the distribution of the standard deviation with $B = 360$ m. 69
4.28 Frequencies of each histogram interval as a function of the budget. 70

A.1 Screenshots of the mobile application MARIA. 82

B.1 Visual representation of typical small radiation doses. 84
B.2 Visual representation of moderate radiation doses. 85
B.3 Visual representation of high radiation doses. 86

List of Tables

2.1 Summary of the methods reviewed in the literature. 24

4.1 Results obtained for a detailed example scenario. 55
4.2 Results obtained from executing ten runs in a scenario containing a single source. 57
4.3 Results obtained from executing ten runs in a scenario containing three sources. 59
4.4 Results obtained from executing ten runs in a scenario containing four sources. 61
4.5 Results obtained from executing ten runs in a scenario containing five sources. 63

List of Algorithms

3.1 Covariance Matrix Adaptation Evolution Strategy optimization routine. 38
3.2 Informative Path Planner. 43
3.3 Function `SolveOptimizationProblem`. 44

Acronyms

APP	Adaptive Path Planning
AUV	Autonomous Underwater Vehicle
CBRN	Chemical, Biological, Radiological, and Nuclear
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
CPM	counts per minute
DAG	Direct Acyclic Graph
GMC	Geiger-Müller counter
GNSS	Global Navigation Satellite System
GP	Gaussian Process
IPFN	Instituto de Plasmas e Fusão Nuclear
IPP	Informative Path Planning
RDD	Radiological Dispersal Device
SLAM	Simultaneous Localizaton And Mapping
UAV	Unmanned Aerial Vehicle
UGV	Unmanned Ground Vehicle
USB	Universal Serial Bus

1

Introduction

Contents

1.1 Motivation 2
1.2 Previous Work 9
1.3 Objectives 12
1.4 Outline of the Approach 12
1.5 Contributions 13
1.6 Dissertation Structure 14

This chapter provides the motivation for this dissertation, describes previous work developed by others to address similar problems, defines the objectives which were set out to be achieved, gives an outline of the approach which is used, lists academic contributions by the author, and details the structure of this dissertation document.

1.1 Motivation

In June 2019 the Stockholm International Peace Research Institute (SIPRI) reported that Chemical, Biological, Radiological, and Nuclear (CBRN) threats are rapidly growing along with considerable changes in the political landscape and developments in technology [1]. The term CBRN refers to the malicious use of these types of materials or weapons with the intention to cause significant damage to a given target [2]. Figure 1.1 illustrates the warning signs for each material. The distinction between radiological and nuclear threats derives from the different effects they each have: radiological threats relate to "*illness caused by exposure to harmful radioactive materials*", whereas nuclear threats relate to "*life-threatening health effects caused by exposure to harmful radiation, thermal or blast effects arising from a nuclear detonation*" [2]. As stated by SIPRI and the United Nations Interregional Crime and Justice Research Institute [3], the paradigm of national security is shifting in order to deal with globalized threats, forcing governments to look beyond their own borders to guarantee the safety of their citizens; this paradigm shift has become especially pertinent with respect to CBRN threats. Accidents, misappropriation, illegal transport, and theft of CBRN agents present new challenges in terms of national and international security. Biological and chemical substances can be acquired relatively easily through the healthcare and industrial sectors. This presents a complex challenge to security organizations trying to prevent actors with malicious intents from obtaining dangerous materials. SIPRI also emphasizes that the combination of technological developments and complex global supply chains can increase the ease of producing and deploying such weapons [1].

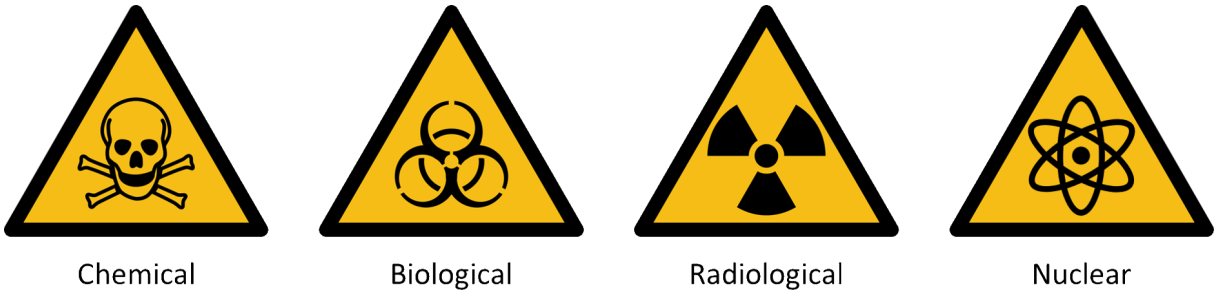

Figure 1.1: CBRN warning signs.

Currently, several commercial CBRN detection technologies are available, although their maturity and suitability for use in a security context requires the input of field specialists [2]. Furthermore, there is a strong focus from organizations to mitigate and address the chemical and nuclear threats [1], such that there is a lack of solutions for radiological incidents. Therefore, this dissertation will concentrate on the proposed solution's applicability in the radiological context, although, in principle, it can also be utilized in other CBRN contexts.

Radioactive materials can be found in numerous scenarios, such as nuclear reactor facilities, the residual ash from thermoelectric power plants, mining sites, medical facilities, and military applications (e.g., nuclear warheads). Hazardous situations can arise as a result of accidents, misuse of work equipment, improper disposal of waste products, or incidents where dangerous material is hidden for use in attacks using nuclear weapons; it is important to understand their origin and the effects that they have. Several situations will be detailed and discussed in the following paragraphs to highlight the need for good radiological monitoring solutions.

There are two nuclear reactor incidents which take the main stage in the public eye: the Chernobyl [4, 5] and Fukushima [6, 7] nuclear disasters. The former occurred on April 26th, 1986, when an explosion and fires at the Chernobyl nuclear plant in Ukraine caused the largest uncontrolled release of radioactive material in the history of the civil nuclear industry. Over the following 10 days, large quantities of radioactive iodine and cesium were released into the air. Most of this material was deposited near the installation, but lighter material was carried by wind currents over Belarus, the Russian Federation and Ukraine and, to some extent, over parts of Europe. Between 1986 and 1990, approximately 530,000 registered recovery operation workers received an average dose of 120 millisievert (mSv), ranging from 20 to 500 mSv [4]. For reference, the global average annual dose from natural background radiation is 2.4 mSv [8]. It should also be noted that receiving a high radiation dose over a short period of time (up to a few hours) has more acute and immediate effects when compared to receiving the same dose over a long period of time [9]. Thus, on average, the recovery operation workers received over 10 times the natural background radiation dose, reaching up to 40 times the natural background radiation dose for those which were most exposed. As a result of the radiation exposure during the recovery operation, 134 workers received very high doses and suffered from acute radiation sickness. Of those, 28 workers died in the first three months following the accident [4, 8]. Furthermore, an increase in the rates of thyroid cancer and radiation induced cataracts has been observed in the workers and populations around the Chernobyl nuclear plant. The nuclear reactor site is strictly monitored and controlled, and a sarcophagus has been constructed around the damaged facilities, as shown in fig. 1.2a.

(a) The Chernobyl No. 4 nuclear reactor as of 2006. Credit: Carl Montgomery, licensed under CC BY 3.0.

(b) The Fukushima Daiichi nuclear power plant, 14 March 2011. Credit: Digital Globe, via New York Times.

Figure 1.2: Aftermath of the two largest nuclear accidents in human history.

The Fukushima Daiichi nuclear disaster occurred on March 11th, 2011 when a magnitude 9 earthquake off the east coast of Japan resulted in a devastating tsunami. As a consequence of the tsunami, the Fukushima Daiichi Nuclear Power Station, located along the shoreline, lost its core cooling capacity which caused severe damage to the reactor's core and led to a nuclear accident [7]. Explosions at the power station on the 12th, 14th, and 15th of March resulted in the release of substantial amounts of radioactive materials into the environment (see fig. 1.2b). Furthermore, it is estimated that the radioactive cloud and ground deposits caused people living in the Fukushima prefecture to be exposed to 10 mSv. Over the first 19 months after the accident, the average effective dose for workers at the power station was approximately 12 mSv, and around 160 workers received doses of more than 100 mSv.

There is still a third well-known nuclear incident which posed a considerable threat to civilian populations: the Three Mile Island accident [10]. This accident occurred on March 28th, 1979, in Pennsylvania, USA, and was "*the most serious accident in U.S. commercial nuclear power plant operating history*" [10]. A combination of mechanical failures in the cooling system and problems related to human-computer interactions led to the partial meltdown of one of the reactors and subsequent releases of radioactive materials into the surrounding environment. Following investigations by several organizations, it is estimated that the maximum dose received at the site boundary was less than 1 mSv. No negative health effects could be directly correlated to the radioactive releases from the nuclear reactor [10]; nonetheless, this event led to sweeping changes across the nuclear industry and the safety regulations. In the aftermath of this nuclear accident, the damaged reactor was never restarted, as illustrated in fig. 1.3, and the remaining operating reactor is scheduled to be permanently shut down by September 30th, 2019 [11].

Figure 1.3: The Three Mile Island nuclear power plant. The cooling towers for the damaged and deactivated reactor are on the left, and the cooling towers for the remaining active reactor are on the right.
Credit: Richard Hertzler/LNP/LancasterOnline, via New York Times.

The story of these nuclear incidents exposes a clear need for radiological monitoring solutions which are adequate both for emergency situations, as is the case of the immediate aftermath of a nuclear disaster, and for long term purposes, e.g. to monitor how the radioactive materials have spread and deposited throughout the environment.

In addition to nuclear disaster sites, there is the case of abandoned or decommissioned mining sites, particularly uranium and radium mines. In the northern region of Portugal, namely the districts of Viseu and Guarda, uranium deposits were discovered as early as 1907 [12]. Extensive mining activities took part during the 20th century; however, activity ceased in 2001 [12]. The closed mines left a legacy of contaminated soil and environmental issues due to the radiological impact of mining by-products [13,14]. Since 2001, the Empresa de Desenvolvimento Mineiro[1] has been responsible for maintaining the safety of the mining sites (that is, keeping civilians out of harm), as well as their environmental clean-up and monitoring [14]. In spite of this costly legacy (which amounted to 49 million euros in 2017 [14]), new issues still arise: near the town of Retortillo, Spain, on the banks of the Yeltes river, a new open-pit uranium mine was opened up in 2018 [15]. As with many mining activities, large amounts of water are required for processing the extracted materials and, as such, the water from the Yeltes river was being used at this site. However, the Portuguese authorities were concerned with the long-term and widespread effects which this mine could have, given that the Yeltes river eventually feeds into the Douro river, which crosses the Portuguese territory into the Atlantic ocean (see fig. 1.4). More importantly, the Portuguese authorities and local populations remain worried about radiological contamination of the river water and underground water sheets, as well as soil contamination from radioactive dust and all the other dangers which it brings [15]. Once again, as mentioned above, CBRN threats have large-scale and potentially cross-border effects, even when there is no malicious intent.

Figure 1.4: Location of the uranium mining site near the town of Retortillo, 40 km away from the Portuguese border. Credit: Global Media Group, via Diário de Notícias.

[1] Mining Development Company

Finally, dangerous situations can emerge when radioactive materials—utilized in nuclear medicine equipment and industrial machines—are improperly transferred or disposed of, or when they are obtained for use in terrorist attacks. A prime example for the origin of such radioactive materials is an orphan source: a sealed source of radioactive material contained in a small volume "*which is not under regulatory control, either because it has never been under regulatory control or because it has been abandoned, lost, misplaced, stolen or otherwise transferred without proper authorization*" [16, 17]. Several orphan source incidents have occurred since the 1980s:

- In September 1987, a private radiotherapy institute in Goiânia, Brazil, abandoned a cesium-137 teletherapy unit when they moved to a new premises [18]. At the time, the source had an activity of 50.9 TBq. Shortly after, two people took this unit, without knowing what it was, and tried to dismantle it to sell for scrap. This ruptured the source capsule, which contained cesium chloride salt (which is highly soluble and easily dispersible), thus contaminating the environment; moreover, the parts of the source assembly were later sold to a junkyard owner. Due to the fascination with the blue glow of the material, fragments were distributed among numerous people. As a result, 249 people suffered radiation contamination, of which 49 were hospitalized and 4 died. In addition, the environment was severely contaminated and required a large clean-up operation, during which an aerial survey with a helicopter was performed. Known as the Goiânia accident, this became one of the most disastrous orphan source incidents to ever occur.

- In the autumn of 1999, a private radiotherapy company based in Bangkok, Thailand, was in possession of several teletherapy devices—without the proper authorization from the Thailand Office of Atomic Energy for Peace—, when it decided to relocate the teletherapy heads to an unsecured location [19]. In January 2000, several people were able to obtain, and partially disassemble one of the teletherapy heads. Even though the object had warning labels and symbols, the individuals were not aware of the dangerous nature of the radioactive materials within. Later, on 1 February 2000, two people brought the teletherapy head to a junkyard in Samut Prakarn, Thailand, to completely disassemble and sell the separate parts. During this process, unbeknownst to anyone around, the cobalt-60 source fell out of its casing and was left at the junkyard. In the following weeks, several people were exposed to large doses of ionizing radiation, resulting in the hospitalization of 10 people, 3 of which died within two months. When the source was finally recovered on 20 February, it was still inside its capsule and had an estimated activity of 15.7 TBq. This accident was one of the reasons which led the IAEA to design a new ionizing radiation warning symbol (ISO 21482, see fig. 1.5), which is intended to clearly warn of the presence of dangerous levels of ionizing radiation, and is recommended for use on "*IAEA category 1, 2 and 3 sealed radiation sources (dangerous sources that can cause death or serious injury)*" [20].

- In February 2010, a Gammacell 220 machine used for research purposes by the University of Delhi was sold to a metal scrap dealer in Mayapuri, New Delhi, India [22]. This machine housed 112 cobalt-60 slugs, contained in 16 pencils (7 slugs per pencil), with a total activity content of 0.69 TBq. A month later, in March 2010, workers at the scrap yard dismantled the machine and divided the

Figure 1.5: Left: the trefoil warning symbol used on containers for radiological materials, highly recognized in the nuclear industry. Right: the ionizing radiation warning symbol launched on 15 February 2007 by the IAEA and the ISO [21], designed to clearly warn anyone of the immediate danger.

cobalt sources into eleven pieces, two of which were moved to a nearby shop and another was kept in the wallet of one of the workers. Then, in April 2010, several individuals were hospitalized due to radiation exposure; subsequently, an investigation and search for the radioactive materials took place. By the middle of April, all of the cobalt sources had been recovered by emergency personnel, and they were transported to the nearest authorized waste disposal facility for safe storage. Furthermore, the clean-up operation collected 400 kg of radioactive contaminated soil from affected locations. As a result of this accident, 7 people suffered from radiation injuries and 1 person died due to radiation sickness.

- Between 1997 and 2006, almost 300 orphan sources were found in Georgia, most of which were remnants from the Cold war [23–26]. To keep the description short, there were two major incidents in that time. The first occurred in 1997 at the Lilo Training Center, about 25 km east of Tbilisi [23]. Sealed radiation sources had been abandoned at this site some years prior to the incident. In April–August 1997 several soldiers developed skin lesions, and an investigation team was assembled to determine the locations and types of the sources which caused the injuries. Hundreds of sources were discovered in numerous locations (see fig. 1.6), including cobalt-60, cesium-137, and radium-226 sources, although the cobalt source was very weak; most of the sources were sealed, but at least 10 of them were missing their lead containers. It was concluded that the main cause of the accident, which resulted in the hospitalization of 11 soldiers, was the improper and unauthorized abandonment of 12 cesium-137 sources. The second major incident occurred in December 2001, near Lia [25, 26]. Three woodcutters drove approximately 50 km east of Lia to collect firewood, when they stumbled upon two cylindrical containers; strangely, the snow around the containers had melted within a 1 m radius. The men proceeded to spend the night in the forest, keeping warm with a campfire on one side and the containers on the other. Three hours later, they began presenting symptoms of acute radiation syndrome, which lasted throughout the night. In the following days, each one of them continued suffering from several symptoms, and by the end of December 2001

they were hospitalized. On 29 December 2001, 27 days after the initial exposure, the Georgian authorities successfully located the “heating devices”; however, recovery and containment of the devices by a team of specialists was only possible on 2 February 2002. It was later determined that the containers were actually components of Soviet-era radioisotope thermoelectric generators, containing strontium-90 sources with an activity of 1295 TBq. The three men suffered serious injuries due to the radiation exposure, and were treated for several years, ultimately resulting in the death of one of them in May 2004.

Figure 1.6: Location where one of the cesium sources was found at the Lilo Training Center. Photo from [23].

Besides the terrible direct consequences of orphan sources upon unsuspecting individuals, there is another concern surrounding orphan sources: their utilization for Radiological Dispersal Devices (RDDs), colloquially known as “dirty bombs” [24, 25]. The basic idea behind these weapons is to disperse radioactive material over a large area by use of conventional explosives. Here, orphan sources would be one way to obtain the radioactive material. However, the radiological effects of a dirty bomb are not expected to be fatal in most situations [27, 28], considering that the more dispersed the radioactive material is, the weaker the dose it delivers. Nonetheless, there are fears that a RDD could be used for psychological terrorism, causing widespread panic because of the fear of radioactivity [28].

In all of the situations described above, it is important to have fast, reliable, and efficient inspection tools, both in the immediate aftermath of a radiological incident and for long term monitoring. Several solutions have already been developed and are discussed in section 1.2. In the case of surveying missions and source localization tasks, efficient path planning solutions can have a big impact, especially with the

advent of autonomous vehicles and robots. By mounting a suite of radiological sensors (and possibly other types of sensors) on a mobile autonomous or remote controlled platform, human operators' safety is vastly improved. Furthermore, by developing path planning tools for these mobile platforms, mission times can be reduced and hotspots can be quickly detected. Thus, based on the motivations presented in this section, this dissertation proposes an informative path planner optimized for radiological hotspot detection. Nonetheless, the methods presented in the chapters ahead need not be limited to applications in radiological monitoring, and can also be applied to chemical and biological threats—a possible theme for future work.

1.2 Previous Work

As was mention in section 1.1, several solutions already exist to address the problem of radiological monitoring. Some of them are commercial solutions and others are developed by research institutions for public safety.

One such solution is the REWARD project, which stands for Real Time Wide Area Radiation Surveillance System [29]. The goal of the REWARD project is the development of "*portable, intelligent radiation detectors that can determine the flux and energy of the incoming radiation, as well as their own location*" [29]. The main features of the REWARD system are: a real-time system with wide area coverage; radiological detectors using novel solid-state technologies, namely cadmium zinc telluride (CZT) detectors for gamma radiation and silicon based detectors for neutron radiation; scalability in terms of complexity and costs; and portability and adaptability to any type of environment. The main applications of the REWARD system include scenarios such as "*nuclear terrorism threats, lost radioactive sources, radioactive contamination, or nuclear accidents*" [29]. Additionally, one important aspect of the REWARD system is that the radiological measurements are georeferenced, that is, associated to a precise geographic location. Since the sensing units are small and have low energy requirements, this allows them to be portable and opens up the possibility for precise radiological hotspot localization [29].

Another solution is provided by Seagull Environmental Technologies [30], a company based in the United States of America. This company provides support for a number of scenarios, mostly involving assessment or removal of hazardous materials (e.g., asbestos and lead). However, of interest to this dissertation is their involvement in emergency responses, where they list their expertise in chemical spills, accidental or intentional releases of CBRN agents, oil spills, counter-terrorism incidents, and federal disasters [30]. It can thus be seen that the application of these solutions is not limited to radiological inspection, but also applies to situations involving chemical and biological agents.

The SEAGULL project [31,32] (unaffiliated with Seagull Environmental Technologies) is developed by the Portuguese Air Force in collaboration with several research institutes, with the goal of developing intelligent systems to support maritime situational awareness based on Unmanned Aerial Vehicles (UAVs). It proposes the creation of an intelligent maritime surveillance system by equipping UAVs with different types of optical sensors. Optical sensors such as cameras (visible, infrared, multi and hyper spectral) can contribute significantly to the generation of situational awareness of maritime events such as de-

tection and georeferencing of oil spills or hazardous and noxious substances, object tracking systems (e.g., lifeboats, debris), recognizing behavioral patterns (e.g., high-speed vessels, atypical navigation patterns), and monitoring indicators of good environmental health.

Gamma-Ex [33] is a project which uses UAVs to perform monitoring and recognition of chemical and radiological patterns in scenarios where biological, chemical, and radiological threats are identified.

Additional solutions for radiological monitoring are developed by Orano, including robots and instruments such as MANUELA (Mobile Apparatus for NUclear Expertise and Localisation Assistance), DORICA (Drone for Onsite Radiological Investigations, Characterizations and Assessments), and RIANA (Robot for Investigation and Assessments of Nuclear Areas) [34]. MANUELA is described as a portable real-time radiological and topographical mapping device [34], that is, it is an instrument which measures radiological intensity and associates each measurement with a real-world position. DORICA performs a similar role, however the sensors are mounted on a drone, allowing for the inspection of high walls and of hazardous areas without exposing a human operator to high radiation doses. Finally, RIANA is an Unmanned Ground Vehicle (UGV) with a modular design and the capability to make radiological measurements or even acquire samples [34].

Many of the existing solutions described in this section require expensive and dedicated hardware components to achieve their goals. However, to map geographic positions to radiological intensity levels a simpler solution can be developed, because the only required components are a radiation detector and some location measuring system (such as Global Navigation Satellite System (GNSS)). The need for such a solution led to the development of MARIA, detailed below.

Mobile Application for Radiation Intensity Assessment

MARIA, a Mobile Application for Radiation Intensity Assessment, is an Android application for handheld devices, such as smartphones or tablets, designed to build a radiological intensity map of a given scenario. MARIA was developed out of a need to acquire and record georeferenced radiological intensity data, and has been in development at the Instituto de Plasmas e Fusão Nuclear (IPFN) since October 2017. The application software for the initial version of the application was developed by the author, and has since seen contributions from other collaborators at the IPFN; development is still in progress. The application combines location data from the device GNSS receiver and/or telecommunications antenna with radiological intensity readings from a Geiger-Müller counter (GMC) transmitted through the device Universal Serial Bus (USB) port, using a serial-to-USB converter. An algorithm is then applied to the combined location and intensity data to create a radiological heatmap of the scenario, displayed as an overlay on Google Maps™ for visualization, as illustrated in fig. 1.7. This allows a user to quickly identify hotspots and dangerous areas. Furthermore, since the measurement data is recorded and can be viewed again at another time, MARIA can monitor radiation levels not only in space, but also in time. Hence, MARIA can be utilized to detect when the activity of a given radioactive source has decreased; as an example, relevant authorities can take advantage of this feature to characterize an area as radiologically safe for the general public in the aftermath of a nuclear event.

Figure 1.7: A heatmap representation of a measurement session made in the field with MARIA, obtained via the Google Maps™ API. The red spot coincides with a structure used for mining operations.

In its current state, MARIA is only apt to measure, visualize, and record radiological data, and has no user guidance system. A more detailed specification of MARIA can be found in Appendix A. The informative path planner developed in this dissertation opens the possibility to enhance MARIA's capabilities; for example, users can be guided along an efficient path. Furthermore, since the application runs on smartphones and the GMC is light weight, the platform can be transported on a quadcopter. Some initial tests have already been performed with a DJI Phantom 3 Standard, namely in [35], illustrated in fig. 1.8.

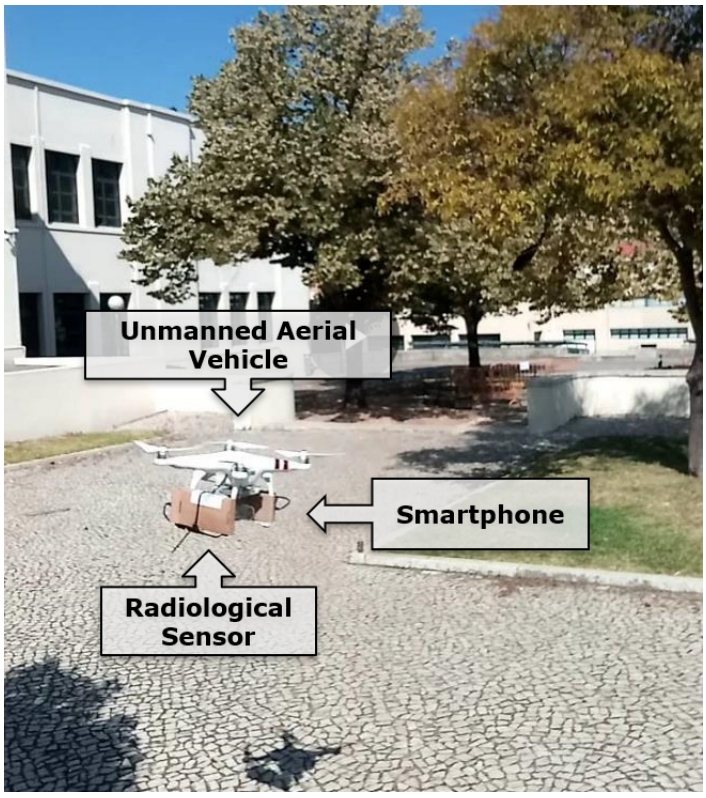

(a) DJI Phantom 3 Standard with MARIA.

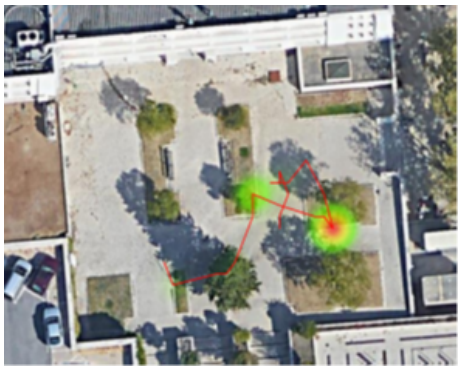

(b) First drone run.

(c) Second drone run.

Figure 1.8: Quadcopter setup (a) and results from flights performed at the IST Alameda campus (b and c).

1.3 Objectives

The goal of this thesis is to develop and implement a path planner for mobile agents (robots or humans) equipped with sensors which measure information about a continuous scalar field. The path planner should determine the optimal path such that the sensor measurements maximize the information gained about the field being analyzed. In order to achieve that goal, the objectives are:

1. Formulate the optimization problem: this involves defining an objective function which can be optimized to determine the best path for information gain, as well as selecting an optimization approach which is suitable for solving the optimization problem.

2. Implement and analyze the selected optimization approach: this has to be done in order to understand its behavior, advantages, and disadvantages.

3. Obtain a solution in the form of an adaptive path planner.

4. Quantify the quality and robustness of the results: by analyzing the results and comparing them to existing solutions, the strengths and weaknesses of the proposed solution can be identified.

1.4 Outline of the Approach

The proposed approach is an informative path planner based on an iterative algorithm which maximizes a utility function in a receding horizon framework, through the utilization of B-splines, Gaussian Processes (GPs), and mutual information. This approach is loosely based on the work presented in [36].

The initial step involves generating a virtual path which acts as a guide for the path which is optimized and traveled. The idea behind this path is to enable defining a start point and an end point for the path planner, as well as ensuring a minimal amount of coverage of the workspace.

Thereafter, the algorithm successively maximizes the defined utility function. A GP is trained on the observation data obtained in the environment, as a tool to estimate the underlying scalar field, in addition to quantifying the uncertainty of the estimated field. The field uncertainty arises as a consequence of uncertainty in the observations and regions of the workspace which have not yet been explored. The GP mean and covariance matrix are combined in order to define the utility function as a modified form of mutual information [37, 38]. The mean of the GP is included in the utility function to reward exploring regions of the workspace where the field's magnitude is large. Furthermore, multiplicative penalty functions are utilized to enforce workspace, budget, and progress constraints, so that the path remains within the workspace, does not exceed the defined budget, and stays within a maximum distance of the virtual path, respectively. Finally, the path is parameterized as a B-spline which is defined by a set of control vertices; consequently, the path can be planned in continuous space and does not require a discretization of the workspace. To tie everything together, the B-spline determines the positions of future observations from which the GP mean and covariance matrix are calculated. When a solution has been found, the agent travels the solution path, making new observations.

1.5 Contributions

The author of this dissertation has also co-authored several research papers related to the topic of radiological threats.

- The first is entitled “Remote inspection with multi-copters, radiological sensors and SLAM techniques,” published in *EPJ Web of Conferences*, Volume 170 (2018) and presented at ANIMMA 2017: Advancements in Nuclear Instrumentation Measurement Methods and their Applications. This work developed a Simultaneous Localizaton And Mapping (SLAM) algorithm which simultaneously makes a 3D reconstruction and estimates a radiological heatmap of a given scenario. Two of the results are illustrated in fig. 1.9; here, a point-like source was placed in the corner of a scenario, and a drone carrying a Structure Sensor and a GMC executed a spiral trajectory through the scenario; then, after some post-processing, a 3D reconstruction of the scenario was made, shown in fig. 1.9a, and each point of the reconstruction was mapped to the respective radiological intensity levels, shown in fig. 1.9b.

- The second paper is entitled “MARIA - Mobile Radiation Intensity Mapper,” and was submitted to the Portuguese committee for URSI (International Union of Radio Science) call for papers in 2017. The paper developed a method to estimate the radiological intensity at all points of a grid given a set of radiological intensity measurements.

- The third paper is entitled “Radioactive Hot-spot Detection Using Unmanned Aerial Vehicle Surveillance” (in press), will be published in *EPJ Web of Conferences*, and was presented at ANIMMA 2019. This work developed a radioactive source characterization algorithm which quantifies the number of sources, their position, and their activity, given a set of radiological observations in a scenario.

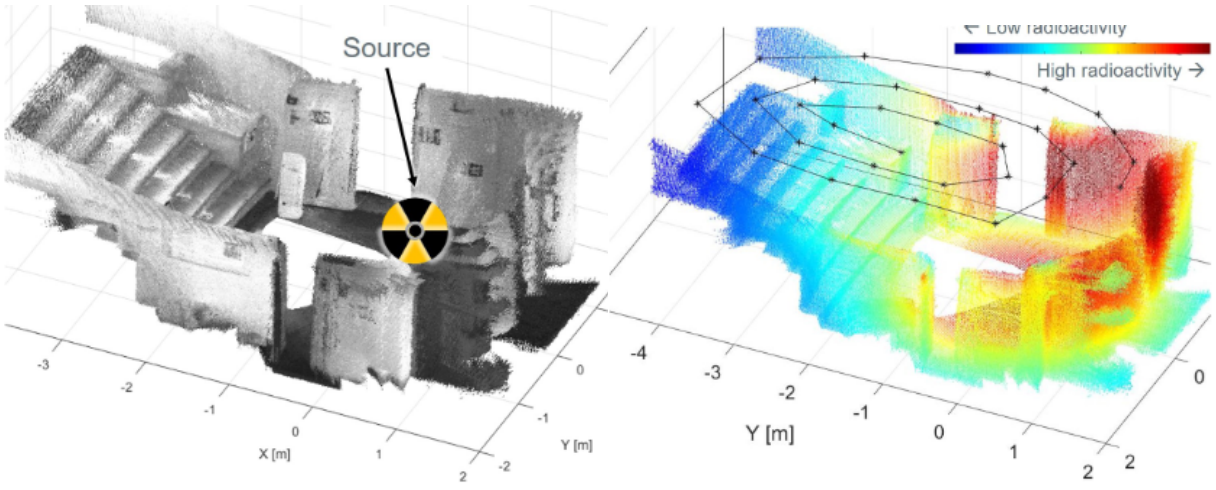

(a) 3D reconstruction. **(b)** Radiological intensity map.

Figure 1.9: Results from the paper Remote inspection with multi-copters, radiological sensors and SLAM techniques.

1.6 Dissertation Structure

The rest of this dissertation is organized as follows:

- Chapter 2 makes a detailed review of the state of the art of path planning, most importantly path planning approaches for complete area coverage and path planning approaches for maximum information gathering.
- Chapter 3 presents a problem formulation for the case in study, accompanied by the assumptions made about the test scenarios, and then followed by a description of the approach which was implemented to solve that problem.
- Chapter 4 describes the results obtained for this dissertation, starting with a technical overview of the radiological heatmap simulator which was developed, succeeded by a description of the metrics which were designed to evaluate the performance of the proposed solution, and ending with an analysis of the results obtained through simulation.
- Finally, Chapter 5 presents the main conclusions of this dissertation, describing the work accomplished and areas to explore in future work.

A list of all the bibliographical references cited in this dissertation can be found in the annex under References. The annex also includes Appendix A, which contains a more detailed description of MARIA, and Appendix B, which gives examples of radiation doses.

2

Background

Contents

2.1 Concepts of Information Theory 16
2.2 Offline Path Planning 18
2.3 Online Path Planning 20
2.4 Comparison of Approaches 24
2.5 Radiological Units 27

This chapter presents a synthesis and discussion of the relevant literature which was reviewed in order to form a solid knowledge base for the development of the dissertation. It is organized as follows: in section 2.1 some commonly used information theory concepts will be defined, to contextualize many of the methods used in other works; in section 2.2 the literature relating to complete coverage offline path planning is reviewed; in section 2.3 the literature relating to online path planning is reviewed, with some emphasis given to informative path planning; in section 2.4 the pros and cons of the studied approaches are listed and compared; and finally, section 2.5 introduces commonly used radiological units of measurement.

2.1 Concepts of Information Theory

In this work we will follow the conventions and definitions for information theory described in [37]. These will be useful to understand the concepts used in the informative path planning literature [36, 39–46], and were necessary for the development of the thesis work.

Entropy is a measure of the uncertainty of a random variable. If we define the probability mass function $p(x) = \Pr\{X = x\}$ for a discrete random variable X with domain $\mathcal{X}$, then the entropy $H(X)$ of X is defined by (2.1). In (2.1) the result is expressed in bits, and we will use the convention that $0 \log_2 0 = 0$ justified by the limit (2.2) and the intuition that adding terms which have a probability of zero does not change the entropy, since they do not add any information [37, p. 14]. Intuitively, entropy can be interpreted as a measure of the average amount of information required to describe a random variable.

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log_2 p(x) \tag{2.1}$$

$$\lim_{x \to 0} x \log_2 x = 0 \tag{2.2}$$

We can also define the joint entropy $H(X, Y)$ of a pair of discrete random variables (X, Y) with a joint distribution $p(x, y)$ as (2.3), and the conditional entropy $H(Y|X)$ as (2.4).

$$H(X, Y) = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log_2 p(x, y) \tag{2.3}$$

$$H(Y|X) = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log_2 p(y|x) \tag{2.4}$$

Relative entropy (also known as Kullback-Leibler distance, cross entropy, or information divergence) is a measure of the distance between two distributions. The relative entropy $D(p||q)$ is a measure of the inefficiency of assuming that the distribution is q when the true distribution is p, and is defined by (2.5).

$$D(p||q) = \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{p(x)}{q(x)} \tag{2.5}$$

In (2.5) we will use the following conventions: $0 \log_2 \frac{0}{0} = 0$, $0 \log_2 \frac{0}{q} = 0$, and $p \log_2 \frac{p}{0} = \infty$ [37, p. 19].

Mutual information is a measure of the amount of information that one random variable contains about another random variable. Given two random variables X and Y with marginal probability mass functions $p(x)$ and $p(y)$ and a joint probability mass function $p(x, y)$, the mutual information $I(X;Y)$ is defined by (2.6).

$$I(X;Y) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x,y) \log_2 \frac{p(x,y)}{p(x)p(y)} \tag{2.6}$$

We can rewrite the definition (2.6) in entropy terms as (2.7), which can be interpreted as the reduction in uncertainty of X due to the knowledge of Y. A graphical representation of the relationships between entropy, joint entropy, conditional entropy, and mutual information is illustrated in fig. 2.1.

$$I(X;Y) = H(X) - H(X|Y) \tag{2.7}$$

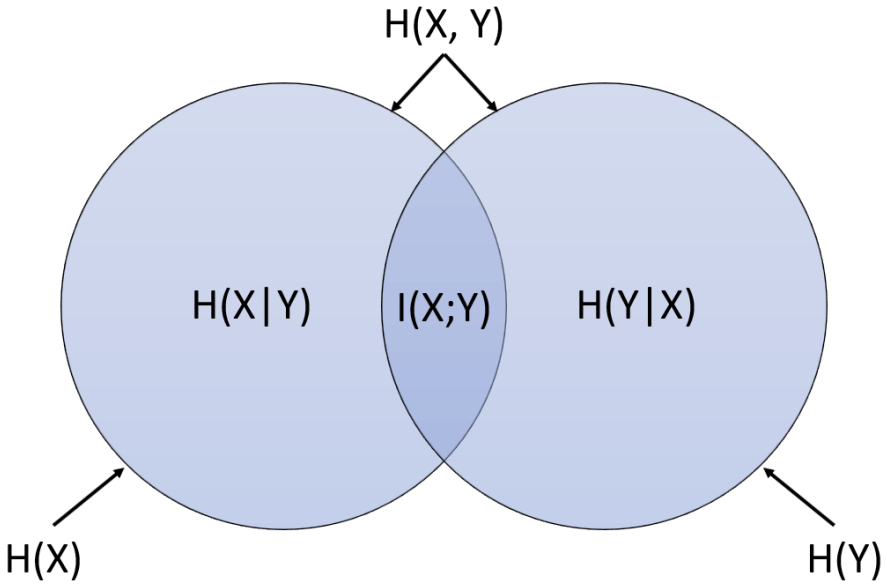

Figure 2.1: Visualization of the relationships between entropy and mutual information. Adapted from [37].

Finally, we define Fisher information. Consider a parametric probability density function $f(x, \theta)$ for which we wish to estimate the parameter θ from a sample size of n. Fisher information is a measure of the amount of information about θ that is present in the sample set and is defined by (2.8). The notation $E_\theta[\cdot]$ denotes the expectation with respect to the density $f(\cdot;\theta)$.

$$J(\theta) = E_\theta \left[\frac{\partial}{\partial \theta} \ln f(X;\theta) \right]^2 \tag{2.8}$$

The concepts of mutual information, Fisher information, and especially entropy come up repeatedly in the informative path planning literature [36, 39–46], and so it was useful and necessary to gain some familiarity with their meaning and mathematical description. For that purpose, the definitions, conventions, and notation used in [37] were followed.

2.2 Offline Path Planning

The objective of complete coverage path planning is to determine a path which passes within a maximum distance from all points of an area or volume of interest while avoiding obstacles [47]. The most commonly used method is known as the boustrophedon cellular decomposition [48] and, in essence, this method consists of:

1. dividing the workspace into smaller regions (cells) by performing an exact cellular decomposition;
2. building an adjacency graph for the resulting cellular decomposition;
3. finding a path which visits each node (cell) in the adjacency graph;
4. computing the explicit robot motions within each cell, in order to follow a boustrophedon path (i.e., sweeping back and forth across the cell).

The resulting path shape is illustrated in fig. 2.2. There is a consensus in the literature that, when using this method for area coverage, the main optimization goal is to minimize the number of turns [47–51]. Thus, in the illustrated example, the path in fig. 2.2a contains a smaller number of turns and therefore has a smaller cost that the path in fig. 2.2b.

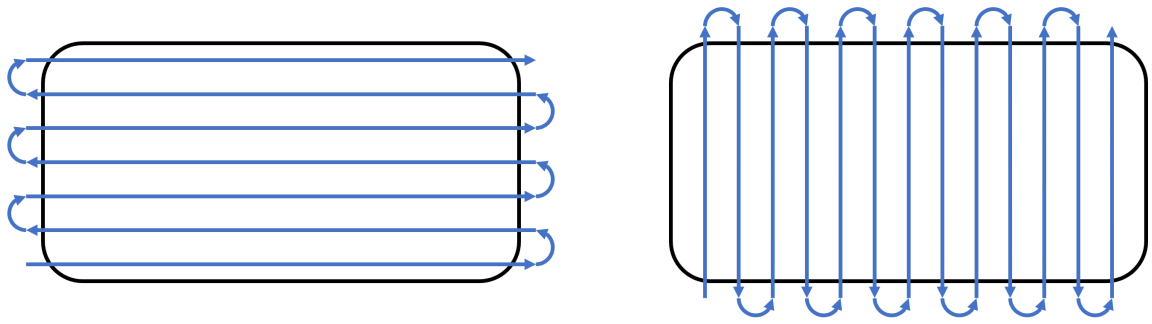

(a) Horizontal Sweep Direction

(b) Vertical Sweep Direction

Figure 2.2: Two examples of a boustrophedon path for complete area coverage. Adapted from [49].

The boustrophedon method was extended to optimize the sweep direction by Huang [49], who addressed the problem of how to achieve optimal sweep coverage of a polygonal planar region which may contain polygonal obstacles. The solution proposed by the author consisted of performing a decomposition of the coverage region into cells, determining the optimal sweep direction for each cell, and then using dynamic programming to combine cells into larger subregions with a single sweep direction, as illustrated in fig. 2.3. Thus, by allowing for different sweep directions in each subregion, instead of using a single sweep direction for the whole coverage region, the proposed method leads to better results, i.e., a path with fewer turns. However, the described algorithm is rather limited in the complexity of the shape of the coverage region that it can decompose in reasonable time, presenting an exponential time complexity. This is due to the initial decomposition, which may divide the coverage area into many subregions (fig. 2.3, left), as well as the combination process, which looks at all subregion recombinations.

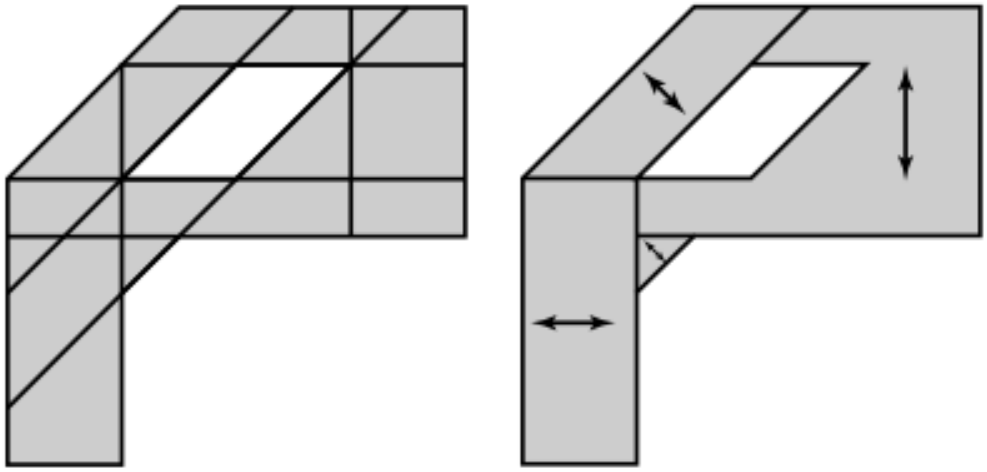

Figure 2.3: Decomposition of the coverage region into subregions with optimized sweep directions. The left figure illustrates the initial decomposition, and the right figure is the final solution. Taken from [49].

The boustrophedon method was further improved upon by Luo et al [52] to address the vicinity problem of obstacles, that is, ensuring complete area coverage with a particular emphasis on the regions closest to obstacles. A result of their proposed solution is illustrated in fig. 2.4; notice how the path generated by their proposed solution does not pass "dangerously" close to some obstacle corners, as well as providing a more uniform coverage around them.

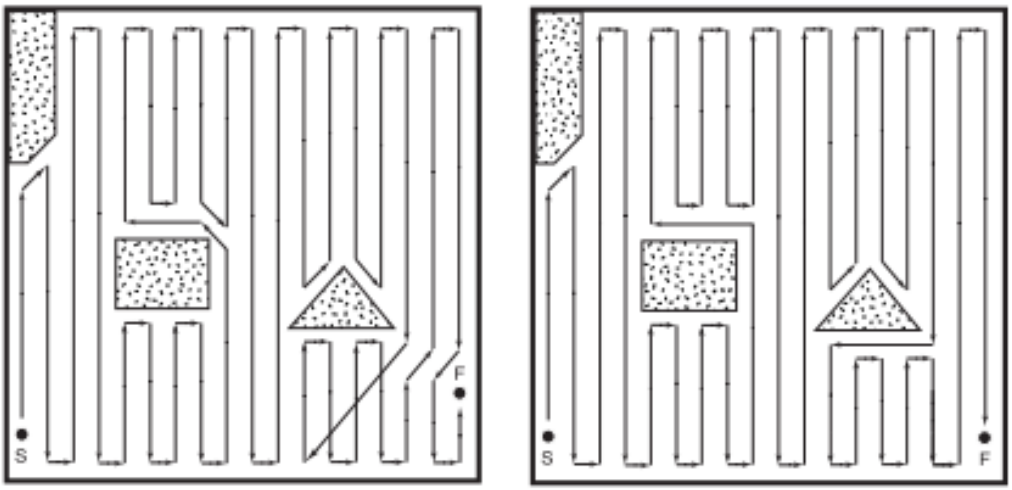

Figure 2.4: Complete coverage path planning solution which avoids obstacles while covering every point in their vicinity. The left figure illustrates the planned path before the vicinity problem is solved, and the right figure is the solution after the vicinity problem is solved. Taken from [52].

Maza and Ollero [50] addressed the problem of area coverage for multiple cooperative robots, dividing the workspace into convex cells, assigning each cell to a robot, and performing a boustrophedon path in each cell. Although their work was only adapted to simple scenarios represented by convex polygons without any obstacles, it presented interesting results relating to the cooperative nature of the study. In particular, their solution takes into consideration the different sensing capabilities of each robot, e.g., assigning a larger coverage area to a robot with a larger energy budget, as illustrated in fig. 2.5.

Their solution also has the flexibility to deal with situations where a robot is lost, triggering a reconfiguration of the areas and sweep directions assigned to the remaining operational robots. Furthermore, the work was focused on applying the proposed methods for use with UAVs, which is one of the proposed applications for the thesis work. However, the authors only demonstrated their results in simulations.

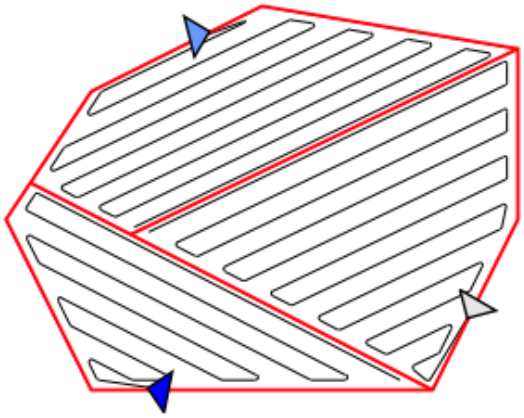

Figure 2.5: Solution to the complete coverage path planning problem applied to cooperative UAVs. Taken from [50].

2.3 Online Path Planning

Online path planning consists of approaches which actively plan a path as the agent moves through the environment. Although these approaches are inherently more complex, they also allow the agent to react to the environment, to avoid collisions and correct errors encountered during motion [53]. This section details some of the existing solutions, starting with adaptive path planners, and then specifying informative path planners, which can be considered a subclass of adaptive path planners.

Adaptive Path Planning

Zhang et al [54] proposed a recursive receding horizon path planning algorithm for UGVs. The goal of their work was to produce a path in a 2D area to reach a certain goal position. Their proposed solution represents the map by a grid, and divides the map into four areas: the implementation area, the observation area, the explored area, and the unknown area. For each observed (past or present) map cell a terrain value is associated indicating the difficulty to traverse that cell. The implementation area represents the positions adjacent to the robot's current position, and, after the robot moves to one of those positions, the map is updated and a new path plan is made. The observation area comprises all map cells which can be currently observed. The explored area is a superset of the observation area, and represents all the cells which have already been observed by the robot. Finally, the unknown area represents all the cells in the map which have not yet been observed. An example of the different map areas is illustrated in fig. 2.6. As it can be seen in fig. 2.6, the path is actually split into two parts: one part (s'') which is optimally planned within the observation area, and another part (s') in the unknown area which connects the end of the path to the goal position by a line segment. Thus, when the robot

moves it makes new observations, expands the explored area, and replans the path taking into account the new observations.

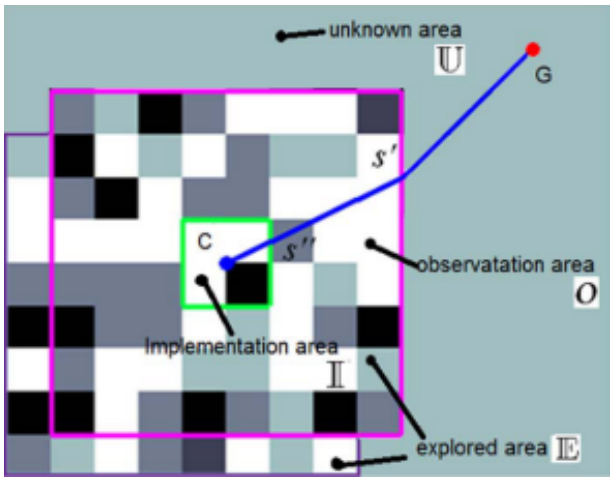

Figure 2.6: The different map areas used by the path planner proposed by Zhang et al. The robot's current position is represented by point C, and the goal position is represented by point G (red). The different shades of gray for the observed cells represent the difficulty to traverse the terrain, from easy (white) to difficult (dark gray) to obstacle (black). Taken from [54].

Furthermore, if the planner is unable to calculate a path towards the goal within the observation area (e.g., due to an obstacle or very rough terrain), then the planning is extended recursively to include the explored area. This behavior is illustrated in fig. 2.7. Thus, with this recursive behavior, the solution proposed by the authors is able to avoid situations where the robot gets stuck or deadlocked. However, if the explored area is large, this operation can be time consuming, since the algorithm searches all available nodes in the explored area to connect the robot to all available nodes on the edge.

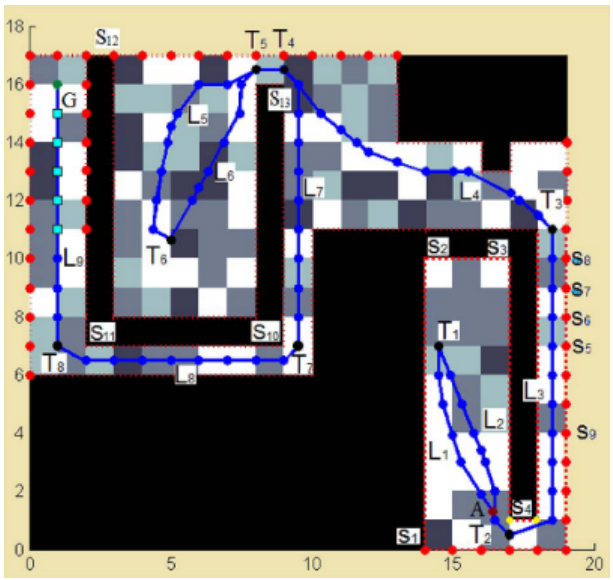

Figure 2.7: Result of the path planner presented by Zhang et al. The robot started at position A (lower right corner) and had to reach the goal (green) at position G (upper left corner). The traversable area is represented by the squares in shades of gray, where white squares are "easy" to traverse and dark gray squares are "difficult" to traverse. The traveled path is represented in blue. At points T_1 and T_6 the planner detects that forward progress is blocked, and so the recursive planner is activated, resulting in some backtracking. Taken from [54].

For the purposes of the thesis objectives the proposed solution is not directly applicable, since it is an algorithm which produces a lowest cost path with the objective of reaching a goal position. However, the authors propose some interesting ideas for Adaptive Path Planning (APP), which can be extended for applications in the subject of the thesis.

Informative Path Planning

Informative Path Planning (IPP) can generally be seen as a subclass of APP, since the general idea is to calculate a path adaptively based on a goal such as entropy reduction or information gain.

Paull et al [42] addressed the problem of coverage path planning for a fixed-wing UAV prioritizing sensor coverage. The authors simulated and tested an online path planner guided by information theory. More specifically, their solution performs a hexagonal cell decomposition of the workspace, then builds a Direct Acyclic Graph (DAG) out of the hexagonal cells (see fig. 2.8). Subsequently, each cell is assigned an average entropy, with which branch entropy for each branch in the DAG is computed. The authors then define information gain through the expectation of the mutual information of observations, and use this, together with the branch entropy, to find an optimal heading. Although their solution produces better coverage results, illustrated in fig. 2.9, the generated path has considerable overlap.

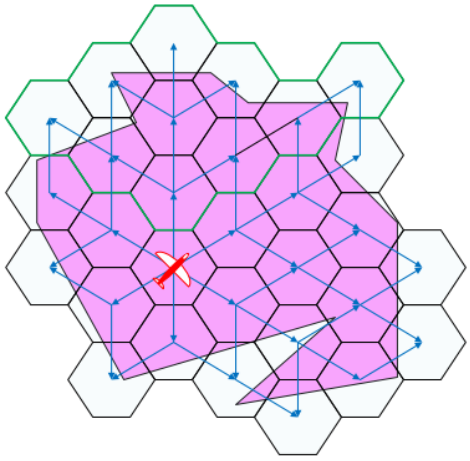

Figure 2.8: Hexagonal tiling over the coverage area (pink) used to generate a directed acyclic graph, represented by the blue arrows. Taken from [42].

Kim and Eustice [55] proposed a perception driven navigation solution for solving a visual SLAM area coverage problem. More explicitly, their work attempted to execute SLAM navigation using photographs as observations, and matching the photographs to build the image of a ship's hull. This work was motivated by the need to perform routine inspections of ship hulls using Autonomous Underwater Vehicles (AUVs). The problem is challenging because the quality of the SLAM navigation performance and the efficient area coverage represent competing objectives. As such, the proposed method consisted of balancing the choice of exploring new areas, thus improving the coverage result, with the choice of revisiting already explored areas for the purposes of loop closure, in order to reduce the pose uncertainty. To achieve that balance, a reward framework is used, such that exploring an area is given a certain utility,

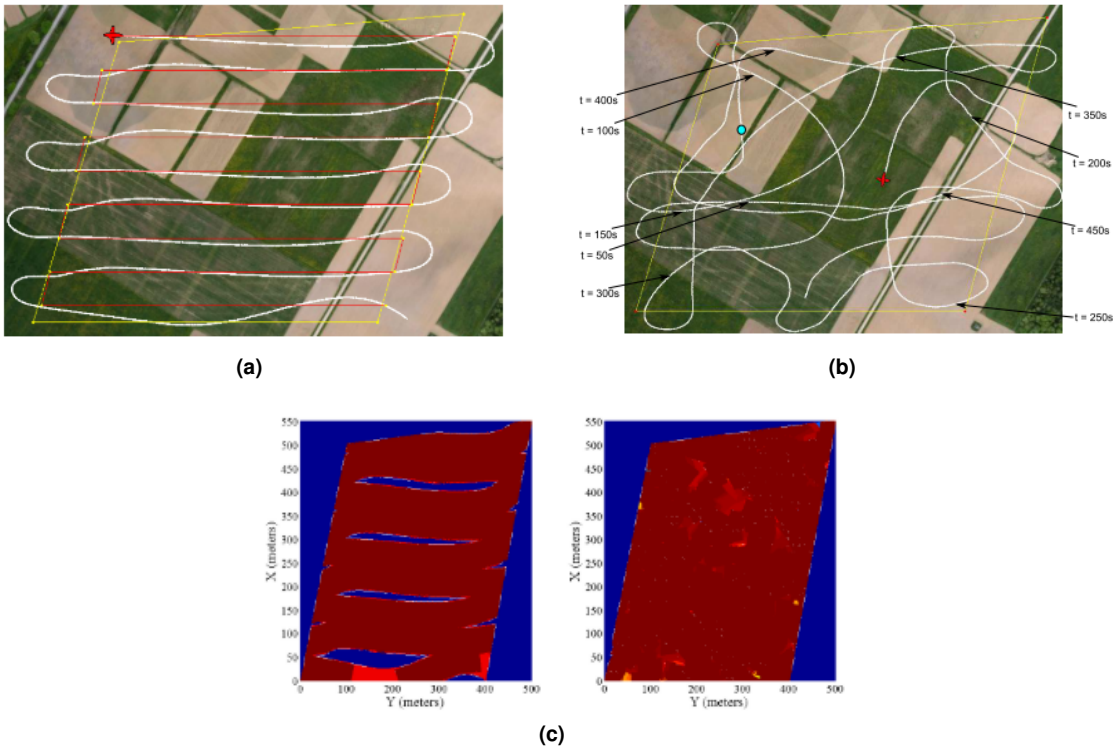

(a) (b)

(c)

Figure 2.9: Resulting paths from hardware in the loop simulations. Figure (a) represents a boustrophedon path, and figure (b) represents a sensor driven path generated online. Figure (c) compares the coverage of both paths. Taken from [42].

as well as a penalty for increasing the uncertainty in the pose; similarly, revisiting a position is also given a utility value.

Paull et al [41] proposed a very similar solution to [42]. This work does however demonstrate the application of similar methods to other situations, namely non-convex and non-homogeneous workspaces. Furthermore, they demonstrate that their algorithm produces structured solutions for regular spaces, as illustrated in fig. 2.10. Additionally, the authors considered the cost of changing heading in this algorithm, by including that cost in the cost function. This was not done in [42].

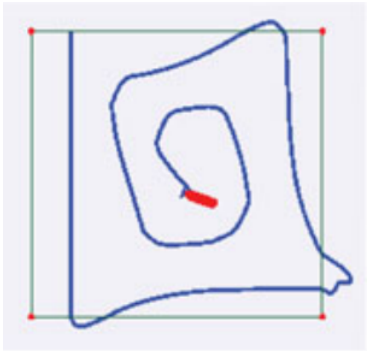

Figure 2.10: Spiral path generated in a square and homogeneous environment by the sensor-driven path planner proposed by Paull et al. Taken from [41].

The work done by Hitz et al [36] introduces some interesting ideas. Namely, they parameterize the path by a set of control points, such that the path is defined by a B-spline with those points. This allows

the path to be planned in continuous space while not adding too much computational burden. Another concept which was appealing was the division of the budget into two parcels: one for the global path and another for the local path. The global path is planned offline in such a manner that it covers the workspace uniformly. Then, as the robot is traveling and acquiring sensor data, a portion of the local budget is used to adapt the path locally. In this way, if the robot passes through an information rich area (which is generally not known a priori) it can use the local budget to explore that area in more detail. Nevertheless, the enforcement of a minimal required progression along the global path ensures that an adequate coverage target is attained. Finally, the proportion of the budget assigned to the global and local paths is controlled by a parameter λ, which allows configuring the planner to search for more uniform coverage solutions or more informative coverage solutions. Results are illustrated in fig. 2.11.

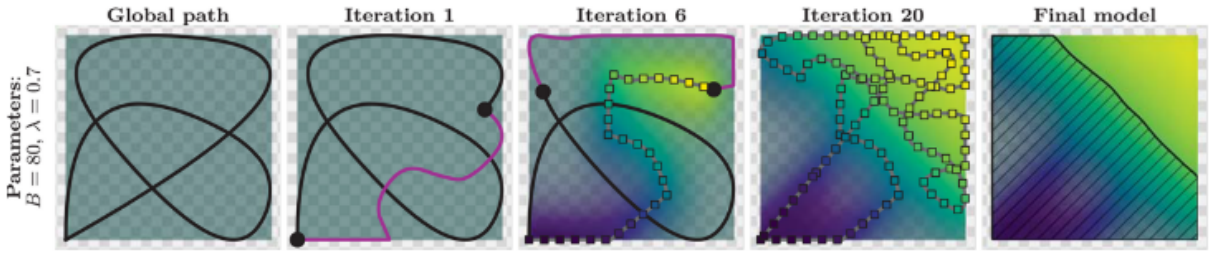

Figure 2.11: Simulation results of the adaptive continuous-space informative path planner proposed by Hitz et al. Illustrated are the global path (black), the locally planned path (pink), the already traveled path (gray) and measurements made (colored squares). More interesting areas are shaded in yellow, while less interesting areas are shaded in blue. Taken from [36].

2.4 Comparison of Approaches

Table 2.1 presents a summary of the more interesting methods reviewed in the literature. These methods represent good candidates to be combined, adapted, and improved upon in the thesis.

Table 2.1: Summary of the methods reviewed in the literature.

Method	Pros	Cons
Optimal line-sweep-based decompositions for coverage algorithms [49]	• Determines the optimal sweep direction for different regions • Optimal sweep direction means fewer turns in the resulting path	• Algorithm has exponential time complexity due to the process of division of the coverage area and subsequent recombination of subregions
Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms [50]	• Adapted for multiple cooperative robots • Takes the capabilities of each robot into consideration when calculating the decomposition • Has the flexibility to reconfigure in case a robot is lost • Time complexity is quadratic in the number of robots and linear in the number of polygon vertices	• Only applies to convex polygonal spaces • Assumes robots are on the boundary of the coverage area • Authors do not propose any method for determining each robot's capabilities, only that the capability is represented by a number between 0 and 1, and that the sum of all robots' capabilities is unity

A recursive receding horizon planning for unmanned vehicles [54]	• Path is planned in the observation area taking into account sensory data about the environment • If there is no path towards the goal within the observation area (e.g., due to obstacles), the planner recursively searches within the explored area (including the observation area), allowing for backtracking and finding alternative routes • Receding horizon approach reduces computational cost (smaller configuration space for the path) and gives flexibility • Solution was tested with a wheeled robot in rough terrain and was able to reach the intended goal	• Solution is just for planning a path to reach a goal position • Requires quantification of the difficulty to cross terrain (between 0 and 1, where 1 represents an impassable obstacle), in order to calculate path cost; no method is proposed to determine the values automatically, would probably require a considerable computation using computer vision • Following the previous point, the presented results had the terrain values inserted manually, with a seemingly arbitrary number of discrete values to assign (in one case 5, in another case 8)
Sensor-driven area coverage for an autonomous fixed-wing unmanned aerial vehicle [42]	• The solution is subdivided into a macroscopic component and a microscopic component • The macroscopic component, represented by a directed acyclic graph, seeks to ensure that the all areas of the map are visited • The microscopic component seeks to achieve complete coverage in each node of the graph • The solution is adaptive and data-driven, meaning it can achieve complete coverage without the need for conservative planning used in boustrophedon techniques • All planning is done online • The planner optimizes over the vehicle's heading, meaning the path is essentially planned in free space	• The resulting path may have considerable overlap
Perception-driven navigation: Active visual SLAM for robotic area coverage [55]	• Solution balances exploration, for which the reward is more coverage, with loop closure, which reduces the uncertainty in navigation • This balance is achieved by attributing a reward to each action • In essence, the proposed perception-driven navigation compares the two propagated uncertainties from revisiting and exploring, and chooses the smaller one whenever the exploration uncertainty exceeds a certain threshold • More emphasis can be given to robot navigation uncertainty or to area coverage performance in the reward calculation through a parameter α	• The proposed method is computationally heavy, requiring image matching, updating of possibly large information matrices (one each for odometry, camera, and perception-driven navigation) and state tracking • Results use real data, but are demonstrated in simulations and offline • The combination of the two above points leads to the conclusion that an online implementation would require considerable optimization in order to be viable

Sensor-driven online coverage planning for autonomous underwater vehicles [41]	• Similar to [42], the solution assumes two components, one aiming to visit all nodes in a directed acyclic graph representing the full coverage area, the other aiming to achieve complete coverage within a cell (node) • Authors consider not only the information gain of exploring more areas, but also the cost of making turns, resulting in a smoother control signal • The planner converges to a structured path for regular shaped workspaces (a spiral path is produced in the case of a square workspace) • The planner optimizes over the vehicle's heading, meaning the path is essentially planned in free space	• Although the solution gives good results, the authors tend to make comparisons with worst case boustrophedon paths
Information-driven sensor path planning by approximate cell decomposition [40]	• The proposed solution takes into consideration the expected entropy reduction of actions, the cost of different sensor configurations, and the cost of distance travelled • By adjusting the weight of the distance travelled component, the planner can explore more "secluded" areas, or execute a faster mission • The solution also takes into consideration the geometry of the robot, thus avoiding paths which the robot can't realistically travel	• Planner requires prior knowledge of the workspace, namely the positions and geometry of obstacles and targets
Adaptive continuous-space informative path planning for online environmental monitoring [36]	• Paths are planned in continuous space through a finite number of control points, from which a B-spline is calculated • Planner divides the budget into two portions, one assigned to the global path and the other assigned to the local path • By adjusting the proportion of the global and local path budgets, we can balance more area coverage with higher resolution of "interesting" areas • The path is allowed to exit the workspace, to account for robot dynamics, but this action has an associated cost in order to minimize its occurrence • The global path ensures that the robot explores most areas, while the local path, which is adapted at each iteration, makes use of sensor information to explore more informative regions	• The global path needs to cover the workspace uniformly, which requires an initial offline path planning solution • The results were all demonstrated using regular workspace shapes (rectangles and squares) • The authors refer that it is unclear to what extent the planner can handle obstacles or highly non-convex environments

2.5 Radiological Units

For the purpose of this dissertation, it is important to understand some radiological units of measurement [56–58]. The most fundamental unit for measuring radioactivity is the becquerel (abbreviated Bq), which corresponds to one nuclear disintegration per second. Most radiological sensors essentially count the number of particle interactions that are detected (although some have the capacity to measure the energy of the particles simultaneously). The detectors which will be simulated in Chapter 4 and which the author has access to are GMCs: these are relatively simple devices which utilize a Geiger-Müller tube to detect ionizing particles and count the number of particles detected in a given time interval. No information about the energy of the particles is measured by a GMC and, practically speaking, it can only detect beta particles, X-rays and gamma rays. Therefore, GMC readings are typically given in counts per second, which are equivalent to Bq, or counts per minute to better approximate the mean radioactivity in a given area.

However, the effects of ionizing radiation do not only depend on the number of particles absorbed, but also the type of particles (alpha, beta, gamma, ...) and their energy. When measuring absorbed dose as a result of radiation exposure, the unit which is used is gray (abbreviated Gy), which corresponds to one joule per kilogram of absorbed radiation. This measure applies to absorbed dose for any type of body, biological or non-biological. In situations where the radiation is absorbed specifically by biological tissue, such as in the case of nuclear industry workers or emergency personnel operating in a radiological disaster site (see section 1.1), the most commonly used measurement unit is the sievert (abbreviated Sv). In terms of dimensions, sievert and gray are equivalent (they both have dimensions of J/kg), but sievert are related to gray by an additional quality factor Q to account for the type of radiation absorbed: $Q = 1$ for beta and gamma radiation, $Q = 20$ for alpha radiation, and (typically) $Q = 10$ for neutron radiation. With this quality factor, the sievert is essentially a measure of biological damage, also described as a unit of equivalent absorbed dose through the relation $\mathrm{Sv} = Q \times \mathrm{Gy}$. Some examples of absorbed doses can be found in Appendix B.

To summarize: becquerel is a measurement unit for radioactivity which corresponds to one disintegration per second, gray is a measurement unit for absorbed dose which corresponds to one joule per kilogram of absorbed radiation, and sievert is a measurement unit for equivalent absorbed dose when referring to radiation absorbed by biological tissue, which multiplies gray by a quality factor to account for the biological effects of different types of radiation.

3

Proposed Solution

Contents

3.1 Problem Formulation 30
3.2 Assumptions 35
3.3 Selected Approach 37

This chapter formally describes the problem which must be solved, along with the assumptions which are used and the approaches which were selected as viable solutions. Section 3.1 defines the problem variables and the objective function which is to be optimized in order to find the optimal path. Section 3.2 describes the assumptions which are used for the scenarios and the radioactive sources. Finally, section 3.3 describes the approach which was selected to solve the given problem.

3.1 Problem Formulation

Given a workspace $\mathcal{W} \subset \mathbb{R}^n$ which is path-connected and may be non-convex, we would like to estimate an unknown scalar field $f : \mathcal{W} \mapsto \mathbb{R}$. This scalar field may represent a plethora of physical quantities, such as radiological intensity, concentration of certain chemicals, or a temperature distribution.

In order to estimate f, an agent must navigate through $\mathcal{W}$ while transporting a sensor which can measure f; the agent can be a human operator or a robot/vehicle which operates on the ground or in the air. Within $\mathcal{W}$, the agent makes a set of observations $\mathcal{O} = \{\mathbf{o}_1, \mathbf{o}_2, \ldots, \mathbf{o}_K\}$ utilizing the transported sensor. Each observation $\mathbf{o}_k$ is composed of the instant at which the observation was made t_k, the agent's position at that instant $\mathbf{q}_k$, and the measured field value m_k, as expressed by (3.1).

$$\mathbf{o}_k = [t_k,\ \mathbf{q}_k,\ m_k]^\mathsf{T} \tag{3.1}$$

Additionally, the sensor measurements are subject to uncertainty due to noise or stochastic behaviors, so that the measured field value is given by (3.2), where $f(\mathbf{q}_k)$ is the field value at position $\mathbf{q}_k$.

$$m_k = f(\mathbf{q}_k) + \varepsilon_k \tag{3.2}$$

In order to reach the observation positions the agent must follow a path. This path is divided into two parts: the first is $\mathcal{Q} : [0, B_{\text{spent}}] \mapsto \mathcal{W}$, representing the path which has already been traveled, consuming a budget B_{spent}, and the second is $\mathcal{P} : [0, B_{\text{local}}] \mapsto \mathcal{W}$, representing the path which has been planned with a budget B_{local}. Moreover, quadratic B-splines will be used to model $\mathcal{P}$. A B-spline is defined by a set of control vertices $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_M\}$, with $\mathbf{c}_m \in \mathbb{R}^n$ for $m = 1, 2, \ldots, M$. With this parameterization $\mathcal{P}$ can be planned in continuous space by manipulating the positions of the control vertices in $\mathcal{C}$. Thus, at a time t each observation $\mathbf{o}_k$ will follow $t_k \leq t$ and $\mathbf{q}_k \in \mathcal{Q}$, whereas $\mathcal{P}$ will determine the positions of future observations. Furthermore, the agent is subject to a total budget B which defines the maximum path length and therefore imposes dynamic constraints on $\mathcal{P}$ through the condition $B_{\text{local}} \leq B - B_{\text{spent}}$.

In this context, we define a belief about the field f, denoted $\hat{f}_k$, which is updated when new observations are made. The belief $\hat{f}_k$ is modeled by a Gaussian Process (GP) which is trained on the data contained in $\mathcal{O}$. The utilization of a GP enables modeling a wide range of field functions, regardless of the underlying physical process [36, 59]. New observations translate into a reduction in the uncertainty of our belief, which is quantified by the GP's covariance matrix. This reduction in uncertainty is directly associated with a gain of information. The concept of reducing the uncertainty of one random variable due to the knowledge of another random variable is known as mutual information [37] (see section 2.1).

Following this reasoning, we understand maximizing information gain as being equivalent to maximizing the mutual information of new observations based on the GP which models $\hat{f}_k$. Additionally, in the case of a radiological intensity field, we are interested in making observations close to the radioactive source, since these contain more information about the source [60] and, consequently, about the field that it generates. Therefore, the mechanism which determines $\mathcal{P}$—the path planner—should take into account not only the expected information gain, but also the expected measurement values, so as to increase the number of observations in high intensity regions.

Our objective is to determine the path $\mathcal{P}$, based on the observations $\mathcal{O}$ and subject to the budget B_{local}, which maximizes the information that is gained about the field f which is present in the workspace $\mathcal{W}$. In the literature this is known as an Informative Path Planning (IPP) problem (see section 2.3). Given that the information gain is inherently dependent on the current belief $\hat{f}_k$, which is dependent on the current set of observations, the path planner must recalculate $\mathcal{P}$ when new observations are made. Therefore, we define the set of N future observation positions $\mathcal{X} = \{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N\}$ with $\mathbf{p}_n = \mathcal{P}|_{nd_{\text{obs}}}$ for $n = 1, 2, \ldots, N$ and $N = \lfloor B_{\text{local}}/d_{\text{obs}} \rfloor$, where d_{obs} specifies an inter-observation distance which determines the equally spaced observation positions along $\mathcal{P}$; the notation $\mathcal{P}|_l$ is used to indicate the position at a length l along $\mathcal{P}$. Note that $\mathcal{X}$ is the variable to be optimized in the IPP problem but that it is conditioned by the shape of $\mathcal{P}$. A simple toy scenario is represented in fig. 3.1.

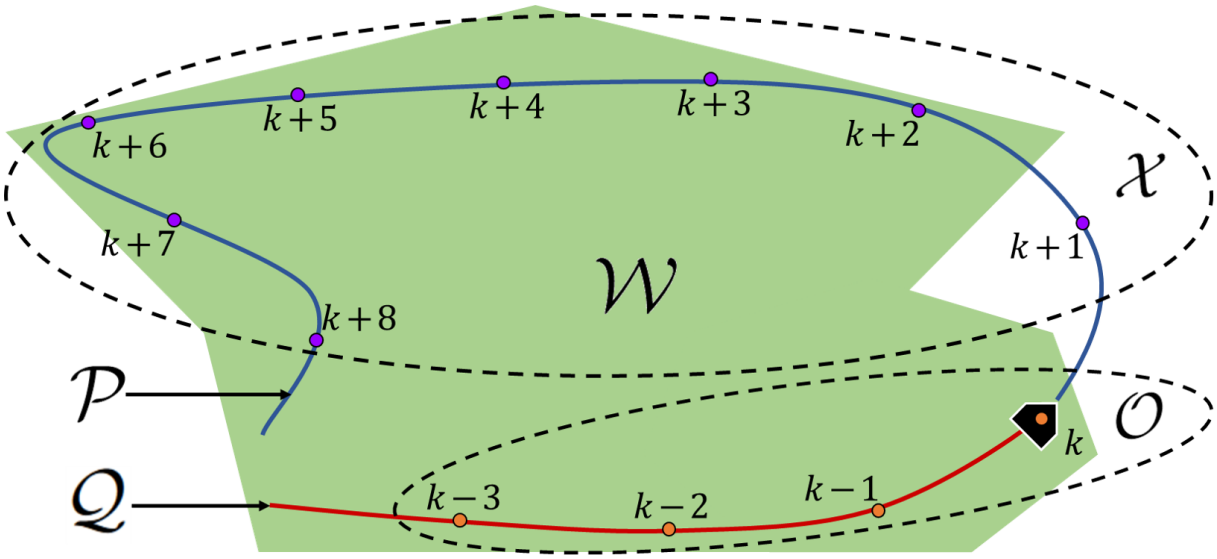

Figure 3.1: Schematic representation of the problem variables. The workspace $\mathcal{W}$ is shown in green, the traveled path $\mathcal{Q}$ is represented by the red curve, and the path $\mathcal{P}$ is represented by the blue curve. The agent (black pentagon) has made k observations, represented as orange circles which constitute the set $\mathcal{O}$. Based on those observations, the set of future observation positions $\mathcal{X}$ is calculated, represented by the purple circles.

In order to achieve the objective described in the previous paragraph with an iterative online path planner, it is convenient to define an additional parameter $\mathcal{G}$ which represents a global path. A fraction $\beta \in (0, 1)$ of the budget is used to determine the length and shape of $\mathcal{G}$, such that $\mathcal{G} : [0, \beta B] \mapsto \mathcal{W}$. The role of $\mathcal{G}$ is to guide $\mathcal{P}$ in three ways: first, it defines a goal position for each iteration (the start position is equal to the agent's position at the start of the iteration), second, it guarantees that a minimum amount of coverage is achieved, and third, it drives the path planner out of local minima (i.e., regions where the field

magnitude is large). In order to accomplish the described guiding effect, $\mathcal{G}$ must be within a maximum distance d_{horizon} of all points in $\mathcal{W}$. This is the horizon distance within which the informative path planner optimizes $\mathcal{P}$ at each iteration; therefore, the maximum distance requirement guarantees that $\mathcal{P}$ can reach every point in $\mathcal{W}$, if necessary. Additionally, a highly desirable property for $\mathcal{G}$ is that it has some degree of periodicity or, alternatively, that it returns to previously explored regions. This property allows the informative path planner to execute an initial exploration of $\mathcal{W}$ and then, upon returning to the explored regions, the path planner can determine which areas are most interesting to explore in greater detail (e.g., around a detected radioactive source).

Thus, the first task for the informative path planner is to determine the shape of $\mathcal{G}$. Then, at each iteration, the informative path planner resolves the optimal path $\mathcal{P}$ between the agent's current position and a preset position along $\mathcal{G}$ using a portion $B_{\text{local}} = (2-\beta)d_{\text{horizon}}$ of the budget (B_{local} is designed so that the agent's final position coincides with the endpoint of $\mathcal{G}$ when the budget is completely spent). Subsequently, the agent travels $\mathcal{P}$ up to a distance $d_{\text{travel}} \leq d_{\text{horizon}}$, making new observations which are added to $\mathcal{O}$. Using a distance $d_{\text{travel}} < d_{\text{horizon}}$ allows the agent to stay in interesting regions for more iterations by not requiring that it end each iteration strictly on $\mathcal{G}$. The preset goal positions are equally spaced along $\mathcal{G}$ by a distance $d_{\text{global}} = \beta d_{\text{travel}}$.

The optimization problem which is solved by the informative path planner described thus far is defined in (3.3)–(3.15). An explanation for each equation is given in the following paragraphs.

$$\text{maximize} \quad u(\mathcal{X},\mathcal{O})g_{\text{budget}}(\mathcal{P},B_{\text{local}})g_{\text{boundary}}(\mathcal{P},\mathcal{W})g_{\text{global}}(\mathcal{P},\mathcal{G}) \tag{3.3}$$

$$\text{subject to} \quad \mathcal{X} = \{\mathcal{P}|_{nd_{\text{obs}}} : n = 1,2,\ldots,N\} \tag{3.4}$$

$$N = \lfloor B_{\text{local}}/d_{\text{obs}} \rfloor \tag{3.5}$$

$$\mathcal{P} = \text{BSpline}(\mathcal{C}) \tag{3.6}$$

$$\mathcal{C} = \{\mathbf{c}_m \in \mathbb{R}^n : m = 1,2,\ldots,M\} \tag{3.7}$$

$$u(\mathcal{X},\mathcal{O}) = \frac{1}{2}\log\left(\det\left(\mathbf{R}_{\text{interest}}\left(\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}\right) + \mathbf{K}_{\mathcal{O}}^{\mathcal{X}}\right)\right) \tag{3.8}$$

$$\mathbf{R}_{\text{interest}}(\boldsymbol{\mu}) = \begin{bmatrix} r(\mu_1) & 0 & \cdots & 0 \\ 0 & r(\mu_2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & r(\mu_N) \end{bmatrix} \tag{3.9}$$

$$r(\mu) = \max\left[1, \mu\left(1 - \exp\left(-\frac{1}{2}\frac{\mu-\mu_{th}}{\mu_{th}}\right)\right)\right] \tag{3.10}$$

$$g_{\text{budget}}(\mathcal{P},B_{\text{local}}) = \Gamma(B_{\text{local}} - |\mathcal{P}|) \tag{3.11}$$

$$g_{\text{boundary}}(\mathcal{P},\mathcal{W}) = \Gamma\left(\int_0^{|\mathcal{P}|} \bar{\phi}(\mathcal{P}|_l,\mathcal{W})dl\right) \tag{3.12}$$

$$g_{\text{global}}(\mathcal{P},\mathcal{G}) = \Gamma\left(d_{\text{horizon}} - \left\|\mathcal{P}|_{d_{\text{travel}}} - \mathcal{G}|_{(i+2)d_{\text{global}}}\right\|\right) \tag{3.13}$$

$$\Gamma(x) = \begin{cases} \exp\left(-\frac{1}{2}\frac{x^2}{c^2}\right) & \text{if } x < 0 \\ 1 & \text{if } x \geq 0 \end{cases} \tag{3.14}$$

$$\bar{\phi}(\mathbf{r},\mathcal{W}) = \begin{cases} -\inf\limits_{\mathbf{r}'\in\mathcal{W}} \|\mathbf{r}-\mathbf{r}'\| & \text{if } \mathbf{r} \notin \mathcal{W} \\ 0 & \text{otherwise} \end{cases} \tag{3.15}$$

The objective function is defined in (3.3) as the product of a utility function u with three multiplicative penalty factors, g_{budget}, g_{boundary}, and g_{global}. The utility function is formally defined in (3.8), and it quantifies the informativeness of a set of potential observation positions $\mathcal{X}$ based on the previous set of observations $\mathcal{O}$ using a modified form of mutual information [36, 38]. The set $\mathcal{X}$ is defined by (3.4) as the N positions on the planned path where new observations will be made, N being specified by (3.5). The path is expressed by (3.6), where BSpline is a function which builds a B-spline curve from the set of control vertices $\mathcal{C}$ defined in (3.7). The mean vector $\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}$ and the covariance matrix $\mathbf{K}_{\mathcal{O}}^{\mathcal{X}}$ in (3.8) are extracted from a GP which is trained on the observation data contained in $\mathcal{O}$. The GP is essentially defined by its covariance function, also known as the kernel. The kernel which was used in this work was the squared exponential kernel [59], which has the form (3.16), where the lengthscale l controls how "curvy" the function can be, and the amplitude σ^2 is a scale factor.

$$k(r) = \sigma^2 \exp\left(-\frac{r^2}{2l^2}\right) \tag{3.16}$$

Additionally, we introduce the term $\mathbf{R}_{\text{interest}}\left(\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}\right)$, which is defined by (3.9) and (3.10). As indicated in (3.9), $\mathbf{R}_{\text{interest}}$ is a diagonal matrix whose diagonal entries depend on the mean vector $\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}$; this dependence is encoded by (3.10). Some sample graphs of the function r with different values of the threshold parameter μ_{th} are illustrated in fig. 3.2. The graphs show that this function has a constant value of 1 for $\mu < \mu_{th}$, and then quickly approaches a linear relation as μ increases. Therefore, in the case where all the elements of $\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}$ are less than μ_{th}, $\mathbf{R}_{\text{interest}}$ is simply the identity matrix and (3.8) becomes the exact form of mutual information; in this situation, maximizing (3.8) produces results which strictly reduce the uncertainty of $\hat{f}_k$ without regard for the field values (i.e., the path planner becomes a complete coverage path planner). However, when there are elements of $\boldsymbol{\mu}_{\mathcal{O}}^{\mathcal{X}}$ greater than μ_{th}, the diagonal entries of $\mathbf{R}_{\text{interest}}$ become greater than 1 and thus increase the value of the utility function (because the covariance matrix is positive semi-definite); conversely, solutions generated in this situation result from a balance between exploring regions where the magnitude of $\hat{f}_k$ is large and exploring unobserved regions.

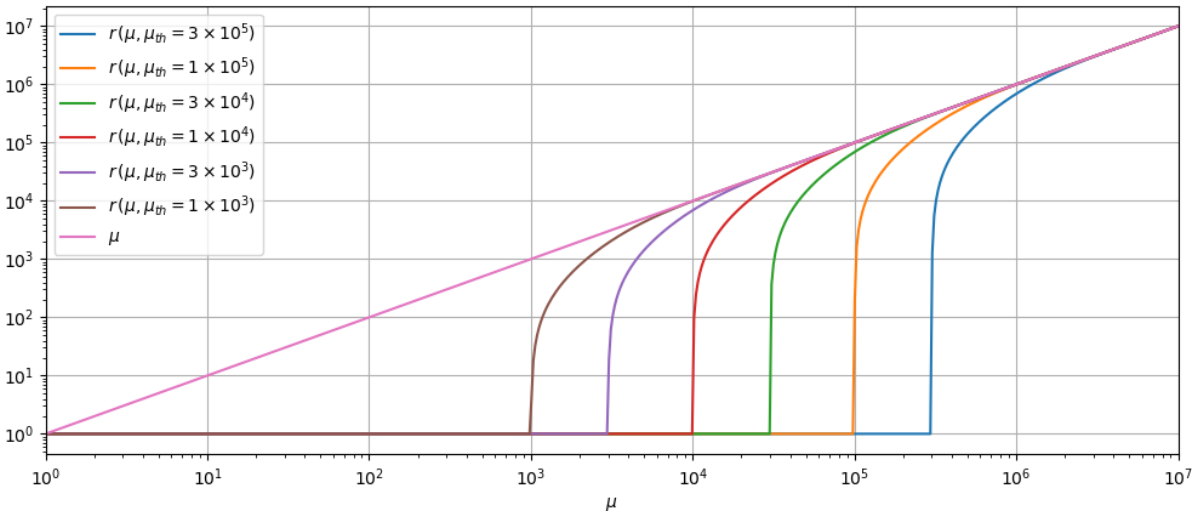

Figure 3.2: Graphs of the reward function r used in the utility function. Each graph corresponds to a different value of the threshold parameter μ_{th}, with the exception of the pink curve which corresponds to the linear function $r = \mu$, represented here for comparison purposes.

For the penalty factors we define a smooth step function Γ in (3.14); depending on the value of the parameter c, the slope of the step function for negative input values is larger or smaller. Some sample graphs of Γ are illustrated in fig. 3.3 for several values of c; it can be seen that smaller values of c lead to a much steeper and sharper step—this equates to having a more severe penalty factor, since smaller deviations outside the desired behavior quickly reduce the utility of a candidate path solution to ≈ 0.

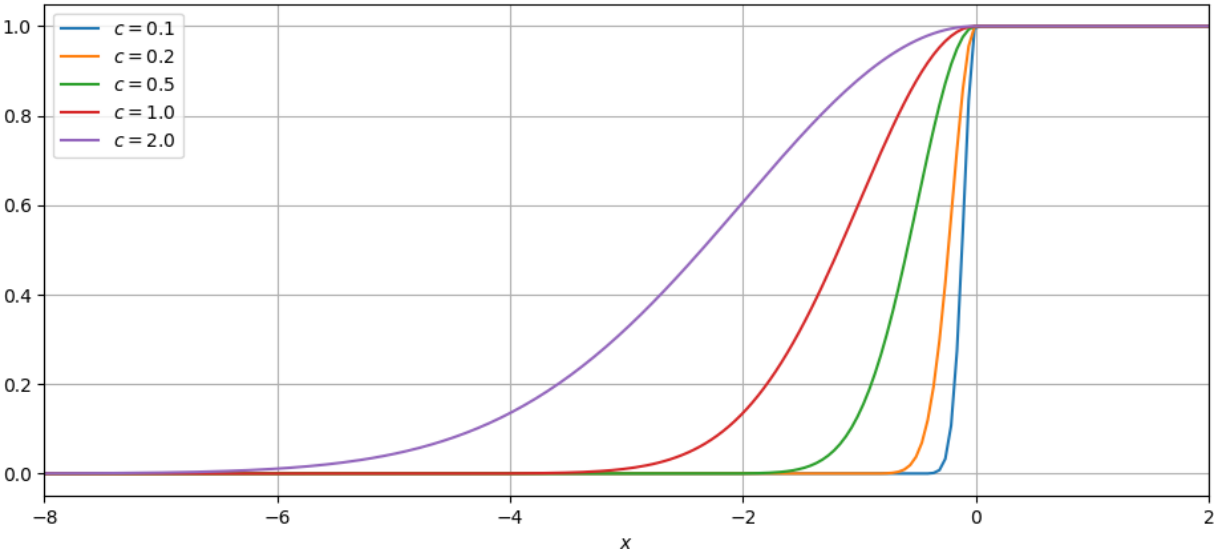

Figure 3.3: Graphs of the smooth step function Γ used in the penalty factors. Each graph corresponds to a different value of the parameter c, where it can be observed that smaller values of c lead to more restrictive penalties (since the penalty is multiplicative).

In the budget penalty factor (3.11) the length of the planned path, denoted $|\mathcal{P}|$, is calculated and subtracted from the budget B_{local}. If $|\mathcal{P}| > B_{\text{local}}$, it follows that $g_{\text{budget}} < 1$ and the utility is reduced.

For the boundary penalty factor (3.12) we define an additional auxiliary function $\bar{\phi}$ in (3.15); for a point $\mathbf{r} \in \mathbb{R}^n$, $\bar{\phi}$ gives the negative of the shortest distance from $\mathbf{r}$ to $\mathcal{W}$, and if $\mathbf{r} \in \mathcal{W}$, this distance is 0. The diagram in fig. 3.4 illustrates how $\bar{\phi}$ is calculated for a given point $\mathbf{r} \in \mathcal{P}$. The boundary penalty factor is calculated by integrating $\bar{\phi}$ for all points along $\mathcal{P}$, where the notation $\mathcal{P}|_l$ indicates the point at a length l along the path, measured from the start of $\mathcal{P}$. In this manner, a path which is completely contained in $\mathcal{W}$ will result in an integral with a value of 0, which leads to $g_{\text{boundary}} = 1$, so there is no penalty; conversely, a path which has a portion outside of $\mathcal{W}$ will result in an integral with a negative value, which leads to $g_{\text{boundary}} < 1$, so the utility is reduced.

The idea behind g_{global} is to generate solutions which guarantee that the optimization problem remains feasible in the next iteration. The term $\mathcal{P}|_{d_{\text{travel}}}$ refers to the end point of the current iteration, i.e., the start point of the next iteration. Similarly, the term $\mathcal{G}|_{(i+2)d_{\text{global}}}$ refers to the goal position along the global path for the next iteration, where i is the iteration index. Hence, the norm in (3.13) equates to the distance between the agent's position at the start of the next iteration and the goal position of the next iteration. If this distance becomes greater than the horizon distance d_{horizon}, then $g_{\text{global}} < 1$ and the solution is penalized. If this penalty were not enforced, then the problem would become infeasible in iterations where the agent's position is further than d_{horizon} from the goal position.

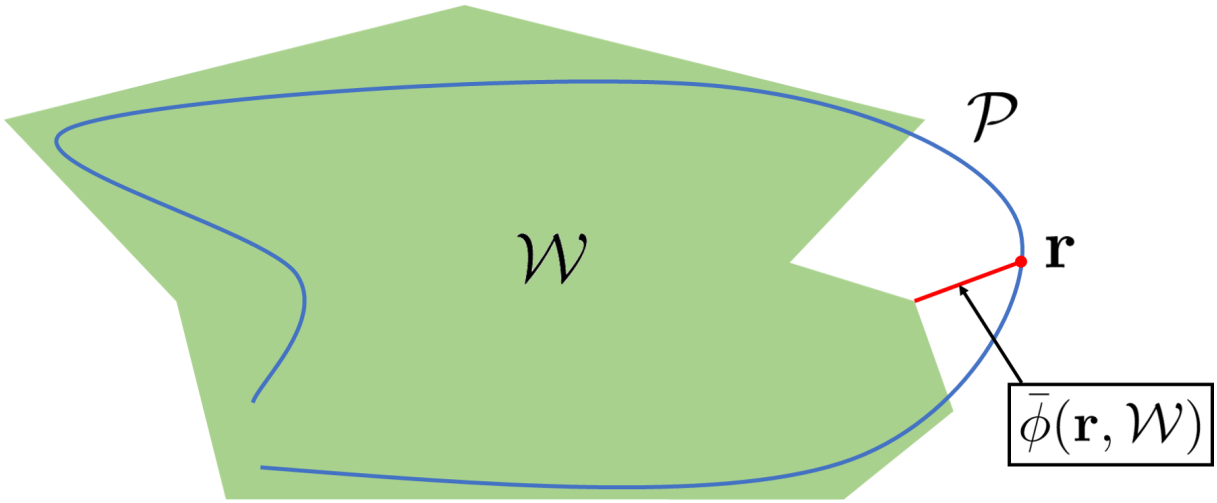

Figure 3.4: Illustration of the signed distance function $\bar{\phi}$: for any point $\mathbf{r}$ outside of the workspace $\mathcal{W}$, $\bar{\phi}$ yields the negative smallest distance from $\mathbf{r}$ to $\mathcal{W}$, represented in red.

In summary, given a mobile agent subject to the budget B, transporting a sensor and operating in the workspace $\mathcal{W}$, we want to determine the optimal path $\mathcal{P}$ such that the information gained about the field $f : \mathcal{W} \mapsto \mathbb{R}$ is maximized; $\mathcal{P}$ is determined by solving the optimization problem defined in (3.3)–(3.15).

3.2 Assumptions

For the implementation of the path planner and the design of experiments a number of assumptions will be made about the behavior of the radiological sources and the environments which will be studied.

First, the radiological sources are modeled as spherical isotropic radiation emitters, with radius r_{source}. As such, the mean radiological intensity (measured in units of becquerel per square meter, Bq/m^2), when observed in a vacuum, decreases proportional to the inverse of the square of the distance to the source, known as the inverse-square law [61]. However, for distances less than or equal to the source's radius, the intensity is assumed to be constant. This behavior is expressed by (3.20).

Second, each source is characterized by an activity $A_n \in \mathbb{R}^+$ (measured in units of becquerel, Bq) and a position $\mathbf{r}_n \in \mathbb{R}^2$. The activity represents the mean number of particles emitted per second by a given source, and is assumed to be constant (i.e., the value of source's half-life is much larger than duration of the survey).

Third, following the previous assumption coupled with the low efficiency of detection in most radiological sensors, the observed number of particles follows Poisson statistics [62]. That is, when an observation is made at a position $\mathbf{r}$, it is taken as a sample from a Poisson distribution with mean $\lambda = R(\mathbf{r})$. Here, $R(\mathbf{r})$ is the mean radioactivity at position $\mathbf{r}$, expressed by (3.21); like the activity, it will be expressed in units of becquerel. Moreover, in the case of a real-world radiological intensity field, a moderate time interval is required in order to make a satisfactory observation, because most radiation detectors rely on counting the number of high energy particles over time. However, for the observations which are simulated in this work, we will consider that the agent makes an instantaneous observation which is a Poisson sample according to (3.17), where H is known as a sensor model [63]. Note that, in

(3.1), the term ε_k is dominated almost entirely by the variance of the Poisson distribution. As such, we will make the approximation $\varepsilon_k \approx \sqrt{m_k}$, where m_k is the measured radioactivity.

$$H(f(\mathbf{q}_k)) \sim \text{Pois}\left(R(\mathbf{q}_k)\right) \tag{3.17}$$

Additionally, for simplicity, the detector which makes the observations is assumed to be spherical with radius r_{detector}, so that detector orientation effects can be ignored and the detector's solid angle can be efficiently calculated. Consequently, in order to convert radioactivity observations into intensity values, we divide the radioactivity values by the cross-sectional area of the detector, expressed by (3.18).

$$I = \frac{R}{\pi r_{\text{detector}}^2} \tag{3.18}$$

Fourth, all physical mediums (air, walls, water, etc.) are assumed to be isotropic and homogeneous; that is not to say that the environment as a whole is homogeneous, since the existence of obstacles may be considered, but each medium on its own is homogeneous. As such, each medium is characterized by an absorption coefficient α, which translates its propensity to absorb ionizing radiation (in the case of vacuum, the absorption coefficient is $\alpha = 0$ m^{-1}). The effect of the absorption behavior on radiological intensity follows exponential attenuation [64], as expressed by (3.19).

$$I(\mathbf{r}) = I_0 \exp\left(-\alpha \left\|\mathbf{r}_n - \mathbf{r}\right\|\right) \tag{3.19}$$

Fifth, different radiological sources are assumed to be independent of one another, such that the intensity observed at a certain position is given by the sum of the intensities of each source. Likewise, the radioactivity observed at a certain position is given by the sum of the radioactivities of each source.

Finally, following all the previous assumptions, the mean intensity and radioactivity observed at a position $\mathbf{r}$ due to a set of N radiological sources are given by (3.20) and (3.21), respectively, where: α_j is the absorption coefficient of medium j, L_j is the length of the line segment $\mathbf{r}_n\mathbf{r}$ which traverses medium j (see fig. 3.5), $\mathcal{J}_{LOS}$ is the set of mediums which are in the line-of-sight between $\mathbf{r}_n$ and $\mathbf{r}$; this includes the air or any other medium, not only solid obstacles such as walls, θ is the angular radius of the detector when viewed from the position of the source, given by (3.22).

$$I(\mathbf{r}) = \sum_{n=1}^{N} \begin{cases} \dfrac{A_n}{4\pi r_{\text{source}}^2} & \text{if} \quad \left\|\mathbf{r}_n - \mathbf{r}\right\| \leq r_{\text{source}} \\ \dfrac{A_n}{4\pi \left\|\mathbf{r}_n - \mathbf{r}\right\|^2} \displaystyle\prod_{j \in \mathcal{J}_{LOS}} \exp\left(-\alpha_j L_j\right) & \text{if} \quad \left\|\mathbf{r}_n - \mathbf{r}\right\| > r_{\text{source}} \end{cases} \tag{3.20}$$

$$R(\mathbf{r}) = \sum_{n=1}^{N} \begin{cases} \dfrac{1}{2} A_n & \text{if} \quad \left\|\mathbf{r}_n - \mathbf{r}\right\| \leq r_{\text{detector}} \\ \dfrac{1}{2} A_n \left(1 - \cos(\theta)\right) \displaystyle\prod_{j \in \mathcal{J}_{LOS}} \exp\left(-\alpha_j L_j\right) & \text{if} \quad \left\|\mathbf{r}_n - \mathbf{r}\right\| > r_{\text{detector}} \end{cases} \tag{3.21}$$

$$\theta = \arcsin\left(\frac{r_{\text{detector}}}{\left\|\mathbf{r}_n - \mathbf{r}\right\|}\right) \tag{3.22}$$

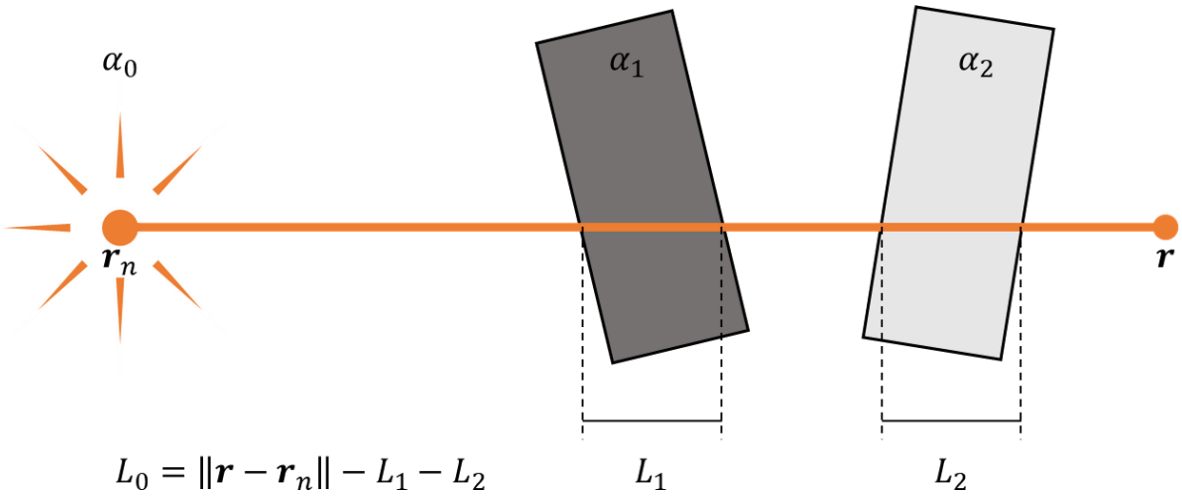

Figure 3.5: Illustration of the terms used in the radiation absorption model. There are three physical mediums: air, obstacle 1 and obstacle 2, with absorption coefficients α_0, α_1 and α_2, respectively. The intensity observed at $\mathbf{r}$ due to the source located at $\mathbf{r}_n$ is influenced by the lengths traversed through each medium, L_0, L_1 and L_2, as well as the absorption coefficients.

One additional assumption will be made about the agent: we will assume that the agent moves with constant velocity, so that the observations are equally spaced along the length of $\mathcal{P}$.

3.3 Selected Approach

In order to solve the optimization problem defined by (3.3)–(3.15), it is necessary to select a solver which is capable of obtaining solutions for that problem. For this purpose, the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) routine will be used [65].

CMA-ES is a generic optimization routine which belongs to the class of evolutionary algorithms. It is a stochastic, gradient-free method which has been successfully applied to high-dimensional and non-linear optimization problems. The steps of the CMA-ES routine are described in algorithm 3.1. The function `CalculateUpdatedParameters` on line 9 involves a considerable amount of complexity and is not within the scope of this work; for a more detailed discussion, the reader is referred to [65]. The simple example in fig. 3.7 graphically illustrates the progression of the CMA-ES routine.

A more interesting objective function for testing CMA-ES is known as Himmelblau's function, which is given by (3.23) and illustrated in fig. 3.6. Himmelblau's function has 4 local minima all equal to 0, whose minimizers are $(x_1, x_2) \in \{(3, 2), (-2.805118, 3.131312), (-3.779310, -3.283186), (3.584428, -1.848126)\}$. Two examples are illustrated in figs. 3.8 and 3.9. It can be seen that, depending on the initial parameters, the solutions converge to different minimizers. Additionally, the example illustrated in fig. 3.10 has the same initial parameters that were used to generate fig. 3.9; however, due to the stochastic nature of CMA-ES, the solutions converged to different minimizers.

$$f(x_1, x_2) = \left(x_1^2 + x_2 - 11\right)^2 + \left(x_1 + x_2^2 - 7\right)^2 \tag{3.23}$$

Algorithm 3.1: Covariance Matrix Adaptation Evolution Strategy optimization routine.

Input: Objective function $f_{\text{objective}}$, initial guess $\mathbf{x}$, population size λ, initial step size σ
Output: Optimal solution $\mathbf{x}^*$

begin
  // Initialize the mean vector to the initial guess.
  $\mathbf{m} \leftarrow \mathbf{x}$;
  // Initialize the covariance matrix to the identity.
  $\mathbf{C} \leftarrow \mathbf{I}$;
  while *has not met exit condition* **do**
    for k **in** $\{1, \ldots, \lambda\}$ **do**
      // Sample a candidate solution.
      $\mathbf{x}_k \leftarrow \mathcal{N}(\mathbf{m}, \sigma^2\mathbf{C})$;
      // Evaluate the fitness of the solution.
      $y_k \leftarrow f_{\text{objective}}(\mathbf{x}_k)$;
    end
    // Update the mean, covariance, and step size parameters.
    $\mathbf{m}, \mathbf{C}, \sigma \leftarrow$ `CalculateUpdatedParameters`$(\mathbf{x}_{1:\lambda}, y_{1:\lambda}, \mathbf{m}, \mathbf{C}, \sigma)$;
  end
  // Return the optimal solution.
  $\mathbf{x}^* \leftarrow \mathbf{m}$;
end

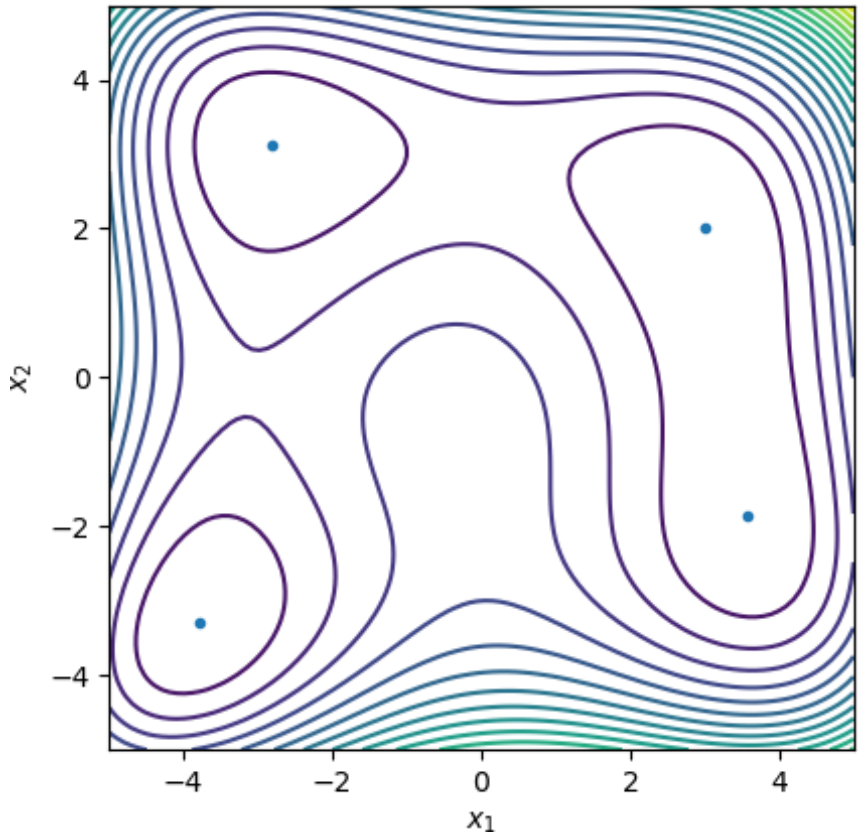

Figure 3.6: Level curves of Himmelblau's function, with the minimizers represented as blue dots.

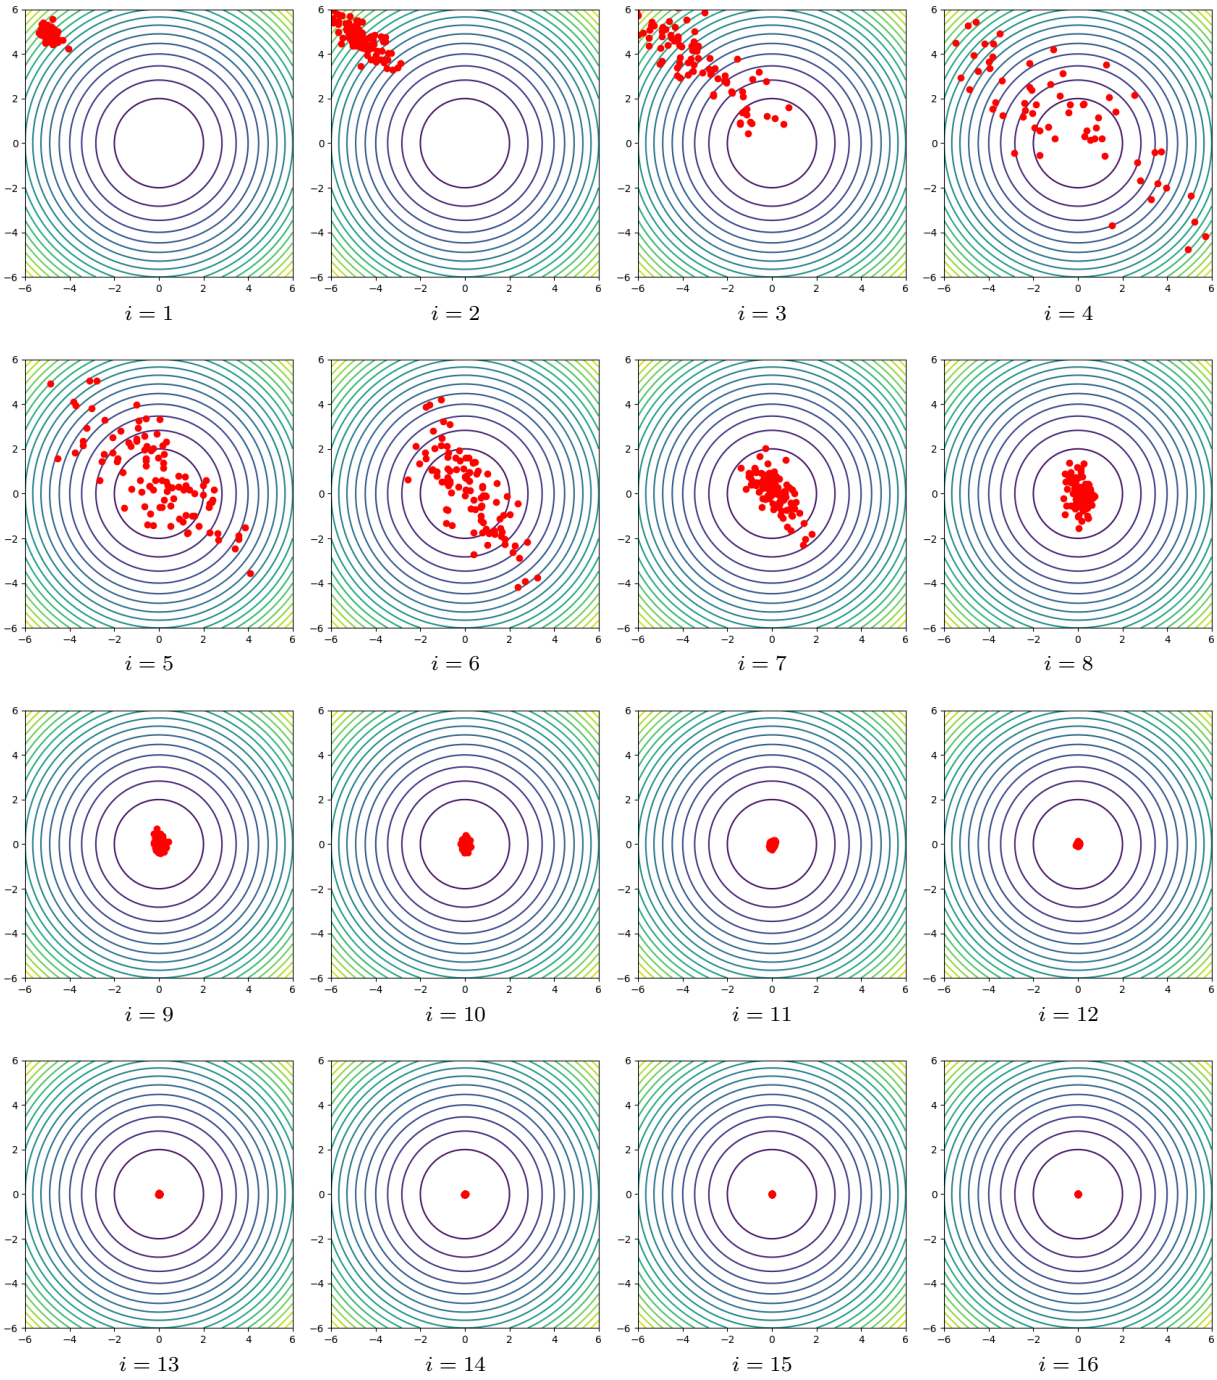

Figure 3.7: Illustration of the evolution of CMA-ES when finding the minimum of the function $f(x_1, x_2) = x_1^2 + x_2^2$. The level curves of the objective function are represented as colored circles, and the global minimum is located at $(0, 0)$. The red dots represent the CMA-ES population (candidate solutions), which has a size of $\lambda = 100$. The initial parameters were $\mathbf{x} = (-5, 5)$ and $\sigma = 0.1$.

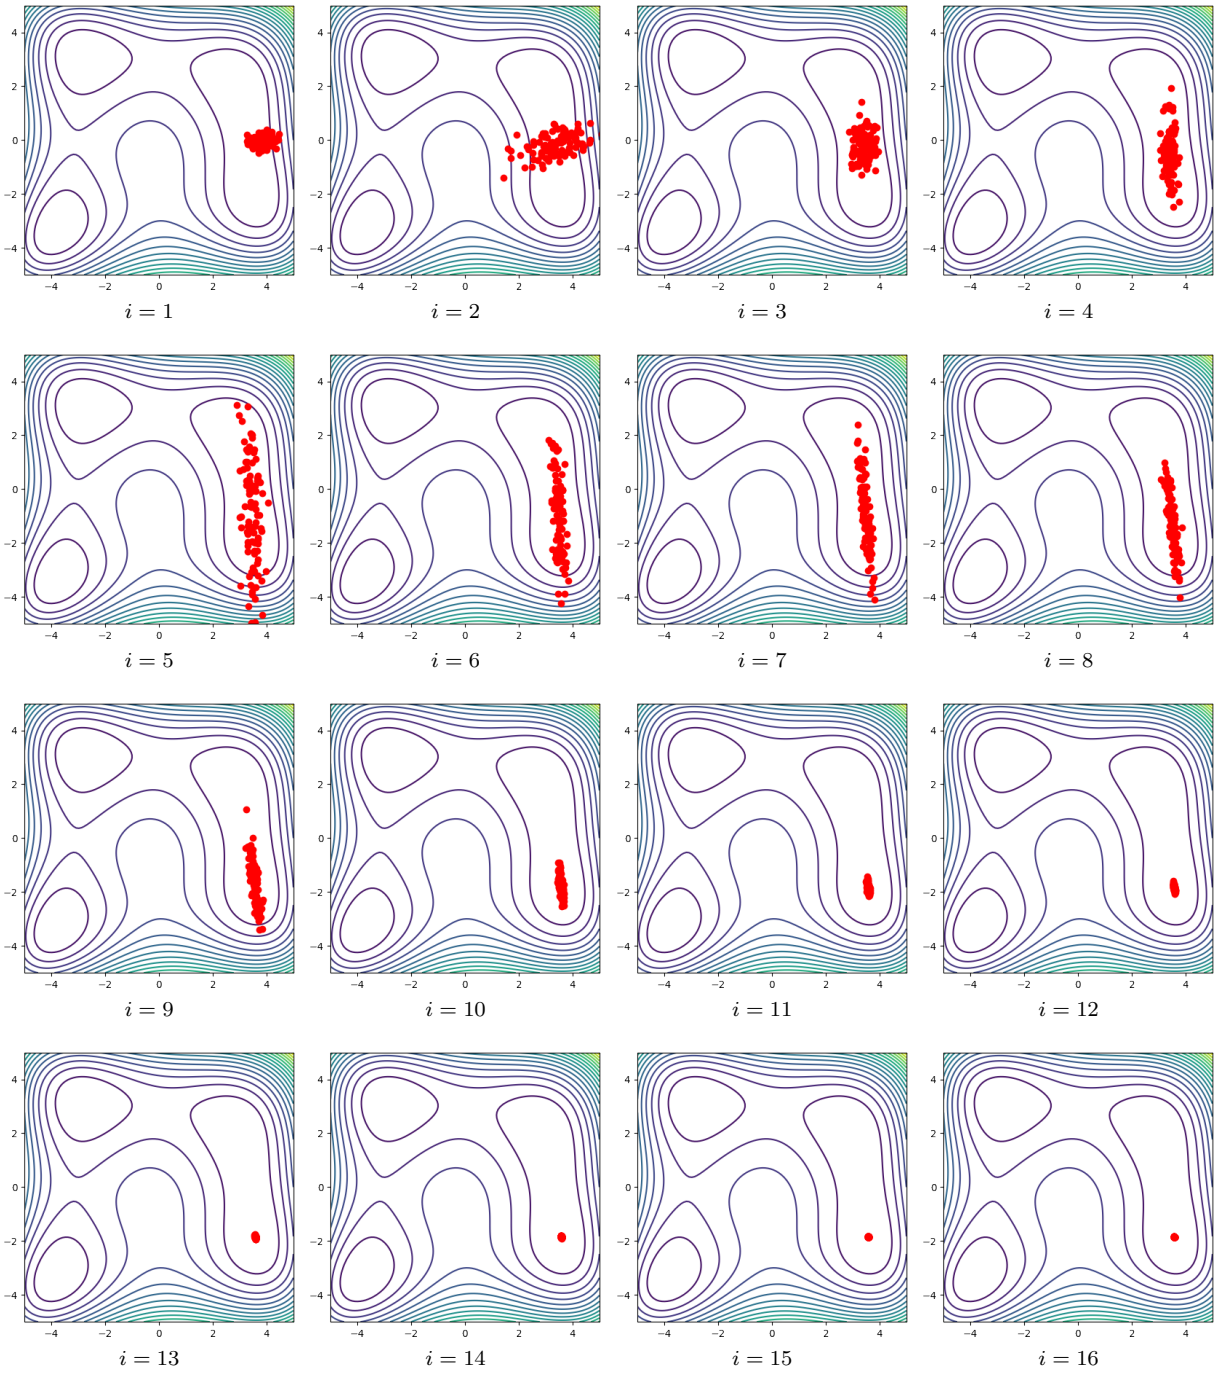

Figure 3.8: Illustration of the evolution of CMA-ES when minimizing the Himmelblau function (3.23). The level curves are represented as colored contours. The red dots represent the CMA-ES population (candidate solutions), which has a size of $\lambda = 100$. The initial parameters were $\mathbf{x} = (4, 0)$ and $\sigma = 0.1$. It can be observed that the solutions "enter" the valley containing the minimizers $(3, 2)$ and $(3.584428, -1.848126)$ and then converge to the latter.

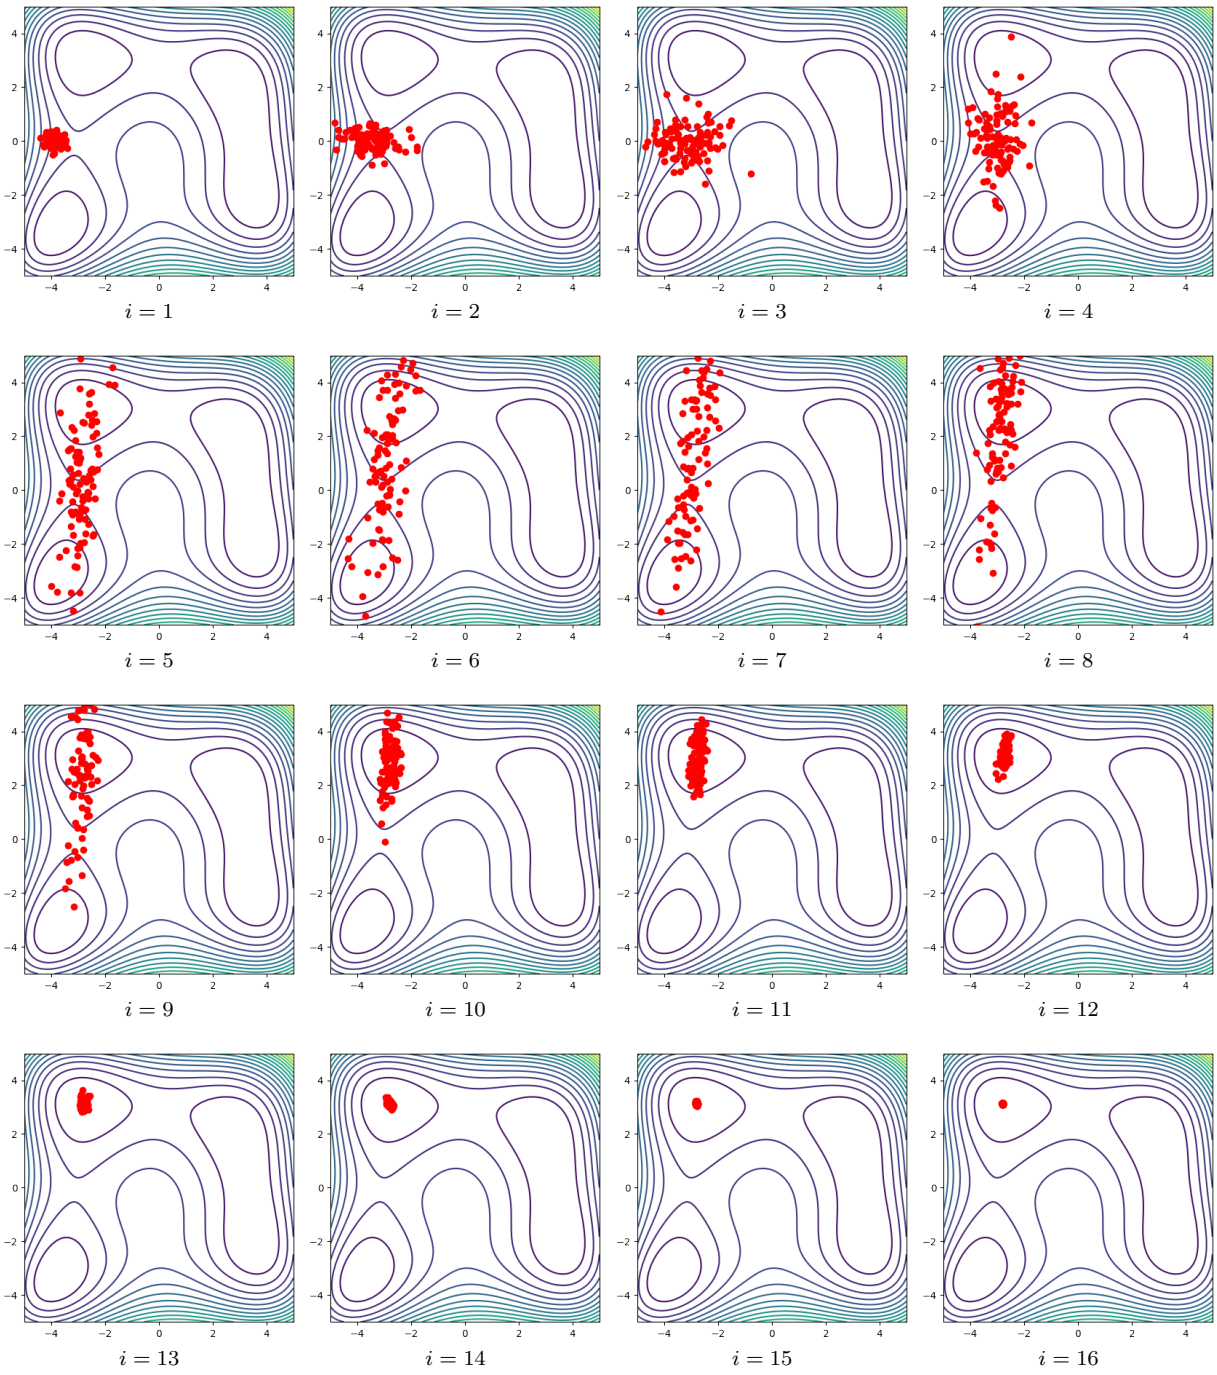

Figure 3.9: Illustration of the evolution of CMA-ES when minimizing the Himmelblau function (3.23). The level curves are represented as colored contours. The red dots represent the CMA-ES population (candidate solutions), which has a size of $\lambda = 100$. The initial parameters were $\mathbf{x} = (-4, 0)$ and $\sigma = 0.1$. In this case the solutions converged to the minimizer $(-2.805118, 3.131312)$.

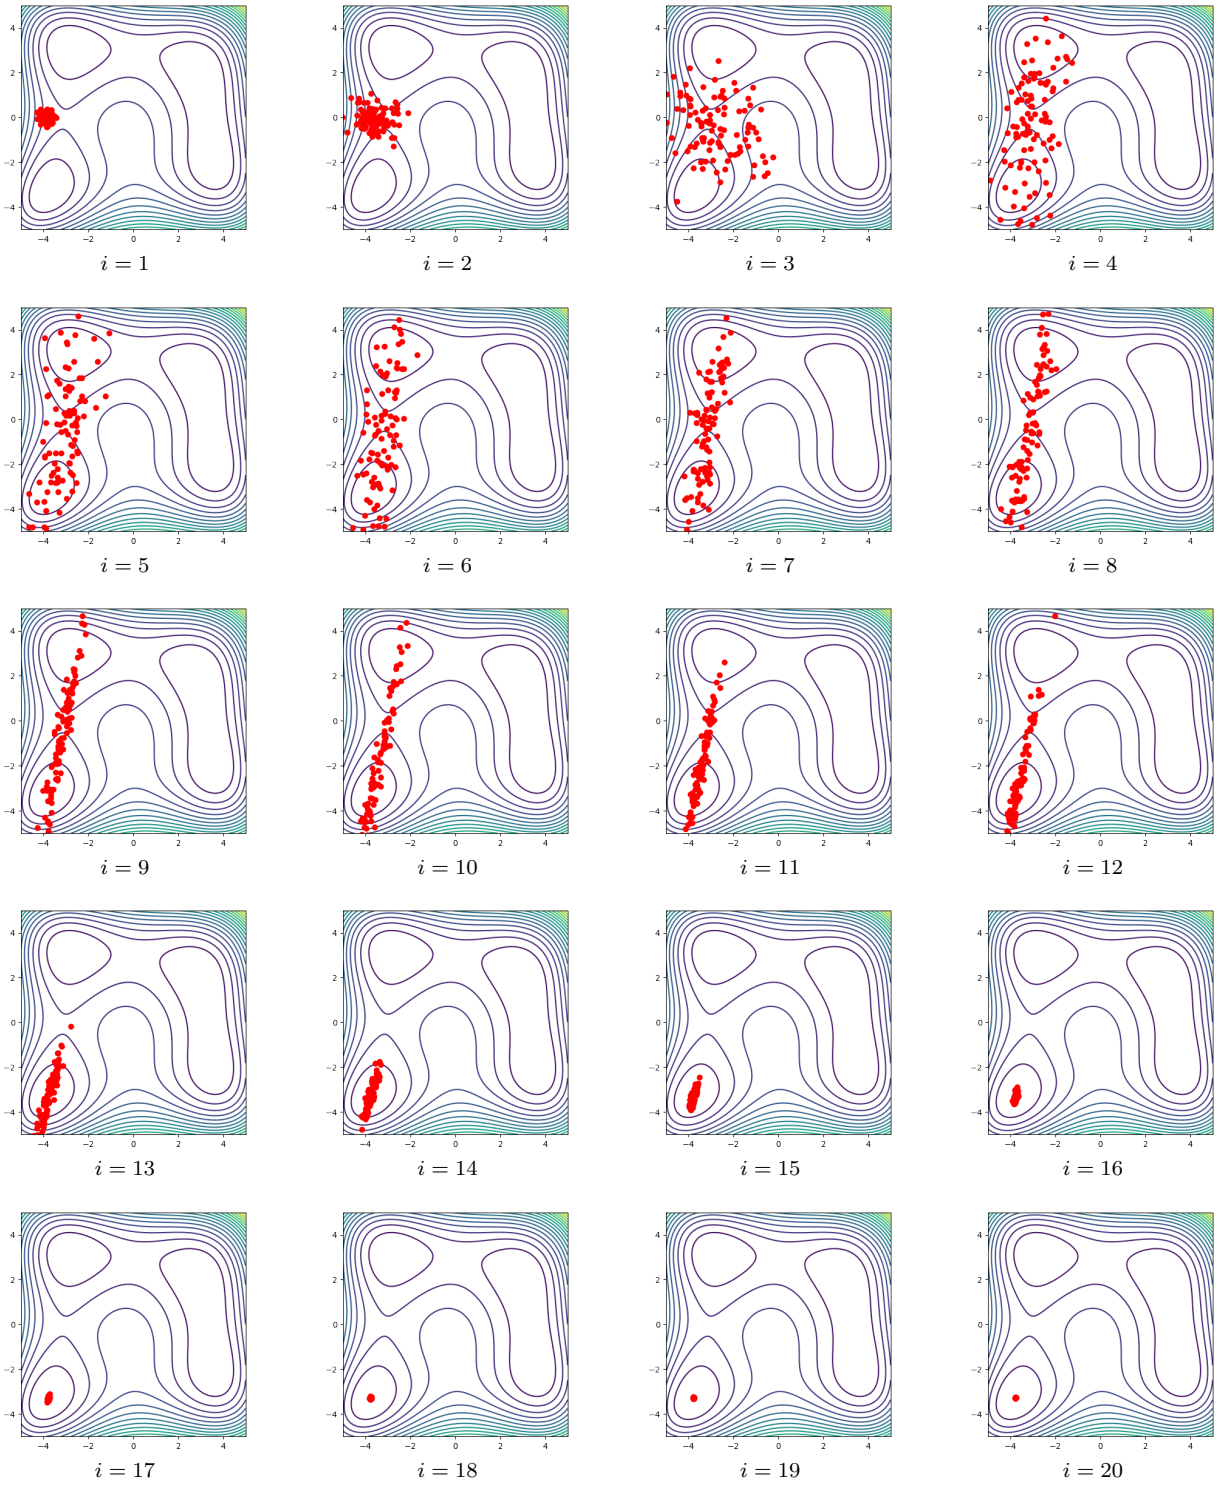

Figure 3.10: Illustration of the evolution of CMA-ES when minimizing the Himmelblau function (3.23). The level curves are represented as colored contours. The red dots represent the CMA-ES population (candidate solutions), which has a size of $\lambda = 100$. The initial parameters were $\mathbf{x} = (-4, 0)$ and $\sigma = 0.1$, the same that were used in fig. 3.9. However, in this case the solutions converged to the minimizer $(-3.779310, -3.283186)$, due to the stochastic nature of CMA-ES.

The informative online path planner algorithm which achieves the goals set out in section 3.1 and was implemented for this dissertation is detailed in algorithms 3.2 and 3.3.

Algorithm 3.2 starts by setting the necessary variables to their initial values (lines 2–12). The function `GenerateGlobalPath` on line 11 generates a global path $\mathcal{G}$ with a predefined shape which is constrained to the workspace $\mathcal{W}$ and has a length equal to B_{global}. The agent's starting position is then set to the start of $\mathcal{G}$. Subsequently, the main loop is started (lines 13–26): the current and next goal positions are calculated (lines 14–17), as well as the available budget (line 18) and travel distance (line 19) for the current iteration; then, the optimization problem is solved (line 20), as described in algorithm 3.3; next, the path and the observation set are updated (lines 21–22), as well as the agent's position (line 23), which is equivalent to traveling the path while making new observations; finally, the remaining variables are updated (lines 24–45), i.e. the spent budget is incremented by the travel distance, and the iteration index is incremented by 1; the cycle is repeated until all the budget is spent (line 26).

Algorithm 3.2: Informative Path Planner.

Input: Workspace $\mathcal{W}$, budget B, budget fraction β, horizon distance d_{horizon}, travel distance d_{travel}, inter-observation distance d_{obs}, previous observations $\mathcal{O}$ (optional)
Output: Path $\mathcal{Q}$, observations $\mathcal{O}$

1 **begin**
  `// Initialize variables.`
2 $i \leftarrow 0$; `// Initial iteration.`
3 $B_{\text{spent}} \leftarrow 0$; `// Initial spent budget.`
4 $B_{\text{global}} \leftarrow \beta B$; `// Global budget.`
5 $d_{\text{global}} \leftarrow \beta d_{\text{travel}}$; `// Global travel distance.`
6 $B_{\text{local}} \leftarrow (2-\beta) d_{\text{horizon}}$; `// Local per-iteration budget.`
7 $\mathcal{Q} \leftarrow \emptyset$; `// Initial path.`
8 **if** $\mathcal{O}$ *not given* **then**
9 $\mathcal{O} \leftarrow \emptyset$; `// Initial observation set.`
10 **end**
11 $\mathcal{G} \leftarrow$ `GenerateGlobalPath`$(B_{\text{global}}, \mathcal{W})$;
12 $\mathbf{p}_{\text{agent}} \leftarrow \mathcal{G}|_0$; `// Agent's initial position.`
13 **repeat**
    `// Set the goal positions for the current iteration.`
14 $d_{\text{goal}} \leftarrow \min(d_{\text{horizon}} + i d_{\text{global}}, |\mathcal{G}|)$;
15 $d_{\text{nextGoal}} \leftarrow \min(d_{\text{horizon}} + (i+1) d_{\text{global}}, |\mathcal{G}|)$;
16 $\mathbf{p}_{\text{goal}} \leftarrow \mathcal{G}|_{d_{\text{goal}}}$;
17 $\mathbf{p}_{\text{nextGoal}} \leftarrow \mathcal{G}|_{d_{\text{nextGoal}}}$;
    `// Set the local budget and the travel distance for the current iteration.`
18 $B_{\text{local}} \leftarrow \min(B_{\text{local}}, B - B_{\text{spent}})$;
19 $d_{\text{travel}} \leftarrow \min(d_{\text{travel}}, B - B_{\text{spent}})$;
    `// Solve the optimization problem.`
20 $\mathcal{P} \leftarrow$ `SolveOptimizationProblem`$(\mathbf{p}_{\text{agent}}, \mathbf{p}_{\text{goal}}, \mathbf{p}_{\text{nextGoal}}, d_{\text{horizon}}, d_{\text{obs}}, d_{\text{travel}}, \mathcal{O}, B_{\text{local}}, \mathcal{W})$;
    `// Travel the optimized path, with observations, up to the travel distance.`
21 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \mathcal{P}|_{d_{\text{travel}}}$;
22 $\mathcal{O} \leftarrow \mathcal{O} \cup \{$`MakeObservation`$(\mathcal{P}|_{k d_{\text{obs}}})$ **for** $k = 1$ **to** $\lfloor d_{\text{travel}}/d_{\text{obs}} \rfloor\}$;
23 $\mathbf{p}_{\text{agent}} \leftarrow \mathcal{Q}|_{\text{end}}$; `// Update the agent's position.`
24 $B_{\text{spent}} \leftarrow B_{\text{spent}} + d_{\text{travel}}$; `// Update the spent budget.`
25 $i \leftarrow i + 1$; `// Update the iteration index.`
26 **until** $B_{\text{spent}} = B$;
27 **end**

Algorithm 3.3 gives a detailed description of the function `SolveOptimizationProblem` used on line 20 of algorithm 3.2. It starts by fixing the positions of the first and last control vertices (lines 2–3), so that the optimized path $\mathcal{P}$ starts at the agent's current position and ends at the current iteration's goal position. The CMA-ES routine is encoded in lines 4–15. The solver is initialized (line 4), to indicate that its initial parameters are set and that it keeps track of its state internally (see algorithm 3.1). Then the main CMA-ES loop starts (lines 5–14). This entails sampling a population of candidate solutions (line 6) and initializing the set of each candidate solution's utility $\mathcal{U}$ to the empty set (line 7). Next, for each candidate solution $\mathcal{C}$ (which is an ordered set of $M-2$ control vertices), the first control vertex is prepended and the last control vertex is appended (line 9), and the objective function is evaluated for the candidate solution (line 10) and its value appended to $\mathcal{U}$ (line 11). The function $f_{\text{objective}}$ on line 10 is the objective function defined in (3.3); notice that the term $\mathcal{G}|_{(i+2)d_{\text{travel}}}$ in (3.13) is equivalent to $\mathbf{p}_{\text{nextGoal}}$. The last step in the main CMA-ES loop is to update the solver's parameters (line 13), which includes recording the best solution found so far. When the exit condition is met (e.g., after a fixed number of generations or when a threshold for the objective function is met), the optimal solution is extracted (line 15). Finally, the local path is parameterized as a B-spline and returned to the main informative path planner algorithm (line 16).

Algorithm 3.3: Function `SolveOptimizationProblem`.

Input: Current position $\mathbf{p}_{\text{agent}}$, current goal position $\mathbf{p}_{\text{goal}}$, next goal position $\mathbf{p}_{\text{nextGoal}}$, horizon distance d_{horizon}, inter-observation distance d_{obs}, travel distance d_{travel}, observations $\mathcal{O}$, local budget B_{local}, workspace $\mathcal{W}$
Output: Optimized path $\mathcal{P}$

```
1  begin
       // Fix the first and last control vertices.
2      C_1 ← p_agent;
3      C_M ← p_goal;
       // Execute the CMA-ES routine to find the optimal set of control vertices.
4      CMAES_InitializeSolver();
5      while has not met exit condition do
6          S ← CMAES_SamplePopulation();
           // Reset the tracked solutions utility set.
7          U ← ∅;
8          foreach C in S do
               // Prepend and append the first and last control vertices to complete
                  the set.
9              C ← C_1 ∪ C ∪ C_M;
               // Evaluate fitness of the solution and add it to the tracked set.
10             U ← f_objective(C, p_nextGoal, d_horizon, d_obs, d_travel, O, B_local, W);
11             U ← U ∪ U;
12         end
13         CMAES_UpdateParameters(S, U);
14     end
       // Get the optimal solution from the CMA-ES solver.
15     C_optimal ← CMAES_GetBestSolution();
       // Convert the control vertices into a B-spline which defines the local path.
16     P ← BSpline(C_optimal);
17 end
```

4

Results

Contents

4.1 Radiological Heatmap Simulator 46

4.2 Quality Metrics 48

4.3 Simulation Results 50

This chapter presents the analysis and assessment of the results of the selected approach. The results consist of sets of runs in different scenarios, such as larger/smaller budgets, or varying numbers of sources; each scenario is simulated ten times, to account for stochastic variations in simulated measurements, as well as fluctuations in the performance of the informative path planner. Additionally, the results are compared to those obtained in the same scenario with a boustrophedon path and with a Brownian motion path. The chapter has the following structure: section 4.1 describes the simulator which was developed in order to test the informative path planner and refine its performance, section 4.2 describes the quality metrics which will be used to compare the results between different scenarios and path types, and section 4.3 presents and discusses the results which were obtained with the simulator.

4.1 Radiological Heatmap Simulator

Development and testing of computer algorithms often requires a configurable environment wherein different situations and conditions can be tested. Furthermore, a computer simulator allows one to perform exhaustive tests with different sets of parameter values, as well as a statistical analysis of the performance of the algorithm. As such, a radiological heatmap simulator was developed using the Python programming language.

The simulator generates 2D spaces containing radioactive sources (a configuration file is used to define the position and activity of the radioactive sources). The simulated scenario is discretized as a dense grid of 600×600 cells so that the intensity for each grid cell can be calculated according to (3.20). In every intensity heatmap generated by the simulator, the background intensity was set to 1000 Bq/m^2—this value was selected so that the average background radiation count is 0.5 Bq (30 counts per minute (CPM), which is within the normal background levels [66]) when measured by a detector with an effective sensitive area of 5 cm^2 [67]. Two simple example scenarios are illustrated in fig. 4.1. The left graph contains a single source located at $(0, 0)$ m with an activity of 30 MBq. The isotropic nature of the simulated source can be seen by the concentric circles of intensity levels. The right graph contains four sources scattered around the scenario, characterized by the following parameters: $(x_1, y_1) = (11.4, -4.0)$ m, $A_1 = 9.7$ MBq; $(x_2, y_2) = (-3.2, -10.5)$ m, $A_2 = 6.1$ MBq; $(x_3, y_3) = (2.8, 2.9)$ m, $A_3 = 5.4$ MBq; $(x_4, y_4) = (-4.4, 9.0)$ m, $A_4 = 7.3$ MBq. This second example demonstrates the additive effect of multiple sources, as can be seen by the large area of high radiological intensity which they create.

Additionally, observations can be simulated in the scenario, which are random samples from a Poisson distribution with a mean given by (3.21), which are then transformed into intensity values using (3.18). Figure 4.2 illustrates a scenario containing three sources where thirty observations at randomly chosen positions were simulated. Due to the stochastic nature of the observations, they will not be exactly equal to the ground truth (i.e., the intensity heatmap). The highlighted rectangle shown on the right side of fig. 4.2 demonstrates this behavior with two observations whose measured intensity was lower than the mean intensity (they have a lighter shade of orange than the surrounding space).

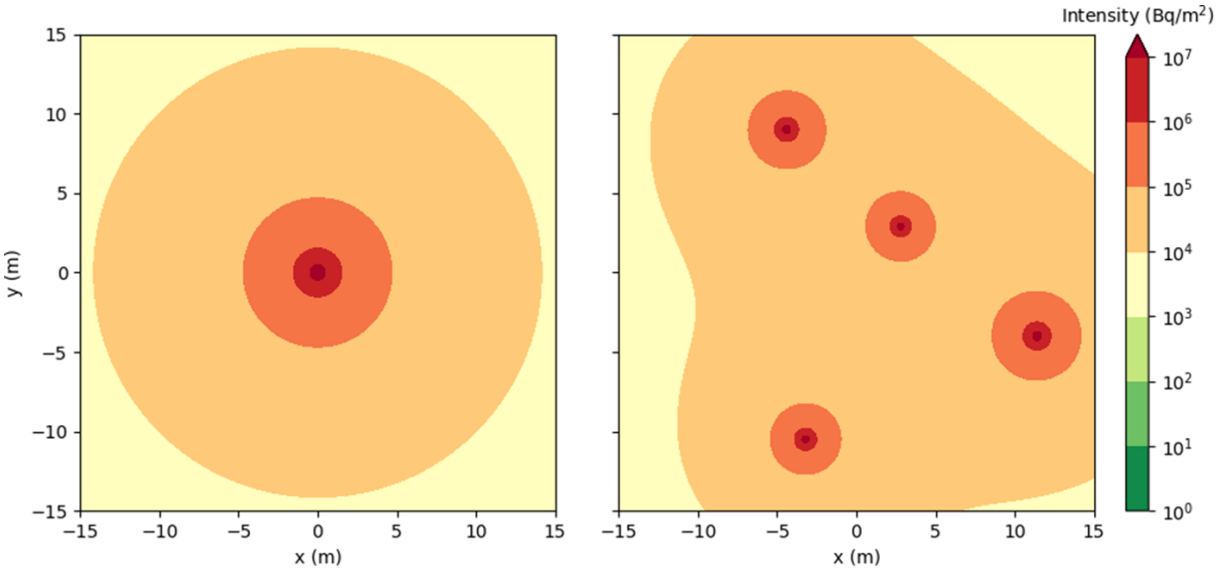

Figure 4.1: Simple toy scenarios generated with the radiological heatmap simulator. The intensity values are represented by the green-yellow-red color gradient, with green being the lowest intensity and red the highest.

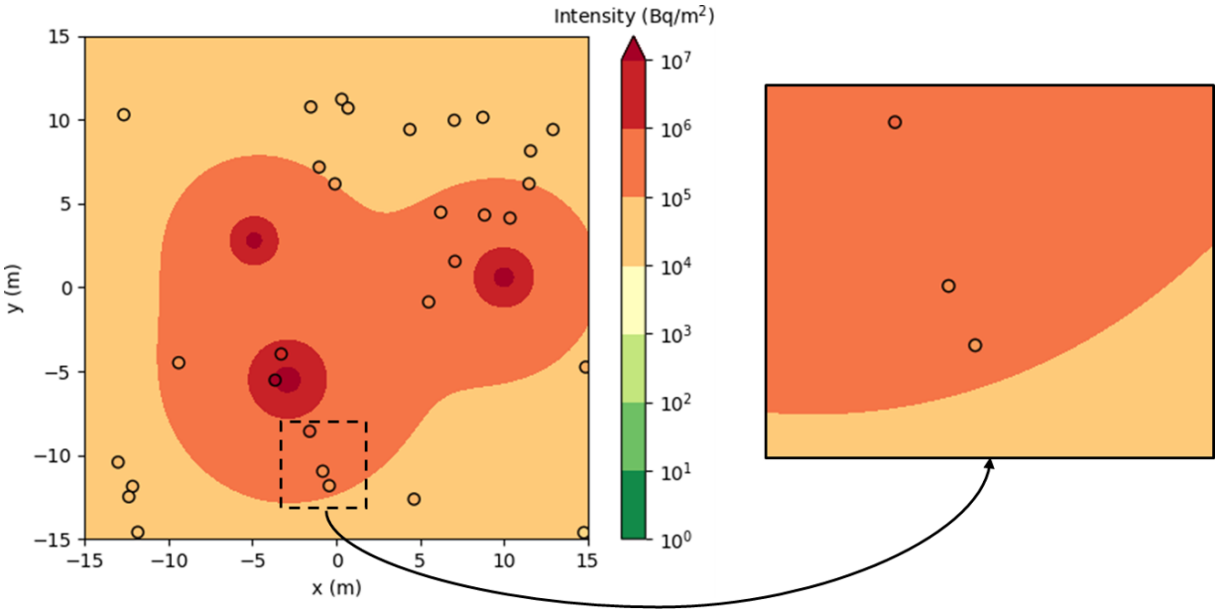

Figure 4.2: Toy scenario containing three sources where thirty observations were simulated. An area is highlighted on the right to illustrate the stochastic nature of the observations.

Note the color scale in figs. 4.1 and 4.2: it is a logarithmic scale, the minimum value is 1 Bq/m^2 (shaded dark green), and all intensities above 10 MBq/m^2 (indicated by the arrow/triangle at the top) are shaded dark red.[1] This scale was chosen for two reasons: first, intensity values larger than 10 MBq/m^2 only occur very close to the source and therefore, visually, there would not be much information gain in the heatmap graphs if the scale was increased; and second, the lower half of the scale (all of the green shaded intervals) is important for illustrating the results of the path planner—although the mean background intensity is 1000 Bq/m^2, it is not uncommon to make observations of 0 Bq a few meters away from any source due to the observation model and the random sampling process (equations (3.21) and (3.17)). Therefore, when the GP is trained on the set of observations and subsequently represented in graphs, these lower intensity values will be illustrated. It is important to represent the ground truth, observations, and GP mean on the same scale so that they can be compared, and that is why the color scale has the described format.

An additional toy scenario is illustrated in fig. 4.3. This scenario showcases the possibility to efficiently simulate shaped, non-point-like sources: by arranging several sources in a certain configuration, as shown on the right, a scenario containing a large L-shaped source is simulated.

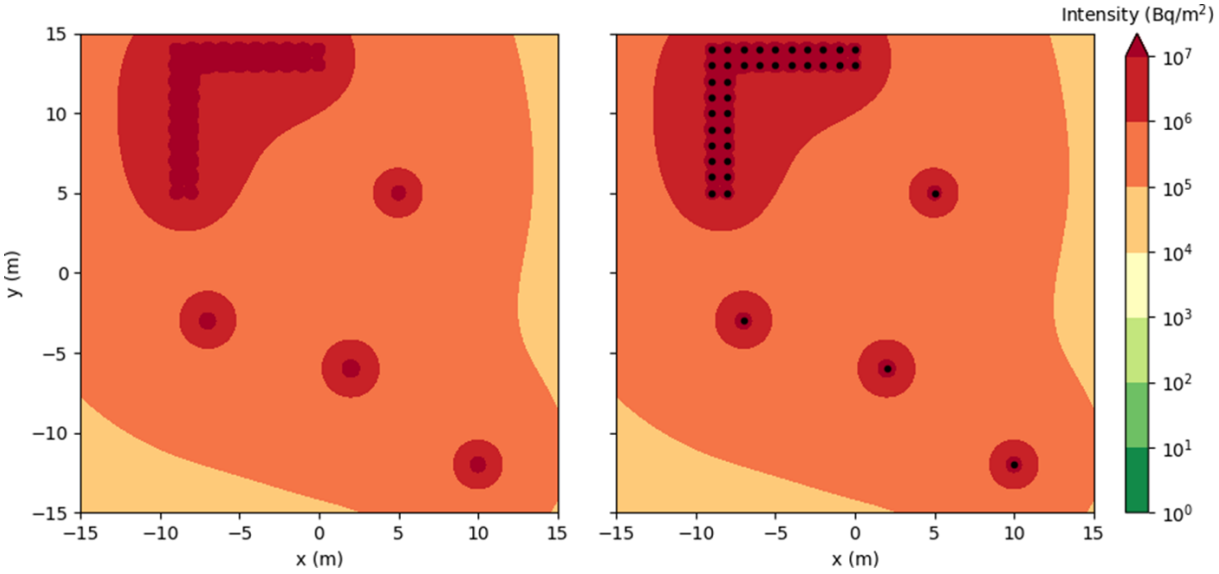

Figure 4.3: A more complex scenario generated with the radiological heatmap simulator. Left: the intensity heatmap is represented; right: the positions of the sources are highlighted as black dots.

4.2 Quality Metrics

In order to compare the informative path planner results with those obtained from boustrophedon paths and Brownian motion paths, three quality metrics will be used.

The first quality metric which will be used in assessing the path planner results is the percentage of observations made within a distance of 3 m from any source, denoted Q_K and given by (4.1). At that distance, the intensity decreases by five orders of magnitude. This quality metric is designed to assess

[1] However, internally, all values greater than or equal to 0 Bq/m^2 are represented.

the effectiveness of the $\mathbf{R}_{\text{interest}}$ term in (3.9), that is, to verify that the informative path planner does in fact result in paths which pass closer to the radioactive sources.

$$Q_K = \frac{1}{K} \sum_{\substack{k=1 \\ \|\mathbf{q}_k - \mathbf{r}_n\| \leq 3, \forall n}}^{K} 1 \tag{4.1}$$

As an example, if the total number of observations along a path is $K = 200$ and the number of observations which were within 3 m of any source was 76, then $Q_K = 0.38 = 38\%$ for that path.

We should expect Q_K to increase for all paths as the number of sources in the workspace increases (as long as the sources are scattered over the workspace and not clustered in one area). Furthermore, we expect that the informative path planner has larger values of Q_K than the other paths, on average. Since the Brownian motion path generates random patterns, it is possible that it makes many observations close to a source just by chance. However, this is not very likely, and so we expect that it will have the worst performance in terms of Q_K.

The two other metrics which will be used to compare the results obtained with the informative path planner to those obtained with the other paths are the error in estimated source position and source activity. The algorithm from [60] can be applied to the set of observations $\mathcal{O}$, giving a maximum likelihood estimate of the source parameters, namely position and activity. Then, since the source parameters are known a priori, the error in estimated position and estimated activity can be calculated, denoted $Q_{\mathbf{p}}$ (expressed by (4.2)) and Q_A (expressed by (4.3)), respectively. These quality metrics were designed to verify whether the informative path planner indeed gathers more information about the radioactive sources, based on the premises stated in section 3.1.

$$Q_{\mathbf{p}} = \|\mathbf{p}_{\text{source}} - \hat{\mathbf{p}}_{\text{source}}\| \tag{4.2}$$

$$Q_A = \frac{|A_{\text{source}} - \hat{A}_{\text{source}}|}{A_{\text{source}}} \tag{4.3}$$

As an example, if there is a source with parameters $\mathbf{p} = (x, y) = (4, 5)$ m and $A = 10$ MBq, and, based on a set of observations $\mathcal{O}$, the estimated parameters are $\hat{\mathbf{p}} = (4.3, 4.6)$ m and $\hat{A} = 8.7$ MBq, then $Q_{\mathbf{p}} = 0.5$ m and $Q_A = 0.13 = 13\%$.

These metrics are calculated separately for two reasons: (1) they operate on completely different scales (i.e., the activity is on the order of 10^6-10^9 whereas the position is on the order of 10^1-10^2), and (2), from experience, we find that the estimation algorithm will often estimate the source's position quite accurately, although it estimates the source's activity with a considerable amount of error. Additionally, $Q_{\mathbf{p}}$ is calculated as an absolute error because it gives a natural interpretation of the distance between the estimated source and the real source, whereas Q_A is calculated as a relative error so that it can be compared between sources with different activities, regardless of the actual value of the activity.

4.3 Simulation Results

The radiological heatmap simulator described in section 4.1 was used to obtain all the results reported in this section. As an introduction to explain the form of the results, fig. 4.4 illustrates an example scenario containing a single radioactive source. The source has an activity of 1 MBq and is located at $(x, y) = (0, 0)$ m. Additionally, as with all the results in this section, the source has a radius $r_{\text{source}} = 1$ cm. The combined effect of the radiation emitted by the source and the background radiation produces the mean intensity heatmap which represents the ground truth. Immediately, it can be seen that the intensity decays very rapidly; in fact, by analyzing (3.20) with $r_{\text{source}} = 0.01$ m, we find that the intensity of a source drops to 1% of its maximum value at a distance of ≈ 10 cm, and it drops to 0.1% of its maximum value at a distance of ≈ 32 cm; at a distance of 10 m the intensity is already below one millionth of its maximum. One possible global path $\mathcal{G}$ is shown on the right side of fig. 4.4. This path extends across a large part of the scenario so that the path planner can explore the whole workspace, if necessary. It has a length of 233.33 m, since it was designed by a path planner with a budget $B = 350$ m and a budget factor $\beta = 2/3$. The path shape was designed with two criteria in mind: (1) to enable the informative path planner to reach all of the workspace, and (2) to have some degree of periodicity so that the path planner can exploit previously explored areas (recall that the global path acts as a guide for the informative path planner, as explained in section 3.1). The global path starts in the bottom right corner, moves up to the top left corner, then down to the bottom left corner, up to the top right corner, back to the top left corner, and so forth, in a loop pattern which rotates $90°$ clockwise at the end of each loop. This particular global path shape is suitable for a reasonably large budget but, as will be discussed further ahead, for smaller budgets a simpler path shape is necessary in order to maintain a periodic behavior.

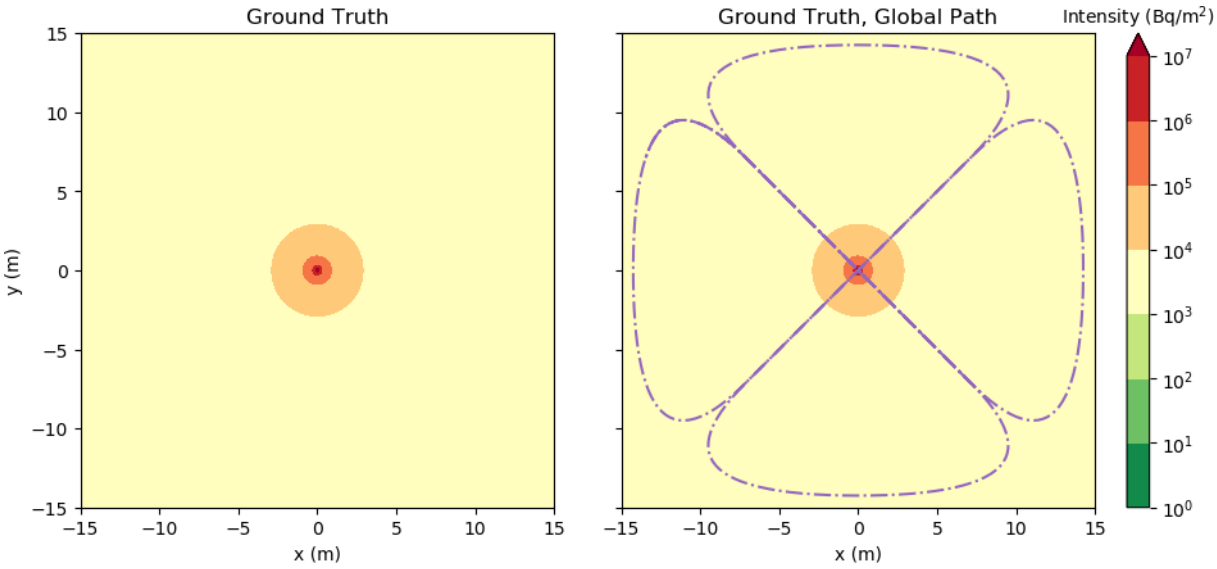

Figure 4.4: A simple scenario containing a single radioactive source at the origin. A global path for the informative path planner is depicted on the right.

The first iteration of the path planner is illustrated in fig. 4.5. The path planner was given a horizon distance $d_{\text{horizon}} = 12$ m, a travel distance $d_{\text{travel}} = 8$ m, an inter-observation distance $d_{\text{obs}} = 2$ m, and an interest threshold $\mu_{th} = 1 \times 10^4$ Bq/m^2.

As indicated by the title and legend, the graph shows the GP mean in the background (green), the global path $\mathcal{G}$ (purple), the planned path $\mathcal{P}$ (gray), the traveled path (orange) and the observations $\mathcal{O}$ (colored circles, labeled as "Field measurements"). Additionally, the agent's current position $\mathbf{p}_{\text{agent}}$ is marked by a yellow dot, and the goal position for the current iteration $\mathbf{p}_{\text{goal}}$ is marked by a red star. In this situation, since d_{travel} is a multiple of d_{obs}, $\mathbf{p}_{\text{agent}}$ will always coincide with the last observation. The GP mean is monochrome green because it is initialized with a constant value of 0 Bq/m^2 (although the GP is utilized as a tool to calculate the utility, it effectively has the same units as the radiological intensity). As noted in section 4.1, the color scale is the same for all graphs, so it will be omitted in the figures depicting iterations of the path planner, along with the legend.

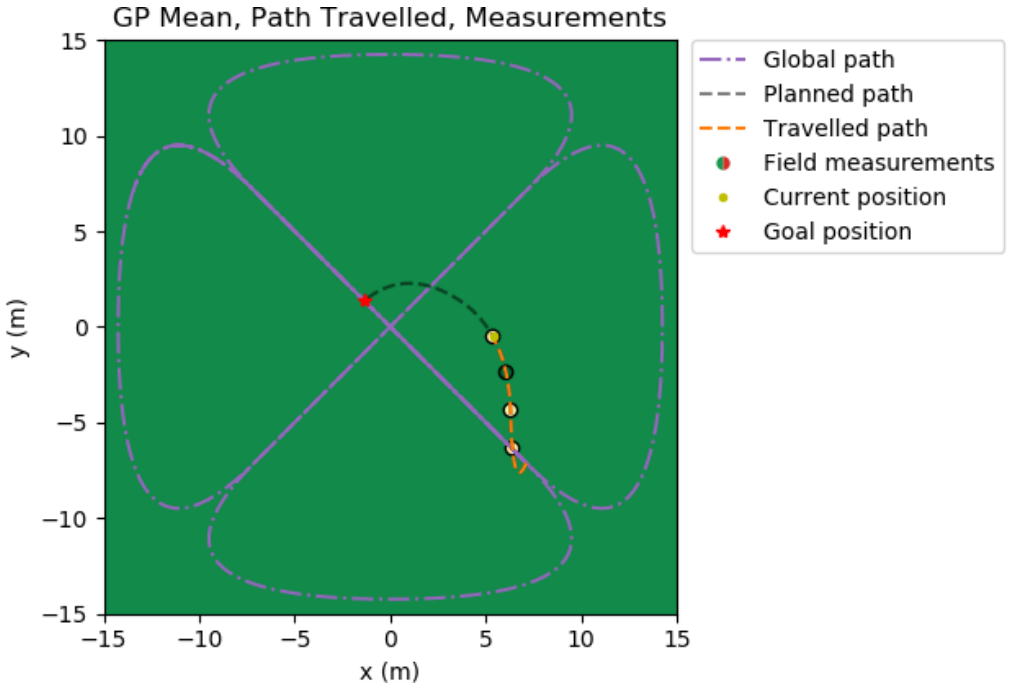

Figure 4.5: First iteration of the informative path planner.

The second iteration is illustrated in fig. 4.6. These graphs demonstrate the influence of the observations on the GP: the observations acquired in the first iteration resulted in two regions with an intensity approximately equal to the background intensity (left), and also produced a region of low uncertainty around the traveled path (right).

The GP standard deviation graph is obtained by evaluating the GP covariance matrix for all points in the grid and taking the square root of the diagonal elements. It is interesting to observe this graph in order to identify regions with high uncertainty, because the informative path planner was designed with uncertainty reduction as a strong component. Likewise, the GP mean is relevant for identifying high intensity regions, which are expected to be explored in greater detail (provided the intensity is above the threshold μ_{th}). As the observation set increases in size and diversity (i.e., observations close to and far from sources, all across the workspace), the GP mean will become more similar to the ground truth, and the GP standard deviation will decrease everywhere (graphically, it will become dark blue).

Similar to the green-yellow-red color scale for the intensity ranges, the standard deviation graphs have a blue gradient color scale, which is also logarithmic. Correspondingly, all uncertainty values larger

than 1000 Bq/m^2 will have the same color (very light blue, almost white). The color scale will be omitted for subsequent graphs since it is the same in all of them.

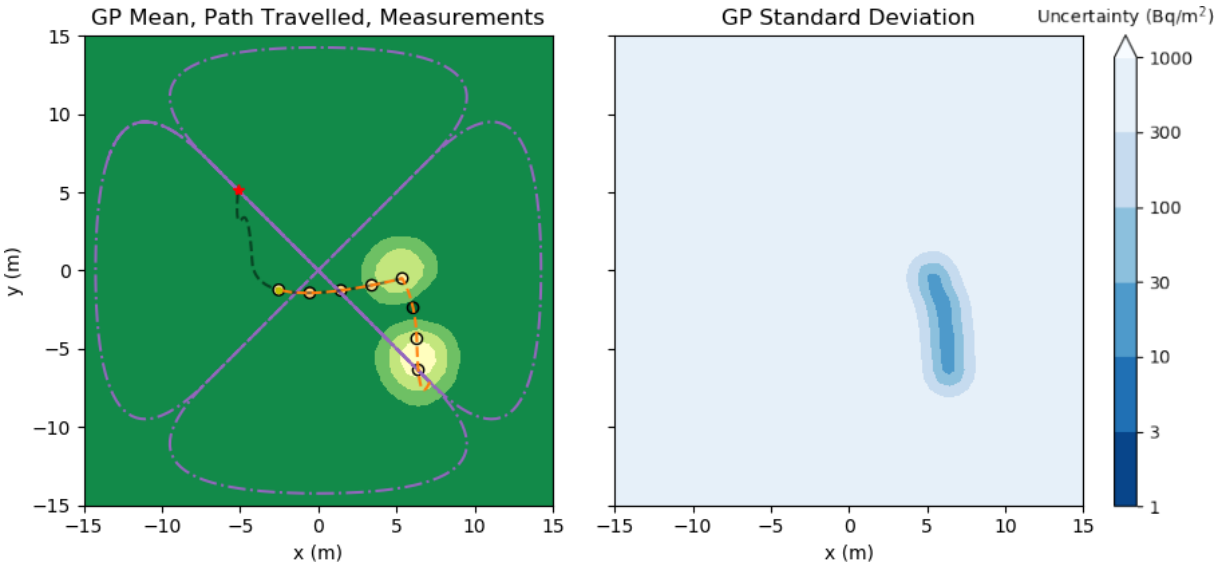

Figure 4.6: Second iteration of the informative path planner.

Iterations 3–5 are illustrated in figs. 4.7a to 4.7c. It is clear that, for these first iterations, the exploration of the workspace is given priority, even though several high intensity observations near the source have already been made; this is because there is large uncertainty everywhere in the workspace. However, at iterations 29, 30, and 38–40 (figs. 4.7d to 4.7h), the path planner gives priority to making observations close to the source, because the region around the source has large intensities and the uncertainty within the workspace has been greatly reduced. The final iteration is illustrated in fig. 4.8a.

This simple example demonstrates how the informative path planner effectively balances exploration of the entire workspace and exploration of high intensity regions. The graphs in fig. 4.8a show how the uncertainty is relatively low everywhere except for the edges of $\mathcal{W}$, and it is very low around the source, where numerous observations were made.

The results can be compared to those obtained from a boustrophedon path, illustrated in fig. 4.8b, and a Brownian motion path (generated as explained in [68]), illustrated in fig. 4.8c. These paths each have a length of 350 m, equal to the budget of the informative path planner. A total of $K = 175$ observations were made along each path. It is evident from the uncertainty graph and the position of the observations that the boustrophedon path results in a very homogeneous coverage. On the other hand, the Brownian motion path has a random coverage pattern as the result of the stochastic generation process: some regions were explored very densely and others were not explored at all. Consequently, few observations were made near the source in either case; specifically, for the boustrophedon path 7 observations were within a 3 m radius, which leads to a value of $Q_K = 4.0\%$, and for the Brownian motion path 11 observations were within a 3 m radius, which leads to a value of $Q_K = 6.3\%$. Meanwhile, a total of 29 observations were made close to the source in the case of the informative path planner, resulting in a value of $Q_K = 16.6\%$. That means that, for this example, Q_K for the informative path planner is over four times greater than Q_K for the boustrophedon path, and over two and a half times greater than Q_K for

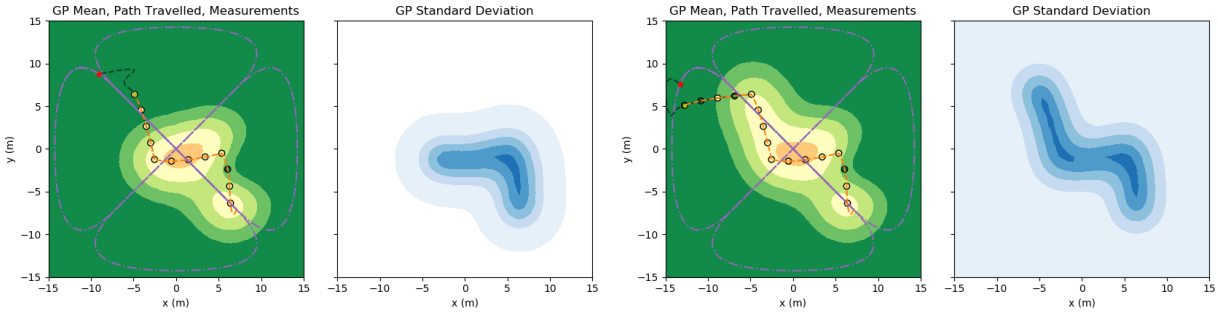

(a) $i = 3$ **(b)** $i = 4$

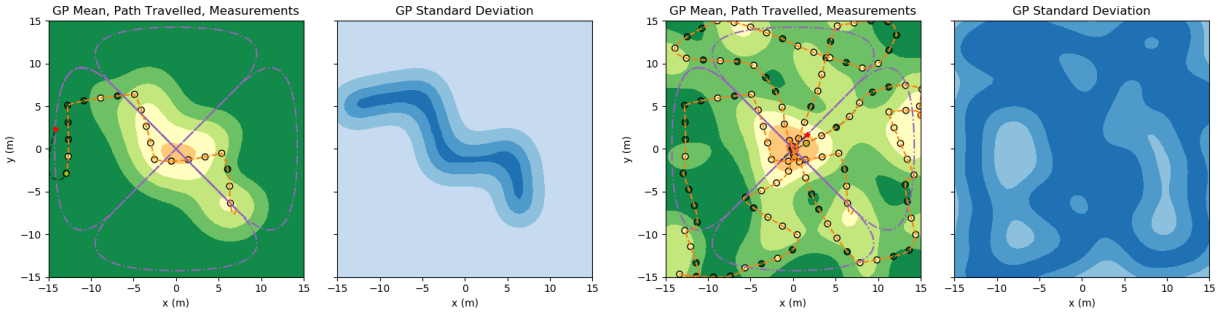

(c) $i = 5$ **(d)** $i = 29$

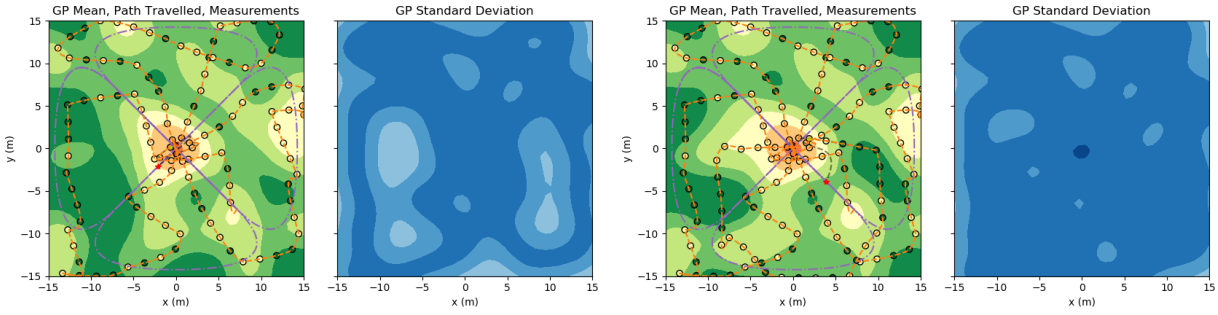

(e) $i = 30$ **(f)** $i = 38$

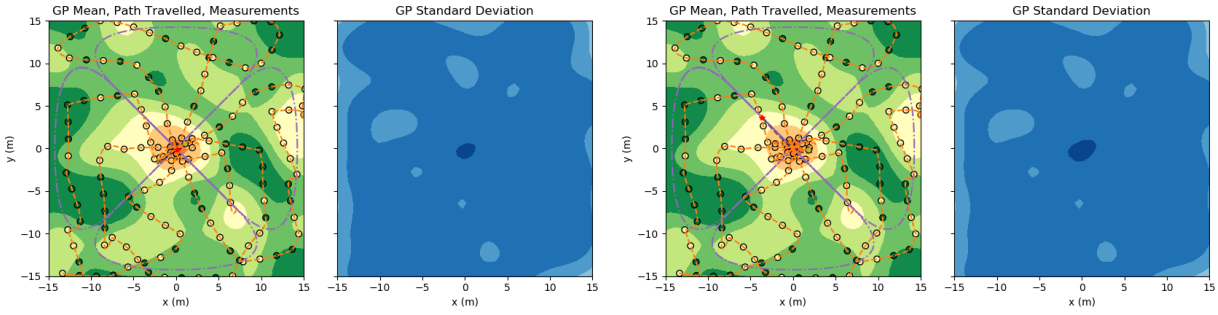

(g) $i = 39$ **(h)** $i = 40$

Figure 4.7: Several intermediate iterations of the informative path planner in a scenario containing a single source at the origin (see fig. 4.4). In the figures **(a)**–**(c)** showing the initial iterations the path planner is exploring the workspace. The remaining figures **(d)**–**(h)** show the iterations where the path planner prioritizes exploiting the region close to the source, resulting in many observations around the origin.

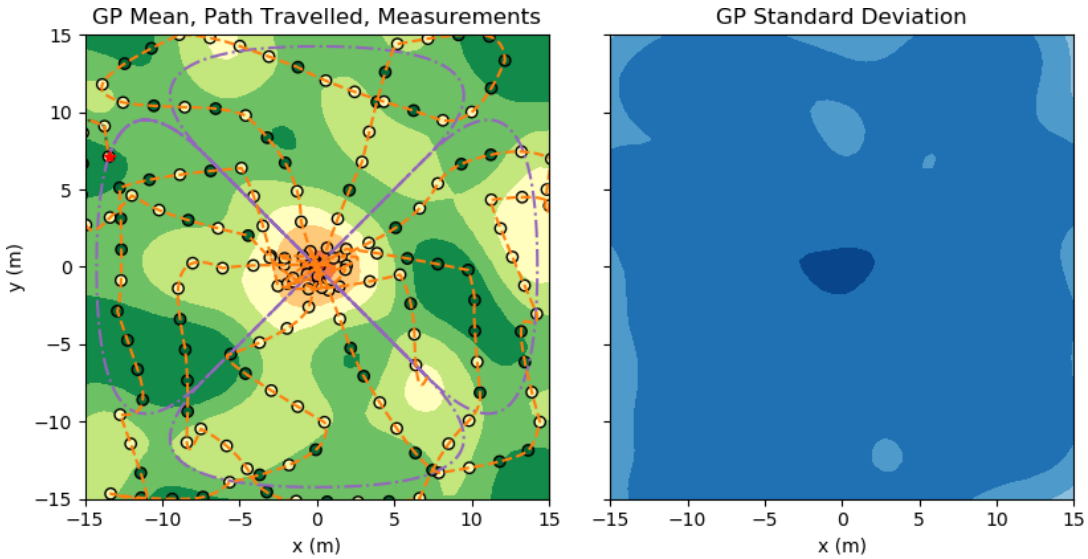

(a) Final iteration ($i = 44$) of the informative path planner.

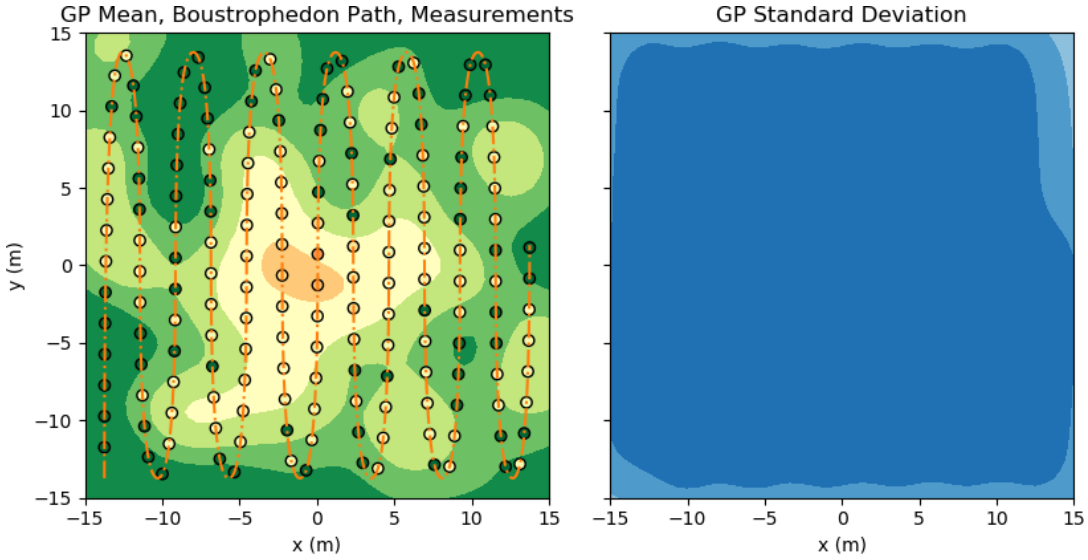

(b) Boustrophedon path simulated for comparison with the informative path planner results.

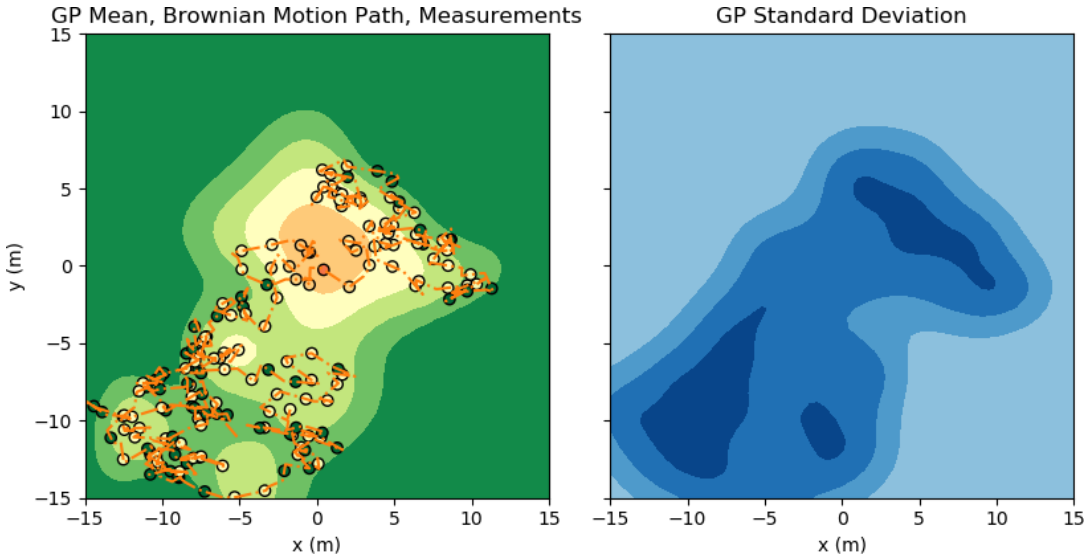

(c) Brownian motion path simulated for comparison with the informative path planner results.

Figure 4.8: Final state of the informative path, the boustrophedon path, and the Brownian motion path.

Table 4.1: Results obtained for the detailed example as explained above.

	Q_K (%)	$\hat{\mathbf{p}}_{\text{source}}$ (m)	$\hat{A}_{\text{source}}$ (MBq)	$Q_{\mathbf{p}}$ (m)	Q_A
Ground Truth	—	$(0.00, 0.00)$	1.00	—	—
Informative Path Planner	16.6	$(0.10, 0.10)$	1.06	0.14	0.06
Boustrophedon	4.0	$(0.82, 0.16)$	0.65	0.84	0.35
Brownian Motion	6.3	$(0.27, 1.33)$	1.59	1.36	0.59

the Brownian motion path. The values of all the quality metrics are presented in table 4.1; additionally, the intensity heatmaps resulting from the source estimates for each path are illustrated in fig. 4.9.

In the example described above the informative path planner performed better in terms of all the quality metrics (see table 4.1). However, due to the stochastic nature of the observations (recall that the measurements are random samples from Poisson distributions), the informative path planner (which adapts depending on the observations), and the Brownian motion (where each step is sampled from a normal distribution), a statistical analysis is necessary. Furthermore, several other scenarios will be tested, namely to assess the informative path planner's behavior when there are more sources with diverse parameters. For each scenario, we will run the informative path planner, boustrophedon path, and Brownian motion path ten times, each; subsequently, we will calculate the quality metrics for each run and take their mean, standard deviation, minimum, maximum, and median, which will be presented.

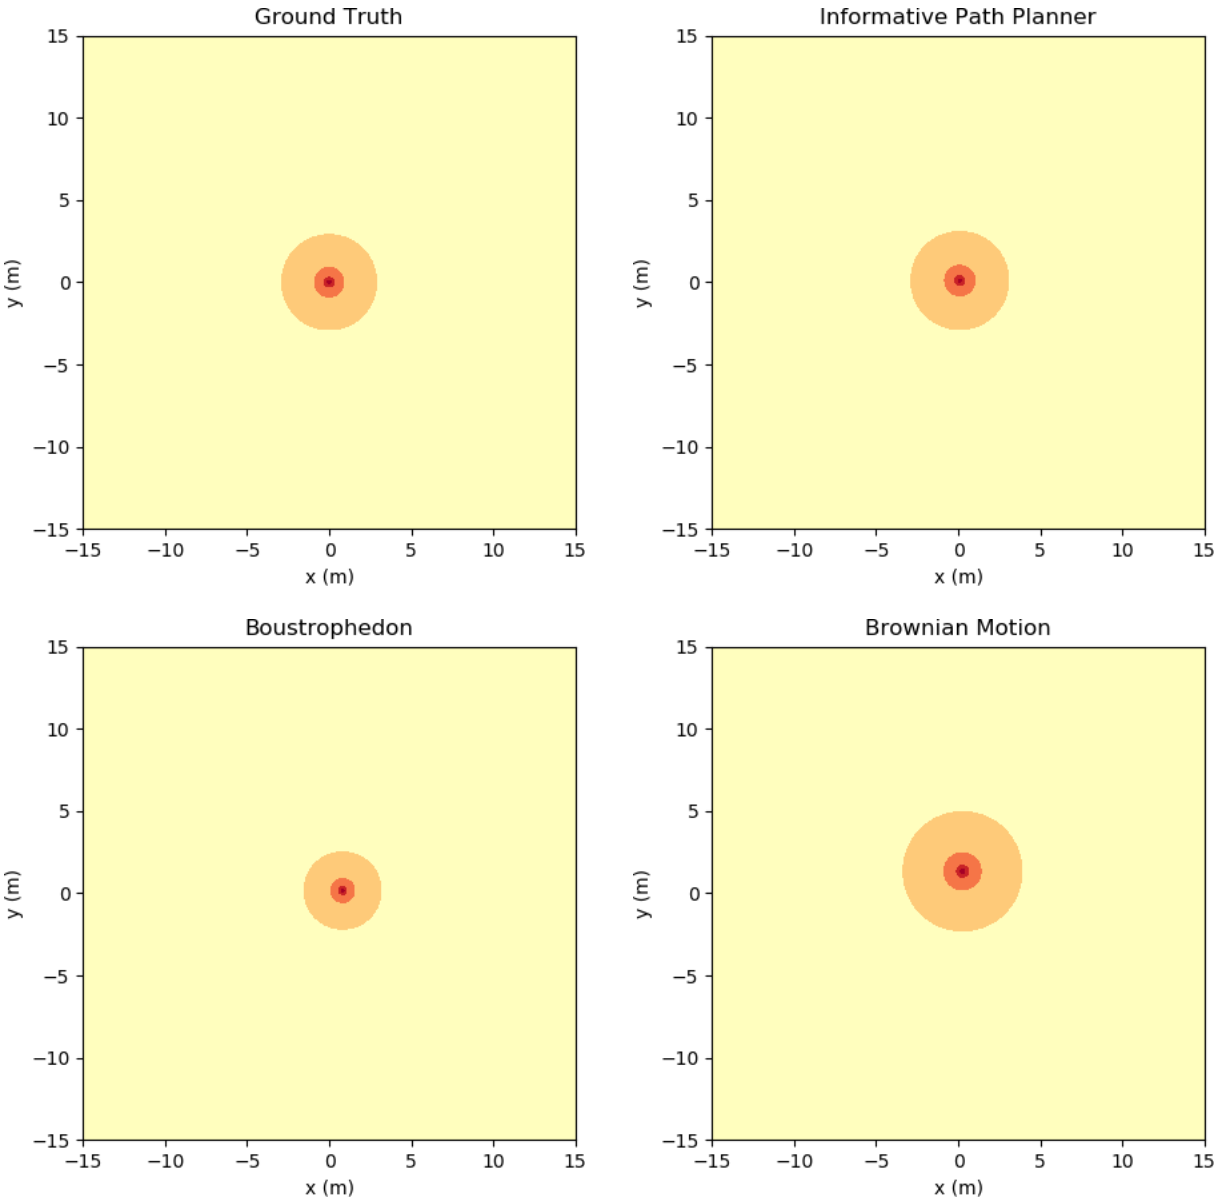

Figure 4.9: Representation of the ground truth and the heatmaps obtained as a result of the source estimates from the simulated paths.

4.3.1 Changing the Number of Sources

We will start by changing the number of sources to assess the path planner's behavior, namely for one, three, four, and five simulated radioactive sources.

One Source

This scenario is identical to the one described at the start of this section, that is, it contains a single source at $(x, y) = (0, 0)$ m with an activity of $A = 1$ MBq. The ground truth is represented in fig. 4.10. The results from performing ten runs are given in table 4.2.

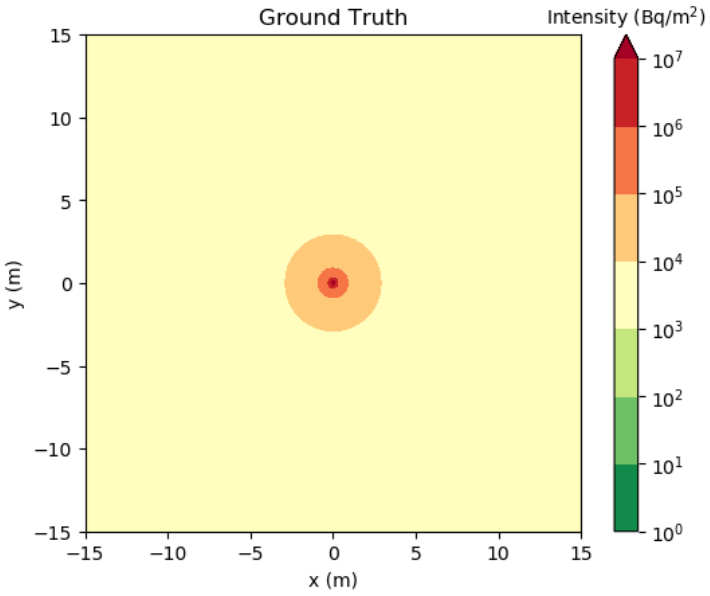

Figure 4.10: The ground truth for the first scenario, which contains a single source at the origin.

Table 4.2: Results obtained from executing ten runs with a budget of $B = 350$ m in a scenario containing a single source at $(x, y) = (0, 0)$ m with $A = 1$ MBq.

	Q_K (%)	$Q_{\mathbf{p}}$ (m)				Q_A			
	$\mu \pm \sigma$	$\mu \pm \sigma$	min	med	max	$\mu \pm \sigma$	min	med	max
Informative Path Planner									
Source 1	15.3 ± 4.5	0.31 ± 0.27	0.14	0.22	1.00	0.40 ± 0.21	0.09	0.42	0.79
Boustrophedon Path									
Source 1	4.0 ± 0.0	1.19 ± 0.55	0.26	1.27	1.94	0.40 ± 0.26	0.02	0.37	0.93
Brownian Motion Path									
Source 1	5.0 ± 6.2	3.07 ± 2.62	0.64	2.39	8.12	0.70 ± 0.26	0.31	0.76	0.98

Three Sources

This scenario contains three sources located at $(x_1, y_1) = (-6.1, 11.3)$ m, $(x_2, y_2) = (5.3, 7.6)$ m, and $(x_3, y_3) = (-10.6, -2.8)$ m, with activities of $A_1 = 10$ MBq, $A_2 = 12$ MBq, and $A_3 = 15$ MBq. The ground truth is illustrated in fig. 4.11. The results from performing ten runs are given in table 4.3. Additionally, selected results from table 4.3 are illustrated in fig. 4.12.

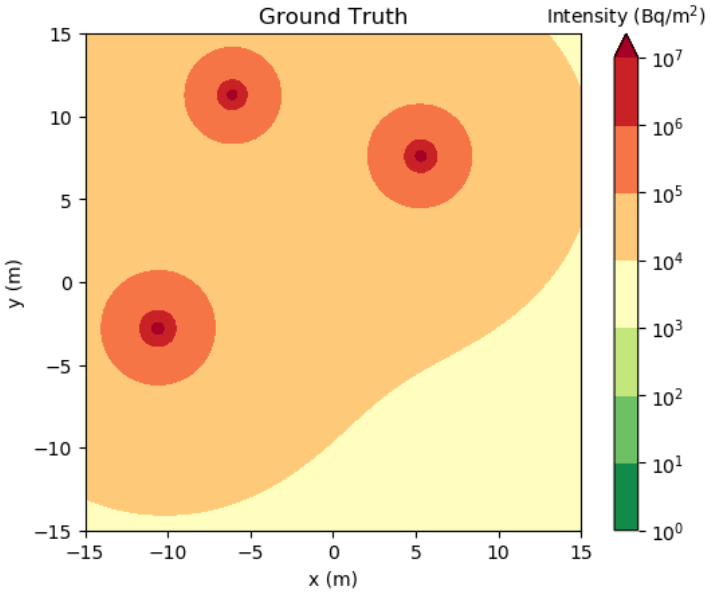

Figure 4.11: The ground truth for the second scenario, which contains three sources.

Table 4.3: Results obtained from executing ten runs with a budget of $B = 350$ m in a scenario containing three sources located at $(x_1, y_1) = (-6.1, 11.3)$ m, $(x_2, y_2) = (5.3, 7.6)$ m, and $(x_3, y_3) = (-10.6, -2.8)$ m, with $A_1 = 10$ MBq, $A_2 = 12$ MBq, and $A_3 = 15$ MBq.

	Q_K (%)	$Q_{\mathbf{p}}$ (m)				Q_A			
	$\mu \pm \sigma$	$\mu \pm \sigma$	min	med	max	$\mu \pm \sigma$	min	med	max
Informative Path Planner									
Source 1		0.44 ± 0.46	0.09	0.30	1.54	0.74 ± 0.16	0.50	0.72	0.97
Source 2	21.2 ± 2.9	1.13 ± 1.13	0.15	0.64	3.13	0.76 ± 0.29	0.24	0.93	0.98
Source 3		0.78 ± 0.57	0.21	0.55	1.78	0.79 ± 0.20	0.46	0.87	0.99
Boustrophedon Path									
Source 1		0.57 ± 0.13	0.33	0.56	0.82	0.15 ± 0.13	0.01	0.13	0.43
Source 2	10.3 ± 0.0	1.12 ± 0.22	0.71	1.12	1.47	0.26 ± 0.15	0.05	0.25	0.47
Source 3		1.06 ± 0.26	0.73	1.04	1.44	0.20 ± 0.14	0.02	0.20	0.42
Brownian Motion Path									
Source 1		13.11 ± 9.52	2.42	12.36	26.26	0.95 ± 0.05	0.85	0.97	1.00
Source 2	7.5 ± 6.6	7.45 ± 5.35	1.77	5.58	19.13	0.90 ± 0.15	0.51	0.96	1.00
Source 3		4.16 ± 5.65	0.09	1.97	18.69	0.79 ± 0.22	0.29	0.87	1.00

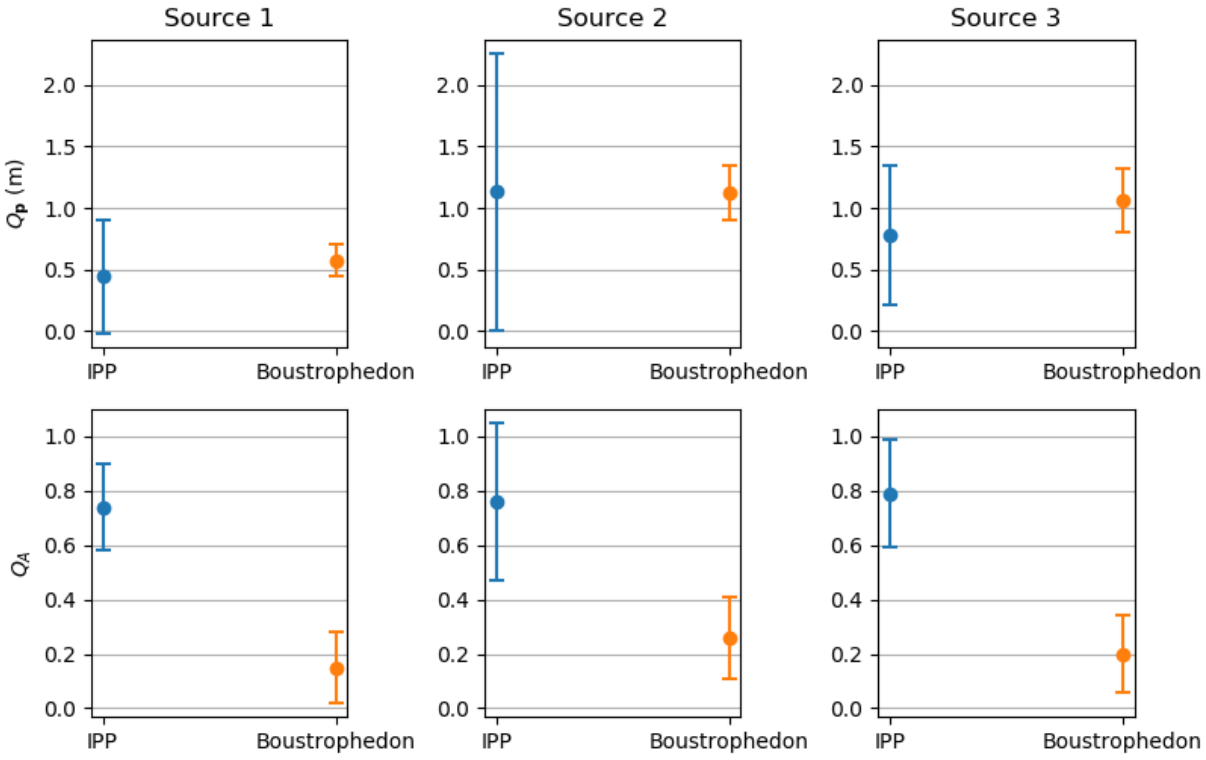

Figure 4.12: Graphs of the mean values of $Q_{\mathbf{p}}$ and Q_A for the informative path planner (blue) and the boustrophedon (orange) paths. The error bars represent a single standard deviation. The Brownian motion path results are omitted because they require a much larger scale which would hide the more interesting comparison between these two paths.

Four Sources

This scenario contains four sources located at $(x_1, y_1) = (-7.1, -9.7)$ m, $(x_2, y_2) = (8.2, -7.4)$ m, $(x_3, y_3) = (8.1, 8.6)$ m, and $(x_4, y_4) = (-2.6, 5.4)$ m, with activities of $A_1 = 6.8$ MBq, $A_2 = 7.6$ MBq, $A_3 = 8.7$ MBq, and $A_4 = 8.7$ MBq. The ground truth is illustrated in fig. 4.13. The results from performing ten runs are given in table 4.4. Additionally, selected results from table 4.4 are illustrated in fig. 4.14.

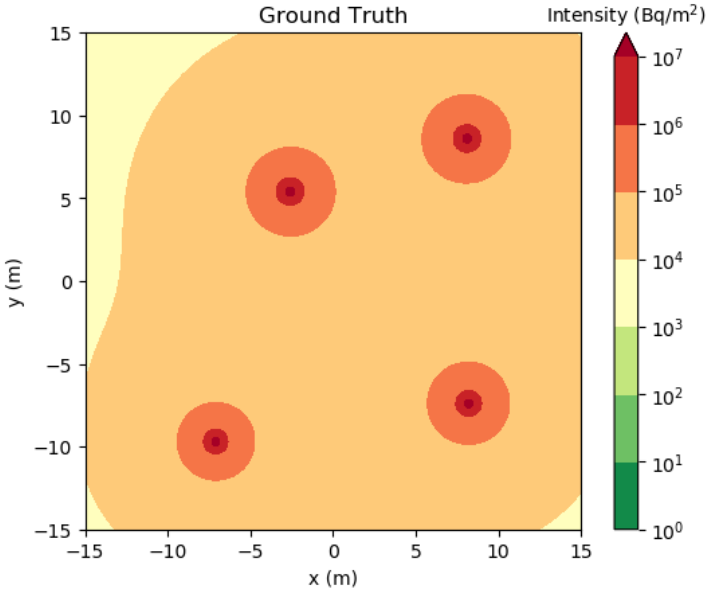

Figure 4.13: The ground truth for the third scenario, which contains four sources.

Table 4.4: Results obtained from executing ten runs with a budget of $B = 350$ m in a scenario containing four sources located at $(x_1, y_1) = (-7.1, -9.7)$ m, $(x_2, y_2) = (8.2, -7.4)$ m, $(x_3, y_3) = (8.1, 8.6)$ m, and $(x_4, y_4) = (-2.6, 5.4)$ m, with $A_1 = 6.8$ MBq, $A_2 = 7.6$ MBq, $A_3 = 8.7$ MBq, and $A_4 = 8.7$ MBq.

	Q_K (%)	$Q_{\mathbf{p}}$ (m)				Q_A			
	$\mu \pm \sigma$	$\mu \pm \sigma$	min	med	max	$\mu \pm \sigma$	min	med	max
Informative Path Planner									
Source 1	30.7 ± 6.3	1.24 ± 1.67	0.09	0.59	5.19	0.81 ± 0.22	0.44	0.94	1.00
Source 2		1.31 ± 1.08	0.17	1.04	2.67	0.87 ± 0.13	0.62	0.92	0.99
Source 3		0.56 ± 0.44	0.06	0.51	1.49	0.78 ± 0.19	0.39	0.86	0.97
Source 4		0.40 ± 0.44	0.10	0.18	1.27	0.70 ± 0.26	0.24	0.84	0.90
Boustrophedon Path									
Source 1	13.7 ± 0.0	0.84 ± 0.65	0.44	0.54	2.20	0.70 ± 0.19	0.28	0.78	0.79
Source 2		0.17 ± 0.09	0.07	0.16	0.37	0.65 ± 0.27	0.09	0.78	0.79
Source 3		0.83 ± 0.11	0.72	0.84	1.06	0.76 ± 0.17	0.37	0.84	0.85
Source 4		0.24 ± 0.30	0.10	0.10	0.92	0.63 ± 0.31	0.04	0.77	0.78
Brownian Motion Path									
Source 1	14.5 ± 7.3	7.08 ± 7.36	0.11	5.27	20.97	0.62 ± 0.32	0.15	0.68	0.99
Source 2		8.91 ± 9.31	0.13	4.76	27.22	0.67 ± 0.32	0.20	0.78	1.00
Source 3		9.35 ± 9.28	0.43	6.60	26.35	0.72 ± 0.28	0.12	0.77	0.98
Source 4		4.41 ± 5.94	0.28	2.44	19.35	0.71 ± 0.22	0.39	0.64	1.00

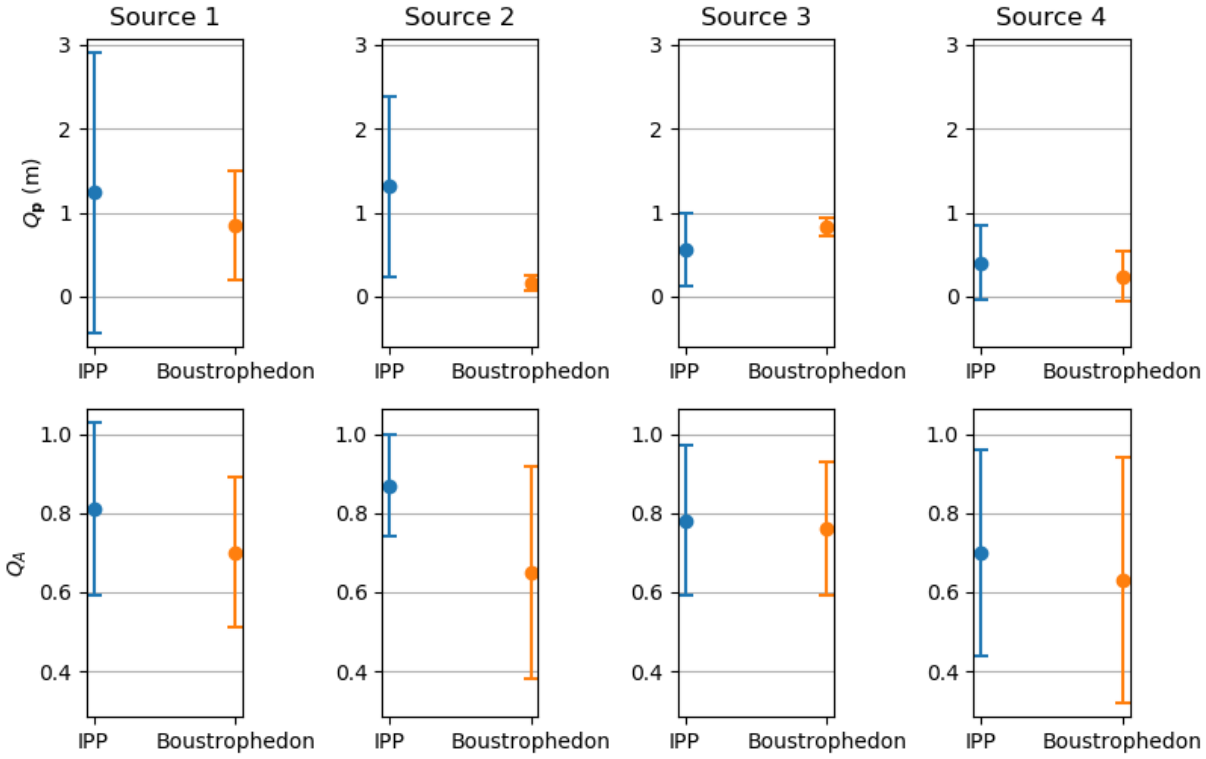

Figure 4.14: Graphs of the mean values of $Q_{\mathbf{p}}$ and Q_A for the informative path planner (blue) and the boustrophedon (orange) paths. The error bars represent a single standard deviation. The Brownian motion path results are omitted because they require a much larger scale which would hide the more interesting comparison between these two paths.

Five Sources

This scenario contains five sources located at $(x_1, y_1) = (2.1, -1.0)$ m, $(x_2, y_2) = (-10.0, -8.1)$ m, $(x_3, y_3) = (7.0, -8.8)$ m, $(x_4, y_4) = (8.9, 11.4)$ m, and $(x_5, y_5) = (-9.5, 6.4)$ m, with activities of $A_1 = 70.0$ MBq, $A_2 = 5.5$ MBq, $A_3 = 0.31$ MBq, $A_4 = 0.92$ MBq, and $A_5 = 9.6$ MBq. The ground truth is illustrated in fig. 4.15. The results from performing ten runs are given in table 4.5. Additionally, selected results from table 4.5 are illustrated in fig. 4.16. Note that, in this scenario, the source parameter estimation algorithm is operating sub-optimally, and was only able to detect source 3 (which is the weakest source) four times in total, two runs from the informative path planner, one from the boustrophedon path, and one from the Brownian motion path. Therefore, it is omitted from table 4.5. This demonstrates the limits of the source estimation algorithm, but also of the proposed solution, since it is attracted to the strong source (source 1) which is located in close proximity, and therefore does not sufficiently explore the region around source 3 so that the sources may be distinguished.

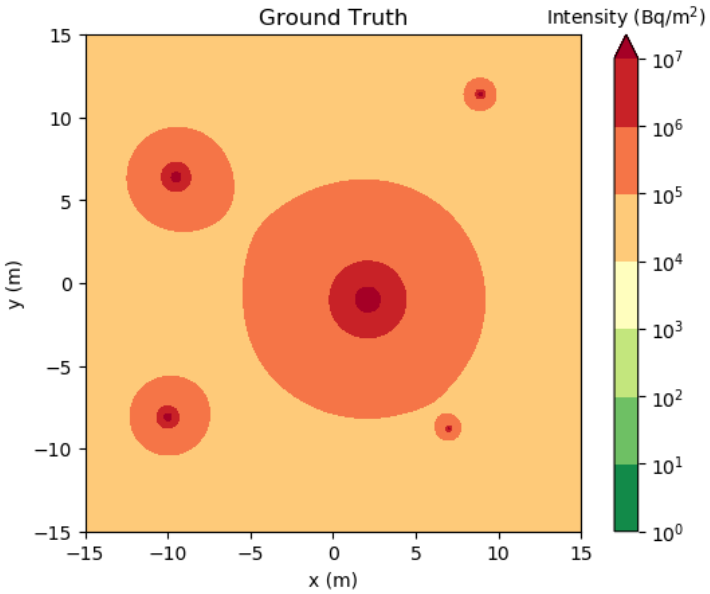

Figure 4.15: The ground truth for the fourth scenario, which contains five sources.

Table 4.5: Results obtained from executing ten runs with a budget of $B = 350$ m in a scenario containing five sources located at $(x_1, y_1) = (2.1, -1.0)$ m, $(x_2, y_2) = (-10.0, -8.1)$ m, $(x_3, y_3) = (7.0, -8.8)$ m, $(x_4, y_4) = (8.9, 11.4)$ m, and $(x_5, y_5) = (-9.5, 6.4)$ m, with $A_1 = 70.0$ MBq, $A_2 = 5.5$ MBq, $A_3 = 0.31$ MBq, $A_4 = 0.92$ MBq, and $A_5 = 9.6$ MBq. Entries for source 3 are omitted because it was not detected in 87% of the runs.

	Q_K (%)	$Q_{\mathbf{p}}$ (m)				Q_A			
	$\mu \pm \sigma$	$\mu \pm \sigma$	min	med	max	$\mu \pm \sigma$	min	med	max
Informative Path Planner									
Source 1	34.7 ± 2.5	0.14 ± 0.06	0.07	0.13	0.25	0.85 ± 0.06	0.72	0.88	0.91
Source 2		1.29 ± 0.89	0.37	1.09	2.92	0.94 ± 0.05	0.84	0.96	0.98
Source 4		1.42 ± 1.04	0.19	0.94	3.39	0.84 ± 0.23	0.22	0.91	1.00
Source 5		0.65 ± 0.75	0.06	0.54	2.67	0.94 ± 0.04	0.89	0.95	1.00
Boustrophedon Path									
Source 1	18.9 ± 0.0	0.15 ± 0.04	0.07	0.16	0.19	0.93 ± 0.07	0.77	0.96	0.97
Source 2		0.72 ± 0.10	0.64	0.67	0.97	0.93 ± 0.09	0.73	0.97	0.99
Source 4		0.49 ± 0.17	0.27	0.45	0.86	0.79 ± 0.30	0.12	0.91	1.00
Source 5		0.23 ± 0.07	0.15	0.18	0.32	0.94 ± 0.05	0.80	0.96	0.97
Brownian Motion Path									
Source 1	15.9 ± 10.1	2.85 ± 2.27	0.20	2.54	6.47	0.87 ± 0.08	0.72	0.88	0.96
Source 2		11.60 ± 8.97	2.44	7.53	28.33	0.77 ± 0.36	0.06	0.94	1.00
Source 4		9.57 ± 6.90	0.11	9.08	19.78	6.59 ± 6.84	0.70	3.92	23.02
Source 5		10.27 ± 9.88	0.53	6.41	28.19	1.65 ± 2.69	0.53	0.83	9.28

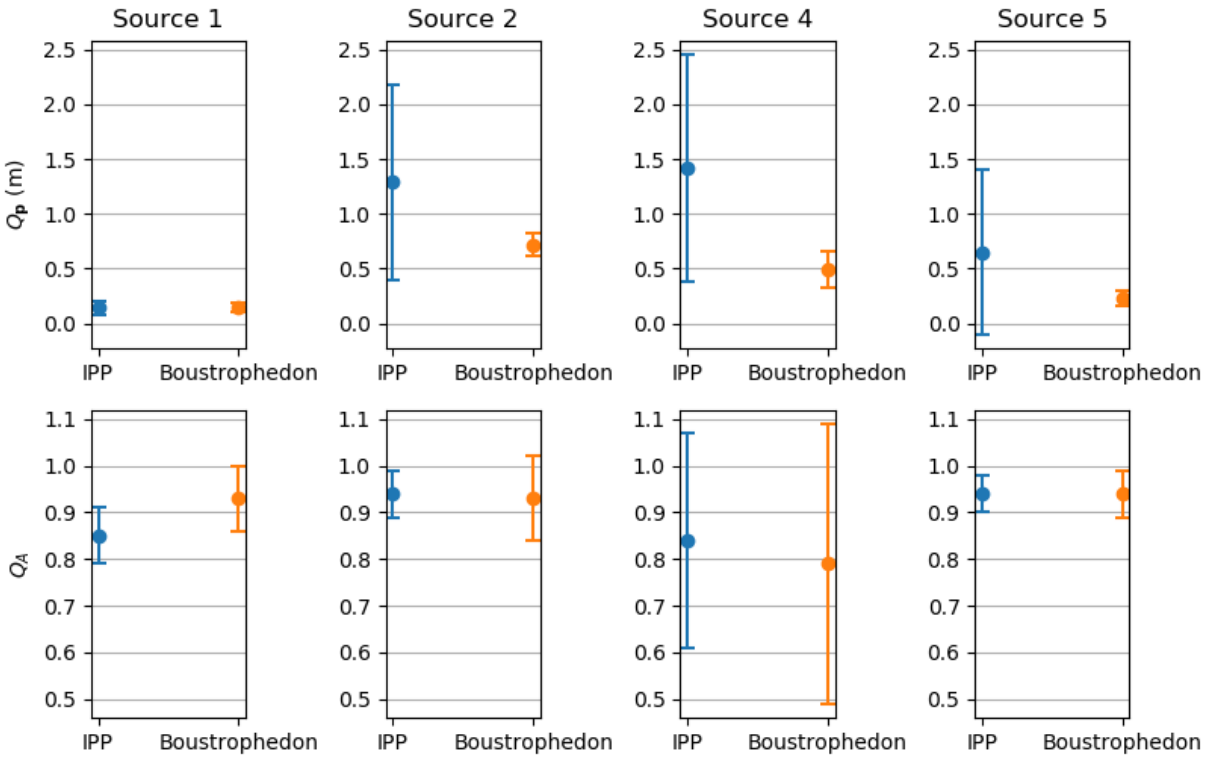

Figure 4.16: Graphs of the mean values of $Q_{\mathbf{p}}$ and Q_A for the informative path planner (blue) and the boustrophedon (orange) paths. The error bars represent a single standard deviation. The Brownian motion path results are omitted because they require a much larger scale which would hide the more interesting comparison between these two paths.

Discussion

A graphic summary of the results of these tests is presented in figs. 4.17 to 4.19. It is interesting to observe that, in relation to Q_K, the boustrophedon path and the Brownian motion path have very similar results, even though the first has a very uniform complete coverage and the latter has a completely random coverage pattern. Furthermore, the informative path planner has, on average, double the number of observations within 3 m of a source compared to the other two paths.

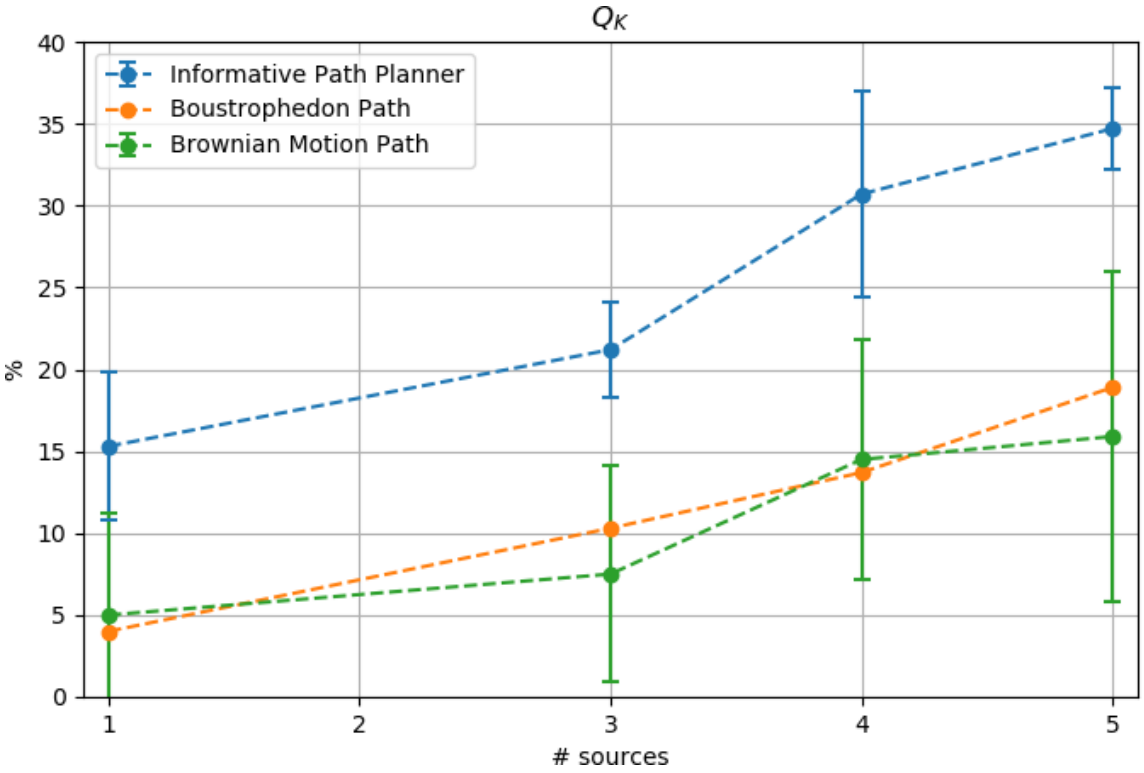

Figure 4.17: Graphical representation of Q_K versus the number of sources.

The shaded graphs in figs. 4.18 and 4.19 capture the full range of $Q_{\mathbf{p}}$ and Q_A, respectively, the circles mark the average for each metric, and the error bars represent one standard deviation. The results of the Brownian motion path were omitted from figs. 4.18 and 4.19 because the maximum values of the quality metrics masked the more interesting comparison between the other paths.

From fig. 4.18 we can see that, for a small number of sources, the informative path planner performs better on average in relation to $Q_{\mathbf{p}}$, especially in the scenario containing a single source. However, the informative path planner results do have a much larger dispersion (except in the case of a single source). In relation to $Q_{\mathbf{p}}$, the informative path planner has the potential to perform better than the boustrophedon path when it detects all the sources in the initial exploration phase, but it also has the possibility of performing much worse than the boustrophedon path when it does not detect one (or more) of the sources and subsequently prioritizes exploiting the regions around sources which were detected. This effect is exacerbated if the activity of the sources is considerably different (e.g. by a factor of 100), because the path planner is attracted to the more intense source.

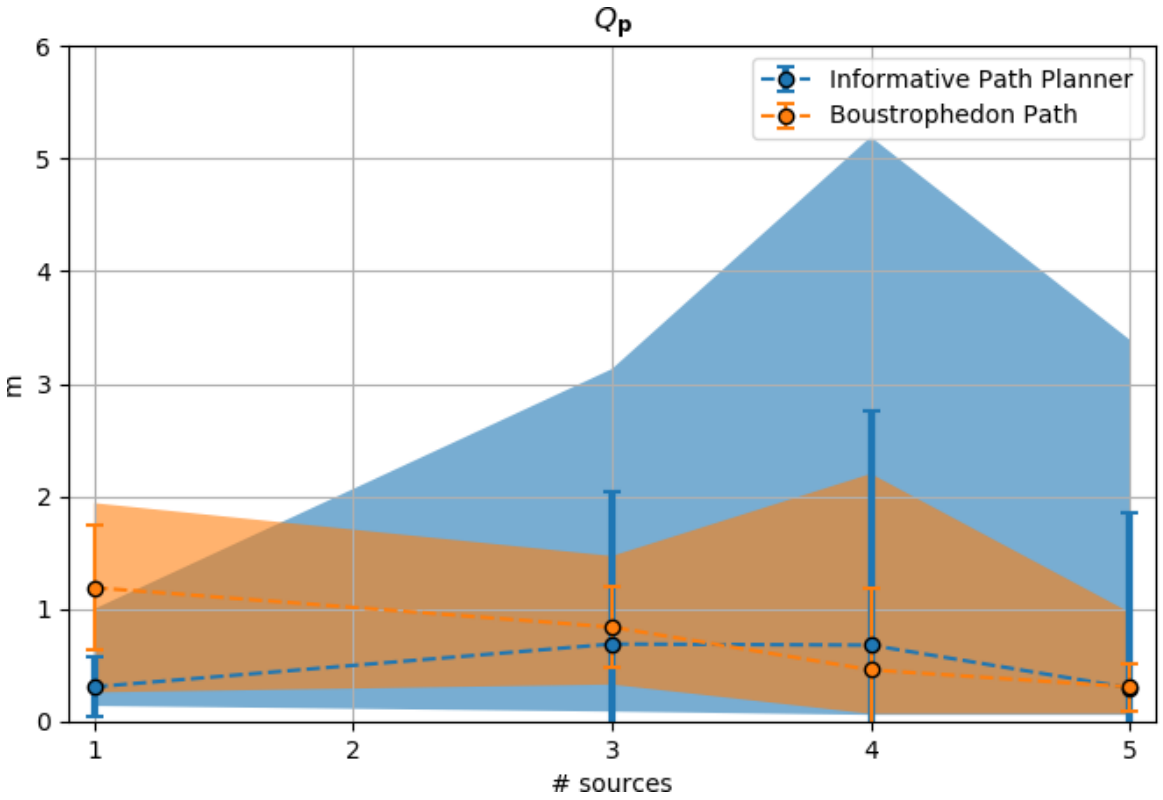

Figure 4.18: Graphical representation of Q_{p} versus the number of sources.

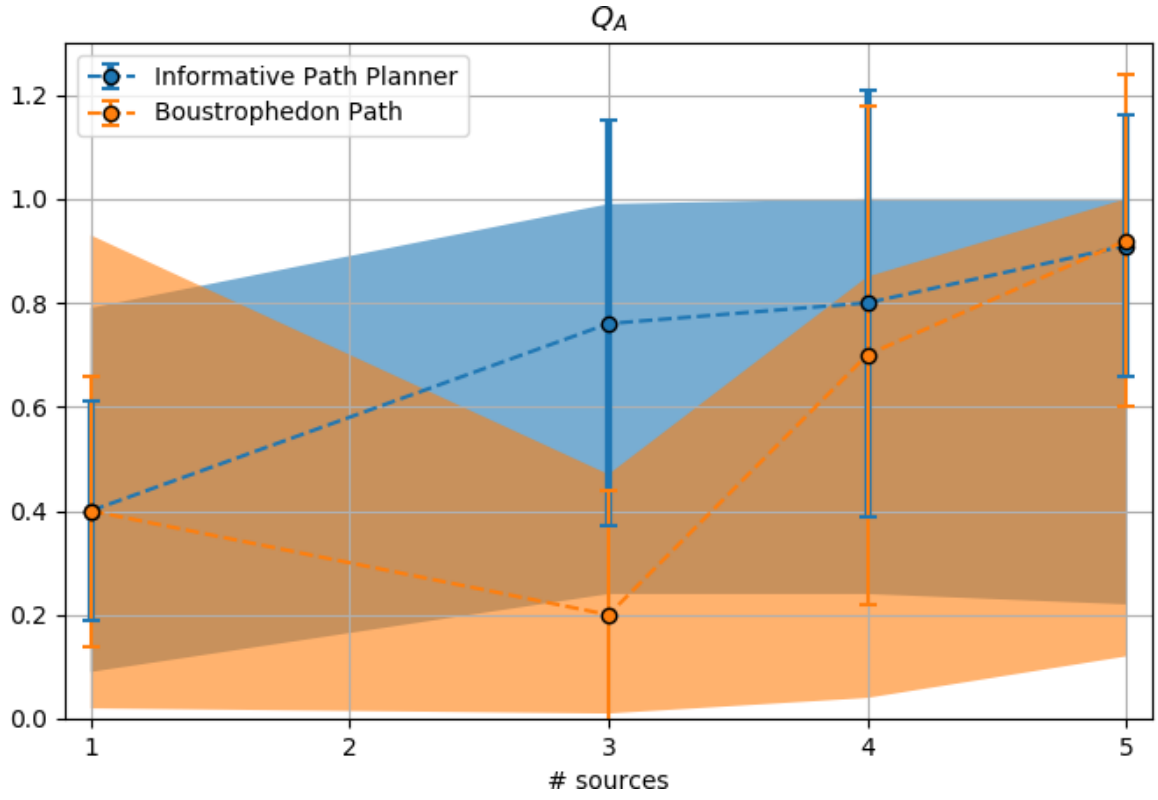

Figure 4.19: Graphical representation of Q_A versus the number of sources.

In terms of the Q_A graphs illustrated in fig. 4.19, the informative path planner consistently performs worse or equal to the boustrophedon path, on average. Additionally, the minimum value of Q_A for the informative path planner results never beats the minimum value for the boustrophedon path. Nonetheless, both paths have large standard deviations, which confirms the consideration stated at the end of section 4.2. However, in practice this may not necessarily be a problem: the informative path planner can be used to accurately locate intense sources in a given scenario, after which a more sensitive sensor can be used to precisely measure each source's activity.

4.3.2 Changing the Budget

Now we will compare the distribution of the final GP standard deviations between the informative path planner, the boustrophedon path, and the Brownian motion paths for several budgets, namely $B = 120$ m, $B = 160$ m, $B = 200$ m, $B = 240$ m, $B = 280$ m, $B = 320$ m, and $B = 360$ m. For this purpose, ten runs were executed for each path and for each budget value, in the scenario illustrated in fig. 4.20, and a GP was trained on the final set of observations for either path. Then, dividing the workspace into a dense grid of 300×300 equally sized cells, the GP standard deviation was evaluated at each cell position. Finally, a histogram was calculated for the standard deviation values, the distribution of which is compared between the paths. Notice that this is similar to calculating the relative areas of each shade of blue in the GP standard deviation graphs. The histograms are illustrated in figs. 4.21 to 4.27. The minimum and maximum of the frequency of each interval were recorded as well, and are represented by a downward pointing arrow and an upward pointing arrow, respectively; arrows which are not shown correspond to a frequency value of 0 (which can not be represented on a logarithmic scale).

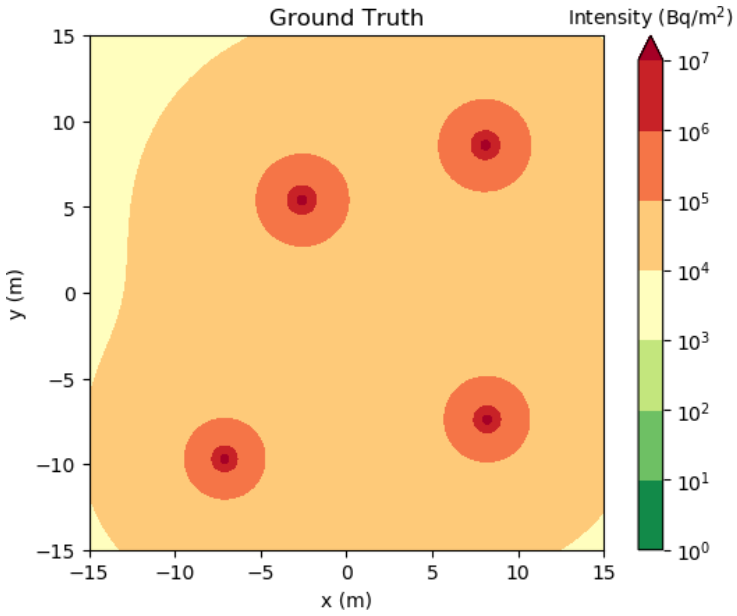

Figure 4.20: The ground truth for comparing the standard deviation distributions between the informative path planner and the boustrophedon path. This scenario contains four sources, located at $(x_1, y_1) = (-7.1, -9.7)$ m, $(x_2, y_2) = (8.2, -7.4)$ m, $(x_3, y_3) = (8.1, 8.6)$ m, $(x_4, y_4) = (-2.6, 5.4)$ m and with activities $A_1 = 6.8$ MBq, $A_2 = 7.6$ MBq, $A_3 = 8.7$ MBq, $A_4 = 8.7$ MBq.

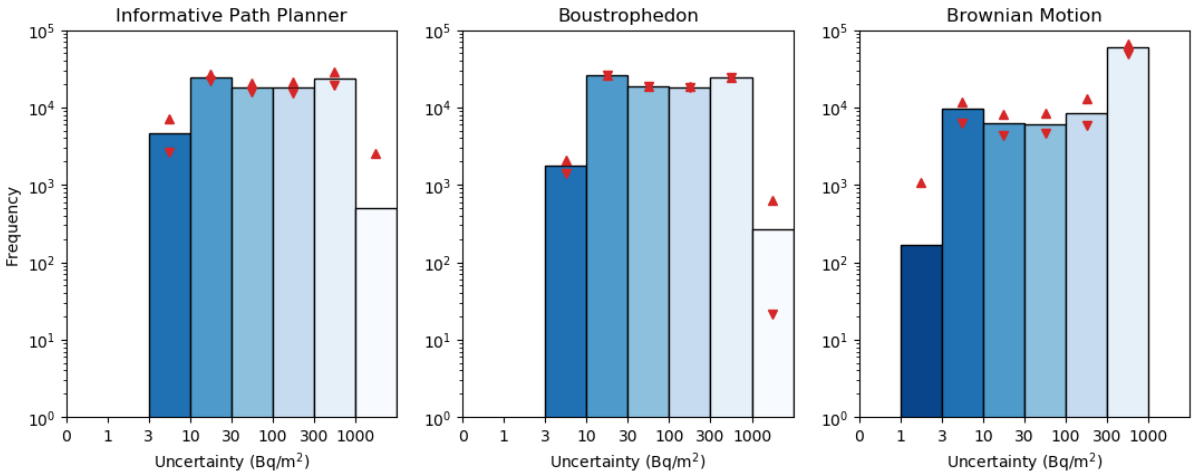

Figure 4.21: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 120$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

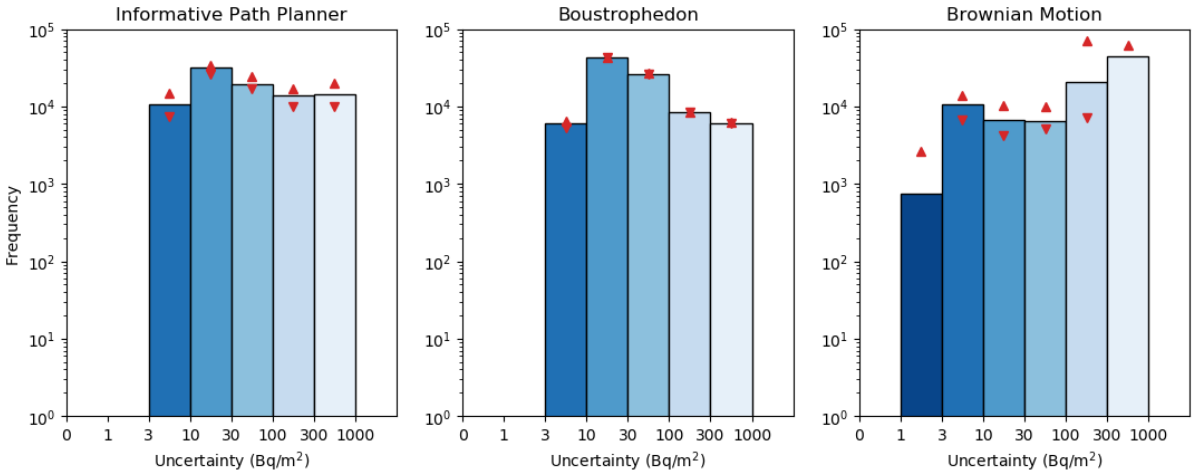

Figure 4.22: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 160$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

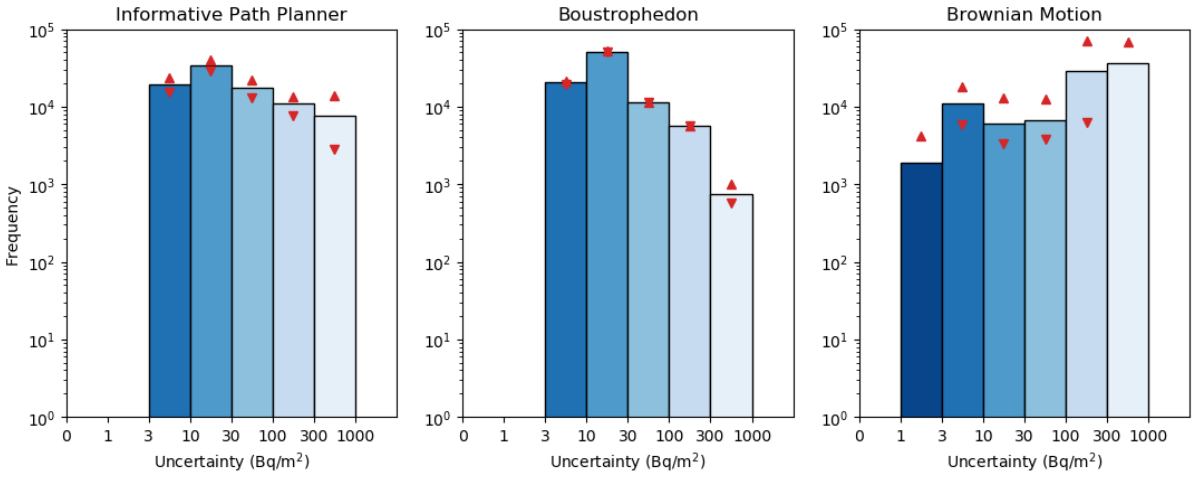

Figure 4.23: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 200$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

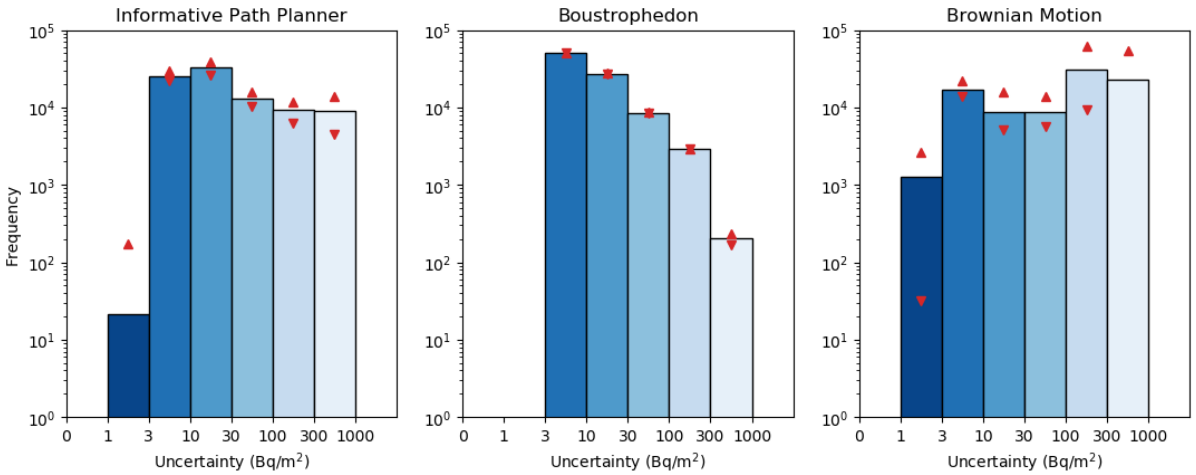

Figure 4.24: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 240$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

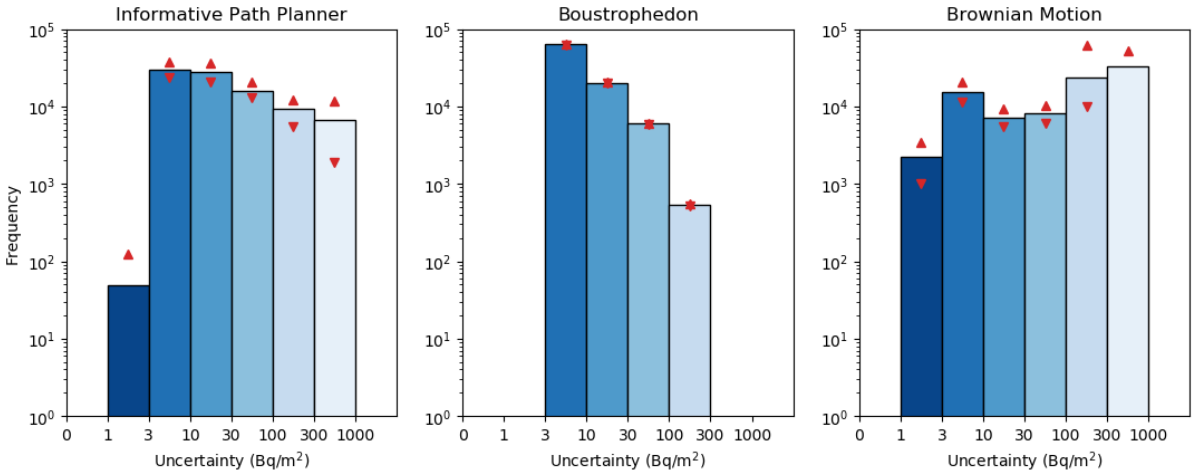

Figure 4.25: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 280$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

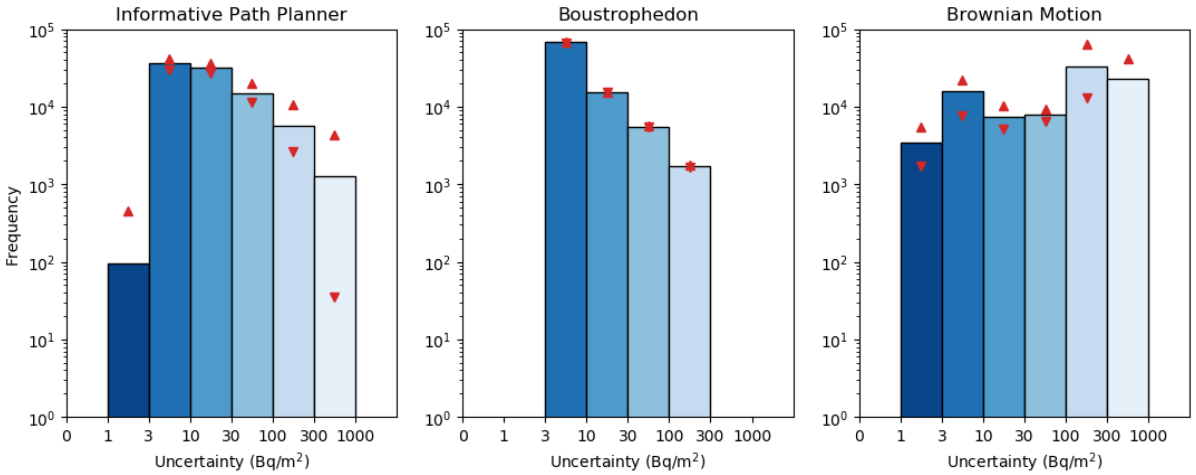

Figure 4.26: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 320$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

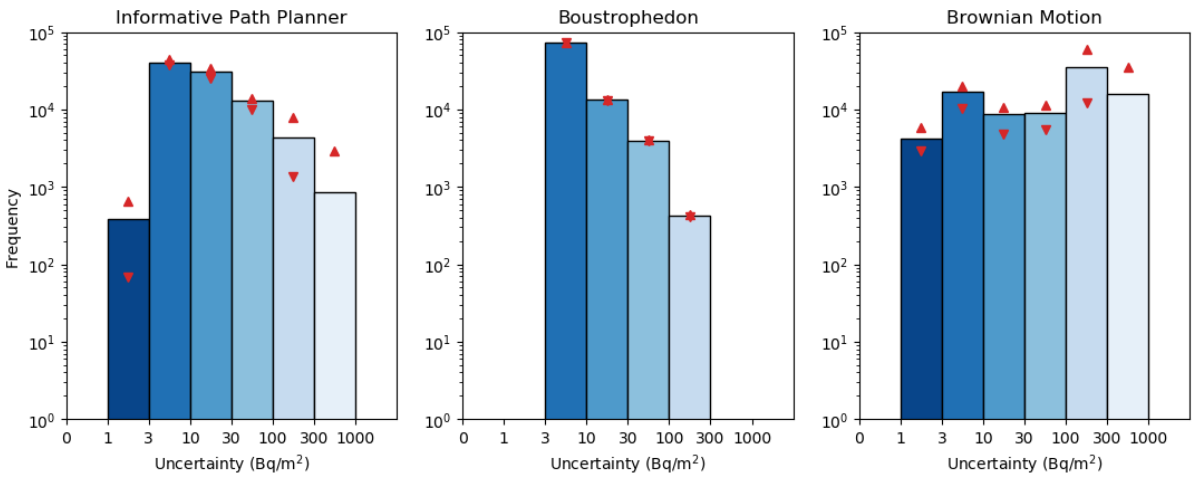

Figure 4.27: Histograms of the distribution of the standard deviation of a GP trained on a set of observations obtained by paths with a budget of $B = 360$ m. The maximum and minimum frequencies are represented by upward and downward pointing red triangles, respectively.

Discussion

From inspection of the histograms, the first observation is that the boustrophedon path is the most consistent (evident from the small difference between minimum and maximum frequencies), followed by the informative path, and lastly the Brownian motion path has the most variability, as would be expected. However, the main point to be made with this analysis is that we can observe the trade-off which the informative path planner makes: although the boustrophedon path has a higher frequency of GP standard deviation points with values less than 100 Bq/m^2 for budgets above $B = 160$ m, it never decreases below 3 Bq/m^2, whereas the informative path planner makes a trade-off between having some regions with very low uncertainty and others with high uncertainty. This is most visible in the histograms in fig. 4.27. The very low uncertainty values arise from the path planner's behavior of passing over high intensity regions very often, thus increasing the certainty of the mean intensity in those regions; naturally, as a consequence, the low intensity regions are not explored in great detail, which explains why there are still some regions with high uncertainty. Moreover, we see that as the budget increases, the frequencies for the boustrophedon histogram are "concentrated" in the interval $[3, 10]$, while the frequencies for the informative path planner histogram tend to shift to the lower intervals and spread out. On the other hand, the Brownian motion path will often explore random regions very densely and leave others completely unobserved, due to its stochastic nature, which leads to high frequencies in all ranges of uncertainty.

The frequencies of each histogram bin as a function of the path budget are presented in fig. 4.28. It is evident from these graphs that the frequencies for the large standard deviation bins decrease as the budget increases, whereas the frequencies for the low standard deviation bins increase.

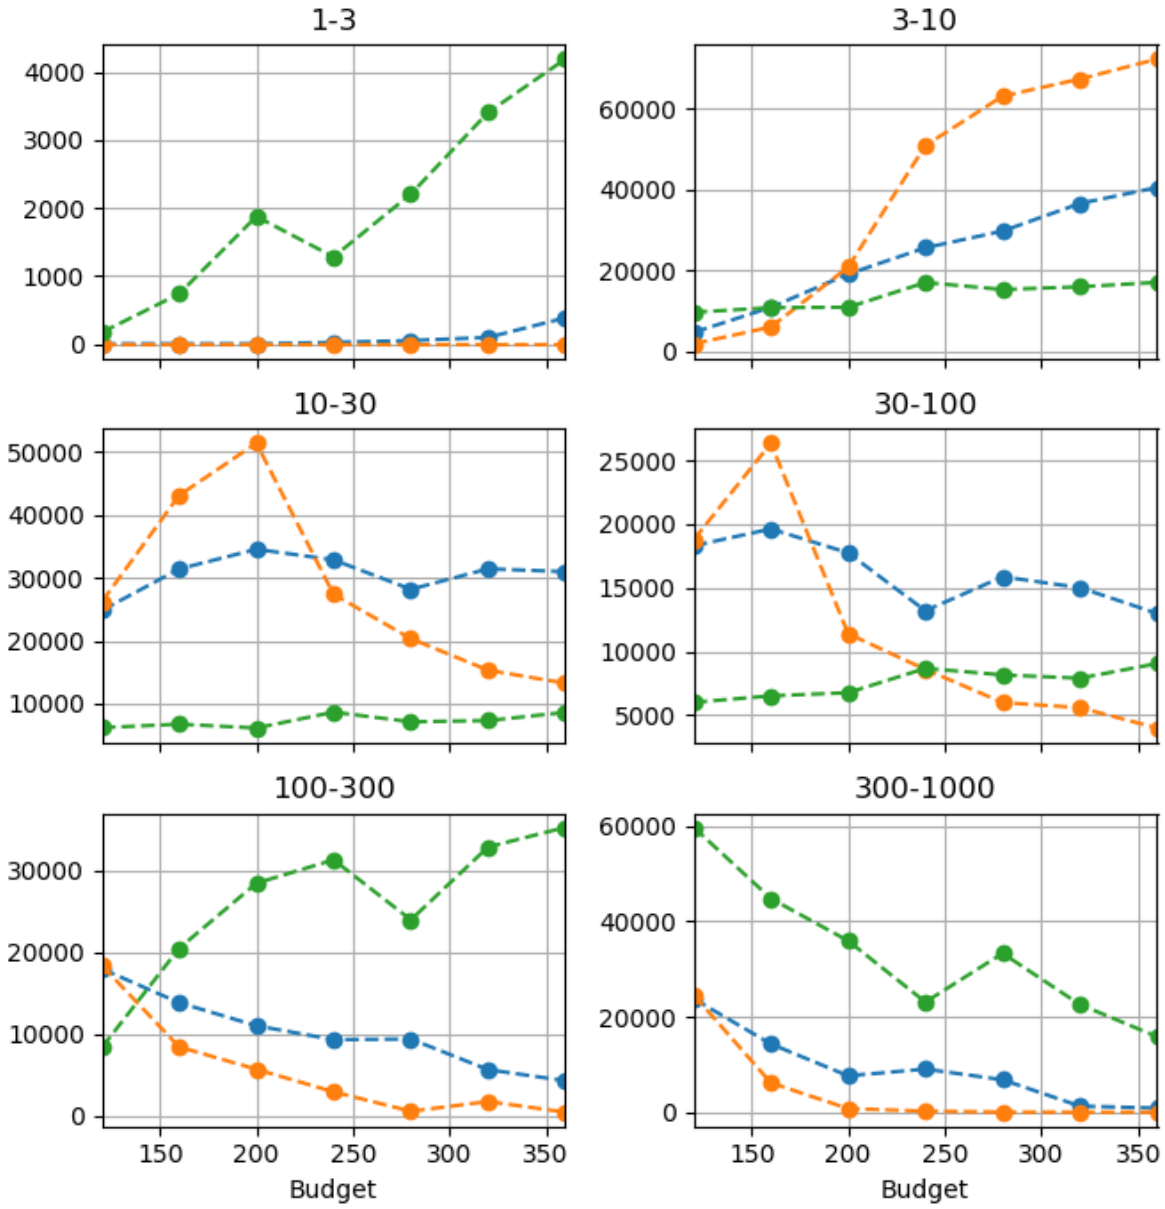

Figure 4.28: Frequencies of each histogram interval as a function of the budget. The informative path planner results are in blue, the boustrophedon path results are in orange, and the Brownian motion path results are in green. The histogram intervals are specified on top of each graph.

5

Conclusions

Contents

5.1 Work Accomplished 72
5.2 Future work 73

5.1 Work Accomplished

CBRN threats are a serious issue in the modern world. Radiological and nuclear threats arise from the use of radioactive materials, and several deadly incidents have occurred since middle of the twentieth century. Sophisticated radiological monitoring technologies can reduce the risk of these threats, both in the prevention phase and in the emergency response phase. Moreover, the development of autonomous robotic solutions reduces the risk toward humans even more by keeping them away from harmful radiation when the monitoring tasks are carried out.

For this purpose, an adaptive informative path planner was designed, implemented, and tested. The path planner is grounded on maximizing mutual information with an additional reward for exploring regions where the field being surveyed has large values. As such, it balances two desirable behaviors: executing a complete coverage path over the workspace being surveyed, and exploring the space around intense radioactive sources in more detail.

The proposed path planner is capable of addressing radiological intensity field survey problems and, on average, outperforms other paths with respect to locating radioactive sources. It is capable of solving problems in finite workspaces with budget constraints to limit the path length and progress constraints to ensure minimal coverage goals. The path planner makes a trade-off between exploration and exploitation, balancing exploring high field magnitude regions versus increasing the total coverage area. When compared to Brownian motion paths, the informative path planner performs better in every metric which was assessed. Additionally, when compared to boustrophedon paths (which are the standard for offline coverage path planning) the proposed solution results in better localization of radioactive sources—especially in scenarios containing a single source. However, with respect to estimating source activity, it performs worse or on par with boustrophedon paths, on average. Moreover, in situations containing sources whose activity differs by two or more orders of magnitude, the informative path planner will prioritize exploiting the region around the more intense source, leading to few observations around the weak source. Nevertheless, considering that satisfactory radioactive source activity quantification requires considerable lengths of measurement time, the path planner can be used to accurately locate sources and, subsequently, a more sensitive radiological sensor can be placed at the estimated source locations for measuring the activity.

Finally, this path planner was applied to radiological field inspection notwithstanding its applicability to surveys of other scalar fields, given the appropriate parameter modification. For example, it could be used for mapping the intensity of WiFi signals.

5.2 Future work

Field experiments were not possible at the time of writing this dissertation due to logistical reasons. This presents the first opportunity for future work, as an important step towards demonstrating the applicability of the proposed solution in real world scenarios.

Several other avenues for future work are evident; however two are probably the most pertinent: (1) to test the path planner in 3D scenarios, that is, to generate paths which can also change in altitude, since in this work we solely considered exploring the xy plane—in principle the algorithm should work, seeing as the problem formulation does not have any specific requirements for 2D workspaces, and the source model already accounts for radiation propagation in three dimensions; and (2) to incorporate agent dynamics in the objective function—one step in that direction would be to generate smooth paths, eliminating the sharp turns which can occur with the proposed path planner. Additionally, different reward functions can be tested, to incentivize the path planner to explore not only in the highest magnitude regions, but to explore all high magnitude regions above a certain threshold, for example. Another interesting opportunity for future work is to incorporate obstacle avoidance capabilities into the path planner; although the radiological heatmap simulator has the ability to simulate rectangular obstacles which absorb the simulated radiation, the development of obstacle avoidance in the path planner was beyond the scope of this dissertation.

References

[1] F. Su, I. Anthony, S. Toprak, E. Dinu, J. Ballard, T. Abe, A. Sellstrom, K. Arie, T. Kim, and D. Santoro, *Reassessing CBRN Threats in a Changing Global Environment*, F. Su and I. Anthony, Eds. Stockholm: SIPRI, 2019. [Online]. Available: https://www.sipri.org/publications/2019/other-publications/reassessing-cbrn-threats-changing-global-environment

[2] Centre for the Protection of National Infrastructure, "Chemical, Biological, Radiological and Nuclear (CBRN) Threats," 2017. [Online]. Available: https://www.cpni.gov.uk/chemical-biological-radiological-and-nuclear-cbrn-threats-0

[3] United Nations Interregional Crime and Justice Research Institute, "CBRN Risk Mitigation and Security Governance Programme," 2004. [Online]. Available: http://www.unicri.it/topics/cbrn/

[4] World Health Organization, "CHERNOBYL at 25th anniversary - Frequently Asked Questions," 2011. [Online]. Available: https://www.who.int/ionizing_radiation/chernobyl/20110423_FAQs_Chernobyl.pdf

[5] ——, "1986-2016: CHERNOBYL at 30 - An update," 2016. [Online]. Available: https://www.who.int/ionizing_radiation/chernobyl/Chernobyl-update.pdf

[6] ——, "FAQs: Japan Nuclear Concerns," 2011. [Online]. Available: https://www.who.int/hac/crises/jpn/faqs/en/

[7] ——, "FAQs: Fukushima Five Years On," 2016. [Online]. Available: https://www.who.int/ionizing_radiation/a_e/fukushima/faqs-fukushima/en/

[8] United Nations Scientific Committee on the Effects of Atomic Radiation, *UNSCEAR 2008 Report: Sources and Effects of Ionizing Radiation*. New York: UNSCEAR, 2010, vol. 1. [Online]. Available: http://www.unscear.org/unscear/en/publications/2008_1.html

[9] United States Nuclear Regulatory Commission, "Biological Effects of Radiation," in *Reactor Concepts Manual*. USNRC, 2003, ch. 9. [Online]. Available: https://www.nrc.gov/reading-rm/basic-ref/students/for-educators/09.pdf

[10] ——, "Backgrounder on the Three Mile Island Accident," 2018. [Online]. Available: http://www.nrc.gov/reading-rm/doc-collections/fact-sheets/3mile-isle.html

[11] J. Fortin, “Three Mile Island Nuclear Power Plant Is Shutting Down,” New York, p. 17, may 2019. [Online]. Available: https://www.nytimes.com/2019/05/08/us/three-mile-island-shut-down.html

[12] R. Pereira, S. Barbosa, and F. Carvalho, “Uranium mining in Portugal: a review of the environmental legacies of the largest mines and environmental and human health impacts,” *Environmental Geochemistry and Health*, vol. 36, no. 2, pp. 285–301, 2014. [Online]. Available: https://doi.org/10.1007/s10653-013-9563-6

[13] A. L. Campos, “Minas desactivadas em Portugal e a sua radioactividade: perigos e alertas,” 2017. [Online]. Available: https://nationalgeographic.sapo.pt/natureza/grandes-reportagens/1205-minas-desactivadas-em-portugal-e-a-sua-radioactividade-perigos-e-alertas

[14] C. Soldado, “A exploração de urânio “foi um erro em Portugal e será um erro em Espanha”,” Lisbon, feb 2018. [Online]. Available: https://www.publico.pt/2018/02/20/sociedade/noticia/a-exploracao-de-uranio-foi-um-erro-em-portugal-e-sera-um-erro-em-espanha-1803726

[15] F. Naves, “Mina de urânio vai utilizar a água de um afluente do rio Douro,” Lisbon, feb 2018. [Online]. Available: https://www.dn.pt/sociedade/interior/mina-de-uranio-vai-utilizar-a-agua-de-um-afluente-do-rio-douro-9142940.html

[16] United States Nuclear Regulatory Commission, “NRC: Orphan Sources,” 2017. [Online]. Available: https://www.nrc.gov/materials/miau/miau-reg-initiatives/orphan.html

[17] International Atomic Energy Agency, *IAEA Safety Glossary: 2018 Edition*. Vienna: IAEA, 2019. [Online]. Available: https://www.iaea.org/publications/11098/iaea-safety-glossary-2018-edition

[18] ——, *The Radiological Accident in Goiânia*. Vienna: IAEA, 1988. [Online]. Available: https://www.iaea.org/publications/3684/the-radiological-accident-in-goiania

[19] ——, *The Radiological Accident in Samut Prakarn*. Vienna: IAEA, 2002. [Online]. Available: https://www.iaea.org/publications/6375/the-radiological-accident-in-samut-prakarn

[20] L. Lodding, “Drop it and Run!” *IAEA Bulletin*, vol. 48, no. 2, pp. 70–72, 2007. [Online]. Available: https://www.iaea.org/publications/magazines/bulletin/48-2

[21] D. Dahlstrom, “New Symbol Launched to Warn Public About Radiation Dangers,” 2007. [Online]. Available: https://www.iaea.org/newscenter/news/new-symbol-launched-warn-public-about-radiation-dangers

[22] R. Kumar, G. Panda, B. Singh, D. Rane, J. Sunil Kumar, and A. Sonawane, “Lessons Learned from the Radiological Accident in Mayapuri, New Delhi,” in *Safety and Security of Radioactive Sources: Maintaining Continuous Global Control of Sources throughout Their Life Cycle*. Vienna: IAEA, 2015, pp. 517–528. [Online]. Available: https://www.iaea.org/publications/10678/safety-and-security-of-radioactive-sources-maintaining-continuous-global-control-of-sources-throughout-their-life-cycle

[23] International Atomic Energy Agency, *The Radiological Accident in Lilo*. Vienna: IAEA, 2000. [Online]. Available: https://www.iaea.org/publications/5968/the-radiological-accident-in-lilo

[24] R. Stone, "The Hunt for Hot Stuff," *Smithsonian Magazine*, mar 2003. [Online]. Available: https://www.smithsonianmag.com/science-nature/the-hunt-for-hot-stuff-76764125/

[25] International Atomic Energy Agency, "Radioactive Sources Recovered in Georgia," 2006. [Online]. Available: https://www.iaea.org/newscenter/news/radioactive-sources-recovered-georgia

[26] ——, *The Radiological Accident in Lia, Georgia*. Vienna: IAEA, 2014. [Online]. Available: https://www.iaea.org/publications/10602/the-radiological-accident-in-lia-georgia

[27] United States Nuclear Regulatory Commission, "Backgrounder on Dirty Bombs," 2018. [Online]. Available: https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/fs-dirty-bombs.html

[28] C. D. Ferguson, "Reducing the threat of RDDs," *IAEA Bulletin*, vol. 45, no. 1, pp. 12–16, 2003. [Online]. Available: https://www.iaea.org/publications/magazines/bulletin/45-1

[29] M. Lozano and A. Fernandez, "Reward Project." [Online]. Available: http://www.reward-project.eu/

[30] Seagull Environmental Technologies Inc., "Seagull Environmental Technologies," 2018. [Online]. Available: http://www.seagullenvirotech.com/

[31] Centro de Investigação da Academia da Força Aérea, "SEAGULL," 2013. [Online]. Available: http://www.academiafa.edu.pt/subPagina-10D00-019.005.003.011-seagull

[32] M. M. Marques, P. Dias, N. P. Santos, V. Lobo, R. Batista, D. Salgueiro, A. Aguiar, M. Costa, J. E. Da Silva, A. S. Ferreira, J. Sousa, M. d. F. Nunes, E. Pereira, J. Morgado, R. Ribeiro, J. S. Marques, A. Bernardino, M. Grine, and M. Taiana, "Unmanned aircraft systems in maritime operations: Challenges addressed in the scope of the SEAGULL project," *OCEANS 2015 - Genova*, 2015. [Online]. Available: https://www.doi.org/10.1109/OCEANS-Genova.2015.7271427

[33] Grupo ISQ, "Gamma-Ex," 2017. [Online]. Available: https://www.isq.pt/projeto/gamma-ex/

[34] Orano DS Communications, "Press Release - Dismantling Robots 2018," apr 2018. [Online]. Available: https://www.orano.group/en/press/dossiers-de-presse/robots-dismantling-and-services

[35] D. F. Macedo, "Radioactive hot-spot detection using unmanned aerial vehicle surveillance," Master's dissertation, Universidade de Lisboa, Lisbon, 2018.

[36] G. Hitz, E. Galceran, M.-E. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017. [Online]. Available: https://doi.org/10.1002/rob.21722

[37] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley, 2006. [Online]. Available: https://www.doi.org/10.1002/047174882X

[38] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, “Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting,” *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012. [Online]. Available: https://doi.org/10.1109/TIT.2011.2182033

[39] B. Zhang and G. S. Sukhatme, “Adaptive Sampling for Estimating a Scalar Field using a Robotic Boat and a Sensor Network,” in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. Rome: IEEE, 2007, pp. 3673–3680. [Online]. Available: https://doi.org/10.1109/ROBOT.2007.364041

[40] C. Cai and S. Ferrari, “Information-Driven Sensor Path Planning by Approximate Cell Decomposition,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 3, pp. 672–689, 2009. [Online]. Available: https://doi.org/10.1109/TSMCB.2008.2008561

[41] L. Paull, S. Saeedi, M. Seto, and H. Li, “Sensor-Driven Online Coverage Planning for Autonomous Underwater Vehicles,” *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 6, pp. 1827–1838, 2013. [Online]. Available: https://doi.org/10.1109/TMECH.2012.2213607

[42] L. Paull, C. Thibault, A. Nagaty, M. Seto, and H. Li, “Sensor-Driven Area Coverage for an Autonomous Fixed-Wing Unmanned Aerial Vehicle,” *IEEE Transactions on Cybernetics*, vol. 44, no. 9, pp. 1605–1618, 2014. [Online]. Available: https://doi.org/10.1109/TCYB.2013.2290975

[43] R. Cui, Y. Li, and W. Yan, “Mutual Information-Based Multi-AUV Path Planning for Scalar Field Sampling Using Multidimensional RRT*,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 46, no. 7, pp. 993–1004, 2016. [Online]. Available: https://doi.org/10.1109/TSMC.2015.2500027

[44] A. A. R. Newaz, S. Jeong, and N. Y. Chong, “Fast radioactive hotspot localization using a UAV,” in *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)*. San Francisco: IEEE, 2016, pp. 9–15. [Online]. Available: https://doi.org/10.1109/SIMPAR.2016.7862348

[45] G. Ferri, A. Munafò, and K. D. LePage, “An Autonomous Underwater Vehicle Data-Driven Control Strategy for Target Tracking,” *IEEE Journal of Oceanic Engineering*, vol. 43, no. 2, pp. 323–343, 2018. [Online]. Available: https://doi.org/10.1109/JOE.2018.2797558

[46] K.-C. Ma, L. Liu, H. K. Heidarsson, and G. S. Sukhatme, “Data-driven learning and planning for environmental sampling,” *Journal of Field Robotics*, vol. 35, no. 5, pp. 643–661, 2018. [Online]. Available: https://doi.org/10.1002/rob.21767

[47] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013. [Online]. Available: https://doi.org/10.1016/j.robot.2013.09.004

[48] H. Choset and P. Pignon, “Coverage Path Planning: The Boustrophedon Cellular Decomposition,” in *Field and Service Robotics*, A. Zelinsky, Ed. London: Springer, London, 1998, pp. 203–209. [Online]. Available: https://doi.org/10.1007/978-1-4471-1273-0_32

[49] W. H. Huang, “Optimal line-sweep-based decompositions for coverage algorithms,” in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, vol. 1. Seoul: IEEE, 2001, pp. 27–32. [Online]. Available: https://doi.org/10.1109/ROBOT.2001.932525

[50] I. Maza and A. Ollero, “Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms,” in *Distributed Autonomous Robotic Systems 6*, R. Alami, R. Chatila, and H. Asama, Eds. Tokyo: Springer, Tokyo, 2007, pp. 221–230. [Online]. Available: https://www.doi.org/10.1007/978-4-431-35873-2_22

[51] H. Choset, “Coverage for robotics – A survey of recent results,” *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113–126, 2001. [Online]. Available: https://doi.org/10.1023/A:1016639210559

[52] C. Luo, S. X. Yang, D. A. Stacey, and J. C. Jofriet, “A solution to vicinity problem of obstacles in complete coverage path planning,” in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, vol. 1. Washington, D.C.: IEEE, 2002, pp. 612–617. [Online]. Available: https://doi.org/10.1109/ROBOT.2002.1013426

[53] Z. Shiller, “Off-Line and On-Line Trajectory Planning,” in *Motion and Operation Planning of Robotic Systems*, ser. Mechanisms and Machine Science, G. Carbone and F. Gomez-Bravo, Eds. Springer, Cham, 2015, vol. 29, pp. 29–62. [Online]. Available: https://doi.org/10.1007/978-3-319-14705-5_2

[54] B. Zhang, L. Tang, J. DeCastro, M. J. Roemer, and K. Goebel, “A Recursive Receding Horizon Planning for Unmanned Vehicles,” *IEEE Transactions on Industrial Electronics*, vol. 62, no. 5, pp. 2912–2920, 2015. [Online]. Available: https://doi.org/10.1109/TIE.2014.2363632

[55] A. Kim and R. M. Eustice, “Perception-driven navigation: Active visual SLAM for robotic area coverage,” in *2013 IEEE International Conference on Robotics and Automation*. Karlsruhe: IEEE, 2013, pp. 3196–3203. [Online]. Available: https://doi.org/10.1109/ICRA.2013.6631022

[56] Health Physics Society, “Radiation Basics,” 2013. [Online]. Available: https://hps.org/publicinformation/ate/faqs/radiation.html

[57] International Atomic Energy Agency, “Radiation in Everyday Life,” 2015. [Online]. Available: https://www.iaea.org/Publications/Factsheets/English/radlife

[58] D. Close and L. Ledwidge, “Measuring Radiation: Terminology and Units,” *Science for Democratic Action*, vol. 8, no. 4, pp. 8–10, 2000. [Online]. Available: https://ieer.org/article/science-for-democratic-action/volume-8-number-4/

[59] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, ser. Adaptive Computation and Machine Learning. The MIT Press, 2006. [Online]. Available: https://mitpress.mit.edu/books/gaussian-processes-machine-learning

[60] M. Morelande, B. Ristic, and A. Gunatilaka, "Detection and parameter estimation of multiple radioactive sources," in *2007 10th International Conference on Information Fusion*. Quebec: IEEE, 2007, pp. 1–7. [Online]. Available: https://doi.org/10.1109/ICIF.2007.4408094

[61] A. Gunatilaka, B. Ristic, and R. Gailis, "On Localisation of a Radiological Point Source," in *2007 Information, Decision and Control*. Adelaide: IEEE, 2007, pp. 236–241. [Online]. Available: https://doi.org/10.1109/IDC.2007.374556

[62] A. Sitek and A. M. Celler, "Limitations of Poisson statistics in describing radioactive decay," *Physica Medica*, vol. 31, no. 8, pp. 1105–1107, 2015. [Online]. Available: https://doi.org/10.1016/j.ejmp.2015.08.015

[63] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, ser. Intelligent Robotics and Autonomous Agents. The MIT Press, 2005. [Online]. Available: https://mitpress.mit.edu/books/probabilistic-robotics

[64] G. W. Nelson and T. D. Reilly, "Gamma-Ray Interactions with Matter," in *Passive Nondestructive Assay of Nuclear Materials*. USNRC, 1991, ch. 2, pp. 27–42. [Online]. Available: https://www.osti.gov/biblio/5428834-passive-nondestructive-assay-nuclear-materials

[65] N. Hansen, "The CMA Evolution Strategy: A Comparing Review," in *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, ser. Studies in Fuzziness and Soft Computing, J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea, Eds. Springer, Berlin, Heidelberg, 2006, pp. 75–102. [Online]. Available: https://doi.org/10.1007/3-540-32494-1_4

[66] United States Nuclear Regulatory Commission, "The Nuclear Regulatory Commission's Science 101: What is a Geiger Counter?" 2019. [Online]. Available: https://www.nrc.gov/reading-rm/basic-ref/students/science-101/what-is-a-geiger-counter.html

[67] International Atomic Energy Agency, *In Situ Analytical Characterization of Contaminated Sites Using Nuclear Spectrometry Techniques*, ser. Analytical Quality in Nuclear Applications Series. Vienna: IAEA, 2017, no. 49. [Online]. Available: https://www.iaea.org/publications/12227/in-situ-analytical-characterization-of-contaminated-sites-using-nuclear-spectrometry-techniques

[68] K. Sigman, "Simulating Brownian motion (BM) and geometric Brownian motion (GBM)," in *Professor Karl Sigman's Lecture Notes on Simulation*. New York: Columbia University, 2013, ch. 7. [Online]. Available: http://www.columbia.edu/~ks20/4404-Sigman/Simulation-Sigman.html

A

MARIA Description

The first version of the application simply read the GMC readings and displayed them on the device screen. The user could start, pause, and stop a measurement, and the data was saved to a file. Currently, the application is in version 2, with many more features (see fig. A.1), such as:

- A splashscreen with information about how the application works, the version number, the IPFN logo which also behaves like a button to open the IPFN homepage, and a button to begin making measurements.

- Automatic detection of a connected sensor; if none in found, then the user is given the option to use a simulator which draws random samples from a Poisson distribution (since, for conventional detectors, the number of radioactive events which are detected in a given time interval follows Poisson statistics almost perfectly [62]).

- An improved tabular view of the past 5 sensor readings, displaying the measured value in particle counts and in counts per minute, as well as the time at which the sensor reading occurred.

- A map view which displays measurements at the location where they were made over Google Maps™; each measurement has a color depending on its value, in a gradient ranging from green (low intensity) to orange (moderate intensity) to red (high intensity); furthermore, the cap for what counts as “low” and what counts as “high” can be configured by the user in the options menu (this

is to avoid situations where all reading may be very low or very high, and thus appear indistinguishable on the screen).

- An options menu to allow the user to configure the reading interval, open the file directory containing all past measurement data files, take a screenshot, switch between map mode and sensor readout mode, and, while in map mode, the user is additionally given the option to load a previous measurement file, which will read the relevant data and display the readings over the map.

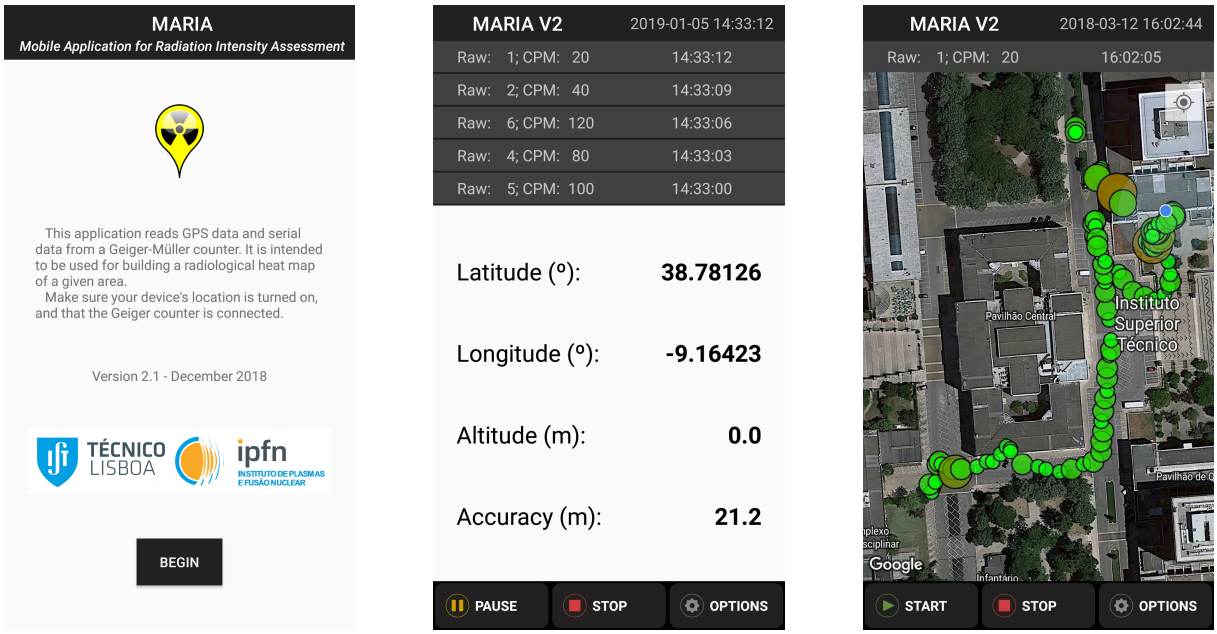

(a) (b) (c)

Figure A.1: Screenshots of the mobile application MARIA. **(a)** The application splashscreen, with the logo, some useful information and a button. **(b)** The application in sensor mode, displaying the past 5 readings and the current positional data. **(c)** The application in map mode. The circles represent measurements, with green being low intensity values and a few orange circles being moderate intensity values. The location is given by the device GNSS module and the measurements are given by a GMC connected through the USB port. The user's current location is represented with the small blue circle.

Since its inception, MARIA has already made some noteworthy achievements. In October 2018 MARIA was selected as one of the 12 finalists (out of 50 candidates) of the 2018 "This is APPStation Challenge," a competition to select the most innovative mobile application organized by Ábaco Consultores. In April 2019 MARIA was selected as one of the 5 finalists of Lab2Market@Técnico, a technology transfer program sponsored by Everis and i-Deals granting 50 hours of business consulting to each of the 5 finalist projects. Additionally, there are plans to integrate MARIA into a project named FRIENDS. FRIENDS (Fleet of Drones for Radiological Inspection, Communication and Rescue) is a collaboration between IST, IPFN, C^{2}TN, LARSyS, ISR, Universidade de Aveiro, IT, and ISCTE-IUL. The goal of this project is to develop a solution for real-time mapping of a scenario and the radiological contamination level of a given area, as well as autonomous and cooperative navigation of a fleet of drones to cover this area in an optimal way. Since project FRIENDS will use a variety of sensors whose data needs to be mapped, it is anticipated that MARIA will have added support for the respective sensors (currently, the application supports the SEN-11345 and PRM-9000 GMCs).

B

Absorbed Radiation Dose Examples

Several examples of absorbed radiation dose values are illustrated in figs. B.1 to B.3.

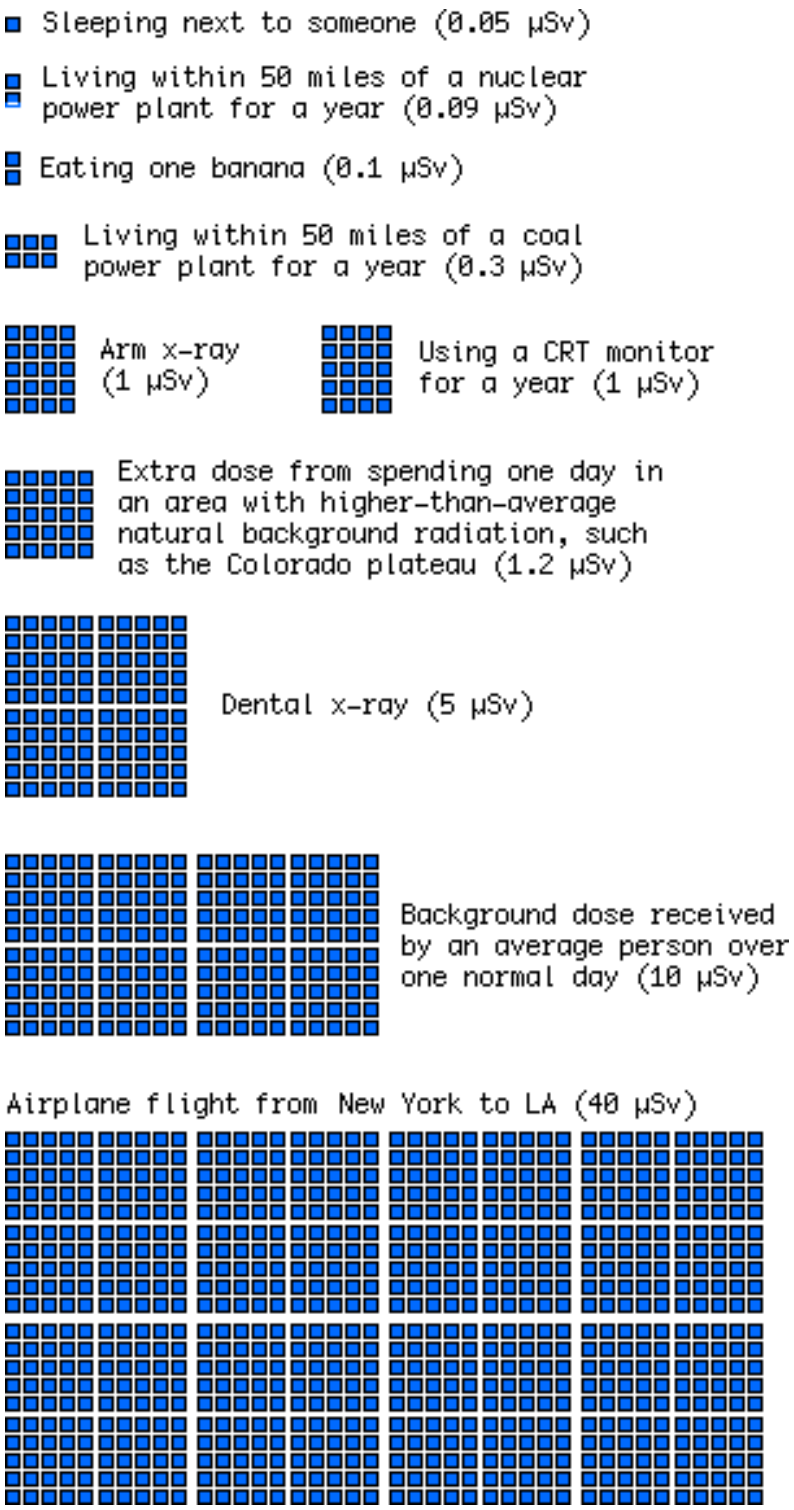

Figure B.1: Visual representation of the typical radiation doses which are absorbed in situations which most civilians might experience.
Credit: Randall Munroe.

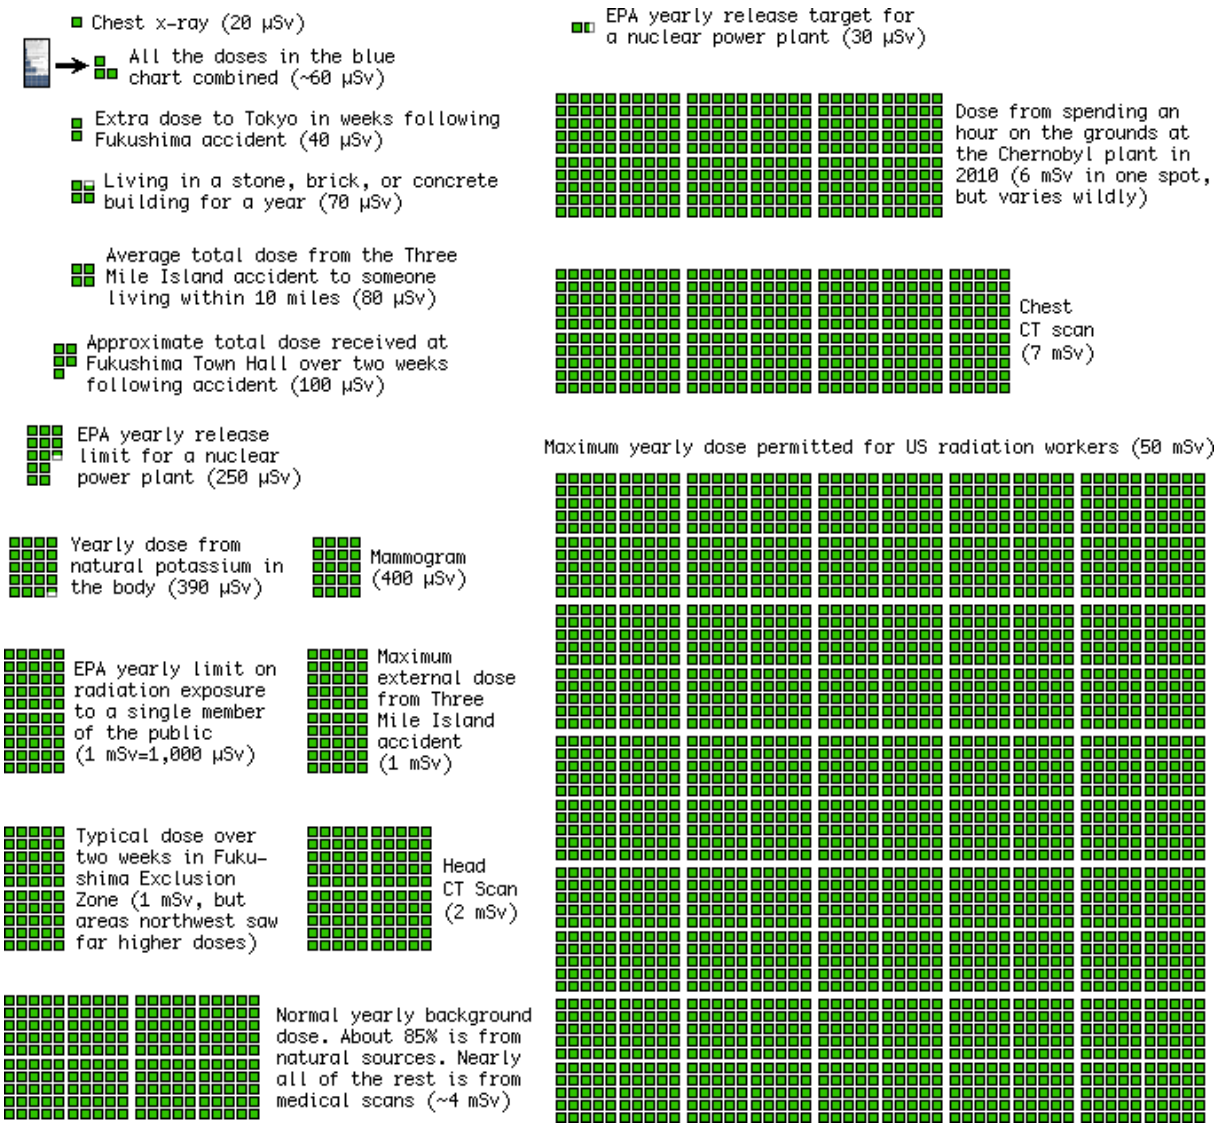

Figure B.2: Visual representation of the moderate radiation doses which may be absorbed in several real-world situations.
Credit: Randall Munroe.

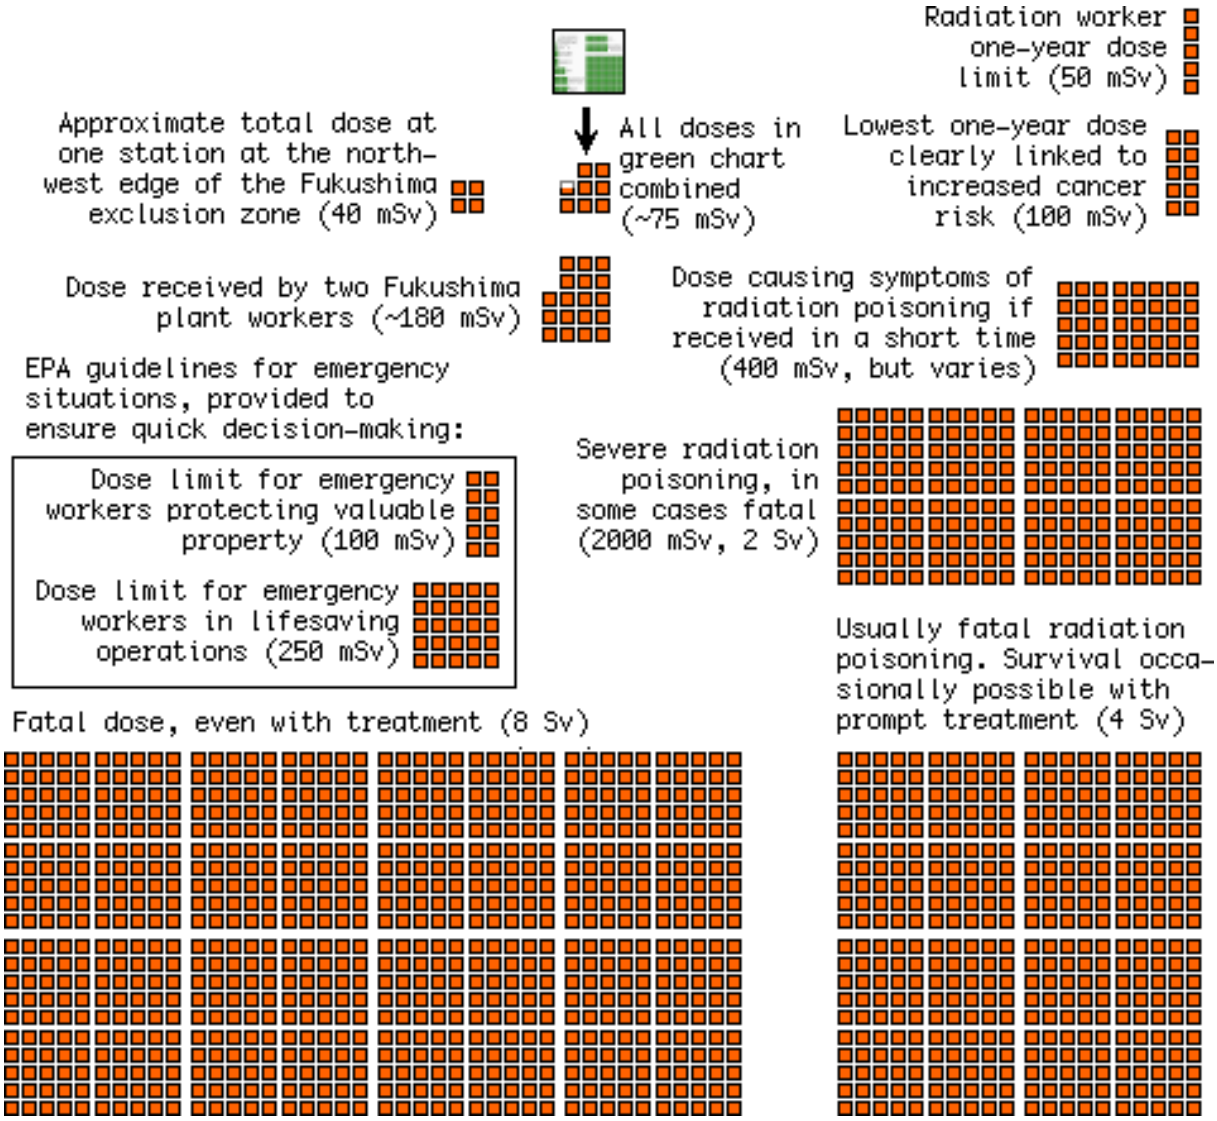

Figure B.3: Visual representation of high radiation doses which humans may be exposed to at radiological disaster sites.
Credit: Randall Munroe.