Emergency rapid mapping with drones

Models and solution approaches for offline and online mission planning

Zur Erlangung des akademischen Grades eines

Doktors der Ingenieurwissenschaften

von der KIT-Fakultät für Wirtschaftswissenschaften
am Karlsruher Institut für Technologie (KIT)
genehmigte

Dissertation

von

Katharina Diana Glock, M.Sc.

Tag der mündlichen Prüfung: 04.09.2020
Hauptreferent: Prof. Dr. Stefan Nickel
Korreferent: JProf. Dr.-Ing. Anne Meyer

Abstract

The increasing availability of unmanned aerial vehicles (UAVs) and the progress in the development of lightweight sensor systems have opened new possibilities for emergency surveillance. After chemical emergencies or large fires, rotary-wing UAVs can provide emergency response units with a quick first assessment of the contaminated area and the affected population. However, unless the affected area is small, the limited flight time and the time-sensitive situation mean that UAVs can only take a limited number of samples. Combined with spatial interpolation methods, these samples allow an estimate of the distribution of contaminants and hazardous substances across the affected area.

This thesis addresses the problem of planning UAV missions that maximize the information gain in emergency surveillance. We study this planning problem both in an offline fashion, where missions are determined before takeoff, and in the online variant where plans are updated while the UAVs are in flight. Our main interest lies in the design of efficient models and solution heuristics that use information about the spatial correlations between observed and unobserved locations to determine high-quality solutions with low prediction errors in time-sensitive situations.

For offline planning, we introduce the generalized correlated team orienteering problem and propose an efficient two-stage solution approach for quickly determining explorative UAV missions. In an extensive study, we confirm the competitiveness of our heuristic in terms of computation time and solution quality compared to state-of-the-art approaches for the team orienteering problem. We introduce new benchmark instances and show the higher information gain of our model compared to related concepts. We demonstrate that it is possible to achieve reliable predictions within computation times that allow the use of these approaches in practice.

In the context of online planning, we study the adaptive mission planning problem, which combines online learning of the surveyed contamination and online planning of UAV missions. We classify a wide range of solution concepts from different disciplines and complement them with new modeling variants. We demonstrate that computationally efficient models yield high-quality solutions within very short computation times, thus enabling their use in emergency surveillance. Furthermore, we show the added value of more accurate but computationally expensive models and solution concepts in case of uncertain initial beliefs and strongly varying processes.

Kurzfassung

Die Verfügbarkeit von unbemannten Luftfahrzeugen (*unmanned aerial vehicles* oder UAVs) und die Fortschritte in der Entwicklung leichtgewichtiger Sensorik eröffnen neue Möglichkeiten für den Einsatz von Fernerkundungstechnologien zur Schnellerkundung in Großschadenslagen. Hier ermöglichen sie es beispielsweise nach Großbränden, Einsatzkräften in kurzer Zeit ein erstes Lagebild zur Verfügung zu stellen. Die begrenzte Flugdauer der UAVs wie auch der Bedarf der Einsatzkräfte nach einer schnellen Ersteinschätzung bedeuten jedoch, dass die betroffenen Gebiete nur stichprobenartig überprüft werden können. In Kombination mit Interpolationsverfahren ermöglichen diese Stichproben anschließend eine Abschätzung der Verteilung von Gefahrstoffen.

Die vorliegende Arbeit befasst sich mit dem Problem der Planung von UAV-Missionen, die den Informationsgewinn im Notfalleinsatz maximieren. Das Problem wird dabei sowohl in der Offline-Variante, die Missionen vor Abflug bestimmt, als auch in der Online-Variante, bei der die Pläne während des Fluges der UAVs aktualisiert werden, untersucht. Das übergreifende Ziel ist die Konzeption effizienter Modelle und Verfahren, die Informationen über die räumliche Korrelation im beobachteten Gebiet nutzen, um in zeitkritischen Situationen Lösungen von hoher Vorhersagegüte zu bestimmen.

In der Offline-Planung wird das *generalized correlated team orienteering problem* eingeführt und eine zweistufige Heuristik zur schnellen Bestimmung explorativer UAV-Missionen vorgeschlagen. In einer umfangreichen Studie wird die Leistungsfähigkeit und Konkurrenzfähigkeit der Heuristik hinsichtlich Rechenzeit und Lösungsqualität bestätigt. Anhand von in dieser Arbeit neu eingeführten Benchmarkinstanzen wird der höhere Informationsgewinn der vorgeschlagenen Modelle im Vergleich zu verwandten Konzepten aufgezeigt.

Im Bereich der Online-Planung wird die Kombination von lernenden Verfahren zur Modellierung der Schadstoffe mit Planungsverfahren, die dieses Wissen nutzen, um Missionen zu verbessern, untersucht. Hierzu wird eine breite Spanne von Lösungsverfahren aus unterschiedlichen Disziplinen klassifiziert und um neue effiziente Modellierungsvarianten für die Schnellerkundung ergänzt. Die Untersuchung im Rahmen einer ereignisdiskreten Simulation zeigt, dass vergleichsweise einfache Approximationen räumlicher Zusammenhänge in sehr kurzer Zeit Lösungen hoher Qualität ermöglichen. Darüber hinaus wird die höhere Robustheit genauerer, aber aufwändigerer Modelle und Lösungskonzepte demonstriert.

Contents

Part I

Foundations

1 Introduction

"Essentially, all models are wrong, but some are useful."
— George Box[1]

In case of emergencies, such as wide-spread fires or accidents involving hazardous gases, the first step toward an effective emergency response is to provide emergency response units with detailed information about the nature and scale of the situation at hand. The quicker this information is provided, the faster decisions can be made concerning the appropriate protective equipment, the warning and evacuation of the potentially affected population, and the effective deployment of response units.

The process of acquiring, combining and analyzing current geospatial information on-demand immediately after an event has occurred is often referred to as rapid mapping. The European Space Agency (ESA) has specifically addressed emergency rapid mapping in its Copernicus program as part of the Sentinel-2 mission. However, up-to-date satellite data of an affected area may only be available within a delay of several hours or even days. This is sufficient for incidents such as floods, where vast areas are affected and need to be monitored. For the majority of emergencies faced by fire services and other emergency response units, however, satellite data are of limited help.

Small unmanned aerial vehicles (UAV) offer a flexible way of providing human decision-makers with an overview of the situation. They are particularly promising for emergencies of a more local scale, i.e., affecting several hundred meters to a few kilometers, compared to the magnitude of natural disasters primarily addressed by the Copernicus program. The degree of contamination can change quickly in space as well as over time and is influenced by wind turbulences and local geographic features such as different terrain heights. In these scenarios, UAVs constitute a major asset for the rapid assessment of the contamination within the immediately affected area.

However, due to the limited flight time of the UAVs and the desire of the response personnel to quickly obtain a reliable overview of the situation, the complete coverage of the affected area is often impossible. In this case, observations can only be made at

[1] Box, G. and N. Draper (1987). *Empirical model-building and response surfaces*. John Wiley & Sons.

a limited number of sampling locations. Taking advantage of the fact that the distribution of gases and contaminants is usually positively spatially autocorrelated, i.e., similar values can be observed at locations close to each other, these samples yield an estimation of the extent and severity of the contamination. This raises the question of how the available UAVs should be deployed to obtain samples that provide as much information as possible.

1.1 Scope of this thesis

Mission planning in the context of emergency rapid mapping revolves around selecting and sequencing sampling locations for one or several UAVs such that the overall information obtained within the limited time frame is maximized. The ultimate goal is to achieve a minimal prediction error at critical locations. This, however, is intractable in practice, as it is impossible to assess the prediction error before the mission. Even after all observations are completed, the actual estimation error cannot be measured, as the "true" distribution is unknown. For this reason, models and planning approaches in this context use a wide range of approximations that have been shown to lead to good estimations, ranging from simple strategies for maximizing spatial coverage to probabilistic models for minimizing prediction uncertainty.

In this thesis, we study offline and online concepts for planning informative UAV missions for emergency surveillance. These two planning variants are depicted in Figure 1.1. Both are initialized with information about priorities in the affected area, which allows the emergency response units to specify locations at which accurate information is particularly crucial, e.g., because a large number of people may be affected. Additionally, these variants use some a-priori information about the spatial interdependencies in the surveyed process, for example, the range of the correlation. This information enables the selection of sampling locations accounting for the fact that observations made nearby will yield similar values, and therefore are less informative compared to observations spread over a larger area.

In offline planning, missions are determined in an a-priori-fashion, independently of the actual observations and based on initial expert beliefs about the spatial correlation. That is, they determine complete sequences of sampling locations for the UAVs to execute. A prediction is computed based on the final set of samples after the mission is completed. Online approaches, in contrast, do not compute entire missions before takeoff. Instead, they maintain some environment representation, which is updated while the mission is still in progress. New sampling locations are determined based on the currently available information. These approaches are also referred to as *adaptive* if they adjust mission based on the data collected during the flight.

Considering the time-sensitivity of real-world applications, our main interest lies in the design of efficient models and solution heuristics for emergency rapid mapping. The main contributions toward this goal can be summarized as follows:

1. We derive new models for representing the information provided by the sampling

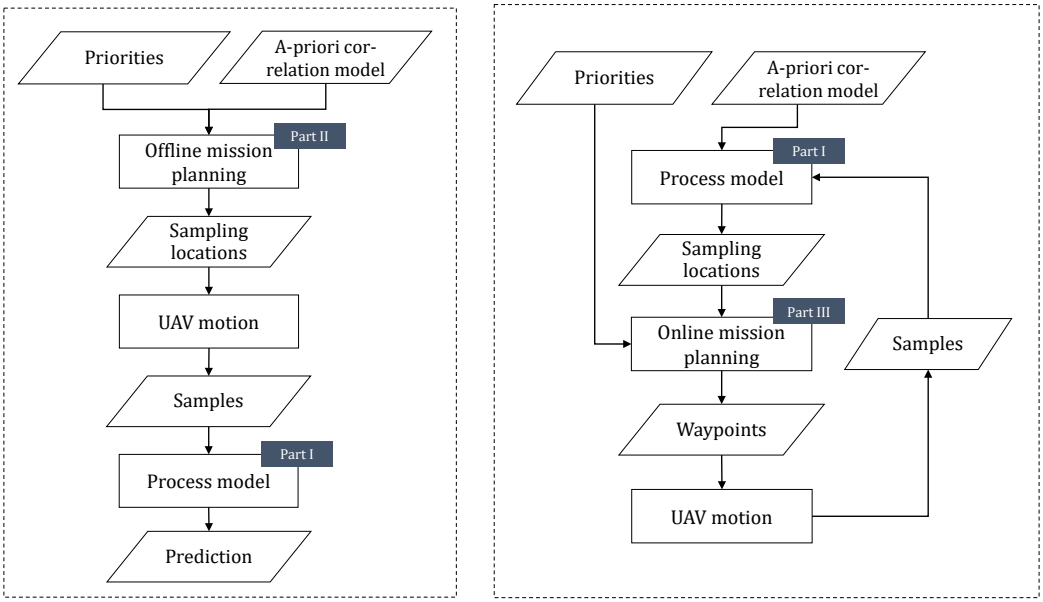

(A) Offline mission planning (B) Online mission planning

FIGURE 1.1: Offline and online mission planning approaches.

locations during a mission.

2. We integrate these models into heuristics tailored toward the efficient planning of multi-UAV missions.

3. We investigate the capabilities and limitations of these models and solution approaches in extensive simulative studies based on newly introduced benchmark instances modeled after a real-world use case.

1.2 Outline

This thesis is divided into four parts that are depicted in Figure 1.2.

Part I - Foundations introduces the technical and theoretical principles on which this work is based. Chapter 2 discusses existing approaches for rapid mapping and situation assessment. The chapter furthermore gives an overview of the technologies that enable UAV-based surveillance systems, e.g., recent progress in sensor technologies. In Chapter 3, we summarize concepts for modeling spatial structure, interpolating spatial processes, and assessing the quality of the predictions.

Part II - Offline mission planning focuses on a-priori planning of informative UAV missions. In Chapter 4, we introduce the mission planning problem for emergency surveillance (MPPES). Chapter 5 gives an overview of related approaches in different fields of research, focusing on modeling variants for incorporating spatial interdependencies within the considered area. Based on this overview, we discuss the gaps remaining in

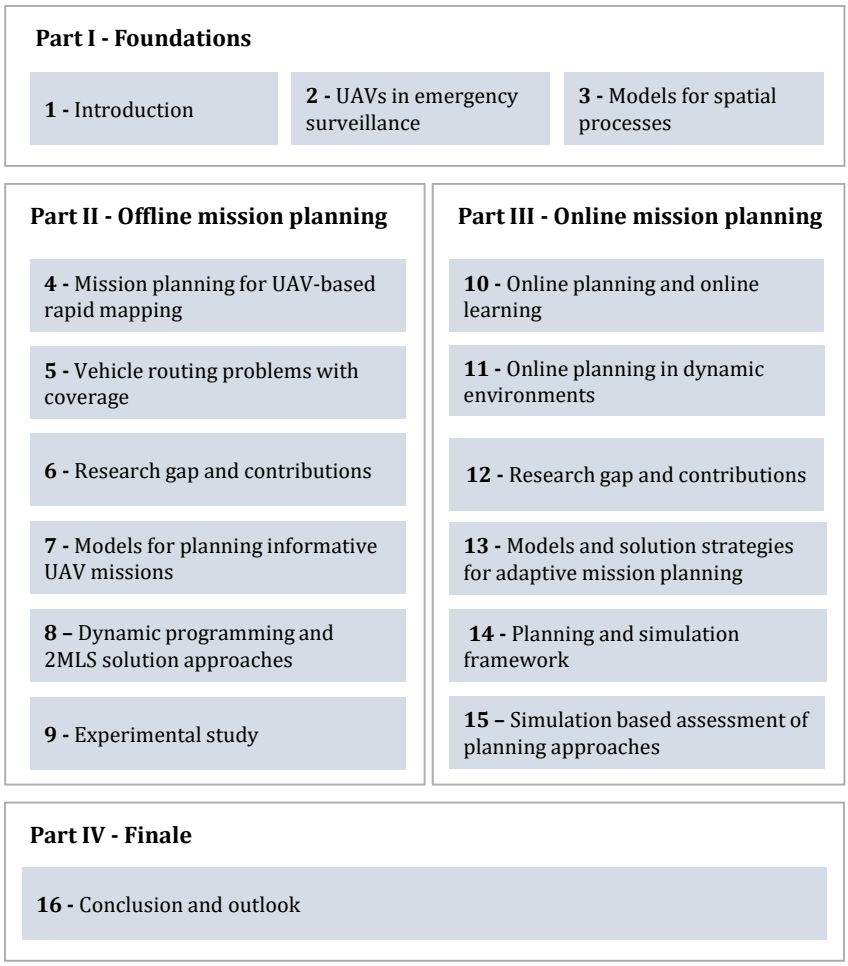

FIGURE 1.2: Organization of this thesis.

literature and summarize our contributions in Chapter 6. The following two chapters are dedicated to our solution approach: In Chapter 7, we propose the generalized correlated team orienteering problem (GCorTOP), a model for solving the MPPES that accounts for spatial interdependencies as well as priorities within a target area. In Chapter 8, we propose an exact approach for benchmarking purposes and introduce a two-phase multi-start adaptive large neighborhood search (2MLS) for quickly obtaining high-quality solutions for larger scenarios. In Chapter 9, we introduce new benchmark instances for the MPPES and assess the performance of our approaches.

Part III - Online mission planning studies the adaptive mission planning problem for emergency surveillance (AMPPES). We introduce the problem and summarize desirable requirements for successful solution approaches in Chapter 10. In Chapter 11, we give a detailed overview of approaches for online planning in dynamic environments proposed in the field of environmental surveillance and discuss methods developed in other disciplines. In Chapter 12, we review the findings with respect to the initially stated re-

quirements. We derive the research gap and summarize our contributions. In Chapter 13, we introduce a unified modeling approach for representing the value of sampling locations in an online setting. Based on this unified model, we adapt existing strategies and propose new methods for solving the AMPPES. These strategies are implemented within a common solution architecture that is developed in Chapter 14. This chapter also introduces a simulation framework for evaluating and comparing the different strategies. The results of these evaluations are presented and discussed in Chapter 15.

Part IV - Synthesis concludes this thesis with Chapter 16. We summarize our main findings and contributions, discuss the results, and suggest future lines of research.

2 Unmanned aerial vehicles in emergency surveillance

This chapter provides provides an overview of the technologies studied within the course of the BigGIS research project[1], which aims at developing new solutions for the effective use of UAVs in emergency services. To introduce the core concepts of emergency rapid mapping, we first establish the role of surveillance and reconnaissance in emergency management in general and fire services in particular. We furthermore summarize existing guidelines for surveying and sampling in the case of widespread hazards. We briefly characterize related technological aspects, specifically concerning the UAV and sensor systems, and conclude with an overview of existing mission planning support.

2.1 Surveillance in emergency management

While the proposed UAV-based surveillance concept itself is new, situation assessment and reconnaissance in general play an important role in emergency management. Established procedures are summarized in this section, which focuses in particular on situation assessment and emergency sampling procedures for fire services.

2.1.1 The emergency management cycle

Emergency management is a term denoting a broad set of organizations and functions that are concerned with the prevention of emergencies as well as response, rescue and reconstruction operations (Haddow et al., 2017). The main purpose of emergency management is best shown at the hand of the emergency response cycle, a general model of how societies and governments deal with emergencies, both before an incident and in the short-term and long-term aftermath. According to this model, emergency management consists of four phases: mitigation, preparedness, response, and recovery. Mitigation includes all actions seeking to prevent an emergency and to reduce the damages if it

[1] `http://biggis-project.eu/` accessed 07.12.2019

cannot be avoided. Preparedness describes those actions taken before an emergency to allow an effective response, such as the training of personnel and the provisioning of resources and equipment. Emergency response refers to the actions taken in the immediate aftermath of an incident. Finally, recovery describes the long-term measures taken to return to a state of normality (ibid).

2.1.2 Rapid mapping

This thesis focuses on the emergency response phase, the main purpose of which is to assist the affected population. These operations start immediately after an incident occurred. Therefore, the first step towards an effective emergency response is to assess the nature and scale of the situation at hand to coordinate the deployment of units and the protection of the population.

For major disasters such as landslides, earthquakes, or floods, dedicated systems have been proposed for acquiring this information using remote sensing technologies. In this context, the process of acquiring and providing information is often referred to as *rapid mapping*, emphasizing the spatial nature of the data that is required. The arguably best-known examples are satellite-based systems such as the Copernicus Emergency Management Service (Copernicus EMS) or the UNOSAT Rapid Mapping Service, which provide users with detailed maps indicating the extent and severity of the crisis (Copernicus EMS, 2018). The target delivery time of these maps is within 9 to 12 hours after an incident, with a first overview provided within 3 hours and detailed maps available after an estimated 5 days (Spruyt, 2017). This delivery time is dependent on the time required for acquiring satellite imagery and for data processing and map preparation. The spatial resolution of the provided data, i.e., the amount of detail that can be captured, varies significantly. Reported resolutions ranging from approximately 1 square meter per pixel (Copernicus EMS, 2016a) to values as high as some hundred square meters per pixel[2] (Copernicus EMS, 2016b).

2.1.3 Situation assessment for first responders

The vast majority of emergencies faced by first responders do not reach the magnitude of natural disasters for which the rapid mapping support described above has been introduced. Nonetheless, reconnaissance is an essential aspect of every operation. It is the first action in the management or command system (the so-called *Führungssystem*) described in the German federal fire service regulations (FwDV 100, 1999). In this work, we focus on operations dealing with large fires and chemical accidents that lead to the release of potentially hazardous substances. While far from being a daily occurrence, the German Federal Office of Civil Protection and Disaster Assistance (BBK) reports that

[2] Different definitions are used in literature to measure the resolution of a sensor system. In this remainder of this work, we understand the resolution as the *ground sampling resolution*, i.e., the size of the ground area captured per pixel.

a total of 99 severe incidents and an estimated 1,000 incidents involving transports of dangerous goods occurred within three years (BBK, 2018b).

When an emergency occurs, the main focus of the response units is to

- detect and identify the specific hazard,
- protect the general population in the affected area and
- provide immediate assistance for those who have come into contact with a hazardous substance, e.g., medical support or decontamination (BBK, 2018a).

These responsibilities require a reliable overview of the nature and distribution of any potentially hazardous substances. The rapid mapping services, which to this date are almost exclusively tailored to large-scale disasters, do not provide data in the speed and accuracy that is required for these smaller incidents. Instead, the BBK maintains several hundred dedicated reconnaissance vehicles that are equipped with sampling sets for a range of biological, chemical or nuclear contaminants (BBK, 2019).

2.1.4 Sampling procedure and guidelines

To support the emergency response personnel in ground-based situation assessment, the BBK has published guidelines and recommendations for sampling in case of chemical, biological and radioactive hazards (Bachmann et al., 2015). They are intended to ensure the effective use of the BBK's reconnaissance vehicles by introducing common standards for "emergency sampling". In addition to instructions addressing specific equipment, these guidelines suggest a basic sampling strategy for first responders, consisting of:

1. Defining the area of interest, i.e., the area suspected of contamination,
2. Identifying potential sampling points where a hazard may be present,
3. Prioritizing these sampling points,
4. Taking samples at a selected set of locations following these priorities,
5. Evaluating the samples, and
6. Decontaminating personnel and equipment, if necessary.

To determine the priorities referred to in step 3, the BBK proposes to distinguish between three general levels, which are summarized in Table 2.1. This classification depends on the likelihood of contamination in an area and the size of the potentially affected population. The purpose of this approach is to ensure that samples are taken at critical locations, i.e., at regions with a high likelihood of contamination where the civil population may need assistance. The guidelines furthermore recommend taking control samples at non-contaminated areas to establish the local background levels of relevant substances and to select representative samples across the designated target area. This is done to ensure that the spread of substances and contaminants can be reliably assessed.

Priority level 1	Areas with very high likelihood of contamination and presence of endangered population
Priority level 2	Areas where contamination is likely
Priority level 3	Areas with moderate risk or uncertain risk of contamination

Table 2.1: Priority levels for emergency sampling according to the German Federal Office of Civil Protection and Disaster Assistance (Bachmann et al., 2015).

2.1.5 Assessment of affected areas and priorities

Emergency services already have tools at their disposal to determine potentially contaminated areas after an incident. The probably best-known example is the WISER service, an online information system developed by the U.S. National Institutes of Health for emergency responders. This service indicates protective ranges and lists protective measures and decontamination procedures for a wide range of hazardous substances (United States National Library of Medicine, National Institutes of Health, 2018). For airborne substances, this service can also estimate and visualize affected areas based on the location of the source and the direction and strength of the wind. An example of the visualization component is given in Figure 2.1, which represents the potentially affected area and recommended actions for a methane spill. This gives a rough indication of affected areas, but the estimate has to be confirmed or corrected in-situ.

2.2 UAV and sensor system

Fire services still largely rely on the ground-based assessment of critical areas, either by using dedicated surveillance vehicles or by performing observations on foot. This, however, is comparatively slow. Furthermore, it requires the availability of qualified personnel and carries the risk of exposing them to hazardous substances.

UAV-based optical remote sensing systems offer an alternative as they can be deployed quickly, require less personnel than conventional surveillance methods, and allow surveying areas that are otherwise inaccessible. Similar to satellite-based rapid mapping services, they can provide information about the extent and the severity of the contamination. In the following, we give a brief overview of the methods and technologies that are combined in such a surveillance system. So far, however, most of these approaches are still in the developmental state.

2.2.1 UAV systems

UAV systems can be grouped into two categories: fixed-wing systems and rotary-wing UAVs. Fixed-wing UAVs fly at higher speeds but need to maintain constant forward motion. Furthermore, they require either a runway or a launcher for take-off and landing.

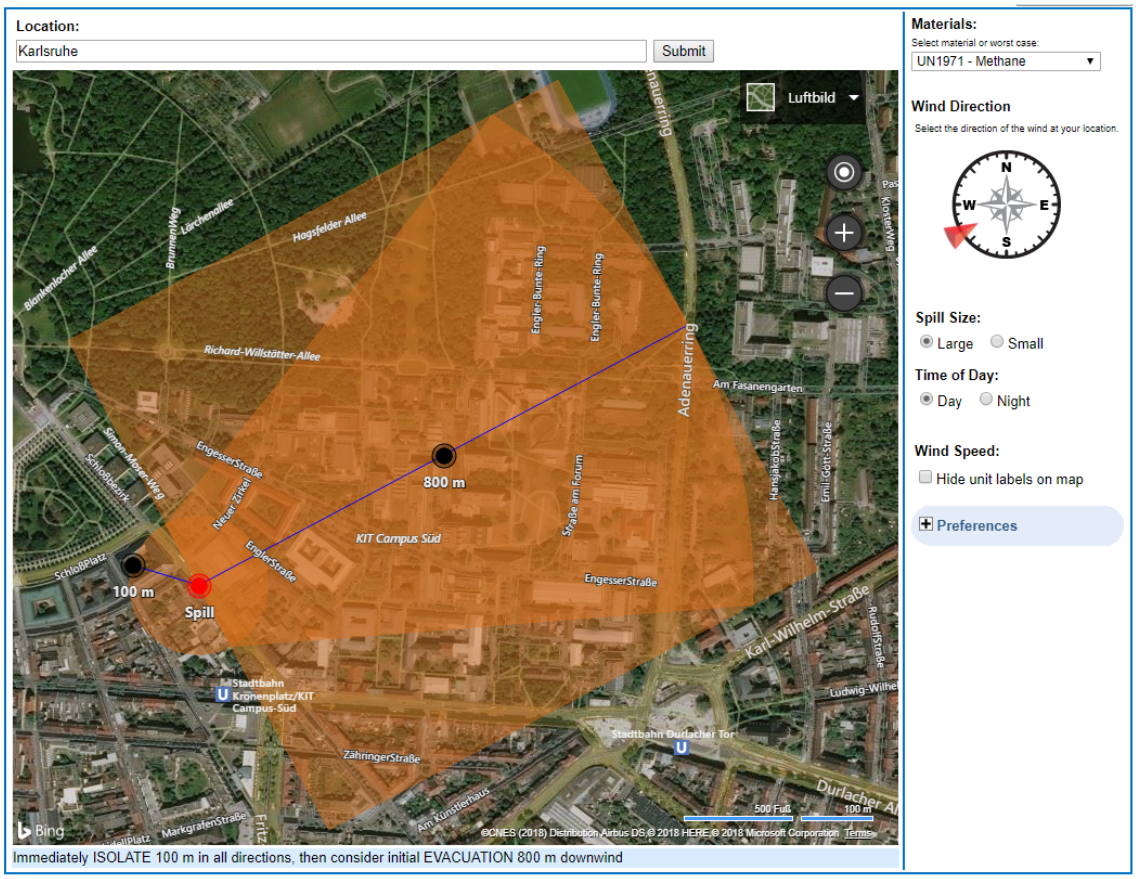

FIGURE 2.1: Screenshot of the WISER system, indicating protective distances for a methane spill depending on source location and primary wind direction. Source: `https://webwiser.nlm.nih.gov/`, accessed 27.05.2018.

Rotary wing UAVs have lower maximum speed and flight duration, but can start and land vertically. They are also able to stay stationary during the flight. This flexibility makes them attractive for emergency surveillance (Boccardo et al., 2015).

For this reason, the BigGIS project focused on rotary-wing UAVs with optical remote sensor systems for surveying an area after an emergency. These vehicles can carry sensors and equipment with a total weight of up to 3 kg. Some systems can fly up to around 40 minutes at a maximum horizontal cruise speed of around 50 km/h above ground, depending on the total payload and environmental influences[3]. In practice, these factors can have a major impact on flight time and cruise speed.

2.2.2 Sensor systems and substance identification

In case of emergencies involving hazardous substances, optical remote sensors are of particular interest. Other than chemical sensors, they allow the provisioning of infor-

[3] compare, e.g., the AiD-MC8 octocopter, `https://www.aidrones.de/`, accessed 05.03.2018

FIGURE 2.2: Prototypical UAV system and hyperspectral camera. Source: Bodo Bernsdorf, EFTAS Fernerkundung Technologietransfer GmbH.

mation without direct contact with a substance, even if these substances are invisible to the human eye. Commercially available systems include, for example, infrared (IR) camera systems that can detect substances such as methane.[4] For emergency surveillance applications, the BBK operates the IR remote reconnaissance device SIGIS 2 for detecting chemical and biological threats (Harig and Rusch, 2011).

In recent years, the emergence of lightweight hyperspectral cameras has allowed their use onboard UAVs (Aasen et al., 2015). A recent overview of available sensor systems in the domain of environmental monitoring is provided by Manfreda et al. (2018). Similar to existing approaches for the remote sensing of chemical agents (Flanigan, 1996; Mayfield et al., 2000), prototypical systems for detecting a wider range of substances using thermal and hyperspectral imaging systems have been studied in the BigGIS project. The sensors used in this project analyze the spectral range of visible light and near IR (wavelength between 450 - 950 nm). As substances reflect or absorb light in specific patterns, this information can be used to identify objects and materials based on their characteristic spectral signature.

A (simplified) example of how substances can be detected this way is given in Figure 2.3, which depicts a proof of concept conducted during the BigGIS research project. Here, an artificially created smoke cloud was blended with additives, i.e., specific substances with known spectral signatures, in this case in the form of chlorophyll. This cloud can be seen at the left-hand side in Figure 2.3a, which depicts the smoke cloud in white while the surrounding vegetation, which similarly contains chlorophyll, is highlighted in red. Two locations are marked in this image, one close to the source and one

[4] https://www.flir.com/browse/industrial/gas-detection-cameras/, 06.02.2018

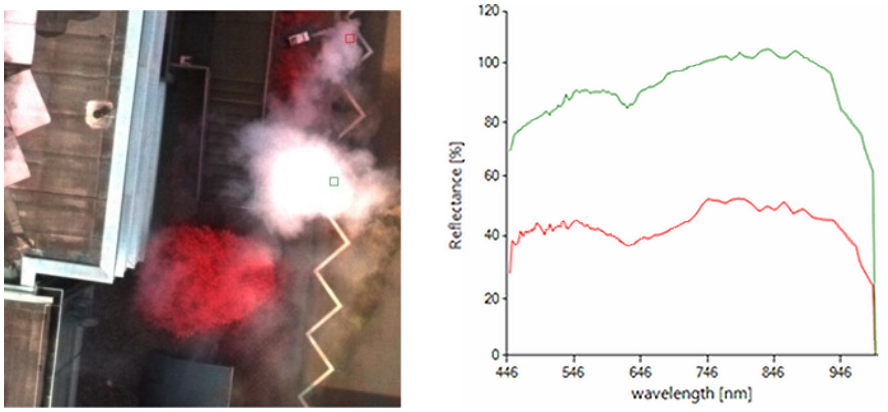

(A) Referenced locations (B) Spectral signatures

FIGURE 2.3: Spectral signatures taken at two points in a chlorophyll-blended smoke cloud, one close to the source (red) and one in the middle of the cloud (green). The sample locations are indicated in the reference image on the left, where the smoke cloud is highlighted in white and the vegetation is indicated in red. The figure on the right gives the observed spectral signatures. Source: Alexandra Bicsan, Bodo Bernsdorf, EFTAS Fernerkundung Technologietransfer GmbH.

in the middle of the cloud. Figure 2.3b gives an extract of the hyperspectral image for both locations, indicating very similar patterns with a first maximum in the visible green spectrum (550 nm), a decrease in the red spectrum (towards 650 nm), and an increase in the near IR.

Problems arise in practical situations due to the intermingled substances with potentially similar patterns and errors in the measurements. Still, these preliminary results indicate that it is possible to detect and identify specific gases using UAV-based hyperspectral imaging systems, thus demonstrating their potential for surveillance applications. Additionally, the UAVs can be equipped with standard RGB-cameras, allowing the user to verify the results.

2.3 Planning and executing UAV missions

In addition to flying under manual control, UAV systems can move autonomously along predetermined waypoints. In the following, we discuss essential flight parameters, which have to be taken into account for planning UAV missions, and briefly review supporting tools that enable this form of automated surveillance.

Flight parameters

The possible speed, maximum height and flight time depend on the deployed UAV platform, i.e., the vehicle itself and the total payload. The effective flight height and resulting ground resolution, i.e., the size of the area represented by one pixel in an image, depend

on the sensor that is used. In practice, the flight height above ground is determined so that a desired target ground resolution can be achieved.

The fundamental relationship between flight height and ground resolution for any given optical sensor is indicated in Figure 2.4. Note that this depiction is idealized, leaving out influences such as lens distortions or tilts of the camera that mean that sensor and covered ground are not strictly aligned. In practice, accurate values need to be obtained from the manufacturers' technical specifications. We denote the flight altitude as h and the ground sampling resolution as *pxr*. The camera is defined by its focal length f with $h \gg f$, horizontal sensor size b and number of pixels per band *pix*. The width of an image in terms of ground distance is denoted as g. We can then express the relationship between flight height and resolution as

$$g = \frac{h \cdot b}{f} \tag{2.1}$$

The total area covered is $A = g^2$ for a quadratic sensor such as the one employed in our use case. The ground sampling resolution, i.e., the covered area per pixel, is

$$pxr = \frac{A}{pix}.$$

The camera depicted in Figure 2.2 achieves a resolution of less than 1 cm^2 per pixel[5], depending on the flight altitude. The image given in Figure 2.3 has a ground sampling resolution of approx. 5 cm^2.

Ground control

During the flight, UAVs are controlled remotely using a ground control software that communicates with the vehicles and shows updates, e.g., concerning position and battery status, to the user. Often used open-source systems are for example QGroundControl[6] or Ardu Mission Planner[7]. Most of these systems offer two basic functions for planning UAV missions:

- Following a manually determined sequence of targets, typically in form of a point-and-click application, and
- Planning a mission for full coverage of a given region, resulting in a grid-like regular pattern.

The ground control systems communicate with the UAV using standardized protocols such as MAVLink[8]. This protocol supports the specification of mission parameters and

[5] `http://cubert-gmbh.com/product/uhd-185-firefly/`, accessed 31.01.2018
[6] `http://qgroundcontrol.com/`, accessed 20.03.2018
[7] `http://ardupilot.org/planner/index.html`, accessed 20.03.2018
[8] `http://qgroundcontrol.org/mavlink/start`, accessed 07.10.2017

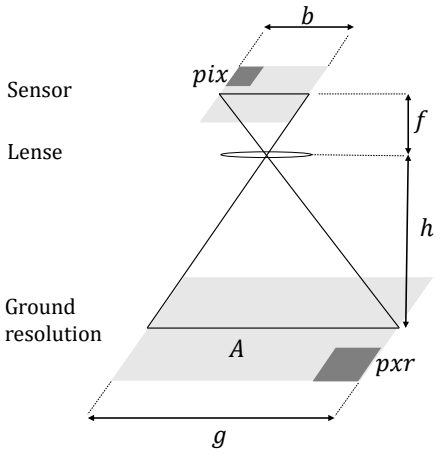

FIGURE 2.4: Dependency between camera specifications and ground resolution.

allows the transmission of data relating to the state of the vehicle such as the current position, speed, and orientation. The mission plan includes directives concerning the orientation of the camera and the number of images to be taken at each location. In an emergency scenario, this setup enables a user to plan and execute a mission by selecting a sequence of target locations without having to remotely control the UAV or the sensor.

FLIGHT CONTROLLER

The term flight controller refers to the on-board hardware and firmware of a UAV. This includes the on-board processor as well as sensors such as the inertial measurement units (accelerometers and gyroscopes), which provide information about the state of the vehicle. It may further include additional modules, for example, GPS or barometers. The main purpose of a flight controller is to control and stabilize the aircraft during the flight. This allows a UAV to fly autonomously, i.e., without continuous input from the user, by following the mission plan transmitted by the ground station. If contact to the ground station is lost, flight controllers typically follow a fail-safe directive, returning the UAV to the launch position or a predetermined home location.

2.4 SUMMARY

In this chapter, we provided an overview of the foundations of emergency surveillance with UAVs. We showed that an initial phase dedicated to reconnaissance and surveillance plays an essential role in emergency management. Currently, emergency services largely rely on ground-based situation assessment, with satellite-based remote sensing systems only being available in case of large-scale disasters. An emerging alternative are UAV-based systems that can be deployed close to the affected area in the immediate aftermath of an incident. Equipped with remote sensor systems, they can be used

to detect and identify substances while avoiding immediate contact with contaminants. We demonstrated that the use of UAV platforms is further enabled by control systems that allow the UAVs to move autonomously through predetermined waypoints without requiring continuous oversight. This makes it possible to deploy UAV fleets with little manual effort, further reducing the demands on the response units.

3 Models for spatial processes

In case of accidents, response units often face hazardous substances that are spread across considerable areas. Hence, the objective of the reconnaissance phase in emergency surveillance is to obtain information about a phenomenon that is inherently spatial. In this chapter, we discuss methods for modeling spatial dependencies and for making inferences based on sampled locations. These concepts play a central role in the mission planning problem introduced in Chapter 4 and are used in the model formulations in Chapters 7 and 13.

The approaches discussed in the following sections have been developed in two distinct fields of research: Multivariate regression and geostatistics. The methods are similar and are occasionally used synonymously. They differ in terminology and sometimes in the details of specific approaches. To establish a common understanding, we cover both approaches: generally applicable regression approaches used in the context of machine learning and the more specific methods for spatial processes developed in the field of spatial statistics.

In this chapter, we first discuss how spatial dependencies within a process can be modeled in Section 3.1. In Section 3.2, we give an overview of approaches that use these properties and models for interpolating and predicting spatial processes.

3.1 Expressing spatial structure

Environmental processes such as the distribution of airborne contaminants follow physical laws but are too complex to be easily described or even understood. A possibility to deal with the complexity is to consider these processes as random and to employ statistical methods to describe and analyze them (Cressie and Wikle, 2011). Representing the affected area as a set of Cartesian coordinates $\mathcal{V} \subset \mathbb{R}^2$, a physical phenomenon can be described as a stochastic process $\{Z(s)\}_{s \in \mathcal{V}}$ over the locations in $\mathcal{V}$, i.e., as a collection of random variables defined over a two-dimensional field (ibid).

The observed phenomena in our case, i.e., plumes or gases, are usually positively spatially autocorrelated, which means that similar values occur at locations close to each other (Stachniss et al., 2009). In other words, they follow Tobler's First Law of Geography, which states that "everything is related to everything else, but near things are

more related than distant things" (Tobler, 1970). This premise is the foundation for the modeling and interpolation approaches that we discuss in the next sections.

3.1.1 Covariance functions and covariance matrices

In all approaches considered in this chapter, the spatial structure of a given phenomenon is expressed by its covariance $\text{Cov}(Z(s), Z(s'))$, which describes the similarity between random variables $Z(s)$ and $Z(s')$ at different locations s and s' with $s, s' \in \mathcal{V}$.

For a stochastic process including a number of locations $\mathcal{V}$, the pairwise covariances are combined in a covariance matrix $\Sigma_{\mathcal{V},\mathcal{V}}$ with each element (s, s') describing the covariance $\text{Cov}(Z(s), Z(s'))$. This matrix comprises the fundamental information and assumptions about the relationship between locations, and therefore about the spatial process itself. Given a sufficient number of observations, it is possible to determine an empirical covariance matrix. However, it is common to use covariance models rather than empirical covariances, as observations are usually not available for all $s \in \mathcal{V}$.

As a consequence, the covariance of a stochastic process is usually modeled using a covariance function $k(s, s')$, also called a *kernel* (Rasmussen and Williams, 2006). It depends on the distance between locations and generally decreases with growing distance (Krause et al., 2008). A process is *stationary* if $\text{Cov}(Z(s), Z(s')) = \text{Cov}(s - s')$, i.e., it only depends on the vector difference between two locations, not the locations themselves. If $\text{Cov}(Z(s), Z(s')) = \text{Cov}(\|s - s'\|)$, i.e., the covariance only depends on the distance, the process is called *isotropic*.

Several covariance functions have been proposed that are used for spatial interpolation and learning. The parameters characterizing a kernel function are usually called *hyperparameters*. To give an example, a popular function is the squared exponential kernel

$$k^{\text{SE}}(s, s') = \nu^2 \cdot \exp\left(-\frac{\|s - s'\|^2}{2l^2}\right) \tag{3.1}$$

where $\|s - s'\|$ indicates the distance between s and s', here and in the following understood as the Euclidean distance. The hyperparameters of this kernel are the so-called *lengthscale* $l > 0$, which determines how smooth the function is, and ν^2 indicating the process variance as a scaling parameter.

3.1.2 The variogram

In spatial statistics, a widely used measure for spatial relationships is the variogram $\gamma(d)$ rather than a covariance function.[1] These two approaches, however, are closely related, as the variogram describes the average variance instead of the covariance. In

[1] The term *variogram* itself is not used consistently in geostatistics literature, and some authors refer to $\gamma(d)$ as the semi-variogram (Bachmaier and Backes, 2011). In this thesis, we only use the term variogram.

other words, it indicates the average degree of dissimilarity depending on the distance[2] $d := \|s - s'\|$ between two locations s and s'.

Formally, with ν^2 as the maximum variance in a process and $\text{Cov}(Z(s), Z(s'))$ denoting the covariance between two locations s and s' at distance d, the relationship between the the variogram and the covariance is (Mitas and Mitasova, 1999):

$$\gamma(d) = \nu^2 - \text{Cov}\left(Z(s), Z\left(s'\right)\right).$$

Given a number of observed realizations $z_s, z_{s'}, \ldots$ of the random process, an empirical variogram can be computed as follows:

$$\gamma(d) = \frac{1}{2|N(d)|} \sum_{(s,s') \in N(d)} \left(z_s - z_{s'}\right)^2 \tag{3.2}$$

In Equation (3.2), $N(d)$ denotes the set of all pairs of locations at distance d, and $|N(d)|$ gives the size of this set, i.e., the number of corresponding pairs. In other words, the empirical variogram gives the average squared difference between observations for all pairs of values at distance d (Rossi et al., 1994).

For subsequent analytical approaches and predictions, these empirical variograms are often replaced by theoretical variogram models for describing the spatial dependence of the random process. They are fitted to the empirical variogram such that the mean squared error is minimal. These variogram models closely mirror the covariance functions discussed in Section 3.1.1. To give an example, the exponential variogram, which corresponds to the exponential kernel in Equation (3.1), is commonly defined as (ibid):

$$\gamma^{\text{SE}}(d) = \nu^2 \left(1 - \exp\left(-\frac{3d}{a}\right)\right). \tag{3.3}$$

This formulation is equivalent to the kernel function in Equation (3.1). In geostatistics, the parameters have been associated with specific characteristics of a spatial process: Similar to Equation (3.1), ν^2 denotes the maximum variance within the process. When describing a variogram, this value is commonly referred to as the *sill*. The parameter a denotes the so-called *range* of the variogram, i.e., the distance at which locations are no longer spatially correlated, with larger values for a indicating that the sill is reached at higher distances. As the exponential variogram converges toward the sill asymptotically, a is usually understood as the "practical" range at which the variogram value reaches 95% of the sill (Bohling, 2005).[3]

[2] This distance is also referred to as the *separation distance* or *lag* in spatial statistics

[3] With $\gamma(d) = \nu^2(1 - \exp(-\frac{3\|s-s'\|}{a})) \overset{\|s-s'\| \geq a}{\geq} \nu^2(1 - \exp(-3)) \approx 0.95\nu^2$

3.2 Spatial interpolation

The main purpose of modeling spatial processes is to allow inferences about unsurveyed locations, based on a finite number of known observations. This section discusses some of the most often used spatial interpolation approaches, both deterministic and stochastic. Afterward, we show how properties of these approaches can be used for assessing the quality of a prediction and for designing sampling strategies.

3.2.1 Deterministic approaches

Deterministic methods interpolate a phenomenon using approximations of the relations between observations. Again, the basic premise behind all of these approximations is that closer objects or points exhibit more similarities than more distant ones, i.e., that the underlying data is spatially autocorrelated. These methods are easily applicable for a vast number of situations but do not include a stochastic component, i.e., they cannot indicate the uncertainty or error in their prediction. Typically, they determine an estimated value $\widehat{Z}(s_0)$ at location s_0 as a weighted average of the known values z_i for $i \in \mathcal{S}$:

$$\widehat{Z}(s_0) = \sum_{i \in \mathcal{S}} \lambda_i z_i. \tag{3.4}$$

They differ in how the weights λ_i and the observations taken into account for the interpolation are determined.

To give one example, in *nearest neighbor interpolation (NN)*, the estimation $\widehat{Z}(s_0)$ equals the value z_s of the location $s \in \mathcal{S}$ that is closest to s_0, i.e., (Li and Heap, 2008)

$$\lambda_i = \begin{cases} 1, & i = \arg\min_{s \in \mathcal{S}} \|s - s_0\| \\ 0, & \text{otherwise} \end{cases} \tag{3.5}$$

When applying *inverse distance weighting (IDW)*, which is one of the most common non-statistical approach, the estimated value $\widehat{Z}(s_0)$ is computed as a weighted average with weights inversely proportional to a power q of the distance between locations (Mitas and Mitasova, 1999):

$$\lambda_i = \frac{\frac{1}{\|s_i - s_0\|^q}}{\sum_{s \in \mathcal{S}} \frac{1}{\|s - s_0\|^q}} \tag{3.6}$$

The parameter $q > 0$ determines how quickly weights decrease with distance, with higher values for q leading to increased weights for close observations and lower values resulting in a smoother interpolated surface.

These approaches and other non-statistical interpolation methods yield quick predictions with low computational effort. Unlike the stochastic methods discussed in the next section, they do not provide information about the prediction error (Li and Heap, 2008).

3.2.2 Statistical interpolation approaches

Statistical interpolation approaches use the covariance models discussed in Section 3.1.1 to make estimations based on a set of observations. In this section, we first discuss Gaussian process regression for spatial processes. We then proceed with a discussion of kriging, which has been developed in the field of geostatistics.

Gaussian process regression

Gaussian processes (GP) are used extensively in machine learning and statistics. This section is based on Rasmussen and Williams (2006), who define a GP, expressed as $Z_{\mathcal{V}} := \{Z(s)\}_{s\in\mathcal{V}}$ as "a collection of random variables, any finite number of which have a multivariate normal distribution" (ibid, p. 13). This means that for any finite subset $\mathcal{V}' \subset \mathcal{V}$, the random vector $Z_{\mathcal{V}'}$ is normally distributed.

A GP is completely defined by its covariance function $k(s, s')$ and mean function $\mu(s) = \mathrm{E}[Z(s)]$. As stated above, for the spatial processes that we are interested in, the covariance function describes the spatial covariance across a random field depending on the distance between locations. The mean function gives the expected value for all $s \in \mathcal{V}$. For simplicity, it is usually set to zero.

The main advantage of a GP in this domain is its capability to make inferences about unobserved locations based on several sampled locations. The GP without taking sample observations into account is referred to as a *Gaussian prior*. As indicated above, the prior is defined by its mean and covariance functions. For simplicity, we denote the prior mean as $\mu_{\mathcal{V}}$, representing a column vector of expected means with one entry $\mu(s)$ for each $s \in \mathcal{V}$. The covariance matrix is denoted $\Sigma_{\mathcal{V},\mathcal{V}}$ and contains an entry $(s, s') = k(s, s')$ for all $s, s' \in \mathcal{V}$. This covariance matrix indicates the essential characteristics of the process. In other words, the prior encapsulates all assumptions or information that one could have about the phenomenon without having performed observations, in particular, assumptions about the kernel function and smoothness of the distribution. In short, we denote the Gaussian prior as

$$Z_{\mathcal{V}} \backsim \mathcal{GP}(\mu_{\mathcal{V}}, \Sigma_{\mathcal{V},\mathcal{V}}). \tag{3.7}$$

If samples are available for a set of locations $\mathcal{S} \subset \mathcal{V}$, the prior belief is updated taking these measurements into account. The result is a Gaussian *posterior*, represented by a posterior mean $\mu_{\mathcal{V}|\mathcal{S}}$ and a covariance matrix $\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}}$ conditioned on set $\mathcal{S}$. We write this as

$$Z_{\mathcal{V}|\mathcal{S}} \backsim \mathcal{GP}(\mu_{\mathcal{V}|\mathcal{S}}, \Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}}) \tag{3.8}$$

An example illustrating the difference between GP prior and posterior distributions is given in Figure 3.1. For simplicity, we use a GP defined over $\mathbb{R}$ rather than a spatial process defined over $\mathbb{R}^2$. Figure 3.1a gives the Gaussian prior. Here, the prior mean is 0. The grey shaded area represents the 95 % confidence interval, indicating the range

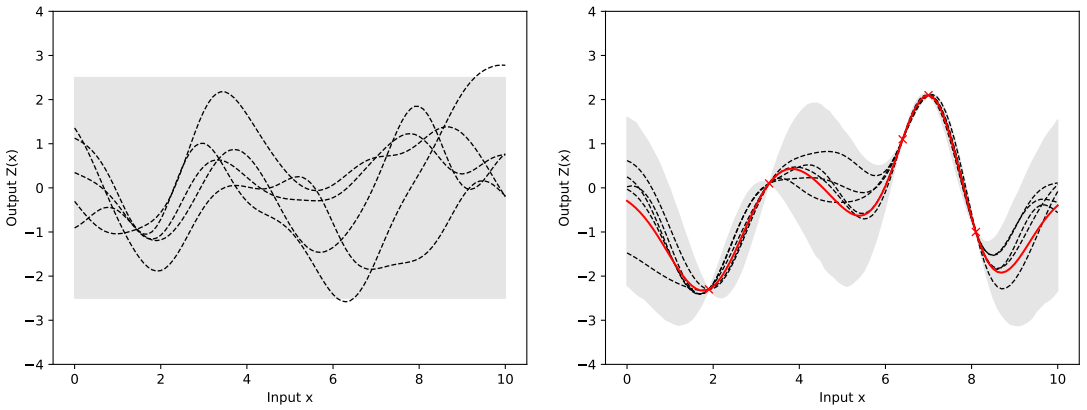

(A) Gaussian prior (B) Gaussian posterior

FIGURE 3.1: Prior and posterior probability distribution. Dashed lines indicate samples drawn from the respective prior or posterior distribution. The grey area gives the 95 % confidence interval. Figure (A) represents the Gaussian prior. Figure (B) gives the posterior distribution, conditioned on five samples indicated in red. The red line gives the posterior mean. Own representation based on Rasmussen and Williams (2006).

in which we expect the output $Z(x)$ to be. Grey lines indicate sample functions that represent possible realizations of this process. The prior covariance determines the shape and amplitude of these functions. For instance, it characterizes the average distance between extrema and the rate at which values change, resulting in similar patterns for all sample functions. Figure 3.1b indicates the Gaussian posterior conditioned on five observations that are marked as red crosses. Again, dashed lines are sample functions, this time drawn from the posterior. The red line represents the posterior mean function.

To derive the posterior, we define $\Sigma_{\mathcal{V},\mathcal{S}}$ as the covariance matrix between all s, s' with $s \in \mathcal{V}, s' \in \mathcal{S}$, and $\Sigma_{\mathcal{S},\mathcal{S}}$ as the covariance matrix between all $s, s' \in \mathcal{S}$. Furthermore, $\mu_{\mathcal{S}}$ designates a column vector of prior mean values $\mu(s)$ for $s \in \mathcal{S}$. The observed values at locations $\mathcal{S}$ are represented as a column vector $z_{\mathcal{S}}$. We can then determine the mean vector and covariance matrix of the GP posterior as follows:

$$\mu_{\mathcal{V}|\mathcal{S}} = \mu_{\mathcal{V}} + \Sigma_{\mathcal{V},\mathcal{S}}\Sigma_{\mathcal{S},\mathcal{S}}^{-1}(z_{\mathcal{S}} - \mu_{\mathcal{S}}) \tag{3.9}$$

$$\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}} = \Sigma_{\mathcal{V},\mathcal{V}} - \Sigma_{\mathcal{V},\mathcal{S}}\Sigma_{\mathcal{S},\mathcal{S}}^{-1}\Sigma_{\mathcal{S},\mathcal{V}}. \tag{3.10}$$

The posterior mean determined in Equation (3.9) gives the expected values $\mu_{\mathcal{V}|\mathcal{S}}$ for all $s \in \mathcal{V}$. The posterior covariance in Equation (3.10) indicates the uncertainty associated with the predictions. The posterior thus incorporates the knowledge obtained using the sampled locations in combination with information about the fundamental characteristics of the process that are formulated in the form of the Gaussian prior. A disadvantage of this approach is its complexity of $\mathcal{O}(n^3)$, resulting from the matrix inversions in Equations (3.9) and (3.10). While several approaches have been proposed that use approximations (see, e.g., Hensman et al., 2013) or exploit structures of partic-

ular covariance functions (e.g., Saatçi, 2011)), the overall complexity of a GP regression remains an issue when fast response times are required (Bui et al., 2015).

Parameter estimation

The Gaussian posterior defined in Equations (3.9) and (3.10) is based on a specific prior covariance matrix, i.e., it depends on the specific parameters characterizing the corresponding kernel function. However, the prior belief about these parameters can be wrong, thus impacting the quality of a prediction. To avoid this, the hyperparameters of a kernel function, here denoted θ, are fitted based on the available samples, i.e., we seek to find the best hyperparameters $\hat{\theta}$, given a set of sampled values $z_{\mathcal{S}}$.

We denote a parametrized kernel function as k_θ. We furthermore denote the covariance matrix defined by k_θ as Σ_θ for $s, s' \in \mathcal{S}$. A means to find the best set of hyperparameters is to optimize the log-marginal likelihood $\log p(z_{\mathcal{S}}|\theta)$, which is defined as

$$\log p(z_{\mathcal{S}}|\theta) = -0.5(z_{\mathcal{S}} - \mu_{\mathcal{S}})^\top \Sigma_\theta^{-1}(z_{\mathcal{S}} - \mu_{\mathcal{S}}) - 0.5 \log |\Sigma_\theta| - 0.5n \log 2\pi \tag{3.11}$$

where $|\Sigma_\theta|$ is the determinant of the covariance matrix and n is the dimension of the process.

The first term on the right-hand side of Equation (3.11) measures the fit to the sampled data $(-0.5(z_{\mathcal{S}} - \mu_{\mathcal{S}})^\top \Sigma_\theta^{-1}(z_{\mathcal{S}} - \mu_{\mathcal{S}}))$. The second term $(-0.5 \log |\Sigma_\theta|)$ is the complexity penalty, which is used to avoid overfitting. Based on this, the fitted hyperparameters $\hat{\theta}$ can be determined as those that maximize the log-marginal likelihood, i.e.:

$$\hat{\theta} = \arg\max_\theta \log p(z_{\mathcal{S}}|\theta). \tag{3.12}$$

We refer to Rasmussen and Williams (2006) for a more detailed discussion but give one example to illustrate the interpretation of Equation (3.11): In a squared exponential kernel (see Equation (3.1)), a higher value for the lengthscale parameter represents a smoother distribution across space. This increases the negative complexity penalty, i.e., the second term. At the same time, it decreases the data fit in the first term, as the posterior cannot follow the sampled data as closely as it is possible with a lower value for the lengthscale. In contrast, low lengthscales imply a high process complexity while increasing the data fit, which means that overfitting is more likely.

Kriging

Kriging refers to a similar prediction technique that has been developed in spatial statistics. Instead of using covariance functions like GP regression approaches, kriging is based on the variogram (Section 3.1.2). Kriging is often referred to as GP regression without further differentiation between the two approaches. However, some differences remain between these concepts. For example, several techniques have been developed specifically for kriging, as it directly addresses spatial data analysis. This section introduces the basic principles behind kriging. Unless otherwise noted, this discussion is

based on Wackernagel (2013) and Cressie and Wikle (2011), where we refer the reader for a more extensive overview.

Several of kriging methods have been developed to address specific properties of the underlying random field: The most common approaches are simple, ordinary, and universal kriging. Simple kriging is based on the assumption that the mean of the underlying random function $Z(s)$ is constant and known. Ordinary kriging assumes a constant, but unknown, mean. Finally, universal kriging can deal with means that are both unknown and not constant, i.e., it can incorporate trends in the data. This way, kriging also allows for variations that can be explained by external factors such as wind. In the following, we focus on ordinary kriging, as it is the most commonly applied and suffices to understand the concept (Curran and Atkinson, 1998).

Note that, often, the observations are restricted to a smaller set of samples $\mathcal{S}' \subset \mathcal{S}$, also called the kriging neighborhood, as samples across larger distances are no longer correlated and thus should have little influence on the result (Oliver and Webster, 2014). For ordinary kriging, this means that the assumption of a constant mean only needs to be fulfilled for the kriging neighborhood rather than the whole process. Formally speaking, it requires local second-order stationarity.

As in Equation (3.4), the predicted value $\widehat{Z}(s_0)$ at location s_0 is calculated as a weighted linear combination of observed values z_s. For each location s_0 to be interpolated, the objective of the kriging approach is to determine weights λ_i for all observed locations such that (a) the estimator is unbiased and (b) the prediction error is minimal. Condition (a) requires that the mean of the estimation must be equal to the mean of the real values:

$$\mathrm{E}\left(\widehat{Z}(s_0) - z_{s_0}\right) = \mathrm{E}\left(\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i} - z_{s_0}\right) = 0$$

Furthermore, the assumption of stationarity in ordinary kriging means that the expected value of the random process is an unknown constant mean and, therefore, that $\mathrm{E}(z_{s_i}) = \mathrm{E}(z_{s_0}) = \mu$ for $i \in \mathcal{S}'$. Together with the above condition and the linearity of expectation, it derives that:

$$\begin{aligned} &\sum_{i \in \mathcal{S}'} \lambda_i \mu - \mu = 0 \\ \Leftrightarrow &\sum_{i \in \mathcal{S}'} \lambda_i = 1. \end{aligned} \tag{3.13}$$

Condition (b) requires that the variance of the estimation error is minimal. The variance

of the estimation error can be determined as:

$$\begin{aligned}
\mathrm{Var}(\widehat{Z}(s_0) - z_{s_0}) &= \mathrm{Var}(\widehat{Z}(s_0)) + \mathrm{Var}(z_{s_0}) - 2\mathrm{Cov}(\widehat{Z}(s_0), z_{s_0}) \\
&= \mathrm{Var}(\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i}) + \sigma_0 - 2\mathrm{Cov}(\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i}, z_{s_0}) \\
&= \underbrace{\mathrm{E}((\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i})^2)}_{\sum_{i,j \in \mathcal{S}'} \lambda_i \lambda_j \mathrm{E}(z_{s_i} z_{s_j})} - \underbrace{\mathrm{E}((\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i}))^2}_{\sum_{i,j \in \mathcal{S}'} \lambda_i \lambda_j \mathrm{E}(z_{s_i})\mathrm{E}(z_{s_j})} + \sigma_0 - 2\mathrm{Cov}(\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i}, z_{s_0}) \\
&= \sum_{i \in \mathcal{S}'} \sum_{j \in \mathcal{S}'} \lambda_i \lambda_j (\mathrm{E}(z_{s_i} z_{s_j}) - \mathrm{E}(z_{s_i})\mathrm{E}(z_{s_j})) + \sigma_0 - 2\mathrm{Cov}(\sum_{i \in \mathcal{S}'} \lambda_i z_{s_i}, z_{s_0}) \\
&= \sum_{i \in \mathcal{S}'} \sum_{j \in \mathcal{S}'} \lambda_i \lambda_j \mathrm{Cov}(z_{s_i}, z_{s_j}) + \sigma_0 - 2\sum_{i \in \mathcal{S}'} \lambda_i \mathrm{Cov}(z_{s_i}, z_{s_0}) \qquad (3.14)
\end{aligned}$$

To determine the best linear unbiased predictor, the objective is to minimize the error variance defined in Equation (3.14) while respecting Equation (3.13). To do so, we introduce a Langrange multiplier 2ψ such that

$$\phi = \sum_{i \in \mathcal{S}'} \sum_{j \in \mathcal{S}'} \lambda_i \lambda_j \mathrm{Cov}(z_{s_i}, z_{s_j}) + \sigma_0 - 2\sum_{i \in \mathcal{S}'} \lambda_i \mathrm{Cov}(z_{s_i}, z_{s_0}) + 2\psi(\sum_{i \in \mathcal{S}'} \lambda_i - 1)$$

and differentiate with respect to the multiplier ψ and weights λ_i:

$$\frac{\partial \phi}{\partial \lambda_i} = 2\sum_{j \in \mathcal{S}'} \lambda_j \mathrm{Cov}(z_{s_i}, z_{s_j}) - 2\mathrm{Cov}(z_{s_i}, z_{s_0}) + 2\psi \equiv 0, \; i \in \mathcal{S}'$$

$$\frac{\partial \phi}{\partial \psi} = (\sum_{i \in \mathcal{S}'} \lambda_i - 1) \equiv 0$$

This leads to the set of kriging equations for determining the weights λ_i:

$$\sum_{j \in \mathcal{S}'} \lambda_j \mathrm{Cov}(z_{s_i}, z_{s_j}) + \psi = \mathrm{Cov}(z_{s_i}, z_{s_0}), \; i \in \mathcal{S}' \qquad (3.15)$$

$$\sum_{i \in \mathcal{S}'} \lambda_i = 1 \qquad (3.16)$$

This means that the weights λ_i depend on the variogram $\gamma(\|s_i - s_j\|) = \nu^2 - \mathrm{Cov}(z_{s_i}, z_{s_j})$ and, therefore, only on the relative distances between sample locations s_i and the interpolated location s_0.

Similar to GP regression, the computational complexity of kriging is $\mathcal{O}(n^3)$, resulting in the same obstacles when facing large data sets or needing quick response times. On the other hand, kriging-based approaches have been shown to perform particularly well for spatial interpolation in terms of prediction accuracy (Li and Heap, 2008).

3.2.3 Measuring prediction quality

The primary advantage of GP regression and kriging over non-statistical approaches is the knowledge about the variability that remains in the process after the prediction. This yields an estimate about the prediction error and thus serves as an indication of the quality of a prediction.

Variability measures

In GP regression, the variability is represented in form of the posterior covariance $\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}}$ determined in Equation (3.10). One possible measure for the variability in the process is the trace $\mathrm{tr}(\Sigma)$ of a covariance matrix, which is defined as the sum of the diagonal elements $k(s,s), s \in \mathcal{V}$. Another measure of uncertainty is the so-called "entropy" $\mathcal{H}(Z_\mathcal{V})$ of a Gaussian process $Z_\mathcal{V}$ in n dimensions (Rasmussen and Williams, 2006):

$$\mathcal{H}(Z_\mathcal{V}) = 0.5 \cdot \log((2\pi e)^n |\Sigma_{\mathcal{V},\mathcal{V}}|), \tag{3.17}$$

where $|\Sigma_{\mathcal{V},\mathcal{V}}|$ denotes the determinant of the covariance matrix $\Sigma_{\mathcal{V},\mathcal{V}}$. In kriging, the variability of the prediction is given in form of the estimation error variance $\mathrm{Var}(\widehat{Z}(s_0) - z_{s_0})$ for all predicted locations s_0, see Equation (3.14). To assess the entire process, this value can be averaged over all interpolation locations (Melles et al., 2011). In general, the lower these measures are, the better is the corresponding prediction.

Performance indicators

Besides these variability measures, a range of performance indicators has been proposed to evaluate the quality of an interpolation. These measures are typically based on the difference between the actual and the predicted values. Among these measures are, for example, the *root mean squared error (RMSE)*

$$\mathrm{RMSE} = \sqrt{\frac{1}{|\mathcal{V}|} \sum_{s \in \mathcal{V}} (\widehat{Z}(s) - z_s)^2}, \tag{3.18}$$

the *mean absolute error (MAE)*

$$\mathrm{MAE} = \frac{1}{|\mathcal{V}|} \sum_{s \in \mathcal{V}} |\widehat{Z}(s) - z_s|, \tag{3.19}$$

and the *mean error (ME)*

$$\mathrm{ME} = \frac{1}{|\mathcal{V}|} \sum_{s \in \mathcal{V}} (\widehat{Z}(s) - z_s). \tag{3.20}$$

RMSE and MAE measure the size of the error, with RMSE giving a relatively high weight to outliers. ME is mostly used in order to measure the bias in the estimated values (Li and Heap, 2008).

3.2.4 Designing sampling strategies

Gaussian processes can be used to design sampling strategies, i.e., strategies for selecting sampling locations that yield as much information as possible about a spatial process (Curran and Atkinson, 1998). This is made possible by the fact that a GP provides knowledge about the variability that remains in the process after the prediction, represented in form of the posterior covariance $\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}}$ (Equation (3.10)). This covariance is independent of the observed *values* $z_{\mathcal{S}}$ and only depends on the sampled *locations* $\mathcal{S}$, which makes it possible to estimate the quality of a set of sampling locations $\mathcal{S}$ prior to actually making the observations.

In the domain of geostatistics, the estimation error variance has been used to determine an adequate sample spacing, i.e., distances between samples that yield an acceptable prediction quality (see, e.g., Curran and Atkinson, 1998; Van Groenigen et al., 1999). In the literature on Gaussian processes, additional variability measures for estimating the information gain yielded by a set $\mathcal{S}$ have been derived. We discuss two measures often used in literature: The first criterion for measuring the quality of a set $\mathcal{S}$ is the *average reduction in variance (ARV)* that is achieved by the selected sampling locations (Das and Kempe, 2008; Krause et al., 2008). The measure is formally defined as:

$$ARV(\mathcal{S}) = \frac{1}{|\mathcal{V}|}(\mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}}) - \mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}})), \tag{3.21}$$

with $\mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}})$ and $\mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}})$ measuring the variances of the GP prior and posterior, respectively.

A criterion based on the concept of entropy (Equation (3.17)) is the *mutual information (MI)* between the selected sensing locations and the interpolated locations $\mathcal{V} \setminus \mathcal{S}$, defined as the reduction in entropy at all unobserved locations (Caselton and Zidek, 1984). The objective of this approach is to select sampling locations that are "most informative" about the remainder of the process. With $Z_{\mathcal{V}\setminus\mathcal{S}}$ as the Gaussian prior over all interpolated locations and $Z_{\mathcal{V}\setminus\mathcal{S}|\mathcal{S}}$ as the corresponding Gaussian posterior, MI is defined as follows:

$$MI(\mathcal{S}) = \mathcal{H}(Z_{\mathcal{V}\setminus\mathcal{S}}) - \mathcal{H}(Z_{\mathcal{V}\setminus\mathcal{S}|\mathcal{S}}). \tag{3.22}$$

The objective of this approach is to select sampling locations that are "most informative" about the rest of the process.

Using these criteria, sensing locations $\mathcal{S}$ can be determined such that ARV or MI is maximal, which means that the variability of the posterior is minimal. Note that the basic complexity of all of these measures is $\mathcal{O}(n^3)$, as computing the GP posterior requires inverting an $n \times n$ matrix with $n = |\mathcal{S}|$ (Rasmussen and Williams, 2006).

3.3 Summary

Environmental models such as the distribution of hazardous substances across an area can be modeled as stochastic processes defined over a two-dimensional field. This makes it possible to use a wide range of methods to model and interpolate these spatial processes. Generally speaking, these approaches are based on the fundamental assumption that close locations will exhibit similar values. These spatial interdependencies inherent in the observed phenomenon characterize the relationship between realizations of the process at different locations.

This property can be used for interpolating the process based on several sampled locations. Stochastic spatial interpolation approaches furthermore provide the possibility of estimating the achieved prediction quality before taking samples. This way, these models can be applied for a-priory selecting informative sampling locations, a property that is exploited in the approaches discussed in the remainder of this work. The main disadvantage of using these probabilistic approaches is their comparatively high computational complexity.

Part II

Offline mission planning

4 Mission planning for UAV-based rapid mapping

In Part I of this thesis, we have discussed the theoretical and practical foundations of this thesis: The role of emergency surveillance, UAV and sensor technologies, and the central concepts for modeling the spatial distribution of hazardous substances. In this chapter, we introduce the central planning problem in a UAV-based approach for emergency surveillance. We summarize the planning requirements and primary objective of such a system and illustrate the planning problem using a representative example.

This part is based on the following paper:

> Glock, K. and A. Meyer. Mission planning for emergency rapid mapping with drones. *Transportation Science 54(2)*, 534-560.

4.1 Mission planning problem for emergency surveillance (MPPES)

As we have shown in Chapter 2, an initial reconnaissance phase is an essential element in emergency management. In the case of large-scale disasters such as floods or earthquakes, this reconnaissance phase revolves around the acquisition of spatial data after an incident, a process which is often referred to as rapid mapping. However, this term is usually understood to apply exclusively to satellite-based surveillance systems. We expand this definition to other remote-sensing based surveillance systems, notably UAV-based systems. This distinguishes it from ground-based approaches, which are state-of-the-art for the vast majority of incidents faced by emergency services.

In this work, we focus on UAV-based rapid mapping approaches for emergencies such as large fires or industrial accidents, where unknown and potentially hazardous substances can be released and may affect a considerable number of people. The concentration of these substances can change rapidly across space, e.g., due to environmental factors such as wind turbulences or because the substances build up at the areas with the lowest altitude. In these situations, obtaining early and reliable information about

the nature and spread of the contamination is crucial to coordinate effective response and rescue operations.

UAVs can provide high-resolution data in the immediate aftermath of an incident without exposing response personnel to contaminants during the reconnaissance phase. These systems can be remotely controlled by an operator or automatically execute missions by following a trajectory of predetermined waypoints. At each of these locations, the UAVs survey the corresponding ground area using, e.g., a hyperspectral sensor system. Each image covers a certain ground area. These images need to be taken at a standstill to avoid blurred results. Therefore, the relevant area of interest can be discretized into a number of candidate sampling locations such that the combination of all images provides a complete map of the area.

Existing tools for planning UAV missions are either based on a manually specified sequence of locations or ensure the complete coverage of a given target area using one or more UAVs. The limited flight times of small UAVs prevent such a complete survey if the affected region is large. In this case, the problem of selecting appropriate sampling locations arises.

We propose a mission planning approach dedicated to planning informative tours for UAVs through a selected subset of candidate sensing locations such that first responders can be provided with a reliable overview of the situation. This planning problem consists of three simultaneous decisions:

1. The selection of appropriate sensing locations,
2. The assignment of the selected locations to UAVs, and
3. The determination of routes through the assigned target locations for each UAV.

The routes have to comply with the flying time constraints of the UAVs. They are furthermore restricted by the specified take-off and landing locations, which may lay outside the specified target area. We call this problem the *mission planning problem for emergency surveillance (MPPES)*.

Following the definition by Irnich et al. (2014), the second and third decisions arise in nearly all types of vehicle routing problems. The first decision is crucial for planning informative missions and mainly characterizes the problem at hand. To increase the informativeness of the missions, this selection has to take into account that the distribution of hazardous substances within a region exhibits positive spatial correlations. This effect means that locations close to one another are likely to be affected to a similar extent, whereas locations further apart do not show such a relation. Hence, samples in close proximity yield less information about the overall distribution of a contaminant than samples taken further apart at locations not correlated with one another. As a consequence, the benefit of any target point is dependent on all other selected sensing locations and cannot be determined in isolation. This interrelationship has to be adequately considered when planning UAV missions.

The priority of sample locations depends on the risk of contamination and the number of individuals that may be exposed to the substances. This ensures that the focus of the

UAV mission lies in areas where an intervention of the response personnel is most likely to be necessary. Jointly considering spatial interdependencies and the target locations' priorities in MPPES is essential to achieve both accurate information in highly prioritized areas and a reliable overview of the distribution of contaminants throughout the entire region.

As it is the case today in the "manual" ground-based reconnaissance process, we investigate the emergency rapid mapping problem in a static setting. In particular, we assume that priorities and the distribution of gases do not change during the mission. These assumptions can be argued, as response teams consider the distribution of the contaminants and the distribution of the population as more or less stable during the considerably shorter time required for the UAV missions compared to ground-based reconnaissance.

4.2 Input data

Before planning the UAV mission, the operator of the system specifies the central planning parameters. At the very least, this includes the following information:

- A rectangular bounding box that encloses the affected region.
- The number and characteristics of the UAVs, especially the available flight time and maximum cruise speed.
- The locations from which these UAVs are deployed and the locations to which they have to return after their mission is finished.
- Priorities assigned to the sampling locations, e.g., based on population data, which can be obtained in a finely grained resolution of 100×100 m for Germany (Statistisches Bundesamt (Destatis), 2018).
- The desired ground sampling resolution of the obtained images.

The target area is discretized into a grid of evenly distributed target locations, each one representing the center of the area that can be covered with one image taken during the flight. These centers make up the set of candidate sampling locations for the UAVs.

4.3 Illustrative example

In Figure 4.1, we illustrate how priorities in the target area and spatial coverage can be taken into account for planning informative UAV missions. The images depict a scenario where a target area of approximately 2.5×2.5 km^2 needs to be surveyed. The colors indicate the priorities, with red representing highly relevant locations, and blue representing lower-ranked ones. Candidate sampling locations for the UAV are indicated at the flight altitude of 120 m, while the route is depicted as a solid line.

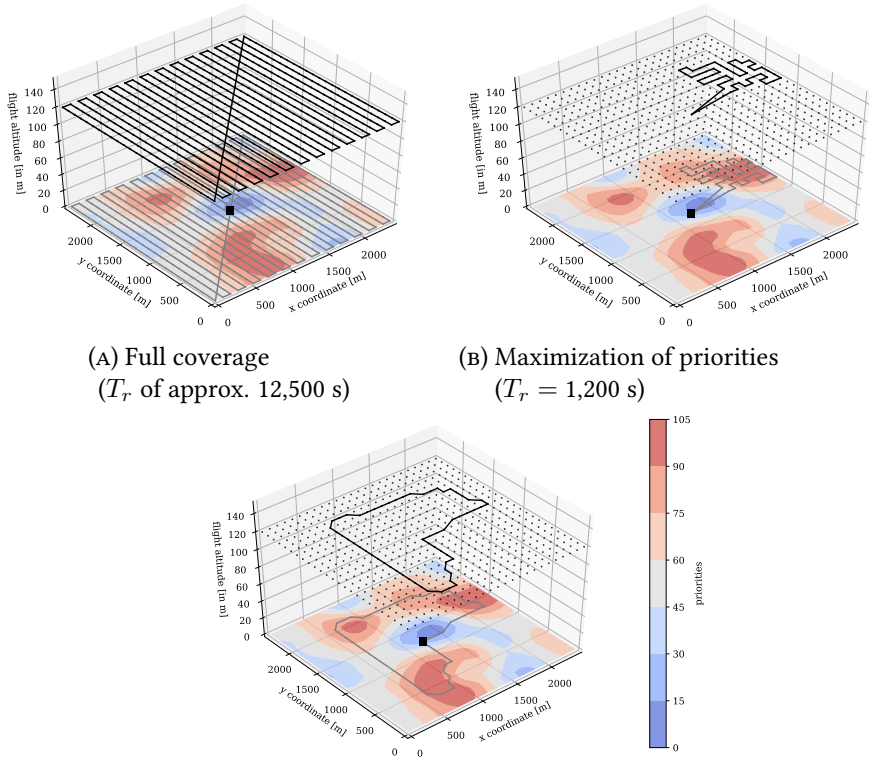

(A) Full coverage (T_r of approx. 12,500 s)

(B) Maximization of priorities ($T_r = 1{,}200$ s)

(C) Emergency rapid mapping ($T_r = 1{,}200$ s)

FIGURE 4.1: Illustration of different mission planning concepts.

Figures 4.1a to 4.1c show three concepts for surveying this area. In all cases, the UAV departs from and returns to a location indicated in the middle of the target region. Figure 4.1a shows a mission pattern that ensures full ground coverage, a concept that is often used in environmental sensing (see, e.g., Stachniss et al., 2009). While providing complete information about the target area, this UAV mission is not feasible in practice due to the long flight time. In contrast, Figure 4.1b represents the vehicle path resulting from maximizing the sum of priorities of the surveyed locations given a realistic flight time limitation. In this solution, the UAV quickly travels to the area with the highest priorities and stays there. Finally, in Figure 4.1c, we depict a mission plan that combines priorities and spatial coverage. This plan offers two advantages: First, it includes all highly prioritized regions. Second, distributing the sample locations across the area allows a more accurate prediction of the distribution of hazardous substances.

4.4 Outline of Part II

In this part, we study the mission planning problem for emergency surveillance (MPPES), which focuses on planning informative tours for surveying a large area to derive accu-

rate information about the extent and the spatial distribution of contaminants in the immediate aftermath of an emergency. To achieve informative tours that yield a reliable overview of the overall distribution of hazardous substances, this approach has to account for the spatial correlation inherent in the observed physical processes, e.g., when considering an airborne substance. Hence, the relative benefit of additional sensing locations depends on the locations' distance to all other selected targets. This means that interdependencies between sampling locations within one vehicle tour as well as between different tours have to be considered during planning, thus introducing new complexity to the planning problem.

In the following chapters, we study models and solution approaches for solving the MPPES. We first summarize the existing work on vehicle routing problems involving the selection of targets under consideration of spatial interdependencies in Chapter 5. We conclude this summary with an overview of the main research gaps that are addressed in this work in Chapter 6. Modeling variants for planning informative vehicle tours are introduced in Chapter 7. Chapter 8 presents two solution approaches, an exact variant based on dynamic programming and a two-phase multi-start adaptive large neighborhood search. We evaluate the models and solution approaches in an extensive simulative study in Chapter 9.

5 Vehicle routing problems with spatial coverage

The mission planning problem detailed in the previous chapter shares similarities with other fields of research: In the literature on vehicle routing, the problem of selecting a subset of candidate locations in the face of tight resource constraints has been considered in the form of the orienteering problem (OP) and the team orienteering problem (TOP). Spatial interdependencies, i.e., relationships between locations that are close to one another, have been addressed in the form of the covering tour problem (CTP) and its variants. In the field of robotics, the problem of routing an unmanned vehicle in a spatially autocorrelated random field is known as informative path planning (IPP).

In this section, we give an extensive overview of the models and solution approaches derived in these three lines of research. We first establish a terminology that allows us to classify and compare these approaches in Section 5.1. In Sections 5.2 to 5.5, we discuss models and approaches originating in these fields in more detail.

5.1 Terminology and classifications

The different problem variants can be distinguished by three major characteristics: the form of coverage, the different types of locations that are involved, and the planning objective.

Nodes and vehicles

Especially in the vehicle routing (VRP) literature, it is common to define a planning problem over a graph, with nodes representing locations such as customers and arcs represent possible connections between these locations. In this thesis, we follow this terminology, using the generic term *node* to refer to any representation of a location that can be part of a planned path or tour. Similarly, we use the term *vehicle* as an umbrella term covering the vastly varying types of mobile resources in different fields.

Spatial coverage

Considering spatial correlations within the observed area is an essential aspect of planning informative tours. From a modeling perspective, this aspect is closely related to the concept of *coverage* that has been introduced in VRP literature, as both concepts imply some degree of similarity or relatedness between nearby nodes. In the following, we use the generic term *spatial coverage* to describe any form of distance-dependent relations between nodes that are considered in a model. These interdependencies can be used to treat nodes that are not selected, but close to a vehicle route, as "visited" for the purpose of determining total profit or fulfilling minimum service requirements. Conversely, without considering coverage, only locations that immediately visited by a vehicle are relevant for determining the objective value or compliance with planning requirements.

We can further distinguish between different forms of coverage that have been discussed in literature:

Complete coverage means that any node can essentially be considered as served by a vehicle as long as at least one vehicle route contains a visit within a given maximum distance.

Partial coverage indicates that any unvisited, covered node only provide partial benefits compared to a solution where they are immediately served.

Cumulative coverage is a more complex form of partial coverage. We use this term to refer to models with partial coverage, in which each additional visit yields additional benefits.

Multiple coverage models, on the other hand, are a generalization of the concept of complete coverage. Here, constraints require a minimum number of vehicle stops near an unvisited node for this node to be considered as covered.

Gaussian process (GP) coverage refers to approaches that use the probabilistic models discussed in Chapter 3 to determine the relationship between nodes.

In all cases, the maximum benefit of a covered node can never exceed the benefit provided if this node were served by a vehicle.

Node types

Nodes differ in the form of coverage that applies to them and in whether they are optional or mandatory:

Type I (optional visit with mandatory coverage) designates nodes that may or may not be included in a vehicle tour, but that must be covered if they are not served directly by a vehicle.

Type II (mandatory visit) refers to nodes that must always be served directly by a vehicle. Coverage aspects are irrelevant for these nodes.

Type III (optional visit without coverage) indicates candidate nodes for a vehicle tour that do not need to be served directly and that cannot be covered. They mainly serve to provide structure for vehicle tours.

Type IV (mandatory coverage) describes nodes that cannot be included in a vehicle tour but that must be covered by one.

Type V (optional visit with optional coverage) involves nodes that do not need to be visited but that provide a benefit if they are either served or covered.

Planning objectives

Generally speaking, the objective of these planning problems is either input minimization or profit maximization. In the following, we classify all problems that focus on profit maximization and do not enforce all nodes to be either served directly or be covered as *(team) orienteering problems*. Problems dealing with the minimization of resource consumption while treating either direct service or coverage of certain nodes as mandatory are considered as variants of the *covering tour problem*. Finally, we denote all problems that plan vehicle tours using statistical means for determining route quality as *informative path planning* problems. Unlike the TOP, these problems do not have predefined profit values that are associated with candidate nodes.

5.2 Covering tour problems

Coverage aspects have been considered in VRP literature in form of the covering salesman problem (CSP), the covering tour problem (CTP), and its multi-vehicle variant (m-CTP). This class of problems deals with the determination of cost-minimal routes such that every node that is not included in a vehicle tour is within a given maximum distance to a node directly served by a vehicle. An overview of model variants and solution approaches is provided in Table 5.1.

The CSP was first introduced and formulated by Current and Schilling (1989) as a variant of the TSP where all nodes have to be within a predetermined maximum distance to a node on the vehicle tour. Afterward, the problem class did not receive much attention until Gendreau et al. (1997) proposed a more general model that distinguishes between nodes that have to be visited, nodes that cannot be visited directly but have to be covered, and optional nodes that do not need to be visited or covered but can be included in vehicle tours to provide coverage for other nodes nearby. In contrast to the TOP, these optional nodes do not yield rewards or benefits. Instead, they serve to provide coverage for other nodes. Both publications propose heuristics based on local search and adapt heuristics dedicated to solving set covering problems to determine the subset of nodes that are included in the vehicle tours. Furthermore, Gendreau et al. (1997) propose the first exact solution approach for this problem class based on a branch-and-cut scheme. The authors note that the performance of their solution approach is highly dependent on

Paper	Model	Objective	Nodes	Coverage	Vehicles	Solution approach	Benchmark instances
Current and Schilling (1989)	CSP	Min. distance	I	Complete	1	Heuristic based on set covering and TSP approaches	One sample problem with 21 nodes
Gendreau et al. (1997)	CTP	Min. distance	II, III, IV	Complete	1	Branch-and-cut, heuristic based on set covering and TSP approaches	Instances with up to 100 nodes of type II or III, up to 500 covered nodes (type IV)
Hachicha et al. (2000)	m-CTP	Min. distance	II, III, IV	Complete	m	Heuristic based on set covering and VRP heuristics	Instances with up to 300 nodes of type II or III, up to 400 nodes of type IV
Golden et al. (2012)	Generalized CSP	Min. travel and visit costs	I	Multiple	1	Two local search heuristics	Adapted TSPLIB instances with up to 1,000 nodes
Naji-Azimi et al. (2012)	m-CTP	Min. distance	III, IV	Complete for multiple commodities	m	Multi-start local search heuristic	Random instances with up to 20 nodes of type III and 100 nodes to be covered (type IV)
Hà et al. (2013)	m-CTP	Min. distance	II, III, IV	Complete	m	Branch-and-cut, metaheuristics	Adapted TSPLIB instances with up to 200 nodes (types II to IV)
Allahyari et al. (2015)	m-CTP	Min. travel and allocation cost	I	Complete	m	GRASP and iterated local search	Random instances with up to 90 nodes and 4 depots

Table 5.1: Related literature on the covering tour problem.

how many nodes can provide coverage, as increasing coverage makes instances easier to solve. The single-vehicle problem was further discussed by Golden et al. (2012) who propose a problem generalization where some nodes need to be covered or served multiple times. The authors refer to this problem as the generalized covering salesman problem. They propose a local search approach where node allocation is heuristically improved by removing and reinserting nodes in vehicle tours taking into account their cost and the coverage they provide. Furthermore, the authors demonstrate that incorporating classic TSP improvement moves yields significant improvements.

The first multi-vehicle variant was discussed by Hachicha et al. (2000) as an extension of the CTP, denoted the m-CTP. Similar to earlier approaches for the CSP and CTP, the authors combine approaches for solving a set covering problem with well known VRP heuristics, notably the savings and sweep heuristics and a route-first/cluster-second approach. Naji-Azimi et al. (2012) address the m-CTP in the context of providing humanitarian aid with multiple commodities and split deliveries. The authors treat coverage as mandatory, but it may be provided by several distinct covering nodes. They demonstrate the limits of exact approaches for coverage problems and propose a local search with added diversification steps that can significantly improve solution quality and runtime. The m-CTP is furthermore considered by Hà et al. (2013), who propose a branch-and-cut approach and a metaheuristic for solving this problem. In their heuristic approach, they first determine covering subsets. Finding tours with minimum distance through these subsets can be considered as a VRP with unit demand. Solutions are improved by local search moves, both classical VRP moves and operators that replace nodes within the tours. Finally, Allahyari et al. (2015) extend the m-CTP to the multi-depot case. In their model, the allocation of an unvisited customer to a vehicle stop incurs a specified cost, which offers incentives for visiting more stops directly. The authors combine a greedy randomized adaptive search procedure and iterated local search for solving this problem.

5.3 Orienteering and team orienteering

The orienteering problem and the team orienteering problem are variants of the TSP and VRP, respectively, in which not all nodes can be visited and the optimization objective is profit maximization rather than minimizing resource utilization (e.g. distance or visit costs). This problem has been widely studied in the VRP literature, and the availability of benchmark instances (Tsiligirides (1984) and Chao et al. (1996a) for the OP, Chao et al. (1996b) and Dang et al. (2013) for the TOP) has enabled the continuous improvement and comparison of solution approaches. Recent surveys (Vansteenwegen et al. (2011) and Gunawan et al. (2016)) provide a detailed overview of this field of research. Due to the popularity of the TOP, a significant body of work has been published. An overview of relevant publications is given in Table 5.2. In this review, we focus on the progress made in the last decade, i.e., in publications from 2007 or later, and include the approaches that are state-of-the-art in terms of solution quality or computational performance.

Most of the work in this area has focused on a wide range of local search and ge-

Paper	Model	Objective	Nodes	Coverage	Vehicles	Solution approach	Benchmark instances
Archetti et al. (2007)	TOP	Max. reward	III	None	m	Tabu search, variable neighborhood search	Chao et al. (1996b)
Ke et al. (2008)	TOP	Max. reward	III	None	m	Ant colony optimization	Chao et al. (1996b)
Vansteenwegen et al. (2009a)	OP, TOP	Max. reward	III	None	m	Guided local search	OP benchmarks, Chao et al. (1996b) (TOP)
Vansteenwegen et al. (2009b)	TOP	Max. reward	III	None	m	Variable neighborhood search	Chao et al. (1996b)
Souffriau et al. (2010)	TOP	Max. reward	III	None	m	GRASP with Path Relinking	Chao et al. (1996b)
Dang et al. (2013)	TOP	Max. reward	III	None	m	Branch-and-cut	Chao et al. (1996b)
Dang et al. (2013)	TOP	Max. reward	III	None	m	PSOiA	Chao et al. (1996b), new instances with up to 399 nodes
Yu et al. (2014)	CorTOP	Max. reward	V	Cumulative	m	MIQP	Structured instances with up to 49 nodes
Vidal et al. (2015)	VRP with profits	Max. reward	III	None	m	Heuristics with implicit customer selection	Chao et al. (1996b)
Keshtkaran et al. (2016)	TOP	Max. reward	III	None	m	Branch-and-cut	Chao et al. (1996b)
Ke et al. (2016)	TOP	Max. reward	III	None	m	Pareto mimic algorithm	Chao et al. (1996b) and Dang et al. (2013)
Ozbaygin et al. (2016)	Maximal CSP	Max. covered demand	V	Partial	1	Branch-and-cut	VRP instances with up to 200 nodes
Archetti et al. (2018)	SOP	Max. reward of covered clusters	V	Complete	1	MIP, matheuristic	New instances based on benchmarks for the generalized TSP
Pěnička et al. (2019)	SOP	Max. reward of covered clusters	V	Complete	1	ILP, variable neighborhood search	Instances by Archetti et al. (2018)

Table 5.2: Related literature on the team orienteering problem.

netic algorithms. Archetti et al. (2007) solve the TOP using generalized tabu search and variable neighborhood search (VNS) algorithms. They show that VNS is superior in both solution quality and runtime, and achieve significant improvements on the previously best-known solutions. Ke et al. (2008) propose an ant colony optimization (ACO) approach and compare four methods for constructing candidate solutions. They achieve further improvements, notably on the largest benchmark instances by Chao et al. (1996b). Vansteenwegen et al. (2009a) propose a guided local search (GLS) approach that combines heuristics tailored to orienteering problems with TSP heuristics to improve the planned routes. They perform extensive experiments on the available benchmark instances for both the OP and the TOP, which highlight the computational efficiency of the proposed approach. The same year, the authors proposed a variable neighborhood search that yielded further improvements (Vansteenwegen et al., 2009b). This algorithm is based on similar local search moves but contains powerful shaking phases to leave locally optimal solutions by removing chains of candidate nodes. Souffriau et al. (2010) introduce a greedy randomized adaptive search procedure (GRASP) with path relinking. The primary search procedure applies local search moves for orienteering problems and the TSP similar to those used in previous publications. Furthermore, new solutions are created based on a pool of elite solutions with high quality.

Dang et al. (2013) introduce particle swarm optimization-inspired algorithm (PSOiA) that extends their previous work on population-based solution approaches. Here, each particle represents a solution to the TOP in the form of a giant tour. Solutions are improved using destruction and recreation operators combined with TSP search moves. Crossover operators create new solutions by merging subsequences from successful tours. Recognizing that most of the recent work consistently achieve high-quality solutions on the benchmark set available until that date, the authors furthermore introduce a new set of larger instances for the TOP with up to 400 candidate nodes.

Vidal et al. (2015) published a study on the VRP with profits, which is a generalization of the TOP and related problems. The authors propose a new neighborhood concept based on exhaustive solutions, i.e., solutions where resource constraints are relaxed such that all visits can be included in vehicle tours. Feasible TOP solutions can be determined in a selection phase by solving a resource constraint shortest path problem (RCSPP) for each vehicle. The authors evaluate the new neighborhood concept by integrating it into three established VRP solution heuristic: a local improvement procedure, an iterated local search procedure, and a hybrid genetic algorithm (HGA). In their evaluation, the authors highlight the trade-off between solution quality and computational effort. The results show that while the HGA achieves the overall best performance of all published approaches up to this date, it computationally expensive. The solution quality of the local search approach is slightly lower. Runtime is significantly reduced in this case.

None of these more recent publications consider the new set of larger instances proposed by Dang et al. (2013). To our knowledge, Ke et al. (2016) are the only ones to report results for these instances since they have been introduced. The authors propose a Pareto mimic algorithm that maintains a population of solutions based on Pareto dominance. Search is based on two new operators. Similar to path relinking (see Souffriau

et al., 2010), a so-called mimic operator is used to build new solutions based on an incumbent one. A swallow operator inserts desirable but infeasible nodes. Finally, local search steps are used to repair and improve.

Exact methods for solving orienteering problems have been less prevalent in literature. Dang et al. (2013) propose a branch-and-cut algorithm strengthened by valid inequalities for symmetry breaking, clique cuts, and bounds on visited customers and profits. The algorithm improves on earlier exact approaches for the set of instances by Chao et al. (1996b), but has not been evaluated on the larger benchmark instances introduced by Dang et al. (2013). More recently, Keshtkaran et al. (2016) propose and compare branch-and-price and a branch-and-cut-and-price schemes. Using the latter algorithm, they can solve more instances among the sets proposed by Chao et al. (1996b) with proven optimality than earlier approaches. However, several of the largest instances with around 60 to 100 customers remain intractable.

5.4 Combining coverage and profit maximization

Coverage aspects in the context of profit maximization have been addressed more frequently in recent years. Yu et al. (2014) and Ozbaygin et al. (2016) propose models that are based on TOP models with additional consideration of coverage constraints such that nodes that are not included in a vehicle tour but are within a specified maximum distance to a visited node provide a positive contribution to the objective function. This contribution is less than the benefit yielded by directly including the node in a vehicle tour.

Yu et al. (2014) introduce the correlated team orienteering problem (CorTOP) as a variant of the TOP with the explicit objective of integrating information about spatial correlations in the model. Their model considers coverage as cumulative, i.e., additional stops within covering distance provide an additional benefit. In this model, the reward associated with a specific node can be divided among the nodes within covering distance. When all covering nodes are included in a vehicle tour, the full reward of an unvisited target node is provided. The authors propose a mixed integer quadratic programming (MIQP) formulation for solving the problem using a commercial solver. However, computation times quickly become impractical even for small instances: The largest instance solved in the single-vehicle case involves 49 candidate locations; this number is reduced to 36 and less for instances with several vehicles. Even for these instance sizes, computation time scales poorly with route duration and number of vehicles. Still, the results of this model are promising with respect to the planning problem addressed in this work due to their consideration of coverage aspects combined with profit maximization. This model provides the basis of the models derived in Chapter 7.

Ozbaygin et al. (2016) propose a branch-and-cut approach for a single-vehicle problem called the maximal covering salesman problem (MaxCSP) with the objective of maximizing covered demand. The authors consider unvisited nodes as covered as long as they are within a given distance to a visited node. Only a fixed percentage of a node's de-

mand can be covered that way. This percentage is independent of the number of visited nodes that may be able to provide coverage, or the distance between the covered and the covering node.

A related problem is the set orienteering problem (SOP) introduced by Archetti et al. (2018), who also introduced a set of benchmark instances. In this problem variant, customers are grouped into clusters and visiting one customer within a cluster allows collecting the full profit associated with this cluster. Clusters can, but do not have to, represent spatial relations. Additional visits within one cluster do not provide additional benefits. Both heuristic and exact solution approaches have been proposed (Pěnička et al., 2019). However, only the single-vehicle case has been considered to this date.

5.5 Informative path planning

The literature in the domain of robotics has its origins in sensor placement problems without consideration of routing constraints. These problems address the question of designing sensor networks for monitoring environmental phenomena. An example in this line of work is the study by Krause et al. (2008), who solve the placement problem by iteratively selecting samples providing the largest incremental benefit with respect to some information measure.

This placement problem has been extended to mobile sensor systems, leading to the informative path planning (IPP) problem that is concerned with planning vehicle trajectories for monitoring a phenomenon that varies in time and space while respecting the vehicles' maximum flight time. The objective is to maximize a measure for the information gain that is achieved by the selected sensing locations. This review focuses on IPP variants with a finite set of target locations. An overview of relevant publications is given in Table 5.3.

Most of the work on IPP addresses environmental monitoring applications, specifically oceanic monitoring using autonomous underwater vehicles, where large areas have to be surveyed in limited time. Similar to team orienteering, these models seek to maximize some measure indicating the benefit of the vehicle tours. In contrast, IPP approaches do not consider profits associated with specific target locations. Instead, they use the variability measures discussed in Section 3.2.3 for determining the information gain achieved by the vehicles with respect to the observed phenomenon. In other words, they seek to determine sensing locations such that the uncertainty remaining in the process is minimal.

These approaches make use of the fact that the posterior covariance matrix can be determined based on the locations at which observations are planned, and is independent of the realizations of the random field (i.e., the actual measurements). This allows to plan optimal tours with regard to these informativeness measures but requires an accurate model for the GP prior distribution, e.g., based on preliminary information or previous samples.

The Gaussian process models that are used in these approaches implicitly introduce

Paper	Model	Objective	Nodes	Coverage	Vehicles	Solution approach	Benchmark instances
Singh et al. (2007)	IPP	Max. mutual information	V	GP	m	Sequential recursive greedy algorithm	Real-world and simulated test data with less than 30 candidate sensing locations
Yilmaz et al. (2008)	IPP	Max. prior uncertainty of selected nodes	III	None	m	Mixed integer linear programming	Real-world data, up to 45×33 locations, tight route length limitations
Singh et al. (2009)	IPP	Max. mutual information	V	GP	m	Sequential recursive greedy algorithm	Real-world and simulated test data, less with 30 candidate sensing locations
Binney et al. (2010)	IPP	Max. mutual information	V	GP	1	Recursive greedy algorithm	Real-world data with 16 candidate waypoints
Binney and Sukhatme (2012)	IPP	Max. avg. reduction in variance	V	GP	1	Branch-and-bound	Artificial data with 25 candidate waypoints
Binney et al. (2013)	IPP	Max. avg. reduction in variance of selected locations	V	GP	1	Generalized recursive greedy algorithm	Simulated data sets, 25 and 33 candidate waypoints resp.; real-world test data with 37 waypoints
Hollinger and Sukhatme (2014)	IPP	Several informativeness measures	V	GP	1	Rapidly-exploring information gathering algorithm	Artificial data with up to 100 candidate locations, real-world instance on continuous fields

Table 5.3: Related literature on informative path planning.

coverage aspects in the objective function: As nearby nodes are similar, observations in close proximity yield less overall information. Hence, the marginal contribution of an observation to the objective value depends on which other locations are visited. At the same time, more measurements always remain beneficial.

Most of the solution approaches discussed in this line of research have focused on optimal algorithms or approximation guarantees to achieve performance guarantees in the face of high sensing costs. Singh et al. (2007) formally define the IPP for multiple vehicles. In their model, they use mutual information (MI) (see Section 3.2.4) to determine the informativeness of a solution. As this objective function is submodular, i.e., the marginal benefit of additional samples decreases with sample size, their solution approach is based on the recursive greedy algorithm proposed by Chekuri and Pal (2005). The authors further improve the running time by decomposing the target area into cells such that locations in different cells can be treated as uncorrelated. The greedy algorithm is run on these cells instead of the full set of candidate locations. Singh et al. (2007) furthermore address the multi-vehicle case by applying the recursive greedy algorithm to a series of single-vehicle problems sequentially, in each step taking into account the information obtained using all previously planned vehicle routes. These concepts are also discussed and evaluated in detail in a later publication (Singh et al., 2009). A version of the recursive greedy algorithm is furthermore used by Binney et al. (2010), who solve an IPP variant with time windows that limit the accessibility of certain areas. The authors also demonstrate how available information, for example, obtained using previous missions, can be incorporated to improve subsequent tours.

A branch-and-bound algorithm for the single-vehicle IPP is proposed by Binney and Sukhatme (2012). Due to the high runtime required for solving even small instances to optimality, the authors furthermore limit the search space. This significantly improves runtime, but problems remain computationally intractable for vehicle routes comprising more than around 15 locations. Binney et al. (2013) extend this model and the recursive greedy approach to a case with time-varying fields. Similar to previous approaches, the algorithm does not scale well for instances with more than a few dozen candidate locations. Hollinger and Sukhatme (2014) propose a rapidly-exploring information gathering algorithm, which iteratively assigns random sampling locations to vehicle routes and expands vehicle paths towards these nodes. This approach is applicable to both discrete and continuous planning problems. However, its performance is highly dependent on maximum route length.

Following a different concept for modeling informativeness, Yilmaz et al. (2008) use the prior uncertainty of a location as an indicator for its relevance, and propose a model maximizing the sum of these values subject to budget constraints. The objective function resembles the classical TOP without considering coverage. Vehicle routes are further constrained by motion and communication constraints. The problem is formulated as a mixed integer linear programming (MIP) model and solved using a commercial solver.

5.6 Summary

In this chapter, we provided an overview of three lines of research related to the UAV mission planning problem derived in Chapter 4 that either consider aspects of spatial coverage or focus on the selection of targets based on their contribution to an objective function. This overview serves as the foundation of the research gap identified in the next chapter.

We introduced a new classification, in which we classified the models proposed in literature based on their planning objective, distinguishing between orienteering problems, covering tour problems, and informative path planning. The problem of selecting subsets of nodes based on reward values associated with each target has been discussed in the form of the (team) orienteering problem. The informative path planning problem focuses on profit maximization under consideration of route lengths constraint. These problem variants are typically based on the sampling strategies discussed in Section 3.2.4. The objective functions introduce spatial interdependencies into the problem, as the marginal benefit of each target depends on all other selected locations in the vehicles' tours. Finally, spatial relationships between nodes are at the center of the covering tour problem and its variants. The focus of this problem class is to determine minimum cost routes such that all specified target nodes are covered by at least one node in a vehicle tour.

We furthermore gave an overview of solution approaches that have been derived for each problem class. The TOP has mostly been addressed using heuristic approaches, which allow the efficient solution of large-scale instances. The same is true for the CTP variants. In contrast, the literature on IPP has typically focused on approximative results that provide approximation guarantees but do not scale as well.

6 Research gap and contributions

In the previous chapter, we have given an extensive overview of the work published in three lines of research that are, in some form, related to the MPPES. In this chapter, we evaluate the models and solution approaches with regard to their suitability for our use case and identify the research gaps that need to be addressed to successfully solve the mission planning problem.

6.1 Research gap

The problem variants related to the MPPES are summarized in Table 6.1. Comparing their characteristics with the requirements for planning UAV missions discussed in Chapter 4, we can derive three major aspects that need to be addressed for solving this planning problem successfully: the joint consideration of coverage and priorities, the derivation of approximative models for spatial interdependencies, and efficient solution approaches for the mission planning problem.

Joint consideration of coverage and priorities The mission planning problem for UAVs combines two practical objectives: Achieving a reliable prediction about the spatial distribution, which requires sampling locations that are well distributed across the area of interest, and obtaining accurate information about prioritized locations. As a consequence, it is necessary to jointly consider each target node's priority and the information it can provide about the surrounding region.

Most publications, however, focus on either profit maximization or the minimization of resource utilization subject to aspects of spatial coverage or spatial correlation. Despite their respective popularity, these two aspects are rarely combined: In the literature on the TOP, the profits collected at each node are generally treated as being independent of the locations visited in the remainder of the tour. In contrast, CTP variants do not involve priorities or utilities, as either direct service or coverage of nodes is considered mandatory. Finally, models derived in IPP literature have explicitly focused on node selection under consideration of spatial interdependencies, but do so without including priorities or benefits associated with individual targets. To the best of our knowledge,

Problem variant	First reference	Vehicles	Nodes	Coverage	Objective	Solution approach
Orienteering problem (OP)	Tsiligirides (1984)	1	III	None	Max. sum of priorities	Heuristics,exact approaches
Covering salesman problem (CSP)	Current and Schilling (1989)	1	I	Complete	Min. distance	Heuristics,exact approaches
Team orienteering problem (TOP)	Chao et al. (1996b)	m	III	None	Max. sum of priorities	Heuristics,exact approaches
Covering tour problem (CTP)	Gendreau et al. (1997)	1	II, III, IV	Complete	Min. distance	Heuristics, branch-and-cut
m-covering tour problem (m-CTP)	Hachicha et al. (2000)	m	II, III, IV	Complete	Min. distance	Heuristics, branch-and-cut
Informative path planning (IPP)	Singh et al. (2007)	1,m^*	V	GP	Min. prediction variability	Exact and approximative approaches
Correlated team orienteering problem (CorTOP)	Yu et al. (2014)	m	V	Partial, cumulative	Max. sum of direct and covered priorities	Branch-and-bound
Maximal covering salesman problem (MaxCSP)	Ozbaygin et al. (2016)	1	V	Partial	Max. sum of direct and covered priorities	Branch-and-cut
Set orienteering problem (SOP)	Archetti et al. (2018)	1	V	Complete	Max. sum of direct and covered priorities	Heuristics,exact approaches
This contribution		m	V	Partial, cumulative	Dynamic programming and heuristic approach	

(*) Multi-vehicle cases are typically solved by sequentially planning single-vehicle tours

TABLE 6.1: Summary of problem variants with profit maximization or coverage constraints.

only the models of Yu et al. (2014) and Ozbaygin et al. (2016) combine profit maximization with coarse models for spatial correlations. Among these two, only the CorTOP by Yu et al. (2014) accounts for the additional benefits that can be gained by making several observations in the vicinity of an unsurveyed location.

Approximative models for spatial interdependencies Considering the relationships between different locations in space is crucial for planning informative UAV missions, which have to account for the information that each surveyed location provides about the surrounding area. The variants discussed in the context of the CSP or CTP consider very coarse forms of spatial coverage. In these models, the benefit of covering a node is independent of the number of or distance to the nodes that provide coverage. In contrast, the IPP relies on Gaussian process models of the observed distribution of hazardous substances. These variants provide sophisticated models for spatial interdependencies, but are computationally expensive ($\mathcal{O}(n^3)$ in the number of sampled locations). This means that the impact of local changes to a solution cannot be computed efficiently – a prerequisite for many successful VRP approaches such as local search based heuristics. Models that are sufficiently accurate but can be evaluated efficiently are missing, even if the CorTOP is a promising first step.

Efficient solution approaches for the mission planning problem To this date, no efficient solution approach exists that is applicable to the mission planning problem. Efficient heuristics designed for TOP do not account for spatial interdependencies between nodes. The heuristics developed for the CTP, in turn, exploit the fact that coverage is complete and all nodes need to be covered. This means that these approaches are not applicable for a model which considers partial and cumulative coverage. For IPP variants, most authors have focused on approximative approaches for the single-vehicle case. While these algorithms scale better than exact solution methods, they do not achieve acceptable computation times for our scenario, with benchmark instances typically limited to less than 100 target locations.

Considering the two models closest to our application, only exact solution approaches have been proposed to this date. The maximal covering salesman problem is solved using branch-and-cut schemes that are tailored to the single-vehicle case with specific coverage requirements (Ozbaygin et al., 2016). The MIQP for the correlated team orienteering problem scales poorly with instance size and the number of vehicles involved (Yu et al., 2014). None of these approaches achieve computation times that are acceptable for our scenario.

6.2 Contributions

By comparing the existing models and applications with the problem statement in Chapter 4, we have identified three points that are relevant for the MPPES, but have received

little attention to this date. In this thesis, we seek to address these gaps. Our contributions to this goal can be summarized as follows:

GCorTOP modeling approach In Chapter 7, we derive new models that account for spatial interdependencies between locations in the target area. These models can be seen as a generalization of the models with partial coverage introduced by Yu et al. (2014) and Ozbaygin et al. (2016). They provide more detailed representations of spatial interdependencies that mirror the ones offered by GP models. Most importantly, the strength of correlation decreases with increasing distance. This way, the information provided by samples at different distances can be modeled more accurately. Furthermore, the proposed models can be combined with priorities that are used to measure the relative importance of obtaining information about specific targets. Compared to approaches derived in IPP literature, these models do not achieve the same level of accuracy in representing the correlations within a process. However, the computational effort is lower, thereby allowing us to integrate these models into local search approaches.

Two-phase multi-start adaptive large neighborhood search (2MLS) To solve large-scale instances with several hundred nodes in reasonable computing time, we introduce an efficient solution concept in form of the 2MLS that incorporates knowledge about spatial correlations for the construction of vehicle routes in Chapter 8. We propose concepts that yield good starting solutions within very short computations times, either by identifying covering nodes or by decomposing the target area. We furthermore design search strategies that specifically emphasize explorative missions, i.e., UAV tours that cover large parts of the target area, while maintaining a balance between exploration and the coverage of high-priority areas. We also show how our proposed models can be integrated into this approach, and use information about the marginal contribution of nodes and bounds to the objective function for speeding up the search. An exact solution approach is introduced for benchmarking purposes.

Benchmark instances for the MPPES and simulative study We study the performance of our approaches as well as their limitations in an extensive simulative study in Chapter 9. The effectiveness of the 2MLS is demonstrated using benchmark instances for the TOP. We furthermore introduce a set of new benchmark instances modeled after a real-world use case. Based on these instances, we demonstrate the performance of the proposed model and solution heuristic compared to other modeling variants. We show the structural changes that result from considering correlations when planning UAV missions and identify the limitations of this approach. We furthermore investigate the applicability of these techniques for UAV based surveillance from a practical point of view.

7 Models for planning informative UAV missions

The fundamental problem in mission planning is to determine informative tours. This requires some measure to quantify "informativeness" in a meaningful way. This measure determines the selection of locations to be surveyed and is crucial for the accuracy of any spatial interpolation approach based on these samples.

In Section 7.1, we formally define the MPPES introduced in Chapter 4. Based on a basic problem formulation, we review stochastic modeling variants in Section 7.2. We derive approximations for spatial interdependencies in Section 7.3 and show how to integrate priorities in Section 7.4.

7.1 Basic problem formulation

Consistent with the notation introduced in Section 3.2, we denote the set of locations within the two-dimensional target area as $\mathcal{V}$ and the set of sensing locations as $\mathcal{S} \subseteq \mathcal{V}$. When planning informative tours, the selection of the sensing locations $\mathcal{S}$ constitutes the main decision. Each UAV m in the set of available UAVs $\mathcal{M}$ is associated with a starting location s_m^{start} and ending location s_m^{end}. In the following, $\mathcal{N}_S$ and $\mathcal{N}_E$ designate the sets comprising vehicle starting and ending locations, respectively. We refer to the set of all locations as $\mathcal{N}$, with $\mathcal{N} = \mathcal{V} \cup \mathcal{N}_S \cup \mathcal{N}_E$. Target locations $i \in \mathcal{V}$ are associated with priorities $u_i \geq 0$ that specify their relevance to the response units. As discussed in Section 4.2, these priorities can, for example, indicate the size of the affected population. A survey at a location i requires a fixed sensing time τ_i to take and process images at standstill. Each UAV mission is limited by a maximum duration of T_m^{max}. The distance between two locations i and j is denoted as d_{ij}. Traveling from a location i to a location j with $i, j \in \mathcal{N}$ requires a nonnegative travel time τ_{ij} that includes the time necessary for acceleration and deceleration. To simplify the formulation, this also includes the sensing time τ_j at the destination location.

Moreover, we define the set of all feasible routes $\Omega = \bigcup_{m \in \mathcal{M}} \Omega_m$. Each route $r \in \Omega_m$ consists of a sequence of locations $r = (s_m^{start}, i_0, i_1, \ldots, i_n, s_m^{end})$ such that the total travel time of this route $\tau_{s_m^{start} i_0} + \sum_{c=0}^{n-1} \tau_{i_c i_{c+1}} + \tau_{i_n s_m^{end}} \leq T_m^{max}$. The set of included locations

$\mathcal{S}_r = \{i_0, i_1, \ldots, i_n\} \subset \mathcal{V}$ corresponds to the sensing locations selected for the tour r. We denote a feasible solution to this problem as ω, where $\omega = (r_1, \ldots, r_m) \in \Omega_1 \times \ldots \times \Omega_m$ represents one element in the Cartesian product of routes, i.e., one combination of routes with one feasible route per vehicle. Finally, we consider a general measure of informativeness $\mathcal{I}$ to determine the quality of a solution.

We can now model the mission planning problem using binary decision variables y_r with

$$y_r = \begin{cases} 1, & \text{if route } r \in \Omega \text{ is selected,} \\ 0, & \text{otherwise.} \end{cases}$$

Then, the set of sampling locations included in a solution is $\mathcal{S} := \bigcup_{r:y_r=1} \mathcal{S}_r$. This notation can consistently be used both for the stochastic informativeness measures of Section 7.2 and for the models introduced in this work (Section 7.3).

The basic problem of planning informative UAV missions can then be stated as follows:

$$\text{(MPPES)} \quad \max \ \mathcal{I}(\mathcal{S}) \tag{7.1}$$

$$\text{s.t.} \quad \sum_{r \in \Omega_m} y_r = 1 \qquad m \in \mathcal{M} \tag{7.2}$$

$$y_r \in \{0, 1\} \qquad r \in \Omega \tag{7.3}$$

Objective (7.1) maximizes total informativeness, while constraints (7.2) ensure that exactly one route is selected per vehicle. Constraints (7.3) set the variable domains.

7.2 Stochastic informativeness measures

One possible approach for measuring the informativeness $\mathcal{I}$ in Equation (7.1) is to use stochastic means that account for the spatial correlation in the surveyed random field. As we have shown in Section 3.2.4, stochastic process models enable us to measure the quality of a set of sampling locations using the achieved reduction in uncertainty in the process. These models are of particular interest for our application as they allow to consider detailed information about spatial interdependencies in the observed distribution of gases.

7.2.1 Model formulations

As discussed in Section 5.5, IPP models typically rely on Gaussian process models discussed in Chapter 3. We focus on the two variants for representing information gain that are most commonly used in this line of work. Binney and Sukhatme (2012) use the average reduction in variance measure as defined in Equation (3.21) in order to determine

solution quality. The corresponding objective can be stated as:

$$\mathcal{I}^{ARV}(\mathcal{S}) = \frac{1}{|\mathcal{V}|}(\mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}} - \mathrm{tr}(\Sigma_{\mathcal{V},\mathcal{V}|\mathcal{S}})) \tag{7.4}$$

We denote the corresponding optimization problem as IPP-ARV.

Singh et al. (2007) propose a formulation based on the concept of mutual information, see Equation (3.22). We refer to this variant as IPP-MI, which is based on the following objective:

$$\mathcal{I}^{MI}(\mathcal{S}) = \mathcal{H}(Z_{\mathcal{V}\setminus\mathcal{S}}) - \mathcal{H}(Z_{\mathcal{V}\setminus\mathcal{S}|\mathcal{S}}) \tag{7.5}$$

7.2.2 Model and solution characteristics

The popularity of these formulations in IPP literature is due to the fact that they have certain properties that can be exploited in exact or approximative solution approaches, notably submodularity and monotonicity (Singh et al., 2009; Krause and Golovin, 2014; Jawaid and Smith, 2015). Submodularity means that the incremental benefit of additional samples decreases with an increased sample size. This can be formally stated as (Singh et al., 2009; Binney et al., 2010):

$$\mathcal{I}(\mathcal{S}_1 + \{i\}) - \mathcal{I}(\mathcal{S}_1) \geq \mathcal{I}(\mathcal{S}_2 + \{i\}) - \mathcal{I}(\mathcal{S}_2) \text{ for } \mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{V},\ i \in \mathcal{V} \setminus \mathcal{S}_2.$$

Monotonicity means that increasing the number of observations cannot decrease the total informativeness. This is written as (ibid):

$$\mathcal{I}(\mathcal{S}_1) \leq \mathcal{I}(\mathcal{S}_2) \text{ for } \mathcal{S}_1 \subset \mathcal{S}_2.$$

MI is submodular and approximately monotonic for sufficiently small $\mathcal{S}$, as shown by Krause et al. (2008). ARV is monotonic (Binney and Sukhatme, 2012). Although it is not always submodular, submodularity has been shown for practically relevant cases, specifically in the context of environmental sampling (Binney et al., 2013).

7.3 Objective functions for approximating spatial interdependencies

The stochastic informativeness measures discussed above provide sophisticated models for spatial interdependencies but are computationally expensive. In the following, we derive objective functions that show similar characteristics regarding the spatial aspects of a solution while avoiding the computational overhead of a GP regression.

Approximating spatial correlations

In statistical interpolation approaches, interdependencies between pairs of locations are described using a kernel function k. Using the correlation rather than the covariance as a normalized measure of spatial similarity, a kernel is a function $k : \mathbb{R}^2 \to [0, 1]$ that decreases with distance and asymptotically converges to 0 (Stachniss et al., 2009). In our approach, we represent these interdependencies using discrete weights w defined over the set of candidate locations $\mathcal{V}$ such that $w : \mathcal{V} \times \mathcal{V} \to [0, 1]$. A weight of $w_{ij} = 1$ for any $i, j \in \mathcal{V}$ indicates perfect correlation, whereas weights close to 0 mean that the observations at i and j are independent, i.e., no information can be inferred about i upon visiting j. Observations at larger distances can generally be assumed to be independent of one another. We can, therefore, make use of local neighborhoods including only pairs of locations between which there is a significant correlation. To this end, we define a covering neighborhood $\mathcal{C}_i$ such that $j \in \mathcal{C}_i \Leftrightarrow w_{ji} \gg 0$ for $j \neq i$.

Figure 7.1 depicts the derivation of weights and of the covering neighborhood given a kernel function $k(i, j)$. Candidate target locations for the UAV missions are determined by discretizing the target area as described in Chapter 4. The weight depends on the pairwise distance between these targets. The size of the covering neighborhood is determined such that the correlation between a location i and any location $j \in \mathcal{C}_i$ is equal to or higher than a predefined minimum value.

Informativeness measures

The discretized weights represent the similarity between different targets. Using these weights, we can estimate the overall information obtained by some sampling locations. A first model for estimating the overall informativeness based on these discretized weights has been proposed by Yu et al. (2014). The weights used in this approximation can be interpreted as the proportion of the information available at i that is obtained upon visiting j. Then, the overall informativeness is modeled as follows:

$$\mathcal{I}^{YU}(\mathcal{S}) = \underbrace{\sum_{i \in \mathcal{S}} \cdot 1}_{\substack{\text{full information at} \\ \text{visited locations}}} + \underbrace{\sum_{i \in \mathcal{V} \setminus \mathcal{S}} \sum_{j \in \mathcal{S} \cap \mathcal{C}_j} w_{ji}}_{\substack{\text{estimated proportion} \\ \text{of unvisited targets}}} \tag{7.6}$$

In other words, the authors estimate the proportion of the overall information that can be collected by a vehicle route. In the following, we refer to the corresponding optimization model as the IPP modeled after Yu et al. (2014), short IPP-YU.

This approximation, however, has drawbacks in practice. One problem is illustrated in Figure 7.2, which gives the optimal routes and objective values on two small graphs that only differ in the modeled degree of similarity between a center and the surrounding targets. On the left-hand side, the corresponding weights are set to $w_{ji} = 0.2$, which means that full information about the graph is only obtained when the center is included in the vehicle tour. On the right hand side, weights are increased to $w_{ji} = 0.3$. In this

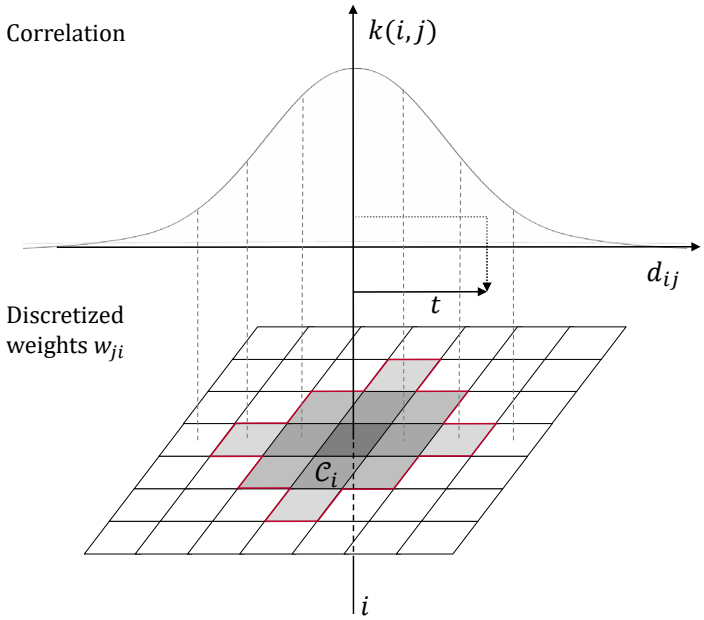

FIGURE 7.1: Derivation of discretized weights relative to a location $i \in \mathcal{V}$.

case, the estimated information gain can be increased by leaving the center unvisited. From a practical point of view and considering the monotonicity of the stochastic models in Section 7.2, this effect is undesirable: In any realistic setting, it is not possible to increase the information available about a location by not performing an observation. Consequently, this model is not a valid approximation for the spatial interdependencies that govern the surveyed area.

Yu et al. (2014) avoid this problem by requiring that

$$\sum_{j \in \mathcal{C}_i} w_{ji} u_i \leq u_i \tag{7.7}$$

holds for all $i \in \mathcal{V}$. In their model, this is achieved by explicitly setting weights w_{ji} such that

$$w_{ji} = \frac{1}{|\mathcal{C}_i|}, \quad j \in \mathcal{C}_i. \tag{7.8}$$

This model, however, introduces a counter-intuitive dependency between the size of the covering neighborhood and the weights. In practice, spatial processes with strong correlations and observations that are similar even at larger distances should be associated with an increased size of the covering neighborhood and increased weights for all $i, j \in \mathcal{C}_i$. This cannot be represented accurately in the approach by Yu et al. (2014): If the neighborhood size $|\mathcal{C}_i|$ increases as more locations are positively correlated, the weights w_{ji} have to decrease in order not to violate Equation (7.7). This means that information

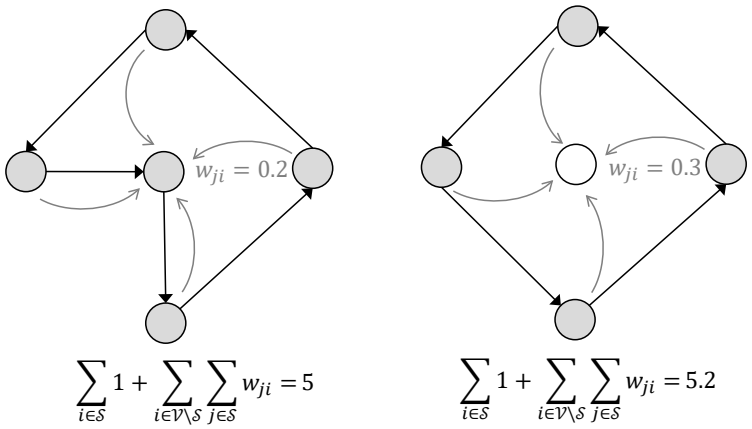

FIGURE 7.2: Illustration of the limitations of the approximation by Yu et al. (2014). Given are five target locations and weights w_{ji} with respect to the center i. The black arrows indicate the optimal path with the objective value given underneath the graph. Visited targets are indicated in grey. The right-hand side illustrates that when weights are increased, an improving solution can be obtained by visiting fewer targets.

about the stronger spatial interdependencies is lost.

We propose to relax the condition stated in Equation (7.7). To avoid that the estimated information gain increases by leaving out target locations, we limit the maximum information gain in a generalized objective function:

$$\mathcal{I}^{GEN}(\mathcal{S}) = \sum_{i\in\mathcal{S}} 1 + \sum_{i\in\mathcal{V}\setminus\mathcal{S}} \min\{1, \sum_{j\in\mathcal{S}\cap\mathcal{C}_i} w_{ji}\} \tag{7.9}$$

We refer to the corresponding planning problem as the generalized IPP after Yu et al. (2014), short IPP-GEN. We model weights using a simple approximation scheme based on inverse distance weighting. We assign a weight $\bar{w} < 1$ to locations at distance d^{min} and specify all other weights relative to $\bar{w}$ as follows:

$$w_{ji} = \begin{cases} \bar{w} \cdot \frac{d^{min}}{d_{ji}}, & j \in \mathcal{C}_i \\ 0, & \text{otherwise.} \end{cases} \tag{7.10}$$

Other approaches based on different kernel functions k are equally possible.

Using such an approximation, we can derive weights w_{ji} that can represent interdependencies of varying strength more accurately. This is illustrated in Figure 7.3: On the left-hand side, we give approximated weights for our approach and the model by Yu et al. (2014) for a comparatively small neighborhood, depending on the distance d_{ji} to an observed location i. The right-hand side gives weights for a situation with stronger correlations, represented by an increase in the size of the covering neighborhood. In the case of the model by Yu et al. (2014), increasing neighborhood size decreases the

weights for all locations. In our model, the stronger correlation can be reflected by increasing $\bar{w}$ without side effects. Note that similar to the stochastic variants, both models are monotonic in the number of observations.

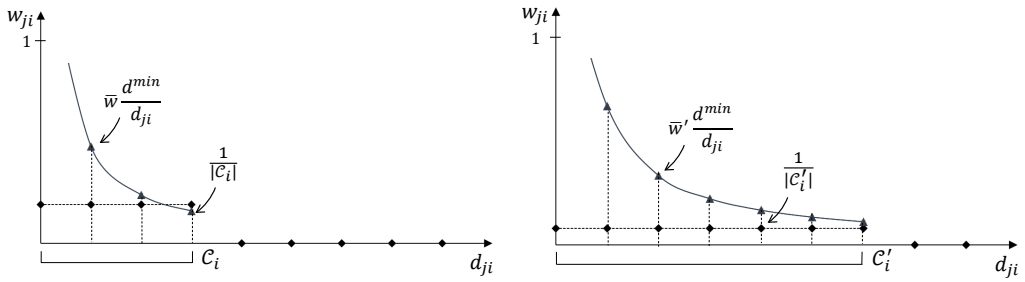

(A) Initial neighborhood size (B) Doubled radius of $\mathcal{C}_i$

FIGURE 7.3: Impact of increased neighborhood size on approximated weights.

7.4 GCorTOP modeling approach

Apart from the spatial interdependencies discussed above, UAV missions should also be based on priorities within the target region that characterize how valuable obtaining information is at a particular location. As the specification of these priorities depends on the user, they do not necessarily have to be spatially correlated. The observations made at these locations, in contrast, are always characterized by spatial correlations in our use case. Hence, some information about a highly prioritized location can also be obtained by making measurements nearby. The objective function needs to adequately account for this effect.

In this work, we combine these aspects: We model relations between target locations based on the approximated weights introduced in the previous section. We then use these weights in the objective function to partially account for the priority of locations that are not visited themselves but are in the proximity of targets included in a UAV mission. This way, the fact that some information can be inferred based on observations nearby decreases the relative importance of performing additional observations at these locations.

Yu et al. (2014) address this aspect by combining the classical TOP, which maximizes direct coverage of priorities, with the approximative model (IPP-YU) in Equation (7.6). This yields a mixed integer quadratic programming (MIQP) formulation, referred to as the correlated team orienteering problem (CorTOP) by the authors, which is based on the following objective function:

$$\mathcal{I}^{CorTOP}(\mathcal{S}) = \sum_{i \in \mathcal{S}} u_i + \sum_{i \in \mathcal{V} \setminus \mathcal{S}} \sum_{j \in \mathcal{S} \cap \mathcal{C}_i} w_{ji} u_i \tag{7.11}$$

Equation (7.11) maximizes the sum of the priorities associated with selected sensing locations $\sum_{i \in \mathcal{S}} u_i$. Unobserved locations are partially taken into account, with their total contribution to the objective depending on the selected targets within their covering neighborhoods. For an unobserved location i, the value of the indirectly inferred information is approximated as $\sum_{j \in \mathcal{S} \cap \mathcal{C}_i} w_{ji} u_i$. Please note that we adapted the formulation of Yu et al. (2014) for reasons of consistency to our notation. The original MIQP formulation corresponds to $\mathcal{I}^{CorTOP}(\mathcal{S}) = \sum_{i \in \mathcal{V}} x_i u_i + \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{C}_i} x_j (x_j - x_i) w_{ji} u_i$ with binary decision variable $x_i = 1$ if $i \in \mathcal{S}$ and 0 otherwise.

This model suffers from the same problems as IPP-YU, see Section 7.3. We derive a new model based on the relaxed variant IPP-GEN. We refer to the corresponding model as the *generalized correlated team orienteering problem (GCorTOP)* based on the following objective function:

$$\mathcal{I}^{GCorTOP}(\mathcal{S}) = \sum_{i \in \mathcal{S}} u_i + \sum_{i \in \mathcal{V} \setminus \mathcal{S}} \min\{u_i, \sum_{j \in \mathcal{S} \cap \mathcal{C}_i} w_{ji} u_i\} \tag{7.12}$$

Equation (7.12) ensures that indirectly collected benefit never exceeds a location's priority. The benefit of each additional observation $j \in \mathcal{C}_i$ with respect to $i \notin \mathcal{S}$ is determined as

$$\min\{w_{ji} u_i, \max\{0, (1 - \sum_{k \in \mathcal{S} \cap \mathcal{C}_i} w_{ki}) \cdot u_i\}\}, \tag{7.13}$$

i.e., the benefit is at most $w_{ji} u_i$, thus imposing a threshold that prevents overestimating a mission's informativeness. This combines the consideration of priorities as proposed by Yu et al. (2014) with the advantages of the relaxed discretization in Equation (7.9).

7.5 Summary

In this chapter, we discussed models for considering spatial dependencies for planning UAV missions, both stochastic and approximative. We first formally defined the mission planning problem for emergency surveillance (MPPES) in a general fashion, i.e., independent of specific models for determining the informativeness or benefit of UAV missions. We discussed models based on Gaussian process (GP) representations that have been proposed in IPP literature. We then derived approximative, discretized models that avoid the computational overhead of stochastic models and allow the consideration of priorities in the objective function. Finally, we demonstrated how these discrete models can be modified to consider priorities and spatial interdependencies in the objective function.

8 Dynamic programming and 2MLS solution approaches

We propose two solution approaches addressing the MPPES: an exact algorithm and a heuristic approach. The exact approach based on bidirectional dynamic programming is introduced in Section 8.1. We then discuss a two-phase multi-start adaptive large neighborhood search (2MLS) that can quickly provide solutions to realistic problem instances in Section 8.2.

8.1 Exact solution approach

We use bi-directional dynamic programming for solving the MPPES exactly. Our algorithm assumes that the objective function $\mathcal{I}(\mathcal{S})$ is monotonic in the number of selected samples, i.e., that additional observations can never decrease the informativeness measure (see Section 7.2.2). We do not impose further requirements on the objective function, which makes this approach applicable to all informativeness measures that we have discussed in the previous section. However, this also means that all UAV routes have to be defined completely to determine their quality.

Generation of vehicle routes

As the benefit of individual vehicle missions cannot be evaluated in isolation, i.e., without considering the routes of all other vehicles, we use a giant-tour representation for modeling the routes of all UAVs at once. Preliminary tests have indicated that this version is faster than a variant where individual missions are created first for each vehicle and the optimal combination of these tours is determined second.

The solution approach is based on the concepts introduced by Righini and Salani (2008). The algorithm maintains data structure $\mathcal{L}^{fw}$ and $\mathcal{L}^{bw}$ comprising the generated labels in the forward and backward direction, respectively. Each label corresponds to a tuple $(\mathcal{S}, m, T_m, i)$, where $\mathcal{S} \subset \mathcal{V}$ indicates the sampled locations, $m \in \{0, \dots, |\mathcal{M}|-1\}$ is the current vehicle, and T_m represents the duration of the emerging tour of UAV m with i as its last reached location. The routes corresponding to a forward label start

at s_0^{start}, i.e., at the starting node of the first UAV, and end at location $i \in \mathcal{N}$. In the backward direction, routes starting at the ending location of the last UAV $s_{|\mathcal{M}|-1}^{end}$ and continue in the reverse direction until i. T_m stores the travel time associated with the current (partial) routes of vehicle m from s_m^{start} to i or backwards from s_m^{end} to i.

A label $(\mathcal{S}, m, T_m, i)$ is extended to a label $(\mathcal{S}', m, T_{m'}, j)$ by adding a previously unvisited location $j \in \mathcal{V} \setminus \mathcal{S}$ and updating the set of visited target locations $\mathcal{S}'$ and route cost $T_{m'}$ accordingly, i.e,

$$\mathcal{S}' = \mathcal{S} \cup \{j\},$$
$$T'_m = T_m + \tau_{ij}.$$

The extension is feasible as long as $T_{m'} + \tau_{j s_m^{end}} \leq T_m^{max}$. In the backward direction, we extend labels along arcs (j, i) starting at possible predecessors j of location i. An extension is not feasible if $T_{m'} + \tau_{s_m^{start} j} > T_m^{max}$.

If the maximum route duration of vehicle m would be exceeded for all $j \in \mathcal{V} \setminus \mathcal{S}$, a label is extended by closing the current vehicle route and opening a new one. In this case, the new label is set to $(\mathcal{S}, m+1, 0, s_{m+1}^{start})$ in the forward direction and $(\mathcal{S}, m-1, 0, s_{m-1}^{end})$ in the backward direction.

Similar to Righini and Salani (2008), we apply resource-based bounding to avoid the generation of redundant labels. In our case, the critical resource is the flight duration. Forward labels are extended as long as

$$m \leq \lceil |\mathcal{M}|/2 \rceil - 1 \equiv m^{fw}.$$

For example, if three vehicles are available, forward extension includes the vehicles indexed 0 and 1. Correspondingly, backward labels are extended while

$$m \geq |\mathcal{M}| - m^{fw} - 1 \equiv m^{bw}$$

If m is not a multiple of 2, the same vehicle is included in forward and backward labels, as $m^{fw} = m^{bw}$. In this case, extensions in any directions are only performed while $T_m \leq T_m^{max}/2$ for all labels with $m = m^{fw} = m^{bw}$.

Dominance tests

As stated in the introduction of this section, we do not assume specific properties of the objective function other than monotonicity. Therefore, dominance checks can only be performed based on the selected locations and routing costs. Consequently, a label $(\mathcal{S}, m, T_m, i) \in \mathcal{L}^{fw}$ is dominated by another label $(\mathcal{S}', m', T_{m'}, i) \in \mathcal{L}^{fw}$ if either

(a) $m' < m, \mathcal{S}' \supseteq \mathcal{S}$
(b) $m' = m, T_{m'} = T_m, \mathcal{S}' \supseteq \mathcal{S}$
(c) $m' = m, T_{m'} < T_m, \mathcal{S}' = \mathcal{S},$

i.e., if at least the same targets can be surveyed using fewer vehicles or if the same candidates can be surveyed within a shorter travel time. Similarly, in the backward direction, a label $(\mathcal{S}, m, T_m, i) \in \mathcal{L}^{bw}$ is dominated by $(\mathcal{S}', m', T_{m'}, i) \in \mathcal{L}^{bw}$ if

$$\begin{aligned}
&\text{(a)}\ \ m' > m, \mathcal{S}' \supseteq \mathcal{S} \\
&\text{(b)}\ \ m' = m, T_{m'} = T_m, \mathcal{S}' \supseteq \mathcal{S} \\
&\text{(c)}\ \ m' = m, T_{m'} < T_m, \mathcal{S}' = \mathcal{S}.
\end{aligned}$$

Optimal solution

To compute a feasible solution, we join forward and backward labels $(\mathcal{S}, m, T_m, i) \in \mathcal{L}^{fw}$ and $(\mathcal{S}', m', T_{m'}, j) \in \mathcal{L}^{bw}$ such that $\mathcal{S} \cap \mathcal{S}' = \varnothing$. If there is an uneven number of vehicles, i.e., $m^{fw} = m^{bw}$, a join is feasible only if $T_m + \tau_{ij} + T_{m'} \leq T_m^{max}$. The objective value $\mathcal{I}(\mathcal{S})$ only depends on the visited target locations. As several solutions may result in the same set of samples, we store already computed objective values in a hash table to avoid the redundant evaluation of expensive informativeness measures. Among all solutions achieving the same informativeness $\mathcal{I}$, we then select the one with the lowest associated cost.

8.2 Two-phase adaptive large neighborhood search (2MLS)

As a generalization of the TOP, the GCorTOP is NP-hard, resulting in high computation times even for small instances when solving these problems exactly. In many applications, heuristic solution approaches are more promising for determining good solutions quickly. For this reason, we introduce a two-phase multi-start adaptive large neighborhood search (2MLS) approach for solving the MPPES.

8.2.1 Overview and algorithm design

In this section, we first present the fundamental ideas and concepts behind our solution approach. We then give an overview of the proposed algorithm. Its constituting elements are discussed in detail in the remainder of this chapter.

Concepts

Existing local search approaches for related problems usually focus on the cost-efficient service of targets, often resulting in narrow tours with visit locations in close proximity. This is in contrast to our goal of surveying as large an area as possible. In contrast, our algorithm specifically promotes explorative UAV routes. This is achieved by a combination of the following concepts:

- We use a multi-start approach combined with a seeding strategy that allows us to quickly generate and evaluate structurally diverse initial solutions.
- The search itself is separated into two phases: Phase 1 is based on reduced problem representations that emphasize spatial exploration, while Phase 2 computes the final solution. Both phases use an adaptive large neighborhood search (ALNS) scheme for improving solutions.
- We design search strategies for the ALNS that encourage explorative routes by accounting for the potential of the surrounding locations instead of considering targets in isolation.

Solution Approach

The solution scheme is summarized in Algorithm 1. Phase 1 consists of a multi-start approach for constructing diverse solutions ω_i^{red} based on aggregated or decomposed problem representations P^{red} (see Section 8.2.3). In each iteration, routes are initialized with seed nodes generated by a k-means++ algorithm (Section 8.2.2) and are completed using an insertion strategy δ^{insert}. The solution is improved using an adaptive large neighborhood search (ALNS) approach and finally transformed to a solution ω_i to the initial problem representation P. In Phase 2, an ALNS is applied to the best solution obtained in Phase 1.

Algorithm 1: Two-phase multi-start adaptive large neighborhood search (2MLS)

Input: Problem P
Output: Solution ω^{final}
`/* Phase 1: multi-start search on simplified problem representations */`
$\omega^{init} = \varnothing, \mathcal{I}(\omega^{init}) = 0$;
for $i = 1$ until the maximum number of multi-starts **do**
  determine problem representation P^{red};
  initialize ω_i^{red} using a k-means++ algorithm;
  $\omega_i^{red} \leftarrow \delta^{insert}(\omega^{red}, P^{red})$;
  $\omega_i^{red} \leftarrow$ `ALNS`(ω_i^{red});
  determine solution ω_i to initial problem representation P based on ω_i^{red};
  if $\mathcal{I}(\omega_i) > \mathcal{I}(\omega^{init})$ **then** $\omega^{init} = \omega_i$;
end
`/* Phase 2: ALNS on complete problem representation */`
$\omega^{final} \leftarrow$ `ALNS`(ω^{init});
return ω^{final};

Adaptive large neighborhood search

The ALNS used in both phases is an adaptation of the framework proposed by Pisinger and Ropke (2007) and is depicted in Algorithm 2. In each iteration, one removal strategy $\delta^{removal}$ and one insertion strategy δ^{insert} are selected from the sets of strategies based on their previous successes and failures. The removal strategy selects nh locations and removes them from the current solution (Section 8.2.5). The removed locations are reinserted by applying δ^{insert} to the partial solution (Section 8.2.6). The search is guided by a reheating-based acceptance criterion that accepts a non-improving move with a probability depending on the number of iterations since the last improvement (Section 8.2.7). The ALNS stops if a convergence criterion is met or $iter^{max}$ is reached.

Algorithm 2: Adaptive large neighborhood search

Input: Problem representation P, starting solution ω^{init}, $iter^{max}$, $iter^{convergence}$, removal strategies $\Delta^{removal}$, insertion strategies Δ^{insert}, neighborhood size limits nh^-, nh^+

Output: Solution ω^{best}

$iter = 0$, $iter^{lastimproving} = 0$;
initialize selection probabilities $\frac{1}{|\Delta^{removal}|}$ for $\delta^{removal} \in \Delta^{removal}$, $\frac{1}{|\Delta^{insert}|}$ for $\delta^{insert} \in \Delta^{insert}$;
initialize $\omega^{cur} \leftarrow \omega^{init}, \omega^{best} \leftarrow \omega^{init}$;
while $iter < iter^{max}$ *and* $iter - iter^{lastimproving} < iter^{convergence}$ **do**
    determine $nh \in [nh^-, nh^+] \cap \mathbb{N}$;
    compute acceptance threshold $\underline{\mathcal{I}}$;
    select $\delta^{removal}$ and δ^{insert} using roulette wheel selection;
    $\omega' \leftarrow \delta^{removal}(\omega^{cur}, nh)$;
    $\omega'' \leftarrow \delta^{insert}(\omega', P)$;
    if $\mathcal{I}(\omega'') > \mathcal{I}(\omega^{best})$ **then**
        $\omega^{best} \leftarrow \omega''$;
        $\omega^{cur} \leftarrow \omega''$;
        $iter^{lastimproving} = iter$;
    else if $\mathcal{I}(\omega'') \geq \underline{\mathcal{I}}$ **then**
        $\omega^{cur} \leftarrow \omega''$;
    update probabilities for Δ^{insert} and $\Delta^{removal}$;
    $iter = iter + 1$;
end
return ω^{best};

8.2.2 Seed nodes

In each iteration of the multi-start approach, a seed for every vehicle tour is determined using a k-means++ variant (Arthur and Vassilvitskii, 2007) that can account for route du-

ration restrictions and starting and ending locations, see Algorithm 3. For each vehicle, the algorithm determines the set of reachable locations $\mathcal{V}_m = \{i \in \mathcal{V} : \tau_{s_m^{start} i} + \tau_{i s_m^{end}} \leq T_m^{max}\}$ and then selects a seed visit $i \in \mathcal{V}_m$ with a probability proportional to the minimum squared distance between i and any $j \in \mathcal{N}$ that is either a vehicle starting or ending location or has already been selected as seed in another vehicle tour. This ensures that the initial seed locations are spread across the area, thereby facilitating the construction of vehicle routes covering large parts of the target region.

Algorithm 3: k-means++ seeding algorithm

Input: Candidate targets $\mathcal{V}$, vehicles $\mathcal{M}$
Output: Seeds $c_m \in \mathcal{V}$ for all vehicles $m \in \mathcal{M}$
initialize $D(i) = \min_{j \in \mathcal{N}_S \cup \mathcal{N}_E} d_{ij}$ for all $i \in \mathcal{V}$;
initialize $C = \varnothing$ as the set of selected seed visits;
for $m \in \mathcal{M}$ **do**
  determine feasible set $\mathcal{V}_m = \{i \in \mathcal{V} : \tau_{s_m^{start} i} + \tau_{i s_m^{end}} \leq T_m^{max}\}$;
  select seed $c_m \in \mathcal{V}_m$ with probability $\frac{D(c_m)^2}{\sum_{i \in \mathcal{V}_m} D(i)^2}$;
  $\mathcal{V} \leftarrow \mathcal{V} \setminus \{c_m\}$, $C \leftarrow C \cup \{c_m\}$;
  for $i \in \mathcal{V}$ **do**
    if $d_{i c_m} < D(i)$ **then** $D(i) = d_{i c_m}$;
  end
end
return C;

8.2.3 Aggregation and decomposition strategies

In Phase 1, we use modified problem representations to quickly determine starting solutions with an emphasis on spatial exploration, i.e., missions covering large areas. For this purpose, we propose two different schemes: The first one is based on spatial aggregation, reducing the number of candidate sensing locations. The second one is based on spatial decomposition to plan individual vehicle tours separately. Both variants provide reduced problem representations, i.e., problem representations P^{red} where the corresponding set of targets $\mathcal{V}^{red}$ is a subset of the targets $\mathcal{V}$. Illustrative examples are given in Figure 8.1.

Grid-based priority aggregation

In this aggregation scheme depicted in Figure 8.1a, the target area is divided into evenly sized grid cells C. Each cell comprises a number of locations $\mathcal{V}_c \subset \mathcal{V}$ such that $\bigcup_{c \in C} \mathcal{V}_c = \mathcal{V}$. For each cell $c \in C$, we determine the priority-weighted center (p_c^x, p_c^y) based on the positions $pos_i = (p_i^x, p_i^y)$ for $i \in \mathcal{V}_c$ such that $p_c^x = \frac{\sum_{i \in \mathcal{V}_c} u_i p_i^x}{\sum_{i \in \mathcal{V}_c} u_i}$ and $p_c^y = \frac{\sum_{i \in \mathcal{V}_c} u_i p_i^y}{\sum_{i \in \mathcal{V}_c} u_i}$. We chose a representative $i_c \in \mathcal{V}_c$ that is closest to this point. The sum of priorities within

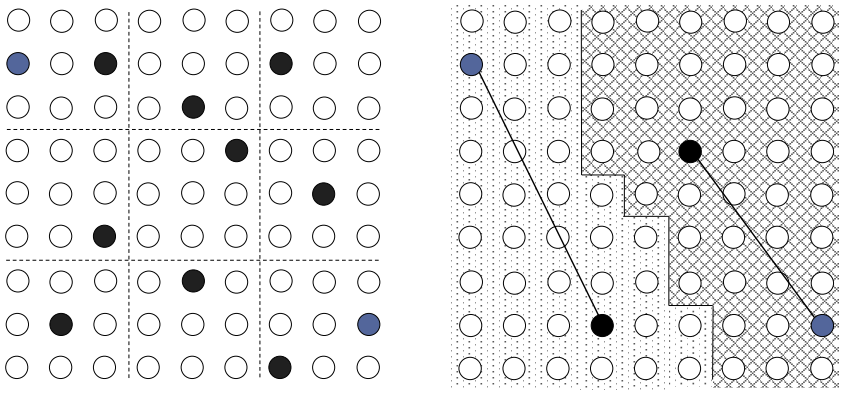

(A) Grid-based aggregation (B) Area decomposition

FIGURE 8.1: Illustration of simplified problem representations. Vehicle starting and ending locations are marked blue. The left-hand side represents a target area with nine representatives in the reduced problem, depicted in black. The right-hand side indicates the same area separated into two non-overlapping single-vehicle instances based on initial seed nodes.

a cell $\sum_{i \in \mathcal{V}_c} u_i$ is assigned to its corresponding representative. This results in a set of target locations

$$\mathcal{V}^{red} = \{i_c \mid c \in C\}. \tag{8.1}$$

An initial solution is constructed considering only these representatives.

In the multi-start approach, we use different levels of aggregation in each iteration to obtain a larger range of candidate solutions. Specifically, we determine the size $|\mathcal{V}_c|$ of the grid cells $c \in C$ such that either 4, 6 or 9 locations are grouped, making exceptions for grid cells located at the border of the target region. Increasing the grid cell size further reduces computation time, as only very few representatives remain. However, we have observed that the solutions rarely represent good starting solutions for Phase 2 of the search approach.

This approach allows us to quickly identify highly relevant regions within the target area and to compute provisional routes for several vehicles. The aggregated priority of each representative point anticipates the priority obtained from locations that are covered by the vehicle tours. Furthermore, any solution to the reduced problem represents a feasible solution with respect to the initial problem representation.

VEHICLE-ORIENTED SPATIAL DECOMPOSITION

The spatial decomposition scheme illustrated in Figure 8.1b separates the target area into independent sections, which are assigned to one vehicle each, and solved as a series of single-vehicle problems. Given the seed routes $r_m = (s_m^{start}, c_m, s_m^{end})$ (see Section 8.2.2), we determine subsets $\mathcal{V}_m^{red} \subset \mathcal{V}$ for all vehicles such that each target location is assigned to the vehicle whose seed tour it is closest to. Ties are broken arbitrarily. This makes it

possible to apply the decomposition scheme to instances with multiple vehicle starting or ending locations.

The result of this decomposition corresponds to a number of single-vehicle problem representations P_m^{red} with $\bigcup_{m\in\mathcal{M}} \mathcal{V}_m^{red} = \mathcal{V}$ and $\mathcal{V}_m^{red} \cap \mathcal{V}_{m'}^{red} = \varnothing$ for any two m, m' with $m \neq m'$. As every single-vehicle problem contains a non-overlapping subset of candidate locations, the combination of routes represents a feasible solution for the original problem. The sizes of the decomposed areas are reasonably well balanced due to the use of the seeding strategy, and the assignments to UAVs can take the vehicles' starting and ending position into account. This increases overall coverage when routes of different UAVs are combined.

8.2.4 Weighted objective function

During the ALNS approach, we employ a hierarchical objective function with the distance traveled by the vehicles as a secondary objective. This ensures that among solutions with similar sensing locations $\mathcal{S}$, the ones with lower resource consumption are preferred. This is particularly relevant whenever it is possible to achieve complete or near-complete coverage of a region, in which case distance minimization becomes more relevant.

However, heavily focusing on distance minimization is detrimental to our interest in obtaining explorative tours, as these typically require long travel times for traversing the entire target area. For this reason, total route length is included in the objective function by multiplying it with a sufficiently small factor of $\varrho > 0$ to adjust its relative weight. Denoting the sequence of locations in each route $r_m \in \omega$ as $(s_m^{start}, i_{m,0}, i_{m,1}, \ldots, i_{m,n}, s_m^{end})$, the planning objective is defined is

$$\max \;\; \mathcal{I}(\mathcal{S}) - \varrho \sum_{r_m\in\omega} (d_{s_m^{start} i_{m,0}} + \sum_{c=0}^{n-1} d_{i_{m,c} i_{m,c+1}} + d_{i_{m,n} s_m^{end}}). \tag{8.2}$$

For each problem instance, we heuristically set ϱ based on the average priority $\bar{u} := \frac{\sum_{i\in\mathcal{V}} u_i}{|\mathcal{V}|}$ and the average distance $\bar{d} := \frac{\sum_{i\in\mathcal{V}}\sum_{j\in\mathcal{V}} d_{ij}}{|\mathcal{V}|^2}$ such that

$$\varrho = 0.01\frac{\bar{u}}{\bar{d}}. \tag{8.3}$$

8.2.5 Removal strategies

The 2MLS includes several removal strategies $\Delta^{removal}$ that iteratively eliminate locations from a solution. The result of this procedure is a partial solution that can then completed again by an insertion strategy (see Section 8.2.6).

The basic removal strategy of our 2MLS is depicted in Algorithm 4. The set of *planned* locations Γ^p contains locations included in a vehicle route, while the set of *unvisited* locations Γ^u contains locations not selected to be visited in the current solution. Note that

we explicitly select locations among the set Γ^u for removal instead of considering all unvisited locations as candidates for re-insertion. First, this reduces the overall neighborhood size and ensures that candidates that are unlikely to be inserted in a tour do not need to be evaluated in the subsequent insertion strategy. Second, this allows us to explicitly create different neighborhoods, e.g., for emphasizing either diversification or more local repair steps.

The number of locations to be removed, i.e., neighborhood size nh, is divided among the planned and the unvisited locations according to their proportion in the initial solution. The planned locations are selected first following a removal heuristic $\delta^{removal-routed}$, resulting in relaxed routes with fewer targets. Then, the free capacity T_m^{free} of each vehicle is computed as the difference between the vehicle's maximum route duration T_m^{max} and the duration of the relaxed route $T_{m'}$. Afterward, the unvisited locations are selected for removal using $\delta^{removal-unvisited}$. A feasibility check ensures that for each removed unvisited location, there is at least one vehicle in the relaxed solution that can reach this location. The removed locations are stored in the set of *open* locations Γ^o.

Algorithm 4: Basic removal strategy

Input: Solution ω^{init} with selected locations $\mathcal{S}$, neighborhood size nh
Output: Open targets Γ^o, remaining planned locations Γ^p, unvisited locations Γ^u

set $\Gamma^p = \mathcal{S}, \Gamma^u = \mathcal{V} \setminus \mathcal{S}, \Gamma^o = \varnothing, nh^p = \lceil \frac{|\mathcal{S}|}{|\mathcal{V}|} \cdot nh \rceil$;
while $|\Gamma^o| < nh^p$ **do**
  select $i \in \Gamma^p$ following removal heuristic $\delta^{removal-routed}$;
  $\Gamma^o = \Gamma^o \cup \{i\}, \Gamma^p = \Gamma^p \setminus \{i\}$;
end
for $m \in \mathcal{M}$ **do**
  compute route length T'_m of the relaxed solution;
  free capacity $T_m^{free} = T_m^{max} - T'_m$;
end
while $|\Gamma^o| < nh$ **do**
  select $i \in \Gamma^u$ following removal heuristic $\delta^{removal-unvisited}$;
  if $\exists m \in \mathcal{M}$ with minimum insertion cost of i no larger than T_m^{free} **then**
    $\Gamma^o = \Gamma^o \cup \{i\}, \Gamma^u = \Gamma^u \setminus \{i\}$;
  end
end

Randomized nearest neighbor removal The randomized nearest neighbor selection strategy seeks to enable local improvements by exchanging targets with others nearby. It iteratively selects nodes in Γ^p at random and removes them as well as their n nearest neighbors among the remaining nodes Γ^u until the maximum neighborhood size is reached. As a result, selected nodes are grouped in clusters, which makes it possible to

replace targets with others nearby. Furthermore, it enables local modifications of routes to reduce total duration.

Sequence-based nearest neighbor removal The sequence-based nearest neighbor strategy removes entire subsequences from vehicle routes. A starting point for each segment is selected randomly from Γ^p. The subsequent targets in the given vehicle route are removed until either the vehicle ending location is reached or the number of removed locations exceeds a given segment length $len^{segment}$. The segment length limit depends on the average route size, i.e., $len^{segment} = \lfloor \upsilon \cdot \frac{|\mathcal{S}|}{|\mathcal{M}|} \rfloor$ with a parameter $\upsilon < 1$. For each segment, we select the closest remaining locations in Γ^u for removal, as these are the most likely candidates for re-insertion.

Route sparsification strategy The route sparsification strategy randomly selects one vehicle $m \in \mathcal{M}$ and removes a large number of target locations within this route. Similar to the sequence-based strategy, these locations are removed in segments of length $len^{segment}$. One location remains between any two consecutive segments. This allows preserving the general shape of the route. Unvisited locations are removed from a plan by choosing one location in each of the removed segment at random, and selecting the closest remaining locations among the set Γ^u for removal.

Priority-delta removal The priority-delta removal strategy selects the locations associated with the lowest contribution to the objective value among the targets in Γ^p. Furthermore, it selects the locations with the highest priority among the reachable unvisited targets in Γ^u. As recommended by Pisinger and Ropke (2007), both the selection of routed and unvisited locations are randomized. To this end, locations in Γ^p are sorted by ascending priority, and locations in Γ^u by descending priority. Given a randomization parameter *det* and a random number $p \in [0, 1]$, the l-th location is selected with $l = \lfloor p^{det} |\Gamma^p| \rfloor$ or $l = \lfloor p^{det} |\Gamma^u| \rfloor$, respectively.

Region-based removal The region-based removal strategy accounts for priorities of other locations within a limited distance d of a target instead of evaluating each location's priority in isolation. The value d is randomized in each iteration. Based on this value, we define the cost-benefit ratio $ratio_i$ of location i as

$$ratio_i = \frac{cost_i^{min}}{\sum_{j \in \mathcal{V}: d_{ij} < d} u_j}, \tag{8.4}$$

where $cost_i^{min}$ corresponds to the minimum detour needed to insert a location i at the best insertion position over all routes of the current solution. If a feasible insertion of location i is not possible, $cost_i^{min}$ is set to ∞. Locations in Γ^p are sorted by ascending $ratio_i$, while those in Γ^u are sorted by descending $ratio_i$. The selection is randomized using a determinism parameter *det* as above.

Weighted center removal The weighted center removal strategy uses the distance to the vehicle tours' weighted centers as an indicator of the cost of inserting a visit. The weighted center is determined as the weighted average of the locations that constitute the tour in a given solution. Irrespective of their associated priorities, this strategy selects the locations in Γ^p that are furthest from the weighted center of the tour that they currently belong to. Among the locations in Γ^u, those that are closest to any weighted center of are removed from a solution. The goal of this strategy is to systematically remove locations that are ill-placed with respect to the remainder of the tour.

Worst angle removal The worst-angle removal strategy selects combinations of target locations that form sharp angles in the route. We assume that these angles indicate room for improvement, as they can be associated with considerable detours. Furthermore, in the case of the MPPES, sharp angles imply that at least two targets in the route are located relatively close to one another, which is rarely beneficial in terms of information gain. For computing the associated angle for each location $i \in \Gamma^p$, we determine the preceding location h and the succeeding location j in the current solution. The angle $\angle hij \in [0, \pi]$ is formed by the arcs from h to i and from i to j, respectively. For the sharpest angles in the given solution, we then remove the three constituting points h, i and j. Unvisited locations are selected randomly from Γ^u.

8.2.6 Insertion strategies

Routes are constructed and re-constructed by iteratively inserting locations into vehicle tours. The basic heuristic for constructing solutions is depicted in Algorithm 5. The heuristic iteratively selects a location i among the set of open locations Γ^o following an insertion strategy δ^{insert}. Each selected location is inserted at its minimum cost position and added to the set of planned locations. If no feasible insertion position is available, a location is left unvisited and added to the set of unvisited locations Γ^u. To prevent the repeated construction of similar solutions, the insertion heuristic is slightly randomized such that a random location is selected with probability $p^{rand} \in [0, 1]$ instead of following δ^{insert}.

Visit insertion relies on the marginal benefit of inserting a target location, i.e., the incremental change of the objective caused by including a location in a route, given the current partial solution. To this end, we distinguish between the remaining priority u_i^r and marginal contribution u_i^m of target i. The *remaining priority* u_i^r reflects the proportion of the location's priority value that has not yet been covered by visiting nearby target locations. The *marginal contribution* u_i^m indicates the incremental change of the objective function when i is included in a UAV route. For the GCorTOP, the remaining priority of a location i is:

$$u_i^r = \max\{0, (1 - \sum_{j \in \mathcal{S} \cap \mathcal{C}_i} w_{ji}) \cdot u_i\}. \tag{8.5}$$

Algorithm 5: Basic insertion strategy

Input: Partial solution ω^{init} with open targets Γ^o, planned locations Γ^p, unvisited locations Γ^u, randomization parameter p^{rand}
Output: Completed solution ω^{final} with selected targets $\mathcal{S} = \Gamma^p$

while $\Gamma^o \neq \varnothing$ **do**
  determine a random number $p \in [0, 1]$;
  if $p \leq p^{rand}$ **then**
    select $i \in \Gamma^o$ at random;
  else
    select $i \in \Gamma^o$ following construction heuristic δ^{insert};
  end
  if $\exists$ feasible insertion position in Γ^p **then**
    insert i at its minimum cost position;
    $\Gamma^p = \Gamma^p \cup \{i\}$;
  else
    $\Gamma^u = \Gamma^u \cup \{i\}$;
  end
  $\Gamma^o = \Gamma^o \setminus \{i\}$;
end

The marginal benefit of a location i can then be determined as:

$$u_i^m = \begin{cases} u_i^r + \sum_{j \in \mathcal{V} \setminus \{i\} : i \in \mathcal{C}_j} \min\{u_j^r, w_{ij} u_j\} & \text{if } i \in \Gamma^o \\ 0 & \text{otherwise,} \end{cases} \tag{8.6}$$

i.e., it comprises the remaining priority of location i plus the priorities of all locations that can be surveyed by i.

Maximum marginal priority insertion The maximum marginal priority insertion strategy greedily inserts the visit with the highest marginal contribution to the objective value. In the case of a tie, the visit with the highest priority u_i is inserted first.

Best marginal priority ratio insertion The best marginal priority ratio strategy seeks to balance marginal benefit and resource utilization. It selects a visit $i \in \Gamma^o$ such that

$$i = \arg\max_{\Gamma^o} \frac{(u_i^m)^q}{cost_i^{min}}. \tag{8.7}$$

The factor $q > 0$ allows to adjust the relative importance of u_i^m and $cost_i^{min}$. Increasing q leads to a higher focus on priorities during construction. Again, ties are broken by the locations' direct priorities u_i.

Weighted center insertion Complementing the corresponding selection strategy, this insertion strategy greedily inserts the visit that is closest to any weighted center of a vehicle tour. Weighted centers are recalculated after each insertion.

Region-based insertion Similar to the corresponding removal strategy, the region-based insertion strategy accounts for the priorities of nearby locations, i.e., locations within a maximum distance d. Locations are selected for insertion such that

$$i = \arg\max_{\Gamma^o} \frac{\sum_{j \in \mathcal{V}: d_{ij} < d} u_j^r}{cost_i^{min}}. \tag{8.8}$$

The objective is to prioritize highly relevant regions early during the (re-)construction of routes. Note that we use remaining priorities u_i^r instead of the marginal contributions u_i^m. Otherwise, we would overestimate the locations' contributions as they are considered in the marginal priority values of all their respective neighbors. Again, the parameter d is randomized to insert locations in different orders during the search.

Cost-greedy insertion The cost-oriented insertion strategy greedily inserts visit i with minimum insertion cost. We use this strategy for diversification during the search, as most other strategies focus on the locations' priorities. Furthermore, in the case of the TOP, it may help to construct and improve routes that include clustered locations.

Orienteering-regret insertion The orienteering-regret heuristic inserts visit locations first that may no longer be feasible if their insertion is delayed. In order to use this strategy, we have to compute two values for each target remaining in Γ^o: the change in the objective value at the current best position for visit i, denoted $\delta_{i,1}^{obj}$, and the change in the objective if i is inserted in its n^{reg}-th best position, denoted $\delta_{i,n^{reg}}^{obj}$, with regret parameter $n^{reg} \in \{2, \ldots, |\Gamma^p|\}$. If the n^{reg}-th best position is not a feasible insertion position, we set $\delta_{i,n^{reg}}^{obj} = 0$. Then, the heuristic selects a location i such that

$$i = \arg\max_{\Gamma^o} |\delta_{i,n^{reg}}^{obj} - \delta_{i,1}^{obj}|. \tag{8.9}$$

If some locations only have few feasible insertion positions left, setting $\delta_{i,n^{reg}}^{obj} = 0$ ensures that these are associated with high regret values. In this case, the strategy prioritizes the locations that offer the highest marginal contribution to the objective, offering additional diversification as locations that are further away from a current route are inserted early. If the n^{reg}-th best position exists for all $i \in \Gamma^o$, the difference between $\delta_{i,n^{reg}}^{obj}$ and $\delta_{i,1}^{obj}$ reflects the change in total route distance due to the weighted objective function. In this situation, the insertion strategy first selects locations where a delay may lead to a high detour. Similar to the cost-greedy strategy, this improves routes in terms of total duration. If no feasible position remains, $\delta_{i,1}^{obj}$ is 0 and the location is added to set Γ^u.

8.2.7 Reheating-based acceptance criterion

As acceptance criterion, we employ a reheating-based scheme that increases the likelihood of accepting a move associated with a decreasing objective value with the number of moves since the last strict improvement of the objective function. Although quite simple, preliminary experiments have shown that this scheme obtains better results than a conventional simulated annealing metaheuristic for the ALNS approach. This is due to the effect that, once an improved solution has been found, subsequent search steps are often able to quickly find further improvements. This way, subsequent search steps focus on intensification after an improving move. The acceptance threshold is only lowered if improvements cannot be found, thereby accepting solutions with decreasing objective values to leave local minima.

Similar to the well-known simulated annealing scheme, a new solution ω^{new} is accepted with a probability that depends on the best solution found so far ω^{best} and the current temperature Ψ. With $\delta^{obj} = |\mathcal{I}^{best} - \mathcal{I}^{new}|$ as the absolute decrease in solution quality, the probability of accepting a move is determined as:

$$P\left(\delta^{obj}\right) = \exp\left(\frac{-\delta^{obj}}{\Psi}\right). \tag{8.10}$$

The temperature Ψ_i in iteration i is increased in non-improving iterations such that $\Psi_{i+1} = \Psi_i \cdot \psi$ with $\psi > 1$. In the beginning and whenever a new best solution is found, Ψ_i is set to a given minimum temperature of $\Psi_0 = 1$.

We initialize heating parameter ψ such that a decrease in solution quality of δ^{obj} will be accepted with probability $p = 0.5$ after κ iterations. To do so, we use the following equation:

$$\psi = \left(\frac{-\delta^{obj}}{\Psi_0 \ln 0.5}\right)^{(1/\kappa)}. \tag{8.11}$$

As proposed by Kilby (2013), we can use Equation (8.10) for computing an acceptance threshold prior to the construction of a solution, which is determined as

$$\underline{\mathcal{I}} = \mathcal{I}^{best} - \delta^{obj} = \mathcal{I}^{best} + \ln(p) \cdot \Psi \tag{8.12}$$

given a random number $p \in (0, 1]$. An insertion step can be stopped preemptively when the maximum achievable priority value $\overline{\mathcal{I}}$ of a solution is less than the required acceptance threshold. We therefore maintain an upper bound on the objective value of GCorTOP, which is computed as follows:

$$\overline{\mathcal{I}} = \sum_{i \in \mathcal{V}} u_i - \sum_{i \in \Gamma^u} (u_i - \min\{u_i, \sum_{j \in \mathcal{C}_i \setminus \Gamma^u} w_{ji} u_i\}) \tag{8.13}$$

The upper bound in Equation (8.13) corresponds to the sum of all priorities reduced by the sum of "lost" priority values, i.e., the proportion of priorities associated with

locations in Γ^u that cannot be covered by any of the locations in Γ^p and Γ^o, respectively.

8.3 Summary

In this chapter, we discussed two approaches for solving the mission planning problem for emergency surveillance. We introduced a bi-directional dynamic programming approach for providing a benchmark for small problem instances. To address realistic instance sizes, we then proposed a two-phase multi-start adaptive large neighborhood search (2MLS). In Phase 1, this approach seeks to quickly determine explorative vehicle tours based on reduced problem representations. To this end, we proposed two different strategies based on spatial decomposition and the aggregation of candidate locations. In Phase 2, these initial routes are further improved based on the full problem representation using an adaptive large neighborhood search approach that iteratively removes several locations from a solution and re-inserts them to improve the objective value. We propose several removal and insertion strategies that enforce desirable characteristics in a solution, e.g., route shape or the priorities of selected targets. Search is guided by a metaheuristic that accepts worse solutions with a probability depending on the solutions' objective values and the number of iterations since the last best solution was found. This allows us to alternately intensify the search to further improve on a recently found solution and to diversify it in case a local minimum was reached.

9 Experimental study

This chapter provides an in-depth evaluation of the proposed models and solution approaches based on an extensive simulative study. It is dedicated to two main issues: (1) The evaluation of the performance of the solution approach considering the achieved solution quality, robustness, and computation times, and (2) the assessment of the applicability of the proposed models with respect to the mission planning problem introduced in Chapter 4.

In Section 9.1, we first introduce the benchmark instances that provide a starting point for this evaluation and derive new benchmark instances for the MPPES based on real-world data and simulated distributions of hazardous gases. In Section 9.2, we propose evaluation criteria for measuring the quality of the modeling variants and solution procedure. The results of this study are presented and discussed in Sections 9.3 and 9.4.

9.1 MPPES benchmark instances

At the core of this evaluation are two sets of benchmark instances that provide insights into the performance of the models and solution approaches. We use the existing benchmark instances proposed for the team orienteering problem (TOP) to compare our solution approach to state-of-the-art algorithms for related problems. To this date, no benchmark instances have been established for IPP. For evaluating the models and algorithms in a realistic use case, we derive new instances specifically tailored to the mission planning problem for emergency surveillance.

9.1.1 Existing benchmark instances

As a result of the attention the TOP has received in literature, a large number of benchmark instances are available for this problem class. An overview of these instances is provided in Table 9.1. Most of the work published in this area uses the instance sets proposed by Chao et al. (1996b) to evaluate proposed solution approaches. However, many of these instances are comparatively easy to solve, either due to the low number of locations in total or due to low route capacities that mean that only a small subset of locations can be feasibly visited. For this reason, Souffriau et al. (2010) and most authors

Reference		Nb.	$\lvert\mathcal{V}\rvert$	$\lvert\mathcal{M}\rvert$
Chao et al. (1996b)	Set 1	54	32	2-4
	Set 2	33	21	2-4
	Set 3	60	33	2-4
	Set 4	60	100	2-4
	Set 5	78	66	2-4
	Set 6	42	64	2-4
	Set 7	60	102	2-4
Dang et al. (2013)		333	100-399	2-4

TABLE 9.1: Overview of team orienteering problem (TOP) benchmark instances introduced by Chao et al. (1996b) and Dang et al. (2013).

of later publications only consider the 157 most "relevant" instances among the 387 instances initially published. These instances are comprised by a subset of the sets 4 to 7 depicted in Table 9.1, leaving out those with tight budget constraints and thus short routes.

To extend the number of available benchmark instances, Dang et al. (2013) introduced a set of larger instances comprising up to 399 candidate target locations. These instances are based on benchmark instances published for the orienteering problem by Fischetti et al. (1998) which are, in turn, either adapted from benchmark instances for TSP and VRP or randomly generated. As the instances by Chao et al. (1996b), the number of available vehicles varies between 2 and 4.

9.1.2 Generation of new benchmark instances

While we can use benchmark instances established for the TOP to compare our solution approach with existing approaches, they do not allow the evaluation and comparison of the proposed modeling variants for the MPPES. Therefore, we introduce new large-scale instances for this planning problem based on the practical use case presented in Chapter 2. An overview of the parameters of the benchmark instances is given in Table 9.2.

In this study, we consider rotary-wing UAVs equipped with optical sensors for detecting hazardous substances. Following the BBK's recommendations, priorities are determined using recent information on population density (see Section 2.1.4). To simulate the distributions of hazardous substances, we generate synthetic autocorrelated random fields. We furthermore consider 5 different target area sizes. For every target area size, we generate 90 instances by varying the assigned priorities, the number of UAVs and maximum flight time as indicated in the table. This leads to a total of 450 benchmark instances.

Target areas and priorities

As the difficulty of the planning problem is primarily determined by the number of candidate targets, we vary the size of the target area between 1.5×1.5 km^2 and 2.5×2.5 km^2.

Parameter	Values
Cruise speed v^{max} (in m/s)	7
Acceleration a^{max} (in m/s^2)	2
Sensing time τ_i (in s)	2
Sensor resolution (pixel)	$1,000 \times 1,000$
Flight height h (in m)	120
Ground sampling resolution pxr (in m^2/pixel)	0.01
Available vehicles	$\{1;\ 2;\ 3\}$
Mission duration T^{max} (in s)	$\{600;\ 900;\ 1,200;\ 1,500;\ 1,800\}$
Target area sizes (km^2)	$\{1.5 \times 1.5;\ 1.5 \times 2;\ 2 \times 2;\ 2.5 \times 2;\ 2.5 \times 2.5\}$

TABLE 9.2: Parameter settings for MPPES benchmarks.

Given the distance of 100 m between targets determined by the sensor equipment, this results in 15 × 15 = 225 to 25 × 25 = 625 candidate target locations. Smaller affected areas either do not require airborne surveillance or can be surveyed in their entirety within the available flight time limit. We do not consider larger areas, as they are not within the UAVs' operation range.

For the MPPES benchmarks, we select target areas within the German state of North Rhine-Westphalia at random. For each area, we determine priorities using the corresponding data on population density published by the Federal Statistical Office based on data collected in the 2011 census (Statistisches Bundesamt (Destatis), 2018). This dataset provides georeferenced information on population density in a resolution of 100 × 100 m, which is sufficient for our use case. Combined with geographical grids provided by the Central Office for Geotopography (2018), it can be loaded in Geographical Information Systems (GIS) such as QGIS (QGIS Development Team, 2017). This allows combining this information with additional map layers, which facilitates the interpretation of these priorities by human operators. Furthermore, this provides an intuitive interface for specifying target regions and UAV starting locations. An example is given in Figure 9.1, which indicates the bounding box specified by a user using the QGIS application on the left-hand side, and the population data corresponding to this area on the right-hand side.

Note that the published data is anonymized such that individuals cannot be identified. This is achieved by modifying values for grid cells with very few inhabitants. We cannot ensure that areas that seem uninhabited in the published data are indeed unpopulated. Therefore, we assign all targets a minimum priority of 1. This does not impact the overall priority distribution but ensures that there are no locations with a priority of zero or negative priority values.

An advantage of this way of determining priorities is its general applicability, as similar data sources are publicly available in many regions. For example, population density estimates on a global scale are available in a spatial resolution of up to 250 × 250 m from the European Commission, Joint Research Centre (JRC) (2015). Thus, we can assume that this procedure can be similarly applied to other regions.

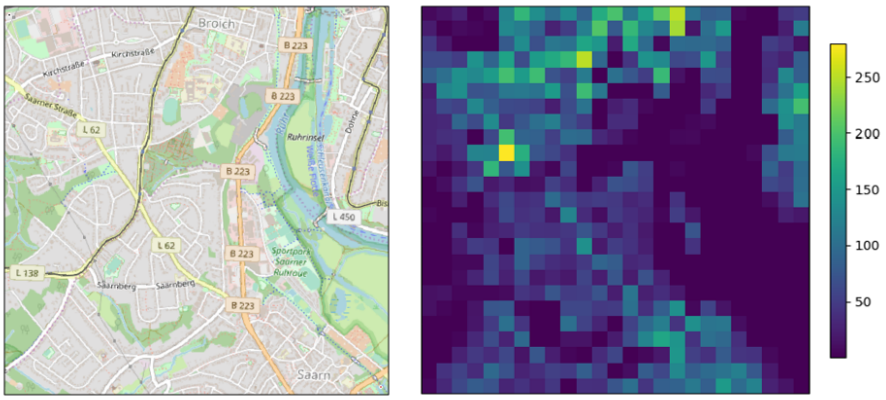

FIGURE 9.1: Generation of benchmark instances based on population density information. Own representation using QGIS (QGIS Development Team, 2017), population data retrieved from Statistisches Bundesamt (Destatis) (2018).

VEHICLES

In our study, we consider UAV systems similar to those used in the BigGIS project (see Section 2.2.1). We generate instances with 1 to 3 vehicles that are identical except for their takeoff and landing positions, which are selected at random among the locations at the border of the target area. To obtain insights on the trade-off between flight time and information gain, we vary the maximum flight time T^{max} between 10 and 30 minutes. We assume a maximum cruise speed v^{max} =7 m/s (approx. 25 km/h), which is realistic when taking payload weight and environmental influences into account. The observations are made at a standstill.

The time required for accelerating and decelerating has a major influence on total travel time between locations in close proximity. For simplicity, we assume that UAVs accelerate and decelerate with a constant horizontal acceleration rate a^{max} =2 m/s^2. Then, travel time τ_{ij} between any two points i and j can be determined depending on traveled distance d_{ij} as follows:

$$\tau_{ij} = \begin{cases} 2\sqrt{\frac{d_{ij}}{a^{max}}} & \text{if } d_{ij} < \frac{(v^{max})^2}{a^{max}} \\ \frac{d_{ij}}{v^{max}} + \frac{v^{max}}{a^{max}} & \text{otherwise} \end{cases} \tag{9.1}$$

We furthermore include a fixed sensing time $\tau_i = 2$ s. This is necessary as the UAVs need to verify their position via GPS, their flight needs to be stabilized at the target location, and several observations are performed and processed before traveling further.

In our experiments, we assume that vehicles are equipped with a hyperspectral sensor with a resolution of 1,000 × 1,000 pixels and a focal length of 12 mm. Given this specification, the flight height and the covered ground area per image depend on the desired target ground resolution, i.e., the covered ground area per pixel in the final image. Based on experiments in the BigGIS project, we consider a ground sampling resolution

of 0.01 m^2/pixel to be acceptable. This corresponds to a covered ground area per image of 100 $\times$ 100 m at a flight altitude of approximately 120 m. Then, candidate sampling locations are located in 100 m distance to one another, each of which represents the center of an image taken by the UAVs.

Spatial distribution of contaminants

The priorities specified above are central to the planning of UAV missions. The quality of the proposed model, however, does not solely rely on whether or not high-priority regions are covered by the vehicle routes. In practice, we use the samples taken along the route to predict the distribution of hazardous substances across the relevant area. Hence, the applicability of the proposed model largely depends on the predictive quality that can be achieved for the distribution of hazardous substances.

To take this into account, we generate spatially autocorrelated random fields that simulate distributions of contaminants across the target area. For this purpose, we define Gaussian process distributions $Z_{\mathcal{V}}$ based on a set of covariance functions. We determine artificial distributions by drawing samples from this distribution (see Section 3.2.2). For each instance, this yields simulated values $Z(i)$ for all $i \in \mathcal{V}$ that follow the specified covariance matrix of the GP prior. We normalize these values such that $Z(i) \in [0, 100]$ for all i, with a mean value of 50 for the multivariate normal distribution.

The covariance functions used are based on the results of Stachniss et al. (2009), who use exponential kernels and Matérn covariance functions for modeling and predicting gas distributions. We alternate between these two kernel functions when generating random fields. This way, we create distributions that vary in their smoothness, some changing more quickly as others. An example for the planning scenarios, defined by their unique set of priorities and simulated distribution, are depicted in Figure 9.2 for a scenario with 500 target locations. Figure 9.2a is generated using the Matérn kernel, resulting in the comparatively rough distribution that is indicated on the right. In contrast, the distribution indicated in Figure 9.2b is based on an exponential kernel, yielding a much smoother distribution.

9.2 Study design and performance measures

As stated at the beginning of this chapter, we evaluate our proposed solution strategy for the MPPES with a particular emphasis on two aspects: The performance of the solution procedure, and the applicability of the models in a realistic setting. In this section, we derive performance measures that allow us to compare and evaluate the proposed strategies and models. Considering solution quality, we use measures established in literature to compare the performance of the 2MLS on standard benchmarks with other published results. For the evaluation of the models, we derive new performance measures that permit a more detailed analysis of the obtained results.

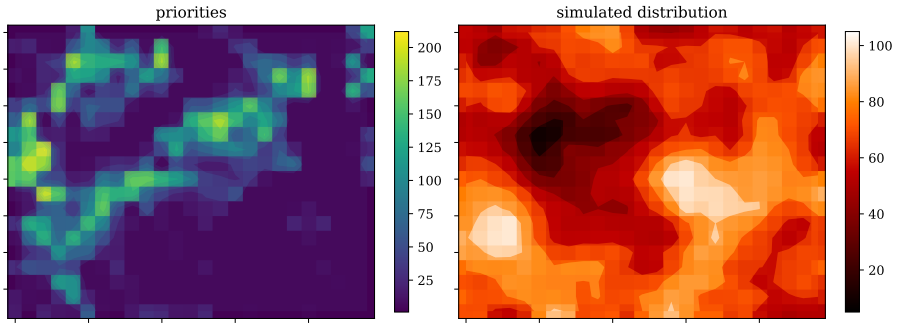

(A) Scenario based on a target area of 2×2.5 km^2 (500 candidate locations) using a Matérn kernel

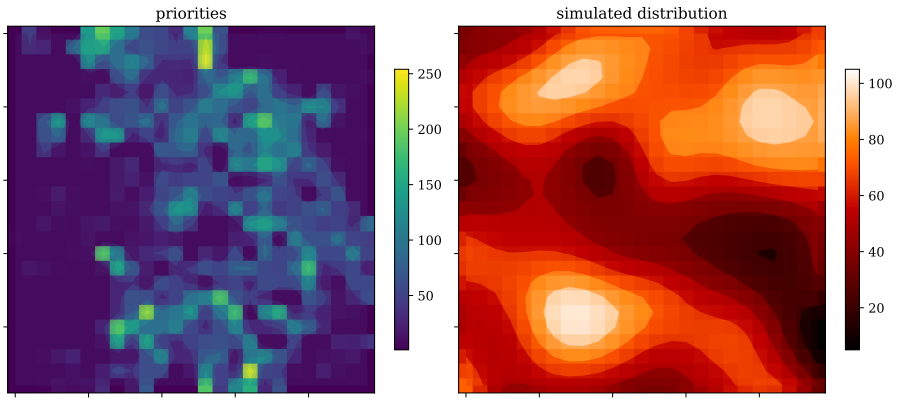

(B) Scenario based on a target area of 2.5×2.5 km^2 (625 candidate locations) using an exponential kernel

FIGURE 9.2: Examples for planning scenarios characterized by priorities and simulated distribution of hazardous substances. Each cell in the images corresponds to one candidate target, represented by the cell center.

9.2.1 SOLUTION QUALITY

To measure the achieved solution quality, we compare the best and the average results out of 10 executions of our algorithm per benchmark instance. This follows established practices in the literature on vehicle routing (see e.g., Dang et al., 2013; Vidal et al., 2015). The best gap compares the best found solution ω^{best} to the best known solution ω^{BKS} as follows:

$$\text{best gap} = \frac{\mathcal{I}(\omega^{best}) - \mathcal{I}(\omega^{BKS})}{\mathcal{I}(\omega^{BKS})} \tag{9.2}$$

The gap between ω^{BKS} and the average informativeness $\mathcal{I}^{avg} = \frac{1}{N}\sum_{i=1}^{N}\mathcal{I}(\omega_i)$

$$\text{avg. gap} = \frac{\mathcal{I}^{avg} - \mathcal{I}(\omega^{BKS})}{\mathcal{I}(\omega^{BKS})} \tag{9.3}$$

serves as a measure of the robustness of the solution approach. The lower this value, the closer the average results are to the best-known solution.

9.2.2 Covered priorities

As we have established in Chapter 4, one objective of our approach is to provide emergency services with information at locations that are particularly relevant to them. Most importantly, this includes locations at which a large number of people may be affected after an incident.

In our model, this relative importance of target locations is considered in form of the priorities u_i. In a straightforward fashion, we could evaluate a solution based on the sum of priorities $\sum_{i \in \mathcal{S}} u_i$. However, doing so leaves aside any information that is obtained about neighboring locations due to the correlations within the distribution of airborne contaminants. We therefore measure direct and indirect coverage using a distance-dependent coverage measure $PCov_d$

$$PCov_d(\mathcal{S}) = \frac{\sum_{i \in \mathcal{V}: \exists j \in \mathcal{S} \text{ with } d_{ij} \leq d} u_i}{\sum_{i \in \mathcal{V}} u_i}, \tag{9.4}$$

which gives the proportion of all priorities within distance d to a sampled location. Note that these measures are independent of how we model informativeness $\mathcal{I}(\mathcal{S})$, thus allowing a comparison of all discussed models.

In case of the MPPES, a UAV may travel further to provide better spatial coverage. In other words, it may reduce the priorities associated with directly served targets $\sum_{i \in \mathcal{S}} u_i \equiv PCov_0$ in order to obtain better overall coverage. To obtain insights into these patterns, we use three parametrizations of the measure. These indicate directly covered priorities ($PCov_0$) as well as those within 100 m and 300 m ($PCov_{100}$ and $PCov_{300}$) respectively. A higher value for the last two measures indicates that a UAV path traverses most of the highly prioritized regions, i.e., that the most relevant target locations are either included in a vehicle route or are close to one.

9.2.3 Prediction quality

A central goal of the MPPES is to provide an accurate overview of the overall contamination across the target area. Keeping this in mind, the quality of planned paths corresponds to the quality of the predicted distributions of contaminants that are determined based on the samples taken during the flight.

To this end, we can use statistical means, specifically, GP regression approaches, which are computationally too expensive for evaluating interim solutions during the search. The theoretical foundation for this analysis is given in Chapter 3, in particular in Section 3.2.2. The planned UAV mission determines a set of sampling locations $\mathcal{S}$. The corresponding observations z_i for all $i \in \mathcal{S}$ are specified by the simulated spatial distribution of contaminants. We condition a Gaussian process on these samples. For fitting

the covariance function of this GP to the sample data, we use the Python `scikit-learn` package (Pedregosa et al., 2011). The GP posterior gives predicted values $\widehat{Z}(i)$ for all $i \in \mathcal{V}$. To measure the quality of this prediction, we measure the deviation of the predicted values from the "true" distribution of contaminants, i.e., the initially generated distribution.

In this study, we use three measures: The *mean absolute error (MAE)*

$$\text{MAE} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} |\widehat{Z}(i) - z_i| \tag{9.5}$$

allows to compare overall accuracy of the predictions.

The *mean error (ME)*, which is defined as

$$\text{ME} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} (\widehat{Z}(i) - z_i), \tag{9.6}$$

indicates whether or not the selected samples lead to a systematic deviation in the prediction, e.g., a systematic underestimation or overestimation of the contamination.

Finally, we also determine the *weighted mean absolute error (WMAE)*

$$\text{WMAE} = \frac{1}{\sum_{i \in \mathcal{V}} u_i} \sum_{i \in \mathcal{V}} u_i \cdot |\widehat{Z}(i) - z_i| \tag{9.7}$$

which accounts for the fact that a reliable prediction is more important to emergency services at locations where more people may be affected. This way, we can evaluate the trade-off between overall prediction quality and the accuracy at highly prioritized locations.

9.2.4 Prediction robustness

The evaluation measures above are based on the assumption that samples taken by the UAVs during their missions are accurate. In practice, this is not always the case, as we cannot avoid measurement errors or inaccurate position estimations. To assess the applicability of our model in real-world settings, we evaluate whether the achieved predictions are robust with respect to inaccurate sample values.

An overview of this approach is given in Figure 9.3. We mimic measurement errors that may occur in practice by adding white noise drawn from a normal distribution with mean 0 and standard deviation σ. A prediction is not computed based on the true data z_i, but on modified observations $y_i = z_i + e_i$, with $e_i \sim \mathcal{N}(0, \sigma)$. This means that the expected value over all measurements does not change, but individual observations may deviate from the exact values. In our experiments, we set $\sigma = k\mu$ with $k \in \mathbb{R}^+$ and μ the mean value of the true data. We evaluate the approach with k = 0.05 and k = 0.1. Due to the way we generated our test sets (see Section 9.1.2), we know that the mean value for

all generated distributions is around 50. Hence, we add a noise with a standard deviation of 2.5 and 5 respectively. This way, the measurement errors are ensured to be within a reasonable range: Observations deviate noticeably from the true values, but the results are not entirely arbitrary, which would prohibit any reasonable prediction.

For evaluating the robustness of the prediction, we compute the prediction quality measures defined in Equations (9.5) to (9.7) and compare the result with those obtained without noise. A high deviation between these two values indicates that the predicted distribution of hazardous substances is sensitive with respect to measurement values. In practice, this would limit the reliability and usefulness of such a UAV-based surveillance system.

9.2.5 Model configurations

To evaluate the performance of our proposed model, we compare three modeling variants in terms of solution quality, runtime and robustness: The TOP, which only considers priorities, CorTOP (Equation (7.11)) as a predecessor to our model with a lower emphasis on covering aspects, and finally, the GCorTOP model (Equation (7.12)).

In case of the latter two models, we have to specify the weights that define the objective function. For CorTOP, these are based on the approximation used by Yu et al. (2014), see Equation (7.8), with the covering neighborhood $\mathcal{C}_i$ of location i made up of its immediate neighbors. For GCorTOP, our model is based on Equation (7.10) with $\bar{w}$ =0.5 and d^{min} =100 m. The radius of the covering neighborhood is limited to 400 m, as the impact of including locations at larger distances in the covering neighborhood is negligible.

9.3 Computational results

The two-phase multi-start adaptive large neighborhood search (2MLS) is implemented in C++. The bi-directional dynamic programming approach is implemented in Python 3.6. All evaluations are performed on an Intel Xeon 2.6GHz processor machine with 119

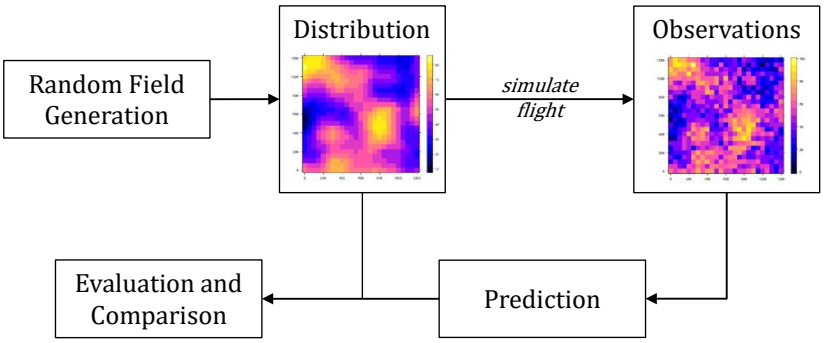

Figure 9.3: Evaluation of robustness with respect to noisy input data.

GB RAM running a 64-bit version of Windows Server 2012. Except for the parameter tuning, each instance was solved 10 times.

9.3.1 Parameter tuning

To determine good parameter settings, we performed a series of preliminary experiments on a subset of 60 randomly selected instances among the TOP and MPPES instances. The idea behind this approach is to find reasonable parametrizations that perform well on a diverse set of instances. For each instance and parameter configuration, we ran the 2MLS three times. The performance of different configurations are compared based on the gap to the best known solution (BKS) averaged over all instances. Detailed results can be found in Appendix A.

Based on the results of the tuning phase, we propose two parametrizations of the 2MLS: a "regular" configuration achieving good results (2MLS), and a particularly fast parametrization denoted as 2MLS-f. This distinction allows us to gain insights into the trade-off between computation time and solution quality. Furthermore, using the fast version, we can evaluate how the algorithm performs in a practical setting where computation times are strictly limited. Table 9.3 gives the parameter settings for the two configurations.

We use the same insertion and removal strategies for both variants, selecting strategies that performed best in the preliminary tests. However, we apply different strategies for Phase 1: We use the vehicle decomposition scheme for 2MLS-f, as it provides better solutions early on during the search. In the regular variant, where we can exploit the available time for further improvements of solutions, we use the aggregation-based scheme. ALNS segment size is set to 200 iterations for the fast version, and 600 iterations for the regular variant. This way, the search reacts more quickly after the first iterations as the total number of iterations is limited. The convergence criterion is set to 100 iterations (2MLS-f) and 300 iterations (regular configuration). Similarly, we use two different configurations for average neighborhood size, with smaller neighborhoods being used for the fast variant.

For the ALNS procedure, except for the parameters explicitly mentioned, we follow the settings proposed by Pisinger and Ropke (2007). Specifically, this includes the strategies' weights and update procedures for the ALNS. Our tests have shown that the overall search is robust with respect to these parameters, as no significant changes have been observed in preliminary tests when varying these parameters.

A note on rounding

Rounding issues have complicated the comparison of results in the past, as solutions obtained using lower precision are occasionally better in terms of total priority, but are infeasible for higher rounding precisions (see, e.g., the discussion in (Ke et al., 2008)). In our implementation, all distances and travel times are rounded to two digits. The same

Parameter	2MLS-f	2MLS
Strategy for Phase 1	Vehicle decomposition	Aggregation
Number of multi-starts	4	4
ALNS limit in Phase 1	100 iterations	100 iterations
ALNS limit in Phase 2	2,000 iterations	2,000 iterations
ALNS segment size	200 iterations	600 iterations
Convergence limit	100 iterations	300 iterations
Neighborhood size limits	$nh^- = \max\{0.05\|\mathcal{V}\|, 10\}$ $nh^+ = \min\{0.2\|\mathcal{V}\|, 100\}$	$nh^- = \max\{0.1\|\mathcal{V}\|, 10\}$ $nh^+ = \min\{0.3\|\mathcal{V}\|, 100\}$
Priority ratio factor q	$\{1, 2\}$	$\{1, 2\}$
Segment length υ	$[0.1\frac{\|\mathcal{S}\|}{\|\mathcal{M}\|}, 0.25\frac{\|\mathcal{S}\|}{\|\mathcal{M}\|}]$	$[0.15\frac{\|\mathcal{S}\|}{\|\mathcal{M}\|}, 0.35\frac{\|\mathcal{S}\|}{\|\mathcal{M}\|}]$
Randomization factor p^{rand}	0.05	0.05
Regret parameter n^{reg}	2	2
Determinism parameter *det*	6	6
Heating parameter κ	400 iterations	400 iterations

TABLE 9.3: Search configurations for computational experiments.

precision is used for calculating the objective value in case the weights in the GCorTOP models make rounding necessary.

9.3.2 2MLS RESULTS ON TOP BENCHMARKS

In a first step, we use the existing benchmarks for TOP in order to compare the proposed two-phase multi-start adaptive large neighborhood search with other approaches, which represent the state of the art with regard to computation time or solution quality:

- Fast path relinking (FPR), proposed by Souffriau et al. (2010), run on an Intel Xeon 2.5GHz processor machine
- Particle swarm optimization-inspired algorithm (PSOiA), proposed by Dang et al. (2013), tested on an AMD Opteron 2.6GHz CPU
- Unified hybrid genetic search (UHGS) with augmented neighborhoods for implicit customer selection proposed by Vidal et al. (2015), run on an Intel Xeon 3.07GHz CPU
- Multi-start local search (MS-LS), again presented by Vidal et al. (2015), with similar augmentations, tested on the same machine
- Pareto mimic algorithm (PMA) by Ke et al. (2016), run on a Intel Core i5 3.2GHz CPU with 4 GB RAM

Due to the different hardware used and the lack of further data, an accurate comparison of computation times is not possible. Except for two newer machines used by Vidal et al. (2015) and Ke et al. (2016), the processors used are of a similar generation as ours and can provide a rough indication of computation times.

Results on instances by Chao et al. (1996b)

The results are displayed in Table 9.4 in aggregated form. For each instance set and solution procedure, the table gives the average and best deviation to the best-known solution over the 10 runs and average computation time.

		FPR	PSOiA	UHGS	MS-LS	PMA	2MLS-f	2MLS
4	avg.	n.a.	0.11 %	0.05 %	1.52 %	0.32 %	3.18 %	1.75 %
	best	0.73 %	0.00 %	0.01 %	0.21 %	0.00 %	0.90 %	0.36 %
	t(s)	8.6	226.7	236.3	15.9	113.1	4.5	22.6
5	avg.	n.a.	0.03 %	0.01 %	0.48 %	0.09 %	1.67 %	0.64 %
	best	0.23 %	0.00 %	0.00 %	0.06 %	0.00 %	0.14 %	0.09 %
	t(s)	2.9	73.9	138.0	3.4	32.3	1.3	12.0
6	avg.	n.a.	0.00 %	0.00 %	0.34 %	0.29 %	1.10 %	0.10 %
	best	0.11 %	0.00 %	0.00 %	0.03 %	0.00 %	0.00 %	0.00 %
	t(s)	2.1	37.3	91.0	2.0	24.8	0.8	7.3
7	avg.	n.a.	0.03 %	0.02 %	1.11 %	0.13 %	1.29 %	0.88 %
	best	0.54 %	0.00 %	0.00 %	0.14 %	0.00 %	0.28 %	0.05 %
	t(s)	6.3	130.3	228.0	9.8	68.6	2.7	22.1
all	avg.	n.a.	0.06 %	0.03 %	0.99 %	0.20 %	2.03 %	1.04 %
	best	0.47 %	0.00 %	0.00 %	0.13 %	0.00 %	0.43 %	0.17 %
	t(s)	5.0	138.4	192.0	9.3	69.3	2.8	18.0

Table 9.4: Summary of results on Chao et al. (1996b) instances.

The best solutions are obtained by PSOiA, UHGS and PMA, all of which are associated with considerable running times. Compared to these approaches, 2MLS comes with a slightly worse solution quality, achieving a best gap of 0.17 % and an average gap of 1.04 % to the best-known solutions. 2MLS-f has a best gap of 0.43 % and an average gap of 2.03 %. However, with an average computation time of 2.8 seconds, it is faster than any of the other approaches. The best gap slightly improves on FPR. As the authors have not published information on their average results, it is not possible to compare the robustness of the algorithms.

Results on instances by Dang et al. (2013)

The instances introduced by Dang et al. (2013) have received considerably less attention in literature. To this date, only two solution approaches are available for comparison: PSOiA proposed by Dang et al. (2013), and PMA introduced by Ke et al. (2016). Both of these approaches are population-based mechanisms with comparatively high computation times.

The computational results are reported in Table 9.5. Note that both previous publications limit their study to a subset of 82 "difficult" instances among the 333 in the large benchmark set. In the same manner, we give average results on these 82 instances as

well as on the entire benchmark set. Note that, while the proposed instances on average include around 171 target locations, this is increased to approximately 240 for this selection. The comparison of results on the entire benchmark set is based on the results reported on the personal website[1] of the authors of the second journal paper.

		PSOiA	PMA	2MLS-f	2MLS
all	avg.	0.13 %	0.13 %	1.62 %	1.03 %
	best	0.01 %	0.00 %	0.42 %	0.19 %
	t(s)	4,379.4	343.9	10.7	32.8
sel. 82	avg.	0.50 %	0.42 %	2.86 %	1.86 %
	best	0.04 %	0.00 %	1.13 %	0.55 %
	t(s)	11,031.0	999.2	25.8	70.5

Table 9.5: Summary of results on Dang et al. (2013) instances, averaged over all instances and the selection of 82 instances considered by Dang et al. (2013) and Ke et al. (2016).

The table shows that PMA obtains the best results in terms of solution quality. 2MLS and 2MLS-f are associated with a higher gap. The best solutions found by 2MLS are close to the BKS, with an average gap of 1.03 % and a best gap of 0.19 %, averaged over all instances. These gaps are slightly higher when considering the 82 more difficult instances but remain below 2 % on average. For 2MLS-f, the average gap of these instances is less than 3 %.

Compared to the two other published approaches, this increase in average gap comes with a considerable improvement in terms of computation time: 2MLS is more than ten times faster than PMA, and 2MLS-f reduces average computation time by a factor of 30. Both variants use less than 1 % of the computation time required for PSOiA.

9.3.3 2MLS results on MPPES benchmark instances

In the next section, we discuss the results on the new benchmark instances proposed for mission planning problem for emergency surveillance. To evaluate the solution quality with respect to known optimal solutions, we study the results on small problem instances and analyze the runtime behavior and quality of the solution approach. The validity of the model itself is evaluated in the final part of this section.

Optimality gap on small problem instances

The exact solution approach based on bidirectional dynamic programming is only tractable for small problem instances. We were unable to solve any of the large-scale instances derived in Section 9.1.2 to optimality within reasonable computation time. To obtain some indication of the limitations of an exact approach and to compare the heuristic results with proven optimal solutions, we generated a set of smaller instances with

[1] `http://gr.xjtu.edu.cn/web/keljxjtu/download`, accessed 23.09.2018

16 to 49 target locations in a rectangular grid. The instances are based on the same parameters as depicted in Table 9.2, except for the maximum duration per vehicle route, which varies between 100 and 250 seconds, and the target area size, which is limited to 0.4×0.4 km^2 to 0.7×0.7 km^2. The computation time for the exact approach has been limited to 10,000 seconds.

Computation times of the exact approach quickly increase with instance size, route duration, and the number of vehicles. In the single-vehicle case, we can only solve solutions to optimality that include around 40 locations or less, unless route durations are very short. Instances with two or more vehicles quickly become intractable for around 20 target locations. These findings are consistent with the performance of the MIQP model proposed by Yu et al. (2014) for solving the CorTOP. Unfortunately, a direct comparison of these approaches is not possible, as Yu et al. (2014) have not published their instances.

Compared with the known optimal results, 2MLS-f has an average gap of 0.1 % with an average computation time of 0.22 s. Meanwhile, the slower configuration of 2MLS has an average gap of 0.0005 % at a computation time of 1.98 s. Both variants can find the optimal solution in 53 out of the 55 cases for which we know the exact solution. Furthermore, for 51 out of 55 instances, 2MLS finds the optimal solution in all runs.

Solution quality and robustness

We now move on to the larger benchmark instances for MPPES introduced in Section 9.1.2. The results on TOP benchmark, particularly the larger ones by Dang et al. (2013), already indicate that even though 2MLS scales better than many other solution approaches, it is still associated with high computational effort for large instances. This is particularly obvious for the 82 difficult instances. These instances comprise on average 240 candidate locations. The largest instances in our study, are almost three times as large. Therefore, the following discussion focuses on the 2MLS-f.

In a first step, we investigate the solution quality achieved by our algorithm. As stated above, exact approaches are limited for instances of this size. For this reason, we cannot use optimal solutions for evaluating the solution quality of the heuristic approach. Instead, we use the best solutions found during our experiments, and consider the average gap relative to this value as an indicator of the overall solution quality.

This gap is given in Table 9.6, which summarizes results based on the three major aspects of target area size, maximum route duration, and the number of vehicles. The minimum average gap is achieved for the smallest instances in combination with high resource availability, i.e., instances with 225 and 300 target locations involving multiple vehicles and high route durations. In these instances, the objective value is generally close to the total sum of priorities $\sum_{i \in \mathcal{V}} u_i$. This means that it is comparatively easy for the solution approach to determine a solution with high coverage, and the solutions found during the search are mostly distinguished by their total duration rather than overall coverage. The effect is exactly opposite in case of very limited resources, i.e., single-vehicle instances with tight route duration constraints. In these situations,

$\|\mathcal{V}\|$	$\|\mathcal{M}\|$	T^{max} (in s)					avg.
		600	900	1,200	1,500	1,800	
225	1	1.8 %	1.1 %	0.5 %	0.3 %	0.1 %	0.7 %
	2	0.8 %	0.1 %	0.0 %	0.0 %	0.0 %	0.2 %
	3	0.2 %	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
	avg.	0.9 %	0.4 %	0.2 %	0.1 %	0.0 %	0.3 %
300	1	2.6 %	2.7 %	1.5 %	0.9 %	0.6 %	1.7 %
	2	1.6 %	0.5 %	0.1 %	0.0 %	0.0 %	0.5 %
	3	1.1 %	0.1 %	0.0 %	0.0 %	0.0 %	0.2 %
	avg.	1.8 %	1.1 %	0.5 %	0.3 %	0.2 %	0.8 %
400	1	1.5 %	1.9 %	1.5 %	1.0 %	0.7 %	1.3 %
	2	0.7 %	0.9 %	0.4 %	0.1 %	0.0 %	0.4 %
	3	0.3 %	0.3 %	0.1 %	0.0 %	0.0 %	0.1 %
	avg.	0.8 %	1.0 %	0.6 %	0.4 %	0.2 %	0.6 %
500	1	1.4 %	1.7 %	1.6 %	1.5 %	1.1 %	1.5 %
	2	0.8 %	1.0 %	0.7 %	0.4 %	0.2 %	0.6 %
	3	0.4 %	0.6 %	0.2 %	0.0 %	0.0 %	0.3 %
	avg.	0.9 %	1.1 %	0.9 %	0.6 %	0.5 %	0.8 %
625	1	2.0 %	2.4 %	2.9 %	2.0 %	1.8 %	2.2 %
	2	1.2 %	1.6 %	1.2 %	0.6 %	0.5 %	1.0 %
	3	0.8 %	1.0 %	0.4 %	0.2 %	0.1 %	0.5 %
	avg.	1.4 %	1.7 %	1.5 %	0.9 %	0.8 %	1.2 %
avg.		1.2 %	1.1 %	0.7 %	0.5 %	0.3 %	0.8 %

TABLE 9.6: Average gap between best found and average solution quality obtained by the 2MLS-f on the MPPES benchmark instances.

efficient use of the available resources is crucial, and exchanging some locations in a solution can have a significant impact on the objective value. Consequently, the highest gap of 2.9 % is observed in the instances with 625 candidate locations and a single vehicle.

Runtime analysis

In this section, we look more closely at the computation times that can be achieved using 2MLS-f. We assume that the target area is defined after arriving at or near the scene of an incident based on a first visual assessment. Our project partners estimate that a few minutes are required after arrival until the UAV missions can be started, for example for determining a feasible take-off location and calibrating the flight controller. At the latest, final routes should be available at the end of this preparation phase in order not to delay the surveillance and rescue operation.

The results for all planning scenarios and vehicle configurations are summarized in Table 9.7. The required running time depends largely on the number of target locations $|\mathcal{V}|$, route duration T^{max}, and number of vehicles $|\mathcal{M}|$. The average computation time over all instances is approximately 56 s, i.e., slightly less than a minute and well within the allowed range in our application. Increasing the number of candidate locations is almost always associated with an increase in computational effort. Whereas 24 seconds suffice for the smallest scenarios involving 225 targets, this is increased to 96 seconds on average for the largest set with 625 locations.

The influence of T^{max} is more ambiguous. The results indicate that increasing maximum route duration can both increase as well as decrease average computation time. This can be explained by the fact that, as routes become longer, they cover larger parts of the target area. In this case, continuing the search does not improve solutions further, as most locations are already included in or in the vicinity of a vehicle route. In other words, routes are already considered "informative" concerning the model at hand and the convergence criterion applies.

This effect becomes particularly obvious when comparing the impact of increasing route duration with respect to different instance sizes. For the smallest instances in terms of target locations, even comparatively short routes can cover major parts of the target area. The highest computation times are reached for T^{max} = 600 s for instances with up to 300 targets. For larger instances, these shorter routes are insufficient, leaving out entire regions of the target area. Consequently, for 400 and 500 target locations, computation times are highest for T^{max} = 900 s and only decrease as T^{max} increases further. In case of 625 target locations, maximum computation times are reached for T^{max} = 1,200 s.

A similar effect can be observed when increasing the number of vehicles. Again, deploying more vehicles allows covering larger areas. Thus, it becomes increasingly difficult to improve on a solution that already covers large parts of total priorities, either directly or indirectly. As a result, computation time almost consistently decreases for the smallest instances when the number of vehicles is raised. Similar to the observations re-

$\|\mathcal{V}\|$	$\|\mathcal{M}\|$	T^{max} (in s)					avg.
		600	900	1,200	1,500	1,800	
225	1	30.2	36.7	33.5	27.1	21.8	29.8
	2	39.3	28.4	16.6	16.3	17.3	23.6
	3	33.1	15.1	14.9	15.0	14.4	18.5
	avg.	34.2	26.7	21.7	19.5	17.8	24.0
300	1	30.9	44.7	51.0	48.8	37.6	42.6
	2	53.1	54.0	39.9	27.6	20.2	39.0
	3	60.5	37.9	21.9	19.0	20.6	32.0
	avg.	48.1	45.5	37.6	31.8	26.1	37.8
400	1	38.3	47.9	55.6	59.7	57.6	51.8
	2	54.8	68.6	58.7	50.3	37.3	53.9
	3	47.6	59.6	43.9	29.4	26.6	41.4
	avg.	46.9	58.7	52.7	46.5	40.5	49.1
500	1	42.3	52.9	69.7	80.0	71.6	63.3
	2	72.1	90.1	97.4	83.6	70.2	82.7
	3	83.2	104.0	71.2	56.5	43.6	71.7
	avg.	65.9	82.3	79.4	73.4	61.8	72.6
625	1	46.2	57.8	77.3	92.6	99.2	74.6
	2	81.4	101.8	134.1	115.2	106.2	107.7
	3	107.7	134.7	121.4	92.9	72.1	105.7
	avg.	78.5	98.1	110.9	100.2	92.5	96.0
avg.		54.7	62.3	60.5	54.3	47.8	55.9

Table 9.7: Average computation time of 2MLS-f on the MPPES benchmark instances.

garding T^{max}, this pattern changes when looking at the largest instances. In these cases, increasing $|\mathcal{M}|$ has a major impact on computation time for smaller values for T^{max}.

The results demonstrate that for the vast majority of instances, computation times are within the maximum delay specified by our project partners. Particularly the impact that increasing T^{max} and $|\mathcal{M}|$ can have on average computation time works to our advantage, reducing computation time for larger and more complex instances.

These observations are consistent with the solution quality analyzed in the previous section: For the scenarios with low gaps, computation time is reduced as the convergence criterion stops search early. Instances with a higher average gap are associated with comparatively high computation times. Here, more potential remains for further improving solutions, thereby avoiding a preliminary stop of 2MLS-f.

9.3.4 GCorTOP model evaluation and comparison

Until now, we have focused on the performance of the solution approach in terms of scalability and solution quality. In this section, we now study the results from a practical point of view. Specifically, we evaluate the applicability of the proposed model itself with respect to the mission planning problem for emergency surveillance. This section aims to determine whether valuable information can be obtained during these UAV missions. As discussed in Chapter 4, this requires the consideration of two aspects: The provisioning of information at particularly critical areas and the reliable prediction of the overall distribution of contaminants within the area.

Covered priorities

In a first step, we investigate the priorities that are covered by UAV missions. In Table 9.8, we report the average percentage of priorities visited directly by the vehicles ($PCov_0$) or that are either surveyed or within 100 m or 300 m of a visited location ($PCov_{100}$ and $PCov_{300}$). The results are grouped by instance size and the number of vehicles.

Focusing on $PCov_0$, i.e., considering only the priorities of locations for which direct observations are obtained, we can see that TOP performs best. This meets our expectations, as this model does not involve any trade-off between spatial coverage and priorities of directly visited locations. Results obtained using CorTOP are similar to these results. The difference to the GCorTOP is larger, and this model yields the overall lowest $PCov_0$ values. This pattern changes when accounting for covered priorities in the vicinity of the vehicle tours. GCorTOP consistently outperforms the two other models with respect to the $PCov_{100}$ and $PCov_{300}$ measures. Even though fewer high-priority target locations are directly visited by a UAV, more of them are at least close to a UAV's sensing location. These effects are obvious when the total number of targets that can be reached is limited with respect to instance size. In these cases, the trade-off between priorities and indirect coverage becomes particularly challenging. This increases the impact of incorporation aspects of spatial coverage and correlation in the objective.

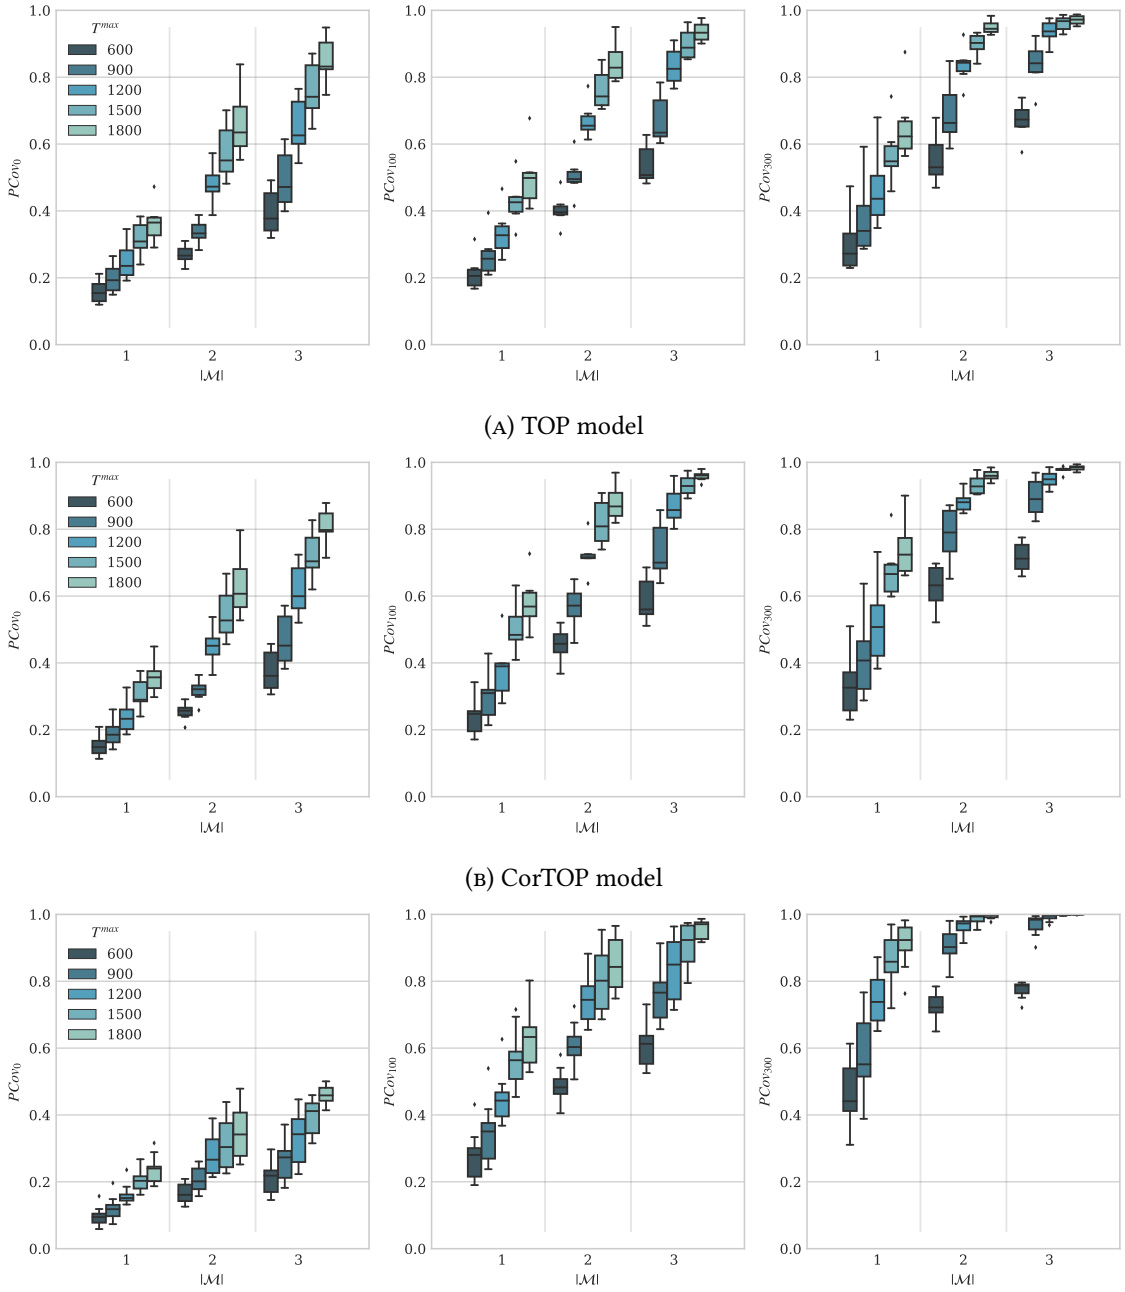

(A) TOP model

(B) CorTOP model

(C) GCorTOP model

FIGURE 9.4: $PCov_0$, $PCov_{100}$ and $PCov_{300}$ on largest scenarios for all three modeling variants, depending on number of vehicles $|\mathcal{M}|$ and route duration T^{max}.

$\lvert\mathcal{V}\rvert$	$\lvert\mathcal{M}\rvert$	$PCov_0$			$PCov_{100}$			$PCov_{300}$		
		TOP	CorTOP	GCorTOP	TOP	CorTOP	GCorTOP	TOP	CorTOP	GCorTOP
225	1	**58.1 %**	56.2 %	46.1 %	71.9 %	79.9 %	**85.2 %**	83.5 %	89.7 %	**95.5 %**
	2	**84.7 %**	79.0 %	62.9 %	92.3 %	94.8 %	**96.7 %**	97.2 %	98.4 %	**99.7 %**
	3	**93.4 %**	88.9 %	75.4 %	97.2 %	98.4 %	**99.1 %**	99.4 %	99.7 %	**100.0 %**
300	1	**42.5 %**	40.3 %	34.8 %	56.8 %	63.1 %	**71.4 %**	73.5 %	78.7 %	**87.8 %**
	2	**69.6 %**	65.9 %	42.7 %	83.2 %	87.2 %	**90.8 %**	93.8 %	95.6 %	**98.3 %**
	3	**83.5 %**	78.5 %	60.0 %	91.5 %	93.9 %	**96.3 %**	96.8 %	98.1 %	**99.7 %**
400	1	**42.9 %**	41.4 %	33.2 %	55.9 %	61.6 %	**70.0 %**	68.5 %	74.4 %	**85.8 %**
	2	**68.4 %**	64.8 %	46.2 %	81.0 %	84.4 %	**88.6 %**	89.3 %	91.4 %	**94.7 %**
	3	**81.7 %**	76.6 %	53.9 %	89.2 %	91.7 %	**93.5 %**	93.4 %	94.5 %	**95.8 %**
500	1	**37.4 %**	36.4 %	29.0 %	49.3 %	56.0 %	**63.1 %**	62.3 %	69.1 %	**80.0 %**
	2	**60.7 %**	57.4 %	45.7 %	75.3 %	79.3 %	**84.2 %**	86.0 %	88.7 %	**93.4 %**
	3	**75.1 %**	70.7 %	56.6 %	85.1 %	87.9 %	**90.9 %**	91.3 %	92.6 %	**95.1 %**
625	1	**25.8 %**	25.0 %	19.9 %	35.0 %	40.2 %	**48.4 %**	47.6 %	53.9 %	**67.1 %**
	2	**46.5 %**	44.0 %	34.9 %	63.7 %	68.8 %	**74.7 %**	78.7 %	83.7 %	**89.4 %**
	3	**63.1 %**	59.8 %	43.4 %	77.6 %	81.7 %	**86.4 %**	87.5 %	90.4 %	**93.8 %**
avg.		**62.2 %**	59.0 %	40.2 %	73.7 %	77.9 %	**82.6 %**	83.2 %	86.6 %	**91.7 %**

TABLE 9.8: Average values for $PCov_0$, $PCov_{100}$ and $PCov_{300}$ depending on target area size, number of vehicles, and modelling approach. Best values for each set are indicated in bold.

More detailed results on the largest instances with 625 candidate locations are provided in Figure 9.4, which summarizes the results over all runs in the form of box plot diagrams. For all models, increasing T^{max} and $|\mathcal{M}|$ has a strictly positive impact on all three measures $PCov_0$, $PCov_{100}$ and $PCov_{300}$. We can furthermore see that GCorTOP obtains much better values for $PCov_{300}$, specifically in the single vehicle case. This comes at the cost of direct priorities, which are low compared with the two other models. In configurations with two and three UAVs, covered utilities quickly exceed 95 % in this modeling variant. For TOP and CorTOP, this threshold is only attained in settings with three vehicles and higher route durations. In other words, the GCorTOP model achieves a high level of indirect coverage at significantly lower resource cost compared to the other two modeling variants.

PREDICTION QUALITY

Apart from the consideration of priorities, the ultimate goal of the UAV-based surveillance system is the provisioning of a reliable overview of the overall contamination across the affected area. Therefore, this section studies the quality of predicted distributions of contaminants that are determined based on the selected sampling locations. These also indicate whether the approximative way of integrating spatial dependencies into the GCorTOP objective is successfully yielding "informative" solutions.

The results with respect to the mean absolute error (MAE), mean error (ME), and

$\|\mathcal{V}\|$	$\|\mathcal{M}\|$	MAE			ME			WMAE		
		TOP	CorTOP	GCorTOP	TOP	CorTOP	GCorTOP	TOP	CorTOP	GCorTOP
225	1	10.67	8.41	**6.00**	0.97	0.82	**0.43**	5.24	**3.95**	4.64
	2	5.45	3.85	**2.06**	0.53	0.48	**-0.03**	1.40	1.87	**1.19**
	3	3.12	1.83	**1.16**	0.49	0.36	**-0.07**	0.47	**0.44**	0.90
300	1	9.62	9.07	**6.14**	1.12	0.62	**0.56**	7.72	6.00	**5.91**
	2	6.37	4.81	**3.84**	-0.14	0.21	**0.07**	2.42	**2.09**	2.69
	3	4.17	2.71	**2.17**	-0.44	-0.07	**-0.02**	1.16	**0.96**	1.46
400	1	12.56	11.22	**7.83**	**1.15**	1.36	1.41	6.37	5.69	**4.92**
	2	6.83	6.44	**3.80**	0.34	0.23	**0.18**	2.24	2.10	**1.93**
	3	4.26	3.71	**2.07**	0.27	0.19	**-0.03**	1.04	1.24	**0.88**
500	1	11.56	10.91	**8.57**	1.64	1.38	**0.83**	6.46	5.53	**4.34**
	2	6.55	6.14	**3.41**	0.46	0.44	**0.32**	2.20	1.93	**1.90**
	3	4.58	4.31	**2.49**	0.44	0.47	**0.18**	1.16	**1.06**	1.08
625	1	13.53	12.64	**9.28**	**0.33**	0.60	0.80	9.52	8.23	**5.69**
	2	7.23	6.49	**5.17**	**1.27**	1.31	1.36	3.28	2.56	**2.27**
	3	4.98	4.42	**2.88**	1.27	1.14	**0.32**	1.56	**1.41**	1.22
avg.		7.03	6.46	**5.09**	0.65	0.64	**0.49**	3.42	2.96	**2.80**

TABLE 9.9: Average prediction quality measures depending on target area size, number of vehicles, and modeling approach. Best values for each set are indicated in bold.

weighted mean absolute error (WMAE) are given in Table 9.9. We can see that with respect to MAE, the GCorTOP model invariably achieves the lowest values, i.e., it yields the most accurate predictions. Compared to TOP, MAE is reduced by approximately 43 % on average. Even compared to the CorTOP, which already incorporates aspects of spatial coverage, MAE is reduced by 33 %.

The prediction bias (ME) is within a reasonable range for all models, and largely negligible for the GCorTOP with multiple vehicles. In the solutions obtained using this model, sampling locations are spread widely across the target area. This reduces the prediction bias that may be introduced by taking several samples in close proximity and thus likely with very similar values.

For all models, the WMAE measure is less than MAE. This is an effect of the integration of priorities. Prediction accuracy tends to deteriorate with increasing distance to the next sampling location. The focus on samples at these prioritized locations ensure that predictions are comparatively good in the corresponding areas. Despite this effect and even though CorTOP and GCorTOP both come with a reduction of the direct priorities $PCov_0$, TOP does not yield the lowest WMAE. This is due to two reasons: First, the overall reduction in MAE that is achieved by incorporating correlations also helps to reduce WMAE. Second, routes planned using TOP are often very narrow, leaving out some areas with high priorities that do not justify a detour when coverage is not considered. This is also demonstrated in the visual comparison at the end of this section.

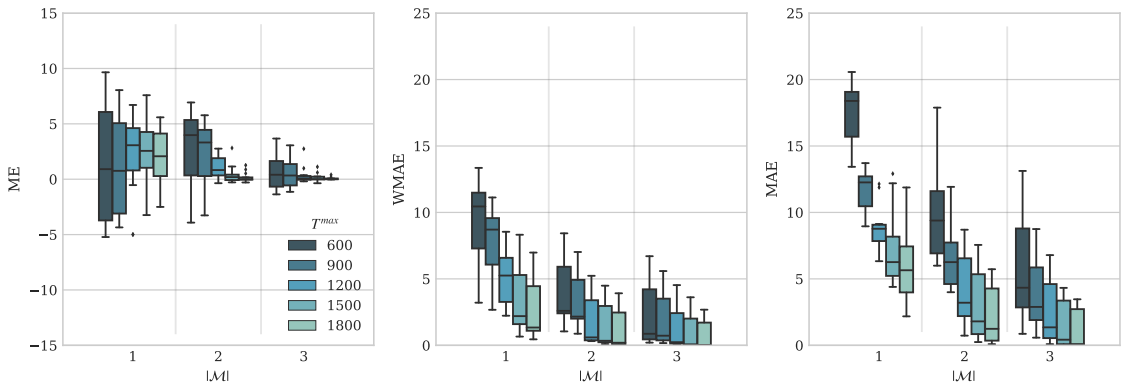

FIGURE 9.5: Comparison of ME, MAE, WMAE for GCorTOP on the largest instances (625 locations), depending on number of vehicles $|\mathcal{M}|$ and route duration T^{max}.

As in the previous section, we look more closely at the largest instances comprising 625 locations. The results are given in Figure 9.5 for GCorTOP. We do not explicitly indicate results obtained using the two other models as they show a very similar pattern. We can observe that both MAE and MAE are unacceptably high for instances involving a single UAV and low flight time budgets. These errors quickly decrease as either flight time or the number of vehicles increases. This is particularly true for WMAE, which decreases more rapidly than MAE. In the case of a single vehicle, we can obtain a WMAE of less than 2 within a flight duration of T^{max} = 1,800 s, which is feasible in practice. If more UAVs are available, this is even possible within much shorter flight times of T^{max} = 1,200 s or less.

Comparing these results with the covered priorities (Table 9.8), high values for $PCov_{100}$ and $PCov_{300}$ are consistently associated with low error measures. A comparison of figures Figures 9.4c and 9.5 demonstrates this in more detail. The covered priorities come close to 100 % given sufficient resources, while MAE and WMAE decrease simultaneously. Conversely, configurations that have a $PCov_{100}$ of 60 % or less yield a MAE of 5 or more, indicating significant deviations between the prediction and the true value. These observations suggest that our approach for modeling spatial coverage in the objective function can successfully approximate the complex spatial interdependencies within the surveyed regions.

From a practical point of view, these are promising results. Even for the largest instance set, prediction quality is improved. While predictions are by no means flawless, this means that for most areas of the target region, predicted values are at least close to the true distribution. As the WMAE indicates, high prediction accuracy can be achieved at least for crucial areas, even given comparatively limited flight times. A visual representation of these results is also given in the last section of this chapter.

Prediction robustness

The measures discussed have been computed without accounting for measurement errors. In this section, we consider the impact these measurement errors can have on prediction quality, thereby evaluating whether the proposed approach is applicable in practical situations. As stated in Section 9.2.4, we modify the observed values by adding noise drawn from a normal distribution. The MAE depending on the noise level is indicated in Figure 9.6. Again, we focus on the largest instances where the impact is most noticeable.

As the figure demonstrates, the measurement errors negatively impact prediction quality. This is most obvious for low values for T^{max}, i.e., the routes containing the fewest samples. In case of s single vehicle, the MAE increases by nearly 20 % for T^{max} = 600 s. For $|\mathcal{M}|$ = 3 and T^{max} = 600 s, it nearly doubles. The results also show that doubling the standard deviation of the noise distribution does not have an equally high influence on the results: Instead, increasing the noise level k from 0.05 to 0.1 comes with much lower increases in MAE than increasing it from 0 to 0.05 in the first place. In other words, while measurement errors have an impact on prediction quality, it does not appear to be highly sensitive with respect to these errors. Furthermore, the model still achieves a major reduction in MAE, even in the case of the highest noise level. This indicates that it is still possible to achieve a prediction of similar quality than in the error-free case, although it may require a higher total flight time.

Visual comparison

In the last step, we give a visual indication of the impact that considering coverage aspects can have on the solutions. To this end, we compare the results for two instances. The first example is depicted in Figure 9.7. We indicate priorities on the left-hand side and the simulated "true" distribution of contaminants on the right. In both figures, high values are indicated in red and low values in blue.

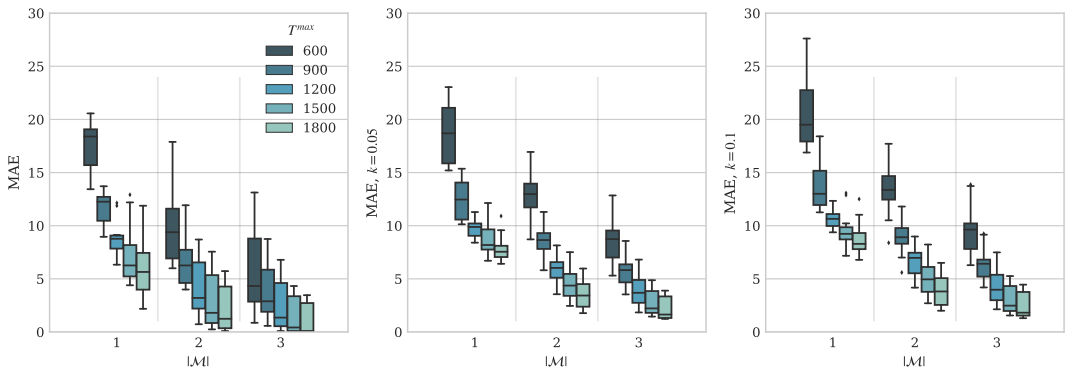

Figure 9.6: Comparison of MAE for different noise levels for GCorTOP on the largest scenarios (625 target locations).

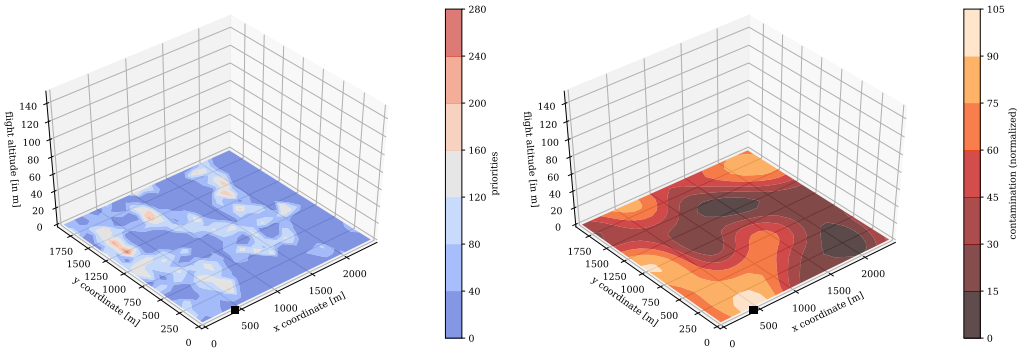

(A) Priorities in the target area (B) Simulated distribution of gases

FIGURE 9.7: Base data of the benchmark instance for example 1.

Exemplary routes obtained using the three modeling variants are given in Figures 9.8 to 9.10. These routes are planned for a single UAV with a maximum flight time of T^{max} = 1,500 s. Flights start and end at position (0, 400). For each solution, we indicate the planned mission on the left-hand side, superposed on the priorities that are the basis for this plan. On the right-hand side, we give the predicted distribution.

The TOP results in the narrow route given in Figure 9.8. We can see that it traverses a region with high utilities but leaves major parts of the target region entirely unexplored. This means that for nearly half the target area, no reliable prediction is available. The mission planned using the CorTOP (Figure 9.9) is broader in comparison to Figure 9.8. While it follows a similar trajectory than this path, sampling locations are better spaced, with more unvisited target locations in between the sampled locations. Finally, the GCorTOP determines a route that differs strongly from the two other models. The UAV travels further to gather more information at the regions with higher priorities that are indicated in the right part of the figure. This significantly improves prediction quality.

The base data for the second example is indicated in Figure 9.11. Compared to the previous example, we can see that the simulated distribution of values is more coarse, i.e., it changes more quickly. Again, results for the three models are depicted in Figures 9.12 to 9.14. All vehicle routes are planned for one vehicle with a flight duration limit of $T^{\max}$ = 1,200 s. The TOP solution in Figure 9.12 clearly mirrors the shape determined by the priorities. It successfully covers almost all of the highly prioritized target locations. Even though, the prediction is comparatively imprecise due to the coarse spatial distribution, which means that prediction quality deteriorates quickly even in short distance to a sampling location.

As the TOP already yields a cyclic shape, these sampling locations are better distributed in space than in the previous example. The CorTOP model leads to a very similar result, both in terms of route trajectory and prediction quality. The GCorTOP further emphasizes spatial interdependencies, yielding a much broader route compared to the two other solutions (Figure 9.14). While prediction quality still suffers due to the

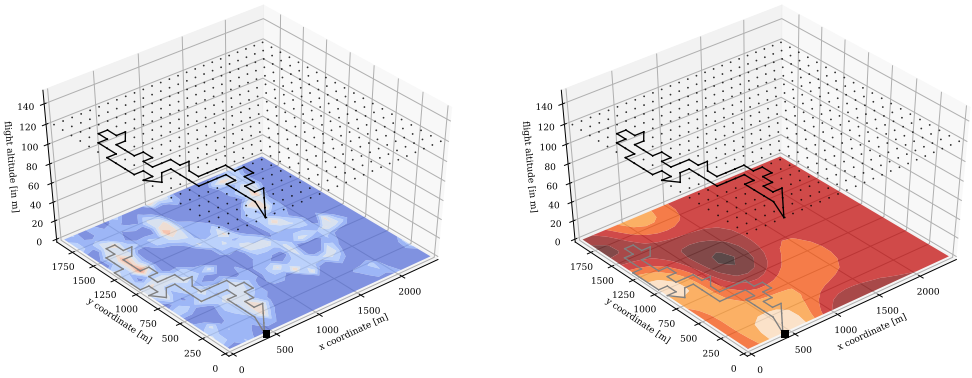

(A) Planned mission (B) Predicted distribution

FIGURE 9.8: TOP results for example 1

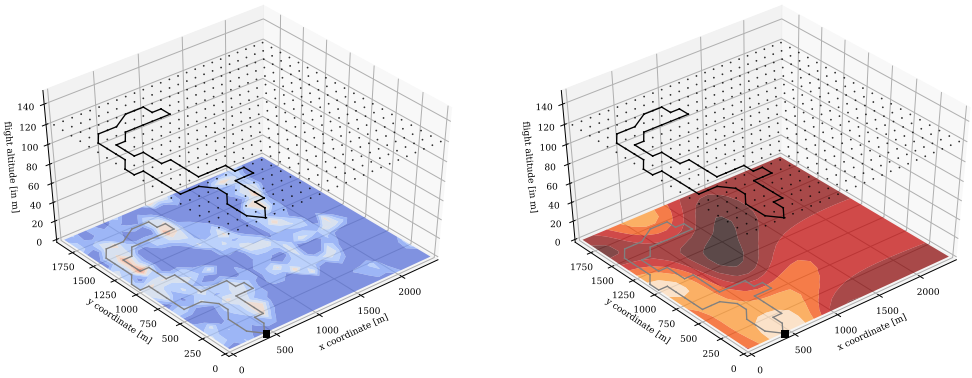

(A) Planned mission (B) Predicted distribution

FIGURE 9.9: CorTOP results for example 1

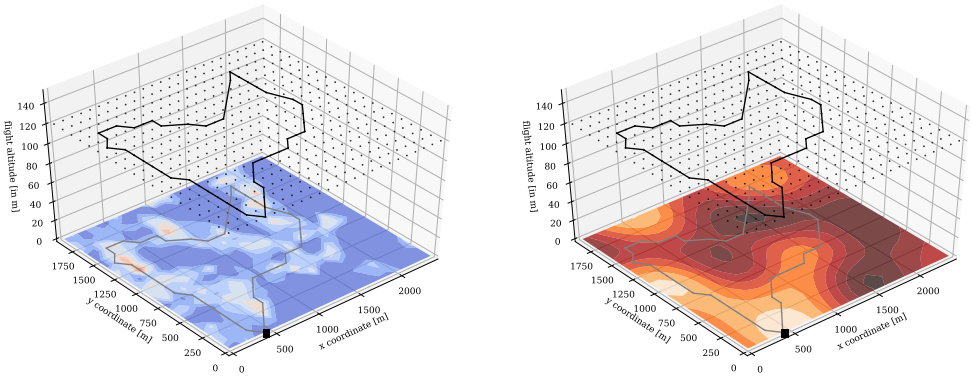

(A) Planned mission (B) Predicted distribution

FIGURE 9.10: GCorTOP results for example 1

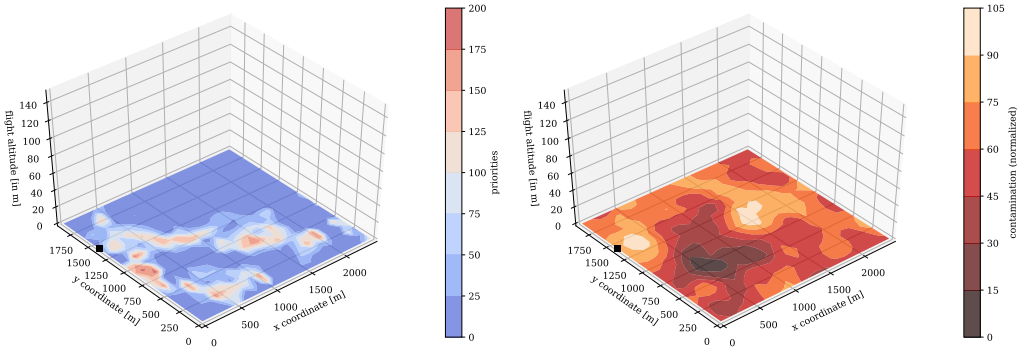

(A) Priorities in the target area (B) Simulated distribution of gases

FIGURE 9.11: Base data of the benchmark instance for example 2

coarse distribution, this is the only model to indicate higher contamination values at the right-hand side of the target area. These are severely underestimated using the two other modeling approaches.

9.4 Discussion

This section summarizes and discusses the effects observed in the computational study. Moreover, we derive insights into the applicability of the proposed solution approaches and models for solving the MPPES in realistic settings.

Solution approaches

The results of the exact approach on small instances with up to 49 target locations showed that instances quickly become intractable if they involve more than 20 targets. The results of the proposed two-phase multi-start adaptive large neighborhood search (2MLS) and two-phase multi-start adaptive large neighborhood search (2MLS)-f on benchmark instances for the TOP highlight the advantage of 2MLS-f compared to the other proposed approaches in terms of computation time, with an average gap of 1.62 % to the best known solution that is achieved within an average computation time of 10.7 s on the largest available instances for TOP.

Considering the new instances for the MPPES with up to 625 targets and 3 UAVs, all except the largest instances involving multiple vehicles have been solved in 100 seconds or less on average. Meanwhile, the gap to optimality for the small instances that can be solved exactly is 0.0005 % for 2MLS and 0.1 % for 2MLS-f. From a practical point of view, the average deviation to the best known solution can be considered acceptable, as the efficiency and scalability of the solution approach are of higher importance in our use case. In this respect, the quick convergence of the proposed approach, particularly when using a vehicle decomposition scheme in the first search space, supports its applicability in practice.

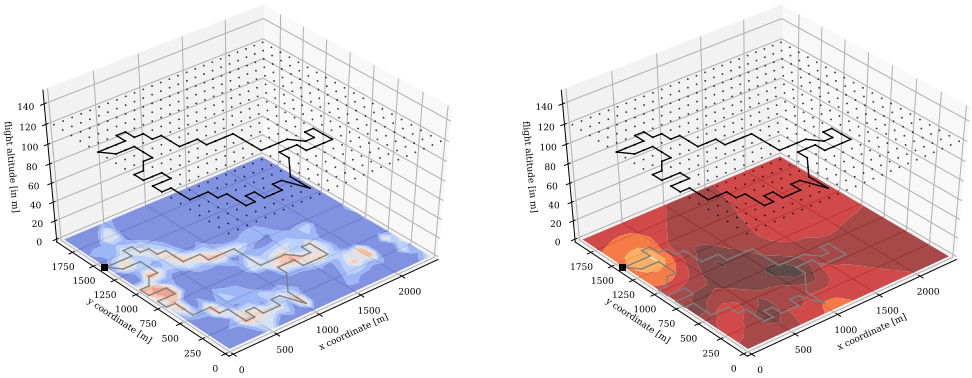

(A) Planned mission (B) Predicted distribution

FIGURE 9.12: TOP results for example 2

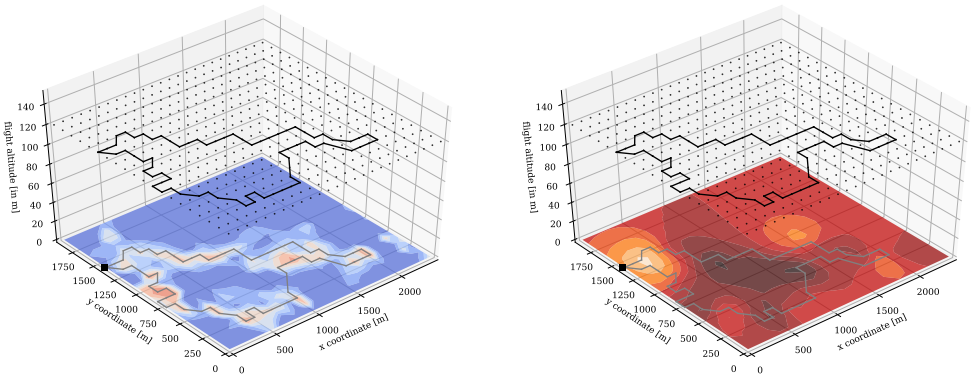

(A) Planned mission (B) Predicted distribution

FIGURE 9.13: CorTOP results for example 2

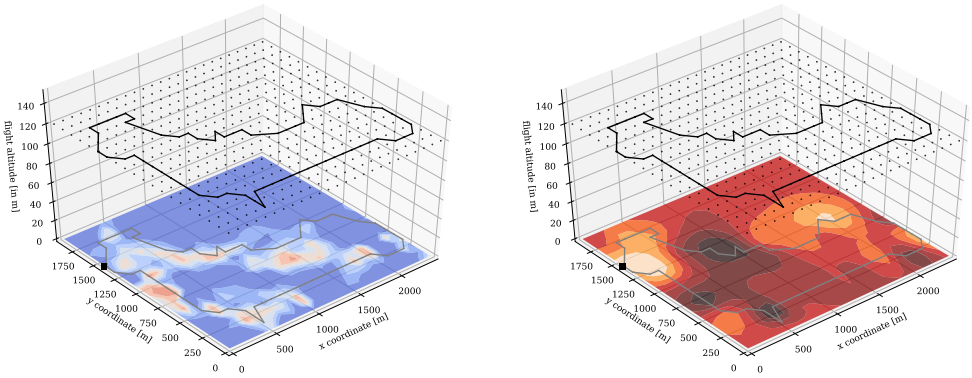

(A) Planned mission (B) Predicted distribution

FIGURE 9.14: GCorTOP results for example 2

Models

To obtain insights into the performance of the GCorTOP model, we compared the results with those obtained using two other related approaches, the TOP and the CorTOP. The results demonstrate the trade-off between maximizing the priorities of targets that are directly included in vehicle routes and maximizing the indirect coverage of those that are in the vicinity of a tour. Increasing indirect coverage is consistently associated with a decrease in the sum of directly included priorities. The models compared in this section deal with this challenge in different ways: TOP emphasizes direct priorities, whereas GCorTOP increases the focus on indirect coverage.

The advantages of the latter approach are demonstrated when looking at the prediction quality. Here, the GCorTOP model consistently achieves the best results. Despite its lower focus on direct priorities, this is even true for the weighted mean absolute error, which focuses on prediction quality at highly prioritized areas. This suggests that our approach for modeling spatial correlations in the objective function is a useful representation of the true spatial interdependencies in the observed phenomenon. As the visual comparison indicates, this leads to the coverage of major parts of the target region, whereas routes planned using TOP and CorTOP models are much more narrow.

9.5 Summary of Part II

In this part, we studied offline planning variants for emergency rapid mapping. In Chapter 4, we introduced the mission planning problem for emergency surveillance (MPPES). The MPPES corresponds to the problem of selecting sampling locations such that the overall information gained at critical locations is maximal within technical limitations of the UAVs. In Chapter 5, we provided a comprehensive overview of approaches that address spatial coverage and priority maximization in vehicle routing. We introduced a new classification to distinguish between different types of spatial coverage and summarized advances from different fields of research: To our knowledge, this is the first detailed overview that combines approaches from vehicle routing with advances made in informative path planning (IPP) in the field of robotics.

Based on this survey, we discussed the capabilities and limitations of existing approaches in Chapter 6. We demonstrated that, while approaches in the VRP literature have lead to efficient heuristic solution approaches for related problems, the models that are used do not provide the level of detail we require for the MPPES. In contrast, advances in literature on IPP offer detailed models, but require a computational effort that prevents the application of these approaches in time-sensitive situations. To fill this gap, we introduced new models and solution approaches in Chapters 7 and 8. Specifically, we proposed the generalized correlated team orienteering problem (GCorTOP) as a flexible means for approximating spatial coverage. This model is more detailed compared to the simple coverage-based approaches proposed in the domain of vehicle routing, but less computationally expensive than the probabilistic approaches currently available.

Moreover, we developed an efficient two-phase multi-start adaptive large neighborhood search (2MLS) for quickly determining high-quality UAV missions. In particular, we combined a first search phase that emphasizes spatial coverage with efficient improvement steps in a second step. We also showed how the GCorTOP model can be integrated into this approach to allow the quick evaluation of search moves. Furthermore, we introduced an exact approach for benchmarking purposes.

In Chapter 9, we evaluated these approaches from a practical point of view. We introduce new benchmark instances for the MPPES based on real-world data on population density and derive measures for assessing the quality of solutions that take priorities and spatial interdependencies into account. Using benchmark instances for the TOP, we showed that the 2MLS achieves high solution quality within short computation times. Extensive experiments on MPPES instances demonstrated that the proposed GCorTOP approach improves over previous modeling variants in terms of both covered priorities and achieved prediction quality. Overall, we demonstrated that the approaches introduced in this thesis can yield high-quality predictions within the limitations on computation time that are imposed in an emergency setting.

Part III

Online mission planning

10 Online planning and online learning

In part III of this thesis, we investigate the planning of information maximizing UAV tours in a dynamic setting, dealing with new information that has been obtained about the surveyed environment. In this chapter, we introduce the corresponding planning problem and summarize requirements that characterize promising solution approaches. We illustrate the combination of online planning and online learning of spatial phenomena in a representative example.

10.1 Adaptive mission planning problem for emergency surveillance (AMPPES)

In Part II of this thesis, we have studied offline approaches for solving the mission planning problem for emergency surveillance (MPPES). The algorithms make decisions based on the assumption that the essential characteristics of the phenomenon are known before the flight. Most importantly, this concerns the spatial correlation of the observed distribution of contaminants, which indicates to which degree samples taken at different locations may yield similar results. Based on this prior knowledge, it proceeds to select "informative" samples, taking both the spatial correlation between surveyed and surveyed locations and priorities in the target area into account.

In this part, we extend the issue of selecting informative sampling locations to the online case, where decisions are made while the mission is still in progress. Again, the overall objective is to provide accurate predictions about the contamination at critical locations, i.e., at locations where a large number of people may be affected or where the likelihood of needing support by the emergency services is high. Addressing this problem in an online setting involves two aspects that differentiate it from the MPPES: Online planning and online learning. Online *planning* refers to the selection of the sampling locations, which is done in an online fashion, i.e., only a few sampling locations are selected to be surveyed at each point in time, and the next sampling locations are determined after completing the previous set. This also facilitates adjusting to new pos-

sibilities or restrictions, e.g., newly available or failing UAVs. Online *learning* means that the information obtained during the flight is used to train and update predictive models that represent the current distribution of gases before the end of the mission. This way, these models provide the best possible prediction at each point in time, and are continuously refined by new incoming samples.

In our study, we seek to gain insights into how these two aspects can be combined to improve the overall information that is gained during the surveillance mission. To this end, we define the *adaptive mission planning problem for emergency surveillance (AMPPES)* as the problem of determining the next sampling locations for UAVs based on the previously obtained knowledge such that the prediction quality at the end of the mission is maximal.

Solution approach to the AMPPES should fulfill three focal properties to be successfully applied within a real-world context:

Adaptivity: Models should be able to adjust to the obtained knowledge about the environment. Most importantly, this involves the utilization of updated environmental beliefs for the selection of the next sensing locations.

Scalability: To support the deployment of UAV fleets in real-world applications, models need to be able to handle an increasing number of UAVs and a large number of candidate sensing locations under consideration of real-time constraints.

Efficacy: Efficacy describes the capability of an approach to achieve the primary mission objective, i.e., a low prediction error at important target locations.

It should be noted that these requirements are conflicting: For example, adaptive models require a higher computational effort compared to less complex online planning approaches due to the need to update predictive models. This is contradictory to the need to reduce computation times to remain scalable if there are several UAVs involved and updates have to be processed often. For this reason, the main focus of this part is to identify strategies that provide a reliable balance between these requirements and to assess the capabilities and limitations of these approaches in the context of emergency surveillance.

10.2 Illustrative example

Figure 10.1 illustrates the combination of heuristic planning strategies and online learning of environmental processes using a simplified example. The figure represents the planned mission and obtained information at three points in time during a mission surveilling a target area of 1.5×1.5 km^2.

In the first snapshot that is shown, a few samples have been collected near the starting location of a UAV. As the collected values are similar, a predictive model trained with these few initial samples initially assumes that the surveyed phenomenon is relatively smooth, i.e., that the observed level of contamination is likely to be stable throughout

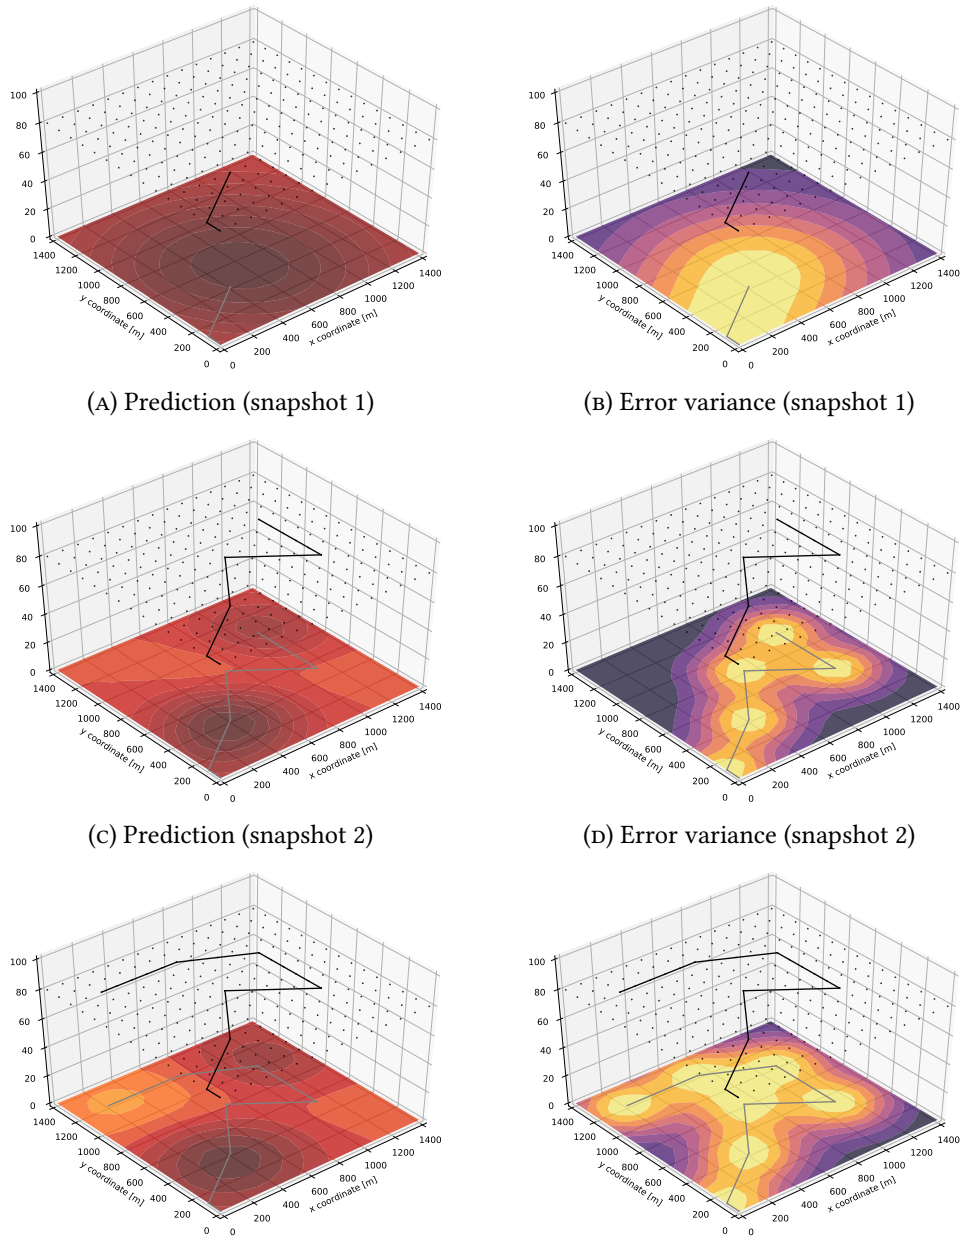

(A) Prediction (snapshot 1)

(B) Error variance (snapshot 1)

(C) Prediction (snapshot 2)

(D) Error variance (snapshot 2)

(E) Prediction (snapshot 3)

(F) Error variance (snapshot 3)

FIGURE 10.1: Evolution of the predicted distribution (left) and variance (right). High predicted contamination is indicated in orange, high error variance is represented in dark purple.

the considered area (Figure 10.1a), and that the prediction error variance is relatively low (Figure 10.1b).

As the UAV travels further, newly obtained samples differ more from the predicted level of contamination than the initial model had assumed. This invalidates the initial assumption of a very smooth process. Most importantly, this means that the initial assumption about the nature of the spatial correlation has been incorrect and that less information can be inferred about unobserved locations further away from the collected samples. Consequently, the estimation error variance after updating the environment model increases much more quickly with distance to the collected samples (see Figure 10.1d relative to Figure 10.1b). This information then motivates the next sampling locations, leading to a more accurate mapping of the process (Figure 10.1e) and an overall reduction in the prediction error estimate, particularly in the center of the target area (Figure 10.1f).

Overall, example 10.1 illustrates the possibility of training a process model in an online fashion and using the obtained knowledge to select the next sampling locations. Most importantly, this means that incorrect beliefs about the correlation can be corrected, which is not possible in offline or non-adaptive planning. This allows, e.g., to increase the spatial density of samples if the surveyed process varies more quickly, and to identify areas where the collected information is insufficient.

However, we can also see that the main outcome of the planning approach is a mission where sampling locations are well distributed across the target area, similar to the effects we have observed in Part II of this thesis. This observation motivates the following chapters, in which we investigate to what extent adaptive planning is advantageous and in which situations simpler strategies are sufficient.

10.3 Outline of Part III

In this part, we study solution approaches for the adaptive mission planning problem for emergency surveillance, i.e., the problem to compute UAV missions in an online fashion, using information obtained during the mission to improve the selection of subsequent sampling locations. Our main concern lies in the consideration of the trade-off between the adaptivity, the efficacy and the scalability of a solution approach. To this end, we first summarize approaches for related problems that have already received attention in literature in Chapter 11. We identify common solution principles within the domain of environmental surveillance and complete this overview with a summary of approaches for solving dynamic problems in related fields.

In Chapter 12, we summarize these findings with particular attention to the planning requirements stated in this chapter. We identify the main research gaps that are currently complicating the implementation of these approaches in practice. In Chapter 13, we introduce a unified model for the AMPPES and derive several modeling and solution variants based on approaches from the literature and the concepts introduced in Part II of this thesis. We also propose new models tailored toward dealing with time-dependencies

in the surveyed data, i.e., phenomena that change over time.

We seek to evaluate and compare these models in an unbiased fashion. Therefore, we implement them within a common solution framework, which is introduced in Chapter 14. This chapter furthermore introduces a discrete event simulation framework that is developed in this thesis to enable the unbiased comparison of solution concepts. The simulative study and the obtained results are discussed and analyzed in Chapter 15. We furthermore summarize the main insights and give recommendations for applying the solution approaches in practice.

11 Online planning in dynamic environments

This chapter reviews the state of the art related to our problem. We do so from two perspectives: In Section 11.1, we summarize approaches addressing related applications in environmental surveillance. Section 11.2 looks into methods for solving dynamic problems proposed in other domains. The findings are summarized in Section 11.3.

11.1 Solution approaches in environmental surveillance

Due to the increasing availability of UAVs, there has been a growing interest in using these vehicles for surveying spatial phenomena, both in environmental applications and in emergency response. At the same time, communication technologies allow the coordinated deployment of multi-vehicle teams and fleets, which has motivated the development of strategies for planning multi-UAV missions in dynamic settings. Possible applications range from precision agriculture (Popović et al., 2017) over radiation mapping (Newaz et al., 2016) up to forest fire surveillance (Casbeer et al., 2006). All of these applications have in common that mobile sensing systems move in a two-dimensional or three-dimensional space to obtain information about the environment that surrounds them.

Numerous solution approaches have been developed addressing these problem settings. However, these approaches are typically tailored toward specific applications. The goal of this overview is to identify common concepts and fundamental principles that constitute the basis of many of these approaches rather than to discuss individual approaches in detail.

To analyze these concepts, we propose a classification that distinguishes between six essential aspects:

1. the addressed mission objective,

2. the coordination mechanism for multiple vehicles,
3. the control scheme (i.e., central decision making or decentralized control),
4. the considered model of the surveyed environment,
5. the proposed planning heuristic (i.e., the logic determining the next sampling location), and
6. the decision scope.

Figure 11.1 gives an overview of these aspects and summarizes the most common problem types. We emphasize that these aspects cannot always be separated. For instance, some planning objectives such as source localization often use specific process models that are intertwined with dedicated solution approaches.

Several dimensions have already been introduced in previous publications, especially addressing the control scheme and modeling and heuristic approaches: Stranders (2010) addresses several application domains in the context of "situational awareness" and proposes unified model-based solution approaches. This survey is closest to the one offered here in its distinction between environment representation and planning approaches. Low et al. (2008) distinguish between adaptive and non-adaptive model-based approaches, while Schwager et al. (2011) study and compare model-based, geometric and potential field approaches. Hutchinson et al. (2017) summarize modeling variants for source location and boundary tracking. Control and communication aspects have been discussed in the survey articles by Ponda et al. (2015). Finally, Otto et al. (2018) cover optimization approaches in several UAV-based application domains.

The previously mentioned publications address several of the dimensions in more detail compared to this work. No consensus or final classification scheme covering more than a few of these aspects has been achieved yet. To our knowledge, our work is the first one in this direction. Due to the versatility of approaches proposed in the literature, we give no detailed overview of solution variants. We rather identify the most common principles used in the design of solution approaches in this line of work.

As the main focus of this work lies in finding sensing locations for rotary-wing UAVs, we do not discuss system dynamics, e.g., maneuverability. In general, we assume that a flight controller is capable of executing the proposed sequence of waypoints. Considering the integration of these aspects, we refer the reader to Strobel (2016) and Ritter (2017). Furthermore, the development and refinement of predictive models for spatial processes is not at the heart of this thesis. We, therefore, give a summary of commonly used models but do not discuss these in more detail.

11.1.1 Mission objectives

The mission objective determines the main purpose of a UAV-based surveillance system. In the context of spatial processes, we distinguish between five types of high-level planning objectives that are often addressed in literature. These objectives differ either in the

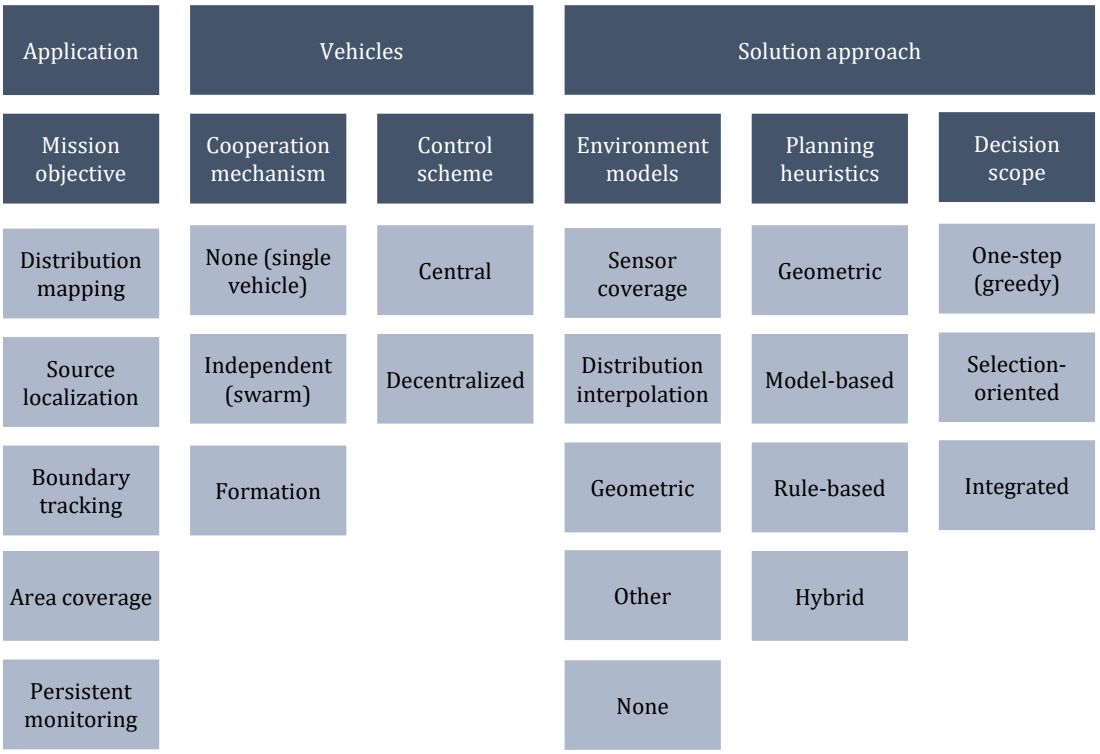

FIGURE 11.1: Online planning approaches for surveying spatial phenomena.

sort of information that is desired or in the way they seek to achieve this information. Figure 11.2 illustrates these applications along with representative mission trajectories. Their main characteristics are summarized below:

Distribution mapping refers to applications where an accurate estimate of the entire distribution of the observed phenomenon is required. This is for example used in the context of agricultural or oceanic applications (Hitz et al., 2017) as well as for gas distribution mapping (Neumann, 2013). Specific to this planning objective is that it is also relevant to obtain information about areas where the surveyed phenomenon is only present on a smaller scale. Distribution mapping can furthermore become relevant for source localization and boundary tracking addressed below, e.g., to identify multiple sources of contamination.

Source localization and characterization applications aim at identifying the source of a gas spill or leakage as quickly as possible. Often, this combines an estimate of the intensity of the leak along with the identification of its position (Kuroki et al., 2010; Ristic et al., 2017). Applications differ in whether the objective is to move a sensor to the source location ("source seeking") or rather to predict this location accurately, independent of the final sensor location ("source term estimation") (Hutchinson et al., 2017).

Boundary tracking variants seek to delineate the extent of a (compact) phenomenon. The objective is to trace, as accurately as possible, the "edge" that is defined by a certain intensity or concentration of the monitored substance (e.g., Casbeer et al., 2006; Euler, 2017). In contrast to distribution mapping, further information about the intensity of a phenomenon, e.g., in less affected areas, is not required.

Area coverage and exploration can be considered as a substitute for surveillance applications where observations spread across a field of action are beneficial (Megerian et al., 2005). In environmental monitoring, it serves to ensure a reliable and broad coverage, often in a setting similar to distribution mapping.

Persistent monitoring seeks to observe a phenomenon that is slowly changing over time, e.g., in oceanic applications (Smith et al., 2011). Instead of a one-time surveillance mission, the target locations are surveyed repeatedly throughout a longer planning horizon.

Other applications for UAV-based surveillance outside the domain of environmental surveillance include strategic patrolling or intruder and object detection and tracking. Strategic patrolling refers to applications where strategically acting entities have to be intercepted by the UAVs. Intruder and object detection requires the detection of other entities, either static or moving, in the target area. Similar to environmental monitoring, many planning approaches for these applications rely on formalized beliefs about a UAV's environment, e.g., the probability of a certain target being present at a given location (Stranders et al., 2013).

11.1.2 Vehicle coordination

Most of the applications addressed in this study rely on the deployment of multiple vehicles with a common mission objective. In this case, coordination mechanisms are needed to determine how a UAV acts and interacts in a fleet. We can distinguish between two forms of coordination:

Independent vehicles can move and act independently from one another. They do not need to remain close to one another and can move freely across the affected area. Using a decentralized control mechanism (see also Section 11.1.3), the group of UAVs is also referred to as a "swarm" (Strobel, 2016).

Vehicles in a formation are constrained by their relative positions. Typically, this means that vehicles have to remain within given distance limits between one another. In most cases, this is due to communication issues. The shape of these formations does not need to be static, as relative positions can change and adjust during the mission (see, e.g., Daniel et al., 2011; Ristic et al., 2017).

For a more in-depth discussion, we refer to the overview presented by Strobel (2016). In our study, we focus on independent vehicles. We include solution approaches targeted

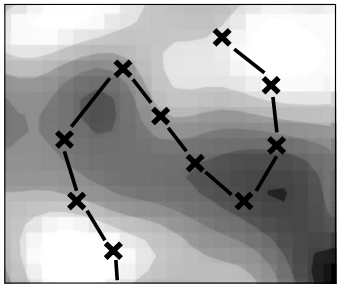

(A) Distribution mapping

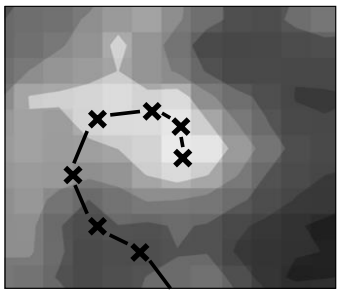

(B) Source localization

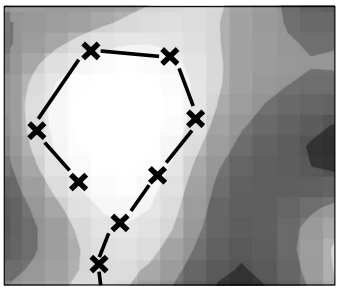

(C) Boundary tracking

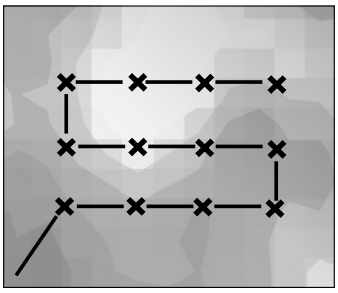

(D) Area coverage

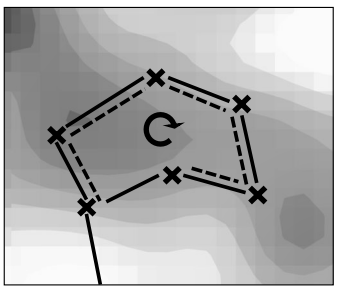

(E) Persistent monitoring

FIGURE 11.2: Applications for online planning approaches.

towards UAV formations in this section if the corresponding planning concepts are transferable to cases with independent vehicles. For example, this includes approaches where a high-level heuristic determines where a formation moves, while low-level coordination mechanisms determine the positions of UAVs within the formation.

11.1.3 Control mechanism

The control mechanism characterizes the underlying architecture that determines the UAVs' actions.

Central control mechanisms run on a central machine with full information on the current system state. The results benefit from the available information and processing power, but require a stable communication with small time lags.

Decentralized control refers to applications without a central instance, no shared or partially shared information basis and without guarantees for arrival reliability of messages passed between the vehicles. Decentralized approaches are more robust, but rely on more complex communication and require processing power on the vehicles themselves, in particular when using computationally expensive process models. We can furthermore distinguish between decentralized control schemes with active or explicit coordination and those with passive or implicit coordination (Stranders, 2010; Ponda et al., 2015). Active coordination relies on messages exchanged between vehicles, with each vehicle communicating its current belief about its environment and its current intentions. Passive coordination is relevant in all cases where each vehicle only obtains information about others through observing the environment, e.g., by monitoring and following other vehicles (Stranders, 2010).

For more information especially concerning decentralized variants, we refer to Ponda et al. (2015).

11.1.4 Environment models

In this section, we summarize models that serve as a representation of the dynamic environment in which the UAVs operate. These models provide the knowledge used by the planning heuristics discussed in Section 11.1.5 to decide on actions and movements. There exist a large variety of concepts for environment representations, and many authors apply highly application-specific models. In the following, we extract principles that are most common in literature. The different types of models are briefly defined below and discussed in more detail in what follows.

Sensor coverage models represent the sensors themselves together with the area that has been covered by previous measurements.

Distribution interpolation variants model the surveyed environmental phenomenon, e.g., the expected contamination, as a spatial process that can be interpolated using the obtained samples. These variants often employ probabilistic models that capture the spatial and temporal correlations governing the process.

Geometric phenomenon representation represent the surveyed phenomenon not as a spatial process, but as a geometric object. As a consequence, they are less interested in the distribution of values but rather in identifying borders delineating the extent of a phenomenon.

Other Other variants are application-specific models or techniques such as artificial potential fields, which are used to plan trajectories based on known positions and obstacles. Hence, we provide only an overview of approaches that have been used several times without claiming completeness.

None Finally, some solution approaches discussed in Section 11.1.5 are model-free, i.e., they use no representation of the environment for planning vehicle missions.

In the following, we introduce the different representations in more detail. We highlight differences between the concepts and give examples for successful applications in literature.

Sensor coverage

Sensor coverage models represent the sensors and their ranges. They indicate areas where information is provided as well as those that are not covered in the current setting. Typically, they model a sensor with a given sensing radius, determining the area in which information can be obtained. In a very basic form, these are, e.g., disc-shaped representations that assume that full information is obtained about the area immediately surrounding a UAV (e.g., Euler, 2017). More complex representations account for the fact that information gain varies with the distance to a sampled location.

The GCorTOP model introduced in Chapter 7 falls into this line of work. Some authors in this line of work also use stochastic model enhancements. An example is given by Lambrou and Panayiotou (2013), who model sensor coverage as the probability of detecting an event, which is decreasing with increasing distance to the sensor location.

Distribution interpolation

In this work, we use the term "distribution interpolation" for summarizing methods that represent the surveyed phenomenon as a process varying in space and potentially over time. This is the most versatile group of methods. Authors use a wide range of techniques that have been derived in other fields of research or apply dedicated variants that are specific, e.g., to oceanic phenomena or gas distribution mapping.

Among the most popular means in this line of work are probabilistic models such as Gaussian processes (GP), which represent the environmental processes themselves as a family of random variables with specific dependencies between one another. They provide, in addition to a prediction of the phenomenon, information about the expected prediction variance and estimation error, which can be exploited for designing strategies that minimize these error measures. We refer to Chapter 3 for a more detailed discussion of probabilistic models. In addition to these techniques, which are widely used in other domains, several dedicated dispersion models have been proposed for modeling gas distributions. Notable approaches are Gaussian plume or Gaussian puff models, which describe the distribution of gases originating from a single source considering influences such as wind and either a continuous or an intermittent release of the substance (Hutchinson et al., 2017). The kernel extrapolation distribution mapping (Kernel DM) algorithm and its extensions use Gaussian kernels to determine predicted distributions and expected prediction variance maps (Neumann, 2013; Reggente, 2014).

These concepts are complemented by models addressing specific distribution patterns. An example is AERMOD which has been developed for modeling atmospheric dispersion (Truong et al., 2016). This approach operates on a larger scale than our use case, e.g., scenarios where areas of several hundred square kilometers are affected. Other dispersion models dedicated to source localization and characterization are reviewed in Hutchinson et al. (2017).

Geometric phenomenon representation

Geometric models for representing the surveyed phenomena are used especially in boundary tracking applications. They try to model distributions using specific geometric shapes such as polygons and fit the parameters of these objects such that they correspond to the obtained observations (Saldaña et al., 2015; Newaz et al., 2016). Some variants of this only use local information rather than modeling the entire phenomenon using one object. In this case, the contour of a phenomenon is for example approximated by a straight line or polygon, which is only valid within a limited distance to a sensor (Casbeer et al., 2006).

Other approaches

We give a brief overview of approaches and ideas that do not fit the categories above but are used in several publications. Motivated by applications such as intruder detection, some publications decide based on some hypotheses about the environment. Information gathered during the mission can be used to confirm a hypothesis or show that it is inconsistent with the observations (Lim et al., 2016). This allows modeling the information gain of each candidate location with respect to the formulated hypotheses.

Originally used for robot motion planning, artificial potential fields (APFs) have been applied to mission planning as well. In APF, the UAVs are represented as objects moving in an (artificial) force field. Areas to be avoided are modeled as repellent areas with

high potential. Prioritized or "informative" target locations are represented as attractive forces with low potential (e.g., Saldaña et al., 2015). The vehicles try to move from high to low potential areas, following a negative gradient in the force field. In environmental surveillance, APFs are often applied when obstacle avoidance or formation control is crucial. To give an example, other vehicles are modeled as high potentials, and desired positions in vehicle formations can be modeled as potential minima at specific angles and distances relative to other vehicles to avoid collisions (e.g., Fiorelli et al., 2006). Due to their close connection with motion control, they cannot be strictly separated from the planning heuristics discussed in the next section.

11.1.5 Planning heuristics

A wide range of planning variants has been proposed in literature. These are usually application-specific and closely tied to the mission objective and process models. Nonetheless, we can identify common features that are used in many of these approaches. Below, we first introduce these fundamental principles. We then introduce specific approaches in more detail. Note that, in all variants, the underlying planning problems are rarely solved exactly. For exceptions in offline planning approaches, we refer to Section 5.5. Some approaches determine optimal solutions for shorter planning horizons, see also Section 11.1.6.

We group planning principles for informative tours into three major categories, which are illustrated in Figure 11.3:

Geometric approaches are essentially segmentation-guided and seek to divide the target space between the available vehicles based on the known current positions and the size and shape of the considered area. They are not used to plan UAV routes but instead aim at achieving a good distribution of sensors. In its most basic form, this is done without considering information about the surveyed process. The partitioning of the area can be static, based on some a-priori partitioning approach. Online versions adjust the partitions while the mission is in progress, i.e., while the UAVs move and collect new information.

Model-based variants use a formalized model representation to evaluate potential targets and to measure a plan's expected quality. If the model and thus the evaluation of targets changes based on previous observations, these variants are typically called *adaptive* in literature. To reduce the number of candidates that need to be evaluated, many model-based approaches use a graph representation of the target area, where nodes denote candidate sampling locations and arcs indicate possible moves of the vehicles. In contrast, many of the other planning variants discussed here can be applied continuously.

Rule-based variants make use of a predefined rule set without explicitly modeling an objective function. Non-adaptive rule-based versions determine, for example, a lawnmower-like pattern or lead to a circular motion of the vehicles (Neumann,

2013). *Adaptive* rule-based versions make decisions based on the current information about the process and a limited set of rules that describe allowed motion patterns.

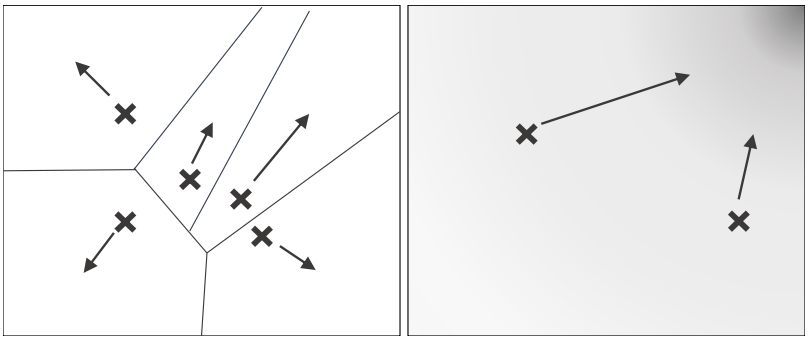

(A) Geometric (Voronoi partitionings)

(B) Model-based: Movement toward high utilities

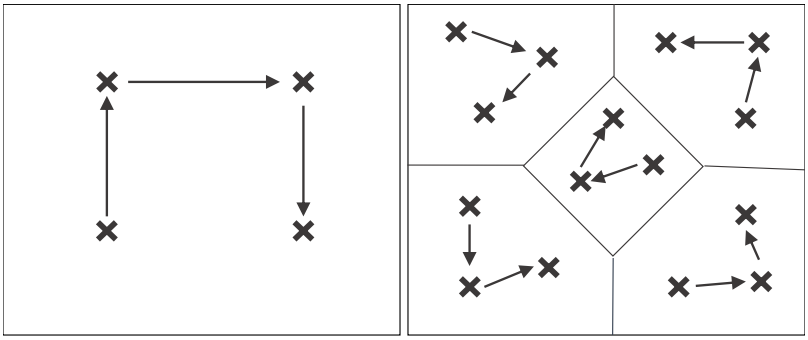

(C) Rule-based: Set of predefined moves

(D) Hybrid: Area partitioning and rule-based traversal

FIGURE 11.3: Representative planning approaches in online surveillance.

A strict separation between these concepts is not always possible, and some publications propose hybrid variants that combine two of these ideas. Examples of hybrid approaches are two-staged approaches where an area is first separated into partitions to reduce the complexity for a subsequent model-based planning step (Stranders et al., 2013). Some authors dynamically switch between concepts, e.g., to balance exploration with a more detailed inspection at some locations (Saldaña et al., 2015).

GEOMETRIC ALGORITHMS

Spatial decomposition techniques have been used to determine an a-priori partition of a target area (e.g., Singh et al., 2009; Xu et al., 2013). Each UAV is assigned to one of these regions. This can lead to inefficient outcomes e.g., in case of vehicle failure (Lambrou and Panayiotou, 2013), unless there are strategies that either adapt the partitions or dynamically switch to other heuristics (see, e.g., Schwager et al. (2017) for a hybridized approach based on different geometry-based heuristics).

Many solution approaches seek to obtain a centroidal Voronoi tessellation of the target area. In this context, a Voronoi tessellation is a partitioning of a plane based on a set of predefined generating points. Each point in the plane is assigned to the generator to which it is closest. The resulting regions are called Voronoi cells. Such a partitioning is a *centroidal* tesselation if all generators also correspond to the center of mass (centroid) of the corresponding Voronoi cell.

Many solution variants apply Lloyd's algorithm, an iterative way to achieve an approximation of a centroidal Voronoi tessellation, rather than relying on a-priori partitions. This algorithm is initialized with the starting position of the available vehicles. Then, the target area is partitioned into regions such that each point in space is assigned to the vehicle closest to it. Each vehicle moves to the current centroid of the corresponding Voronoi cell. Afterward, the partition is updated and the vehicles are again replaced to their next centroid. This is repeated until the computed partitions and vehicle positions remain sufficiently stable. The objective is to achieve and maintain a homogeneous distribution of sensors and sensing locations across the target space. In approaches targeted towards area coverage or persistent monitoring, such as the algorithm proposed by Schwager et al. (2017), the Voronoi decomposition determines the final position, i.e., the vehicles do not move further after a stable distribution has been achieved.

Voronoi tesselations have also been used to achieve other desirable properties in a solution procedure. They are a popular means for decentral coordination because each UAV only requires knowledge about the current positions of its neighbors. This means that the communication effort is low, see, e.g., (Cortes et al., 2004; Kemna et al., 2017; Todescato et al., 2017). Furthermore, as noted by Pavone et al. (2011), an adaptive partitioning scheme has the additional benefit of maintaining a balanced workload between vehicles. Other authors use Voronoi tesselations for collision avoidance, e.g., by specifying minimum distances to the boundaries for the Voronoi cells that limit the vehicle movements (Nguyen et al., 2015). Finally, it is possible to combine these approaches with environment models – current predictions or uncertainties as well as priorities – are weighted Voronoi diagrams, as used for example by Kemna et al. (2017), or power diagrams (Pavone et al., 2011). These approaches enable an adaptive decomposition of the target space based on previous observations.

Model-based heuristics

At the heart of model-based approaches are representations of the surveyed field and phenomenon. Usually, these are based on the environment models introduced in the previous section. Many model-based approaches fall into one of two broad categories: Gradient-based algorithms and utility maximization.

Gradient-based approaches use an estimated concentration gradient to either move along a contour line or towards areas of higher concentrations. In source localization applications, heuristics in this line of research are also referred to as "chemotaxis" algorithms. A more detailed overview of different implementations and variants can be found in Neumann (2013) and Hutchinson et al. (2017). For contour line tracking, sev-

eral specific algorithms have been proposed in literature. An overview can be found in the study by Hutchinson et al. (2017). This includes, for example, strategies that switch direction once a threshold is reached (i.e., the contaminant field is either left or entered too far) rather than directly computing and following the gradient. Vehicle movements in an APF similarly follow a negative potential gradient to move to a potential minimum.

Utility maximization or cost function minimization approaches decide based on some model that represents the benefit of sensing locations or vehicle movements. These variants are often discretized to allow the evaluation and comparison of a finite number of candidate locations (e.g., Nguyen et al., 2015). An example of a non-probabilistic model is given by Lambrou and Panayiotou (2013), who use a cost function that combines the distance to other sensors and obstacles, the boundary of the region, and the distance to the sensing candidate itself.

Given a probabilistic model of the phenomenon, several approaches use information criteria such as mutual information or entropy (see Section 3.2.4). Popović et al. (2017) combine two probabilistic utility functions: entropy reduction, which favors exploration, and an application-specific measure for the number of areas without a sufficient amount of information. They gradually move from the first to the second to encourage intensification as time progresses. Similar to chemotaxis, the literature on source localization often refers to utility-maximizing variants as "infotaxis" strategies (Hutchinson et al., 2017; Ristic et al., 2017).

Rule-based approaches

Similar to model-based and geometric approaches, rule-based approaches (called "expert system" by Kuroki et al. (2010)) exist in adaptive and non-adaptive versions. Adaptive rule-based variants are closely related to model-based variants. The main difference is that in rule-based algorithms, the current assumptions about the environment determine the vehicle motion, but are not used to compare options or to evaluate the impact of the next decisions. Exemplary strategies in this line of work use predefined moves based on the current location of the sensor and the current estimated location of a gas spill (Kuroki et al., 2010). In this example, the UAV moves downwind or traverses the area diagonally. Flight direction changes when the measured concentration is below a threshold value, i.e., when the UAV as likely left the plume. Non-adaptive variants essentially describe deterministic flight patterns, e.g., spiral or lawnmower patterns. We refer to the surveys by Otto et al. (2018) and Cabreira et al. (2019) for more details on flight pattern planning for area coverage. These variants can also be used for detecting a plume or spill before proceeding with more finely grained model-based algorithms (Neumann, 2013). Finally, combinations of these variants are also addressed. To give an example, Brink and Pebesma (2014) use predefined non-adaptive strategies until detecting a smoke plume and then proceed by using a limited set of movement adaptions for tracking the detected object.

11.1.6 Decision scopes in model-based planning

The model-based planning approaches introduced in the previous section can be furthermore characterized by the decision scope that is considered. The decision scope comprises two aspects:

- the length of the planning horizon, and
- the way decisions are made within this horizon.

In other words, the decision scope characterizes the options that are available each time a UAV mission is updated.

The planning horizon describes how approaches account for the effects that a particular decision or sequence of decisions can have on an (expected) future state of the system, typically within a given moving window relative to the point in time where a decision is made. The objective is to make decisions such that the outcome at the end of this horizon is improved with respect to some evaluation criterion. The literature on environmental surveillance distinguishes between greedy approaches, variants with a finite planning horizon, and non-myopic or infinite horizon approaches, depending on the length of the considered horizon (Stranders et al., 2013).

Considering the decision scope, we distinguish between the following three concepts for making decisions within these planning horizons:

Greedy algorithms disregard the future state of the system and evaluate each decision only based on its immediate impact. Gradient-based approaches as discussed in the previous section can generally be considered as greedy.

Selection-oriented approaches separate the model-informed selection of sampling locations from the routing decision. They first select sampling locations without consideration of their spatial order. Often, this selection is done in a heuristic fashion to reduce the computational effort in each step (e.g., Ma et al., 2017). In the next step, a visit sequence is determined such that the overall resource consumption for visiting the selected targets is minimal. A heuristic implementing this concept is, e.g., presented by Popovic et al. (2019).

Integrated variants determine the best possible sequence of sensing locations while explicitly considering the length of the planning horizon (Low et al., 2008). Unlike the previous version, these approaches ensure that the selected target sequence is feasible, i.e., they yield an optimal solution to the resource-constrained information gathering problem with respect to the current environment model. Examples for this concept are the algorithms for solving the informative path planning (IPP) in Section 5.5, which evaluate plans based on the estimated overall information gain at the end of the mission.

Similar to model predictive control (Section 11.2.3), the decisions are often implemented in the form of a receding horizon control strategy: While the approaches determine a sequence of decisions in each step, only the first few actions are implemented in

each step. Then, the vehicle path is updated again based on the new data (e.g., Hitz et al., 2017).

11.2 Relation to other optimization concepts

In Section 11.1, we have summarized approaches in the field of environmental sensing and surveillance that are similar to AMPPES. In this section, we discuss how these approaches relate to concepts for online planning and learning that have been proposed in operations research and machine learning. All of these approaches can be roughly summarized in Figure 11.4. They have in common that decisions are made (a) several times in a dynamic setting, (b) are (partially) informed by beliefs about an expected future state of the system, and (c) that the decisions impact this future state and thereby influence possible future actions. Often, actions and predictions are only valid within a certain planning horizon relative to the point in time where a decision is made.

11.2.1 Dynamic vehicle routing problems

The AMPPES is similar to a dynamic VRP in that information is revealed over time. The main difference is the source of dynamism: In dynamic VRPs, the most common dynamic elements are the requests to be served, travel times, or vehicle availability (Bektaş et al., 2014). Dynamic requests are by far the most common problem variant in this area (Pillac et al., 2013; Psaraftis et al., 2016). In contrast, the main source of dynamism in our case is the acquisition of new samples that change the current benefit associated with the target locations.

The literature on dynamic VRP distinguishes between deterministic and stochastic problems (Pillac et al., 2013). In deterministic problems, the input is dynamically revealed, and no forecasts or probabilities relating to future events are used. This is also referred to as *online* VRP (ibid). In stochastic problem variants, future conditions are described by a known probability distribution that can be exploited. An example is the generation of "dummy" customers, that are included in tours and can be replaced by actual realizations (e.g., Ichoua et al., 2006), or the repositioning of vehicles to cover expected future requests (e.g., Gendreau et al., 2006).

Stochastic orienteering problems, where the profit of nodes is unknown but follows a known distribution, are close to the AMPPES. In contrast to the variants of the VRP where the demand or profit is revealed upon arrival at a location, the profit in our use case depends on the other visits performed by the UAV fleet during the mission. Furthermore, in the case of the AMPPES, no accurate information on the known distribution is available before planning. This prevents the application of solution concepts for stochastic problems.

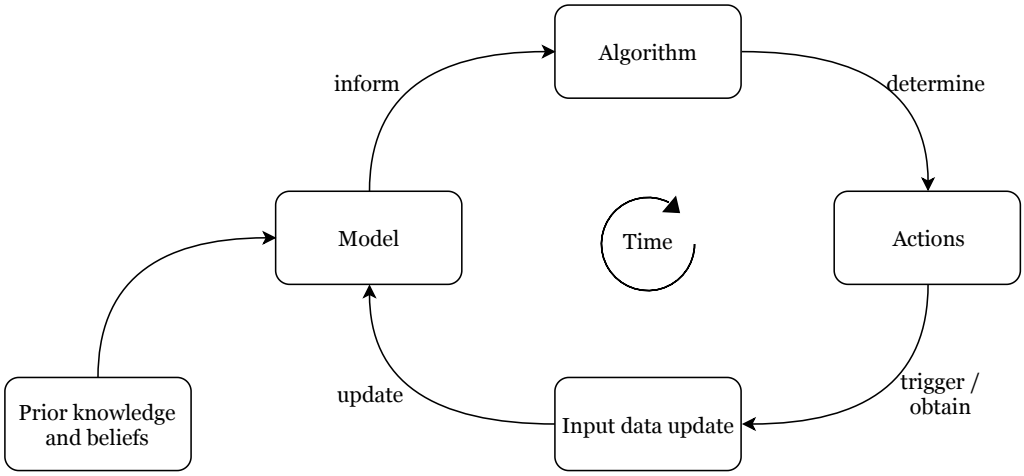

FIGURE 11.4: Generalized update loop in dynamic optimization and planning approaches.

11.2.2 Online optimization with lookahead

The term "lookahead" has been used in recent publications on environmental surveillance in the context of robotics as well as in literature on online optimization. A classification and overview in the field of combinatorial optimization are given by Dunke and Nickel (2016). This study has been extended to the notion of gradual lookahead by Dunke and Nickel (2019). Examples for publication using this term in environmental surveillance are the studies by Popović et al. (2017) and Lim et al. (2016).

The interpretation of what constitutes a lookahead is different in the two areas. Table 11.1 gives a summary of the main differences in these two lines of research, based on the distinction between an informational and a processual component proposed by Dunke and Nickel (2016). In online optimization, a lookahead is defined as a preview on future events (e.g., incoming requests). In the case of gradual lookahead, this preview can be uncertain. Typically, the uncertainty associated with the preview is reduced as the release time comes closer (Dunke and Nickel, 2019). This preview can be used to make better decisions. In the context of environmental surveillance, "lookahead" is often used interchangeably to the term "planning horizon". In these applications, no preview exists in the sense that future inputs are known. Instead, solution approaches update the current belief about the process and use this belief to evaluate potential choices within a specified planning horizon to get a better objective value or minimize resource usage.

From this perspective, the literature on environmental sensing focuses on the processual point of view: The approaches determine possible action sequences over some planning horizon, which are evaluated with respect to an expected final system state. However, the evaluation is not informed by a preview of future events and only depends on past inputs and the considered actions. This is supported by the literature itself, which typically refers to algorithms with lookahead as an alternative to greedy approaches (Stranders, 2010). At the beginning of this chapter, we have therefore referred to these approaches using the term moving planning horizon rather than lookahead.

	Online optimization	Environmental surveillance
Interpretation of lookahead	Preview on future events	Evaluation of the impact of future decisions
Informational component	Future events and associated uncertainty set	Current belief about the surveyed process / value of candidate visit locations
Processual component	Various processing rules for the available data	Moving planning horizons

Table 11.1: Interpretation of "lookahead".

Similarities between the two concepts lie mostly in the implementation of the processing rules. Specifically, all approaches have to consider the following aspects:

- the frequency with which new input data is processed,
- the definition of permissible actions for each update, which are evaluated and compared based on the available information.
- the exploitation of information relating to the uncertainty of the available information.

11.2.3 Model predictive control

Model predictive control (MPC) refers to feedback control mechanisms for dynamic systems where the impact of decisions or control inputs on a future state can be assessed using a model representation of this system. Its origins lie in process control in industrial applications, e.g., chemical plants. Its final goal is often to achieve and maintain a stable system state.

Usually, MPC works in three steps (Seborg et al., 2010): (1) optimal control inputs are determined for finite prediction and control horizons, (2) the first action is implemented, thereby changing the underlying physical system, and (3) when a new input is received or a given time has passed, the horizon is moved forward in time and new control input is determined. Consequently, MPC is also referred to as receding horizon control. This makes it possible to deal with uncertain influences and unexpected system changes.

MPC can be considered to be more closely related to AMPPES than stochastic VRP: Stochastic VRP use information provided externally to make decisions. In the case of AMPPES and MPC, the available information is updated based on the system's own decisions, which result in new measurements or control inputs. This similarity has been noted by some authors, e.g., Indelman et al. (2015). In the context of controlling autonomous vehicles, MPC is used to ensure that vehicles follow the waypoints determined by a higher-level mission planning approach as closely as possible (Strobel, 2016).

11.2.4 Active learning

Active learning, as a discipline in machine learning, refers to a class of supervised learning algorithms that actively requests new (labeled) input data. This can, for example, reduce high labeling and training effort by iteratively selecting only the most "informative" input. These approaches are often used in the context of classification problems. To this end, active learning systems combine two components: one component for training a model on the current training set, the other for selecting unlabeled samples to be added to the training set.

In the context of this study, the selection component is of particular interest. A wide range of measures has been proposed in the literature on active learning to evaluate and compare candidate input samples. Following the surveys by Settles (2009) and Wang and Hua (2011), we can distinguish between several recurring strategies, which are summarized in the following:

Error reduction: Strategies targeted towards error reduction select the sample that minimizes the expected estimation error or prediction variance (see Section 3.2.4). These are typically computationally expensive, as the expected error reduction has to be computed for all candidate samples.

Uncertainty: These variants select the most uncertain samples. In the case of a classification problem, those are the samples that are closest to the classification boundary. Other variants select samples that are most informative, e.g., using measures such as entropy.

Diversity: Given a measure for the distance between two samples, these strategies seek to maximize the diversity of the training set, i.e., the dissimilarity between pairs of selected samples. This is for example based on kernel functions modeling the "distance" between two samples.

Density: Similar to the previous criterion, density-based strategies use a distance measure to identify "representative" samples, i.e., samples with similar characteristics as others, such as their respective location. Examples are clustering techniques, where only cluster centers are selected.

Committee-based: To improve the predictions, some authors use several learning models that are trained simultaneously on the available input data. Committee-based selection strategies make use of these models, e.g., by selecting the samples where the learning components disagree most.

From this point of view, AMPPES can be considered as an active learning problem where the planning heuristic selects sampling locations for training, e.g., a Gaussian process model of the environment. As we have already seen, some methods for deciding on sampling locations are based on similar concepts, for example, the selection of the next sample as the one where the prediction is least certain. The main difference is

that in AMPPES, the resources available for sampling are limited, and more informative samples, i.e., those at larger distances to the already surveyed area, can come at a higher cost. This means that not all subsets of samples can be feasibly included in the tours and that each decision impacts the choice of additional samples in the future.

11.3 Summary

In this chapter, we reviewed the state of the art relating to the AMPPES. In a first step, we identified common concepts that are used in the context of environmental sensing. To this end, we introduced a classification scheme for dynamic approaches and summarized a wide range of approaches in this line of work (Section 11.1). We have shown that, even though the literature applies a variety of environment models and solution approaches, many of these approaches are derived from a few core concepts. In a second step, we summarized solution techniques proposed in different domains, notably dynamic solution variants for the VRP, online optimization with lookahead, model predictive control and active learning (Section 11.2).

In the next chapter, we discuss the applicability of these approaches for the AMPPES, and identify the main research gaps that are addressed in the subsequent chapters. The concepts identified in this overview then provide the foundation of the models and heuristic approaches proposed in Chapter 13.

12 Research gap and contributions

As shown in the previous chapter, numerous concepts for problems related to the AMPPES have been proposed in literature. In this chapter, we review these concepts and identify promising online and adaptive approaches, whose suitability and limitations with respect to the AMPPES we examine in detail in the remainder of this work. We then identify open questions regarding the design of mission planning strategies and summarize our contributions in this direction.

12.1 Discussion of solution approaches

In the previous chapter, we have identified six dimensions characterizing online approaches, which we have summarized in Figure 11.1. Among these dimensions, our main interest lies in the solution approaches for planning UAV missions, which are characterized by (1) the applied environment model, (2) the planning heuristic, and (3) the decision scope.

Considering the environment representations, two concepts are promising for our application: sensor coverage and process interpolation methods. Sensor coverage models are inherently inadaptive in the sense that they do not adjust based on previous measurements. They are computationally inexpensive and are likely to scale well to scenarios where the affected areas are large and spatial coverage is most important. In contrast, probabilistic process models offer the possibility to adapt based on the obtained samples. They provide detailed information about predicted values and remaining uncertainty in the process. However, updating these models to new information itself can be expensive, as is the evaluation of candidate sampling locations.

The remaining modeling variants discussed in the previous chapter are often specific to other applications or, in the case of APF, focus on aspects such as collision avoidance that are not relevant to our use case. Geometric representations of the phenomenon are promising when the primary goal lies in estimating the extent of a phenomenon. For AMPPES, where we are interested in the concentration of contaminants in "hotspots" as well as in less affected areas, these do not provide a sufficient amount of information.

Among the different planning heuristics, geometric and model-based approaches seem promising for our use case. Geometric approaches serve as a substitute for the "true" objective of achieving low prediction errors as they seek to improve the spatial distribution of sensors. As dissimilarities between two locations increase with distance, an even distribution of UAVs is generally advantageous for the prediction quality. These approaches furthermore scale well to large UAV fleets and are stable with respect to vehicle failure.

Model-based approaches explicitly use the information provided by the environment representation to select the next sampling locations. Hence, their performance strongly depends on the information encoded in the underlying model. The gradient-based approaches in this group largely rely on local information, i.e., in our use case, the gradient of the concentration. Often, these approaches assume that the observed field has appropriate features, such as a single source of contamination. If this assumption holds, they can perform well for the stated mission objective.

In the AMPPES, where the informativeness associated with candidate locations depends on factors such as the previously sampled locations, priorities in the target area, and potentially the time since the last observation, utility-based measures offer the possibility to evaluate sets of sampling locations within the entire affected area. In other words, they offer a "global" view on the relative benefit of sampling locations. This can be used to quickly move toward the most relevant locations, which is promising for our use case compared to gradient-based strategies. The computational complexity depends on the applied environment model and tends to be high, especially when using longer planning horizons and more expensive probabilistic interpolation models.

12.2 Research gap

In our use case, we face two particular challenges compared to other applications addressed in literature. First, the available computation time for making decisions is limited. UAVs require the next sampling location within a few seconds after a measurement in order not to delay the mission. Second, there is little time available for executing the mission: The technological limitations restrict the total flight time of the UAVs, and the response units require information within a very short time after arriving at the scene of an emergency.

Many approaches proposed in literature are tailored toward specific applications, and the relative performance of different designs is rarely considered. Our goal is to identify solution strategies that are promising for solving the AMPPES, keeping the limitations stated above in mind. We focus on three major issues that are relevant for designing planning concepts for our use case:

Assessment of the benefit of adaptive models The literature leaves open whether and to which extent adaptive planning improves plans compared to less expensive models. Most authors assume that adaptive planning is better, independently of its cost.

Advantages of using adaptive process models over simple coverage models have been shown for sensor placement problems (e.g., Krause et al., 2008). Other authors highlight the benefits over complete coverage approaches (e.g., Popović et al., 2017). To the best of our knowledge, it has not yet been empirically assessed whether adaptive planning provides benefits over less complex coverage-based approaches in the case of moving sensors.

Furthermore, adjusting to the obtained information is expensive. For example, the trained process models have to be updated and evaluated, which requires a higher computational effort than more simple coverage-based models. In our application, this computation effort might become prohibitively high if plans have to be updated quickly. Consequently, even if adaptive environment models yield better results than other approaches, it is not certain that using these models would be feasible in our context.

Comparison of geometric and model-based approaches As our literature review has shown, geometric approaches and model-based planning strategies constitute the two recurrently used strategies for environmental surveillance. The more efficient geometric approaches are often used in applications where a large number of vehicles is available or when the focus lies on decentralized planning approaches. However, these methods are targeted toward obtaining a rough distribution of sensors over a large area and do not address the problem of planning more finely tuned missions for individual UAVs within this area. Model-based approaches, in contrast, are advantageous in this respect, but are more expensive and depend on the applied models.

It remains unclear which approaches perform better under which circumstances–e.g., number and flight time of the deployed UAVs or size of the affected area. In case several planning concepts are hybridized, it remains an open question what each component contributes toward the quality of a solution. A better understanding of these effects would make it possible to identify promising combinations of planning strategies.

Design of effective decision scopes Larger decision scopes are often applied to account for the impact that decisions can have on future states of the system. For the AMPPES, in which the benefit of observations can usually only be assessed jointly, increasing the decision scope means that this interdependence between sampling locations can be better taken into account.

There are two aspects in which larger decision scopes promise improvements: The first one is the improved selection of samples that, together, provide more information than sampling locations selected sequentially. The second factor is the increased efficiency of the routes as detours can be avoided and "zig-zagged" routes are avoided. However, it is not certain whether these potential effects can be realized. First, all plans depend on the current information that is encoded in the environment model. As these beliefs about the surveyed phenomenon are not necessarily accurate, decisions based on these assumptions may not yield improvements. Second, choices that are optimal for a shorter planning horizon are not guaranteed to also improve the final outcome of a

mission.

As a consequence, it remains an open question what constitutes an effective decision scope for the AMPPES, especially considering the additional computational effort that may be required. Specifically, it is not certain to which degree the joint selection of sampling locations is worthwhile, or whether similar advantages can be achieved through the improved sequences for visiting sequentially selected sampling locations. Likewise, it has not yet been examined what length the considered planning horizon should have to provide an appropriate balance of computing effort and improved decisions.

12.3 Contributions

Our contributions toward closing the research gap are summarized below. We focus on three aspects: The derivation of models and solution approaches for solving the AMPPES, the combination of these approaches in a unified modeling and solution framework, and the comparison and evaluation of designs as a basis for future research in this and related fields.

Models for spatial processes We study modeling paradigms for environmental processes that range from simple covering variants to probabilistic process models. To this end, in Chapter 13, we develop a unified model that comprises these different aspects. We adapt existing modeling variants to our use case and derive new concepts based on our models introduced in Part II of this thesis. Based on these modeling variants, we empirically investigate the benefits of using complex environment models and the performance of less expensive strategies in an extensive computational study in Chapter 15. We demonstrate the high performance of coverage-based models as well as their limitations and show in which situations models offering a high degree of adaptivity yield better results than other approaches. We furthermore demonstrate that the GCorTOP modeling approach proposed in Part II of this thesis achieves a reliable trade-off between an accurate representation of spatial interdependencies and computational complexity for a wide range of scenarios.

Plannung heuristics and decision scope design We combine the models studied in this work with solution strategies for different decision scopes to provide insights into what constitutes an effective heuristic design. We focus on assessing the potentials of three interlinked aspects: (1) the sequencing of targets in longer planning horizons, (2) the modified selection of samples when considering resource constraints within these planning horizons, and (3) the usage of the obtained information in adaptive planning and the degree of adaptivity that should be supported when using longer planning horizons. Similar to the investigation of the environment models, we empirically identify suitable solution procedures and weight up the computational costs against the quality of solutions using a large set of scenarios for the AMPPES.

Architecture and simulation framework An unbiased comparison of different solution strategies can only be obtained in a controlled environment based on a large number of scenarios. Therefore, we integrate all of the models and solution techniques summarized in this chapter in a common solution architecture. This framework is introduced in Chapter 14. Using this architecture, we can freely combine different models, planning heuristics, and decision scope variants. We furthermore implement a discrete-event simulation framework that allows us to assess the performance of these solution concepts in realistic scenarios.

13 Models and solution strategies for adaptive mission planning

In the previous chapters, we have summarized promising concepts for solving the AMPPES. However, we have also shown that there are several open questions concerning the design of strategies. Our goal is to answer these questions and to identify strategies that perform well, especially considering the computational restrictions in our use case.

For this purpose, this chapter presents different solution principles. Our focus lies in model-based planning approaches. First, we formally describe the underlying planning problem in Section 13.1 and give an overview of the concepts studied in this work in Section 13.2. We introduce three heuristics addressing different decision scope variants for model-based planning in Section 13.3. To compare these approaches with less expensive planning strategies, we also adapt a geometric approach to our use case. In Section 13.4, we then introduce a unified modeling approach and discuss several environment representations that can be combined with these planning heuristics, These models are partially based on existing variants, comprising the approaches derived in Part II of this thesis together with new concepts. This enables us to gain a better understanding of the trade-offs between computational efficiency, adaptivity, and performance offered by the various strategies.

13.1 AMPPES problem definition

Following the notation introduced in Chapter 7, we represent the target area as a discrete set of locations $\mathcal{V} \in \mathbb{R}^2$. These are the locations we wish to receive information about, either by surveying them directly or by interpolating the distribution of contaminants based on other samples within the target area. Again, each location $i \in \mathcal{V}$ is associated with a priority u_i that represents the relative importance of obtaining accurate information at i. The position of a location i is denoted pos_i with coordinates (p_i^x, p_i^y).

The set of available UAVs is denoted $\mathcal{M}$. Each vehicle is associated with takeoff and

landing locations s_m^{start} and s_m^{end}, respectively. The maximum flight time of a vehicle is denoted T^{max}. For simplicity, we assume that each sensing operation requires a fixed sensing time τ_i at standstill. Travel distance between all locations $i, j \in \mathbb{R}^2$ is denoted as d_{ij}. The corresponding travel time including the sensing time at the destination as well as the time necessary for acceleration and deceleration is denoted as τ_{ij}. The planning horizon is a finite interval of time $[0, \mathcal{T}]$, with 0 indicating the start of the mission.

The overall mission planning problem consists of selecting sampling locations $\mathcal{S} \subseteq \mathcal{V}$ that can be visited by the available UAVs within their maximum flight time and that maximize the information obtained about $\mathcal{V}$. In case of the AMPPES, decisions are made in an online fashion at time instants $t \in [0, \mathcal{T}]$. We use a receding horizon concept where the optimization is performed for a moving planning horizon with length $\Delta\mathcal{T} \leq \mathcal{T}$, starting at time t. New sampling locations are computed when a vehicle has no target locations left but does not yet need to return to its landing position. Additionally, new targets can be computed after a fixed decision horizon $\Delta\mathcal{T}^{fix} \leq \Delta\mathcal{T}$ has passed since the mission has last been updated.

To formalize this problem, we require additional notations indicating the state of the system while the mission is in progress: We denote the position of a UAV m at time t as pos_{mt}. All locations sampled by the vehicles in $\mathcal{M}$ from the beginning of the mission until time t are comprised in the set $\mathcal{S}_{<t} \subseteq \mathcal{V}$. For the purpose of adaptive planning, we denote the corresponding measurements as $\mathcal{Y}_{<t}$.

The AMPPES at time $t \leq \mathcal{T}$ is defined as the problem of selecting the next sampling locations $\mathcal{S}_t$ such that the estimated information gain $\mathcal{I}(\mathcal{S}_t, \mathcal{S}_{<t})$, which depends on both the already surveyed locations in $\mathcal{S}_{<t}$ and the set of candidate samples $\mathcal{S}_t$, is maximal. A feasible solution to the AMPPES consists of a set of selected targets $\mathcal{S}_t$ that can be surveyed within the moving planning horizon $\Delta\mathcal{T}$ starting at the current positions of the UAVs $pos_{mt}, m \in \mathcal{M}$, at time $t \in \mathcal{T}$ such that all UAVs can reach their destination locations s_m^{end} within their maximum flight time T^{max}.

13.2 Overview of solution approaches

In this chapter, we focus on the three aspects characterizing a solution approach summarized in our literature review, see Figure 11.1: the environment model, the planning heuristic, and the decision scope. We first introduce planning strategies for deciding on the next sampling locations during the mission, i.e., for determining the set $\mathcal{S}_t$ to be surveyed. Second, we propose a range of environment models that offer an indication of the expected information gain of a set of locations. That is, they provide an informativeness measure $\mathcal{I}(\mathcal{S}_t, \mathcal{S}_{<t})$.

In Section 13.3, we introduce four planning heuristics addressing different decision scopes:

Greedy We use a greedy heuristic for maximizing $\mathcal{I}$ as an initial approach for model-based planning. In each step, this selects one sample per UAV to visit next, accounting for all previous sample locations $\mathcal{S}_{<t}$. This strategy is less expensive

than the other two model-based planning approaches and has found frequent use in literature. Furthermore, previous work on online optimization has shown that increasing the decision scope and using more sophisticated strategies is not necessarily better (Dunke and Nickel, 2016). Hence, it provides a useful benchmark for more elaborate approaches.

Sequential selection The sequential selection strategy selects target locations for each UAV in a heuristic fashion. To this end, it selects samples one after the other until the moving horizon $\Delta\mathcal{T}$ is exceeded. In each step, the next sampling location is selected greedily as the one that provides the largest incremental benefit with respect to all previously sampled locations in $\mathcal{S}_{<t}$ as well as those already chosen for visitation at time t, i.e., those that are already included in $\mathcal{S}_t$. For each UAV, a TSP is solved to determine the order of survey locations. This means that the selection of targets is done in a heuristic fashion, while the routing is done optimally. This allows us to assess whether larger decision scopes improve over the greedy strategy due to a more efficient order in which samples are taken.

Integrated planning The integrated planning approach considers sample selection and routing together, determining the optimal combination of sampling locations within the considered moving planning horizon $\Delta\mathcal{T}$. That is, we solve the static MPPES in each iteration. In contrast to both greedy and sequential selection, this approach gives an indication of how much decisions can be improved by a more sophisticated method of choosing sampling locations.

Geometric approach The geometric approach adapts LLoyd's algorithm for moving the UAVs to positions that are distributed as good as possible across the target region. In its basic version, this is independent of the environment representation $\mathcal{I}$. This allows us to measure the benefit of using one of the model-based planning approaches discussed above against a simple but robust strategy that does not consider previous samples or predictions.

We combine these planning strategies with different environment representations. As demonstrated in Chapter 11, previous publications have proposed a wide range of possible models, many of which are highly application-specific. In this work, we propose strategies that represent the most common approaches in literature. As our main interest lies in the assessment of the necessary level of detail and adaptivity, we focus on three broad categories of representations, each one including several concepts for representing the environment and estimating the informativeness $\mathcal{I}$ of samples:

Nonadaptive models We apply nonadaptive models that use information about the targets' priorities together with a nonprobabilistic model of the spatial correlations that is fixed throughout the mission to measure $\mathcal{I}$. This means that they account only for the distances between sampled and unsampled locations. We propose different concepts, with a simple disc-based approach and a more detailed representation based on the GCorTOP model introduced in Part II.

Extensions to time-dependent models We extend the existing nonadaptive models such that they account for the time that has passed between two samples close to one another. These extensions are new and have not yet been discussed in literature. In the context of AMPPES, they allow us to account for the available time during the mission: Longer flight times mean that more samples can be taken, which is considered when evaluating a target's value. Furthermore, these models provide a means to survey time-dependent processes.

Adaptive process models To assess the benefit of adaptive planning, we propose several measures for $\mathcal{I}$ that rely on Gaussian process models for the surveyed phenomenon. Again, we adapt different strategies from literature that vary in the level of detail with which they estimate the information gained by a set of sampling candidates. This facilitates an assessment of the trade-offs between model complexity, information gain, and efficiency.

13.3 Planning heuristics

In this section, we discuss planning strategies for selecting $\mathcal{S}_t$ that we have introduced above. In Section 13.3.1, we derive model-based planning concepts with different decision scopes. This way, we can vary the degree of adaptability, the replanning frequency, and the horizon length to obtain a better understanding of these approaches. In Section 13.3.2, we adapt Lloyd's algorithm, which is the most popular geometric approach, to the AMPPES. This allows assessing the performance of purely geometrically motivated approaches in our problem setting.

The different concepts are illustrated in Appendix B, Figures B.10 to B.12.

13.3.1 Model-based approaches

The model-based approaches discussed in this section solve the AMPPES for arbitrary informativeness measures $\mathcal{I}$. However, this is a difficult problem: Identifying a profit-maximizing set of samples under consideration of tour length constraints is NP-hard. The same holds for the optimization of submodular objective functions, e.g., the mutual information criterion, even without considering sequence-dependent budget restrictions (Das and Kempe, 2008).

To deal with this challenge, model-based approaches for the AMPPES restrict the decision scope, e.g., by limiting the moving planning horizon $\Delta\mathcal{T}$ that is considered at the time of the planning, or by considering only a subset of possible combinations of sampling locations at each point t in the planning horizon. For evaluating these concepts, we implement and compare three variants: a greedy informativeness maximization, a heuristic selection, and an optimal solution strategy for the time-restricted subproblems. These variants are defined in the next sections.

Greedy informativeness maximization

The least expensive model-based approach is greedy utility maximization, i.e., the selection of a single target location s_t that is associated with the highest estimated informativeness, based on the information that is available in the environment model. This can easily lead to "zig-zagged" movements, e.g., because the prediction variance is highest at distances that are very far from the sampled locations. To prevent this, we maintain a distance limit d^{limit} between subsequent sampling locations, i.e., we solve the following optimization problem each time a new sampling point is required for UAV:

$$\text{(AMPPES-greedy)} \quad \max_{s_t \in \mathcal{V}_m} \mathcal{I}(\{s_t\}, \mathcal{S}_{<t}) \tag{13.1}$$

with $\mathcal{V}_m = \{s \in \mathcal{V} : d^{limit} \geq \| \mathit{pos}_s - \mathit{pos}_{mt} \|\}$. AMPPES-greedy is solved for each UAV $m \in \mathcal{M}$ sequentially. When the distance limit is narrow, this closely mirrors a gradient ascent approach, resulting in a local maximum.

The environment model is updated in between the different vehicles to account for the selected sample. This means that s_t is provisionally added to $\mathcal{S}_{<t}$ for the purpose of evaluating $\mathcal{I}$. This way, if several UAVs start at the same position, the second one will move to the second-best target.

The main advantage of the greedy planning approach over other model-based strategies is that it is computationally inexpensive. Its computational effort is quickly dominated by the cost for evaluating $\mathcal{I}$ in case of more sophisticated environment models, as is shown in what follows. However, this approach provides no means to anticipate the impact of a decision. For example, there might exist beneficial targets that can be included at low cost in between the current and the next sampling location of a vehicle. These are left out by the greedy approach, requiring a larger detour at a later point in time.

Sequential selection heuristic

The idea of a larger decision scope is to select several targets in each iteration. In the sequential selection heuristic, this selection is done by iteratively selecting the target with the highest marginal benefit and including it in the planned mission until the moving planning horizon is exceeded. The approach is depicted in Algorithm 6. Similar to the greedy approach, this problem is solved for each UAV separately.

Throughout the algorithm, we maintain a set of candidates that can be included in $\mathcal{S}_t$ for the currently considered UAV $m \in \mathcal{M}$, denoted $\mathcal{V}_m$. This set is initialized as the set of target locations that can be reached from pos_{mt} within the moving planning horizon $\Delta\mathcal{T}$ such that a return to the end location s_m^{end} is feasible for UAV m within its maximum flight duration T^{max}. In each iteration of the algorithm, one location s_t is greedily selected from $\mathcal{V}_m$ and added to the set of selected sampling locations $\mathcal{S}_t$ such that $\mathcal{I}(\mathcal{S}_t \cup \{s_t\}, \mathcal{S}_{<t} \cup \mathcal{S}_t)$ is maximal. This means that s_t provides the largest incremental benefit assuming that the locations in $\mathcal{S}_{<t}$ as well as all already selected sampling

Algorithm 6: Sequential selection strategy

Input: Environment model and informativeness measure $\mathcal{I}$, time instant t, UAV position pos_{mt} and destination position s_m^{end}, previously selected samples $\mathcal{S}_{<t}$, moving horizon length $\Delta\mathcal{T}$, candidate locations $\mathcal{V}$, flight time limit T^{max}

Output: Sampling points $\mathcal{S}_t \in \mathcal{V}$

initialize $\mathcal{S}_t = \varnothing$ and corresponding shortest travel time $T_{\mathcal{S}_t} = 0$;
initialize reachable candidates
$\mathcal{V}_m = \{i \in \mathcal{V} : \tau_{pos_{mt}i} \leq \Delta\mathcal{T} \text{ and } t + \tau_{pos_{mt}i} + \tau_{is_m^{end}} \leq T^{max}\}$;
while $\mathcal{V}_m \neq \varnothing$ **do**
    $s_t = \arg\max_{s \in \mathcal{V}_m} \mathcal{I}(\mathcal{S}_t \cup \{s\}, \mathcal{S}_{<t} \cup \mathcal{S}_t)$;
    add s_t to planned targets $\mathcal{S}_t = \mathcal{S}_t \cup \{s_t\}$;
    update $T_{\mathcal{S}_t}$ as the shortest travel time for visiting $\mathcal{S}_t$ starting at pos_{mt};
    update $\mathcal{V}_m = \{s \in \mathcal{V}_m : s \notin \mathcal{S}_t \text{ and } T_{\mathcal{S}_t \cup \{s\}} \leq \Delta\mathcal{T}\}$;
end
return $\mathcal{S}_t$;

locations $\mathcal{S}_t$ are visited. This way, each iteration accounts for all previously selected targets. After deciding on the next target s_t, $\mathcal{V}_m$ is updated such that it includes only locations that can be added to $\mathcal{S}_t$ while still reaching all selected locations within the moving horizon $\Delta\mathcal{T}$.

Algorithm 6 stops when $\mathcal{V}_m$ is empty, i.e., when there are no candidates left that can be added to the $\mathcal{S}_t$ without exceeding the moving planning horizon. The time required for visiting $\mathcal{S}_t$, denoted $T_{\mathcal{S}_t}$, is determined by solving an open traveling salesman problem, i.e., a TSP without returning to a fixed location. This way, we avoid the zig-zag effect discussed in the greedy approach, as samples close to the current position are visited first even though their current informativeness value might be lower.

The sequential selection strategy results in a heuristic selection of samples for the given planning horizon length $\Delta\mathcal{T}$, as not all possible visit combinations for $\mathcal{S}_t$ are evaluated. In comparison to the greedy approach, it promises two advantages: It anticipates the effect that the first observations can have on the benefit of subsequent targets, and it can improve the resource utilization by changing the order of target locations.

Integrated approach

The last option considered in this work is the integrated approach, which consists of determining the optimal selection of sampling locations. This means that we solve the AMPPES defined in Section 13.1 to optimality for a given planning horizon $\Delta\mathcal{T}$.

For this purpose, we adapt the dynamic programming approach introduced Section 8.1. In contrast to Part II, we model the planning problem as an open planning problem, as the UAVs do not need to return to their destination location s_m^{end} at the end of the moving horizon. Therefore, we only use forward labeling, as the final position

at the end of the considered moving planning horizon is not fixed. Still, the algorithm ensures that all paths built this way allow the UAVs to return to their destination within their flight time T^{max}.

In contrast to the two previous approaches, the integrated consideration of target selection and routing ensures that the best possible solution within the moving horizon is found. Moreover, this planning approach optimizes the routes for all UAVs together rather than considering them one after the other. However, this is computationally expensive, as we have shown in the experimental study in Part II.

Moving and fixed planning horizon

Both the sequential selection and the integrated planning approach determine missions within a moving planning horizon $\Delta\mathcal{T}$. However, it is not necessary to fully execute the planned mission. To this end, we use a fixed decision horizon $\Delta\mathcal{T}^{fix} \leq \Delta\mathcal{T}$. Routes are planned throughout $\Delta\mathcal{T}$, but after a time $\Delta\mathcal{T}^{fix}$ has passed, the previous plans are discarded in favor of an updated mission. This way, we can control the degree of adaptivity as well as the effort that is necessary for updating missions throughout the entire planning horizon.

If the fixed decision horizon $\Delta\mathcal{T}^{fix}$ is small, missions are recomputed very quickly, i.e., new information is quickly taken into account. However, only a subset of the samples that have been selected in the previous iteration is surveyed before the mission is updated. This means that possible benefits due to a good combination of samples can be lost. The main advantage of this approach is that it anticipates the impact of the next few visited target locations on the outcome of the mission, either with respect to their contribution to the overall information gain or concerning the flight time necessary to reach them, while offering a high degree of adaptivity.

In contrast, if $\Delta\mathcal{T}^{fix}$ is large, the planned mission, which is either optimal for $\Delta\mathcal{T}$ in case of the integrated approach or at least includes an optimal sequencing of locations in case of the heuristic strategy, is executed in full or nearly in its entirety. This takes advantage of the more sophisticated approaches that use larger decision scopes but loses the ability to adapt more quickly. Furthermore, the possibly expensive computation of a mission within the moving planning horizon $\Delta\mathcal{T}$ is not repeated as often.

13.3.2 Geometric solution approach

We introduce one geometric solution approach for benchmarking the model-based approaches introduced above. This approach adapts Lloyd's algorithm for determining an even distribution of sensors.

As briefly described in Section 11.1.5, Lloyd's algorithm is an iterative way to construct Voronoi tessellations based on a set of initial generators. Note that even though Lloyd's algorithm can be applied to continuous input data, we need to compare results to model-based variants operating on discrete candidate locations. Hence, we adapt the algorithm for selecting sampling locations based on a discrete set of candidate sampling locations.

This algorithm relies on Voronoi tessellations of the target regions at time $t \in \mathcal{T}$, which depends on the current position pos_{mt} of all UAVs $m \in \mathcal{M}$ at time t. To this end, we determine a partition (i.e., a Voronoi region)

$$V_m = \{i \in \mathcal{V} : \|pos_i - pos_{mt}\| \leq \|pos_i - pos_{nt}\| \text{ for } n \in \mathcal{M}, n \neq m\} \tag{13.2}$$

for each UAV $m \in \mathcal{M}$. UAVs then move toward the centroid c of V_m, which is located at

$$pos_c = \left(\frac{\sum_{i \in V_m} \mathcal{I}(i) \cdot p_i^x}{\sum_{i \in V_m} \mathcal{I}(i)}, \frac{\sum_{i \in V_m} \mathcal{I}(i) \cdot p_i^y}{\sum_{i \in V_m} \mathcal{I}(i)} \right), \tag{13.3}$$

with (p_i^x, p_i^y) as the coordinates of all target locations $i \in \mathcal{V}$.

The position pos_c of centroid c depends on the informativeness measure $\mathcal{I}$ defined over $\mathcal{V}$ that assigns weights to target locations. In the purely geometric variant, this function corresponds to the locations' priorities. In Equation (13.3), we use the more generic expression that allows integrating the different environment models and informativeness measures in Section 13.4. In the most simple case, where $\mathcal{I}(i)$ corresponds to the priority u_i, pos_c corresponds to the priority-weighted center of the partition.

In our implementation, each UAV calls Algorithm 7 to determine its next sampling position s_t. This algorithm first determines the Voronoi region for UAV $m \in \mathcal{M}$ and its centroid c, depending on the measure $\mathcal{I}$. The next sampling location is the candidate in $\mathcal{V}_m$ that is within a maximum distance of d^{limit} with respect to the UAV position pos_{mt} at the time the algorithm is called and that is associated with the lowest detour compared to a direct flight to c.

Algorithm 7: Lloyd's algorithm adapted to the AMPPES

Input: Available UAVs $\mathcal{M}$ with positions $pos_{nt}, n \in \mathcal{M}$, relevant UAV $m \in \mathcal{M}$, target locations $\mathcal{V}$ with informativeness measure $\mathcal{I}(i), i \in \mathcal{V}$, step width restriction d^{limit}

Output: Next sampling point $s_t \in \mathcal{V}$ for UAV m

1 determine Voronoi region V_m using Equation (13.2);
2 determine centroid c of V_m with position pos_c based on Equation (13.3);
3 determine the next sampling position
$s_t = \text{argmin}_{pos_i \in V_m : \|i - pos_{mt}\| \leq d^{limit}} (\|pos_i - pos_{mt}\| + \|pos_c - pos_i\|)$;
4 **return** s_t;

Note that Algorithm 7 does not depend on the UAV flight time T^{max} or moving planning horizon $\Delta\mathcal{T}$: It seeks to incrementally achieve a balanced spatial distribution of UAVs across the target area. As it does not consider UAV movements further, temporal aspects have no impact on the outcome. Instead, UAVs simply return home to their designated landing position when their flight time is exceeded. This means that Algorithm 7 is relatively quick and robust in achieving a reliable spatial distribution, but may not utilize the available mission time once this distribution is achieved.

13.4 Environment models and information measures

Environment models define the informativeness measure $\mathcal{I}$, which has been used as a black box model for evaluating possible sampling candidates $\mathcal{S}_t$ in the planning heuristics discussed in the previous section. They allow us to represent the knowledge obtained about the phenomenon and to assess the information provided by candidate sampling locations, depending on the previously sampled locations $\mathcal{S}_{<t}$. To study the effects of different environment representations, we introduce several variants, ranging from basic coverage models to complex estimations based on probabilistic process models. In the following, we formally describe these environment representations and utility functions, we apply a unified model comprising four elements:

- The **input parameters** that are used to initialize the model.
- The **model data**, i.e., model-specific data structures that are maintained and updated during the mission.
- An **update** component for updating the model data based on obtained samples. This is done each time a UAV completes an observation at some location $s \in \mathcal{V}$.
- An **evaluation** component for determining the value of candidate sampling locations $\mathcal{S}_t$ based on the updated model data.

The modeling variants are discussed in more detail below. For all models, we define the four components introduced above and state the formal definition of the underlying informativeness measure for candidate samples. Illustrative examples for the application of these models to a representative planning scenario can be found in Appendix B with a focus on the evolution of the informative measure as the mission progresses.

13.4.1 Nonadaptive models

In a first step, we discuss nonadaptive models, i.e., models that do not rely on previous observations. They represent strategies that often lead to good results, in particular by favoring an even distribution of measurements.

Direct priority-based model

The first environment model is based on the priorities in the target area. Hence, this model corresponds to an online version of the TOP. As the most basic model, this variant ensures that unsampled locations associated with high priorities in the input data are visited first. Similar to our study in Part II, this model offers an indication of the results that can be achieved without considering spatial interdependencies in the informativeness measure.

Table 13.1 summarizes the relevant components. The model is initialized with information about the targets' priorities $u_i, i \in \mathcal{V}$. The model maintains an additional data

structure for keeping track of the visited locations $\mathcal{S}_{<t}$. The information gain then corresponds to the sum of the priorities of the candidates $\mathcal{S}_t$ that have not yet been surveyed before.

Parameters	Priorities $u_i \geq 0$ for all locations $i \in \mathcal{V}$
Model data	Completed sampling locations $\mathcal{S}_{<t}$
Update	Add sampled location s to the set of completed samples $\mathcal{S}_{<t}$
Evaluate	$\mathcal{I}^{DIRECT}(\mathcal{S}_t, \mathcal{S}_{<t}) = \sum_{s \in \mathcal{S}_t \setminus \mathcal{S}_{<t}} u_s$

Table 13.1: Definition of the direct priority coverage model.

Disc-based coverage model

The disc-based coverage model assumes that each sampled location provides complete knowledge about all surrounding locations within a maximum covering distance d^{cover}. This distance represents the extent of the correlation in the process, i.e., the distance up to which one can assume that observations yield similar results. The assumption of disc-based coverage has motivated several models in literature, e.g., the covering salesman problem (see Section 5.2) as well as online environment representations (see Section 11.1.4). In their most basic variant, these models only consider whether a location has been within the covering distance of a completed observation. In our model, we additionally account for how many observations have been made within d^{cover}. This way, we encourage a balanced distribution of sampled locations across the target area.

The constituting components of the disc-based model are summarized in Table 13.2. As indicated above, we assume that each sampled location provides information about other targets within a fixed radius of d^{cover} that has to be specified by the user. For each location $i \in \mathcal{V}$, we keep track of the number of times n_i a sample was taken at any location within the maximum distance d^{cover} to i. The model represents the value of either sampling a visit i directly or surveying a location within d^{cover} to i as the squared difference between n_i and the maximum coverage count $\max_{j \in \mathcal{V}} n_j$ over all targets, weighted with the location's priority u_i. The estimated informativeness of a set $\mathcal{S}_t$ corresponds to the sum of these values for all $i \in \mathcal{V}$ that are within covering distance of at least one sample $s \in \mathcal{S}_t$.

Using the squared difference $(\max_{j \in \mathcal{V}} n_j - n_i)^2$ between these coverage counts means that the regions that have received the least attention are prioritized as the total number of samples increases. The proposed model furthermore ensures that even if nearly all targets in the considered area have been surveyed or covered, their informativeness is not simply set to 0, as is the case in less complex coverage models. Instead, it offers a means to select further samples in an informed manner, thus ensuring that the available flight time is used in its entirety.

Parameters	Coverage radius $d^{cover} > 0$, priorities $u_i \geq 0,\ i \in \mathcal{V}$
Model data	Target coverage counter n_i for all $i \in \mathcal{V}$, initialized as 0
Update	Set $n_i = n_i + 1$ for all $i \in \mathcal{V}$ within distance d^{cover} to a sampled location s
Evaluate	$\mathcal{I}^{DISC}(\mathcal{S}_t) = \sum_{i \in \mathcal{V}: \min_{s \in \mathcal{S}_t} d_{is} \leq d^{cover}} u_i \left(\max_{j \in \mathcal{V}} n_j - n_i\right)^2$

Table 13.2: Definition of the disc-based coverage model.

An example for $\mathcal{I}^{DISC}$ is given in Figure 13.1, which illustrates the informativeness values of all targets depending on previous samples indicated as black crosses. Note that for this example, the priority of all locations is 1, i.e., only the modeled spatial dependencies impact the informativeness values of different locations. As expected, we can see that the informativeness is low at areas where the mission has provided samples and high in areas further away. The next sample chosen using the greedy heuristic based on this informativeness measure is connected to the last sampling location using a dashed line.

Kernel-based coverage

In the next step, we propose a more refined coverage model compared to the disc-based coverage approach. Specifically, this means that we model the correlation within the target area in more detail. To this end, we use the GCorTOP objective function derived in Section 7.3, see Equation (7.12), which represents the spatial correlations as a distance-dependent kernel function.

In the offline case, the informativeness of a set of samples $\mathcal{S}$ is measured as the sum of

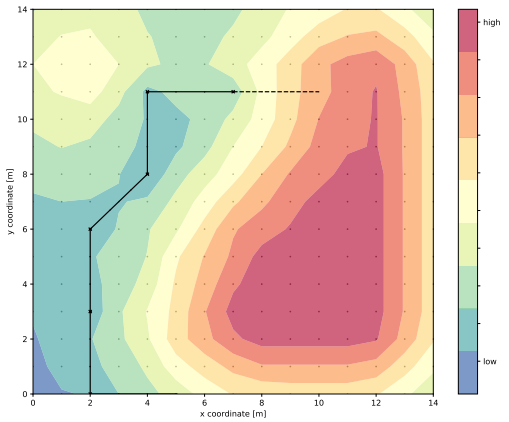

Figure 13.1: Informativeness of sampling locations using $\mathcal{I}^{DISC}$.

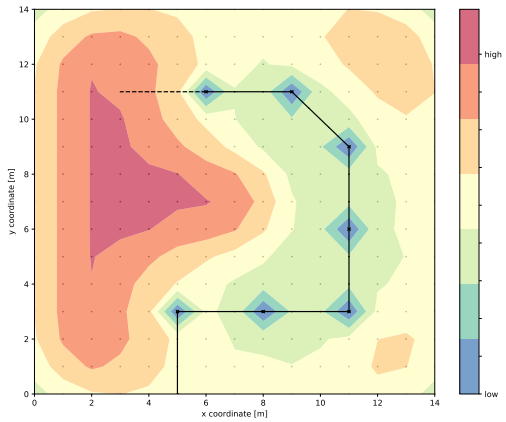

Figure 13.2: Informativeness of sampling locations using $\mathcal{I}^{GCorTOP}$.

the priorities u_i of all directly surveyed locations $i \in \mathcal{S}$ plus the priorities of unsampled locations $j \in \mathcal{V} \setminus \mathcal{S}$, weighted with a factor $w_{ji} \in [0, 1]$ depending on their distance to the next surveyed location $i \in \mathcal{S}$. This objective is formally defined as follows:

$$\mathcal{I}^{GCorTOP}(\mathcal{S}) = \sum_{i \in \mathcal{S}} u_i + \sum_{j \in \mathcal{V} \setminus \mathcal{S}} \min\{u_j, \sum_{i \in \mathcal{S} \cap \mathcal{C}_j} w_{ji} u_j\},$$

with the covering neighborhood $\mathcal{C}_i$ that includes all targets within a maximum distance d^{cover} to location $i \in \mathcal{V}$, i.e., $\mathcal{C}_i = \{j \in \mathcal{V} \setminus \{i\} : d^{cover} \geq \| pos_j - pos_i \|\}$. The weights w_{ji} allow us to represent the correlations using arbitrary distance-dependent kernel functions. Hence, we refer to the coverage model for online planning as kernel-based coverage in contrast to the disc-based model.

The environment model for the online case is defined in Table 13.3. We use the same inverse-distance based approach for determine weights $w_{ij}, i, j \in \mathcal{V}$ as in the offline case, see also Equation (7.10). This means that we specify all weights w_{ij} for $i, j \in \mathcal{V}$ relative to a baseline weight $\bar{w}$ at some distance d^{min} as

$$w_{ji} = \begin{cases} \bar{w} \cdot \frac{d^{min}}{d_{ji}}, & j \in \mathcal{C}_i \\ 0, & \text{otherwise.} \end{cases} \tag{13.4}$$

The benefit of a set of candidate samples $\mathcal{S}_t$ is then determined as the incremental change $\mathcal{I}^{GCorTOP}$, i.e., $\mathcal{I}^{KERNEL}(\mathcal{S}_t, \mathcal{S}_{<t}) = \mathcal{I}^{GCorTOP}(\mathcal{S}_{<t} \cup \mathcal{S}_t) - \mathcal{I}^{GCorTOP}(\mathcal{S}_{<t})$.

Parameters	Coverage radius $d^{cover} > 0$, baseline weight $\bar{w} \in (0, 1]$ and distance $d^{min} > 0$, priorities $u_i \geq 0$ for $i \in \mathcal{V}$
Model data	Weights w_{ij} based on Equation (13.4), completed sampling locations $\mathcal{S}_{<t}$
Update	Add sampled location s to the set of completed samples $\mathcal{S}_{<t}$
Evaluate	$\mathcal{I}^{KERNEL}(\mathcal{S}_t, \mathcal{S}_{<t}) = \mathcal{I}^{GCorTOP}(\mathcal{S}_{<t} \cup \mathcal{S}_t) - \mathcal{I}^{GCorTOP}(\mathcal{S}_{<t})$ $= \underbrace{\sum_{s \in \mathcal{S}_{<t} \cup \mathcal{S}_t} u_s - \sum_{s \in \mathcal{S}_{<t}} u_s}_{\text{incremental value for direct coverage}}$ $+ \underbrace{\sum_{s \in \mathcal{V} \setminus (\mathcal{S}_{<t} \cup \mathcal{S}_t)} \min\{u_s, \sum_{j \in \mathcal{S}_{<t} \cup \mathcal{S}_t} w_{js} u_s\} - \sum_{s \in \mathcal{V} \setminus \mathcal{S}_{<t}} \min\{u_s, \sum_{j \in \mathcal{S}_{<t}} w_{js} u_s\}}_{\text{incremental value for indirect coverage}}$

Table 13.3: Definition of the kernel-based coverage model.

The kernel-based coverage model is more expensive to compute than the disc-based coverage model. Using the weights w_{ij}, however, means that more detailed information about spatial coverage can be used for evaluating targets. For instance, this allows a

better approximation of the benefits of including multiple samples that cover the same locations.

Figure 13.2 provides an example for $\mathcal{I}^{GCorTOP}$. Similar to the example for the disc-based coverage model, it indicates the current informativeness values depending on previously sampled locations, with priorities $u_i = 1, i \in \mathcal{V}$. We can observe that the resulting measures for the targets' priorities are similar to the disc-based model. The main difference between the two measures in this example is the value assigned to targets close to previous samples. This value is higher relative to the remainder of the samples in case of $\mathcal{I}^{GCorTOP}$. This accounts for the fact that additional samples still provide information in these areas, even though unsampled regions should be preferred for the next sampling locations.

13.4.2 Extension to time-dependent models

The models discussed up to this point only rely on spatial information, i.e., they are blind to the point in time at which an observation was performed at some location. We extend them to spatio-temporal models, where information loses value as time progresses. This allows us to address spatio-temporal phenomena, where the surveyed distribution changes in time. Additionally, these models can be applied to balance observations over time as well as over space and to account for the available flight time when evaluating the information gain of samples.

Latency-based coverage model

In a first step, we adjust the disc-based coverage model. Again, we assume that locations provide full information about all samples within a given distance d^{cover}. Instead of counting the number of samples, however, we keep track of the time since the last measurement has been made in the vicinity of a location. We then measure the informativeness of a location relative to the sampling latency, i.e., the time since the last measurement at the location itself or nearby.

The definition of this model is given in Table 13.4. We maintain an additional data structure for storing the last sampling time $\tau_i^{lat}, i \in \mathcal{V}$. Similar to the disc-based coverage model, we use the squared difference between the current time $t \in [0, \mathcal{T}]$ and the last sampling time of a target to penalize long delays between samples.

The latency-based model, similar to the disc-based model, favors the regions in the target area that have received the least attention. These tend to be the locations furthest away from the current UAV positions. This way, we encourage spatial diversification.

Figure 13.3 illustrates the latency-based measures, again using an instance with unit priorities. The UAV mission starts at the lower left side of the represented area and proceeds as indicated using the black line. We can see that the informativeness of sampling locations near the current position of the UAV in the middle of the area is low. However, $\mathcal{I}^{LAT}$ is already high in the area that has been sampled first, thereby encouraging additional samples in this region.

Parameters	Coverage radius $d^{cover} > 0$, priorities $u_i \geq 0$ for $i \in \mathcal{V}$
Model data	Last sampling time $\tau_i^{lat},\ i \in \mathcal{V}$, initialized as 0
Update	Set $\tau_i^{lat} = t_s$ for all $i \in \mathcal{V}$ within d^{cover} to location s sampled at time t_s
Evaluate	$\mathcal{I}^{LAT}(\mathcal{S}_t) = \sum_{i \in \mathcal{V}: \min_{s \in \mathcal{S}_t} d_{is} \leq d^{cover}} u_i \left(t - \tau_i^{lat}\right)^2$

Table 13.4: Definition of the latency-based coverage model.

Kernel-based coverage with devaluation

The GCorTOP approach assumes that a process is time-invariant. We, therefore, need to adapt the idea in situations in which we want to account for temporal relationships. We propose to devalue the information gain provided by samples as time proceeds depending on the length of the planning horizon. For this purpose, the model maintains information about previously sampled locations and sampling times in the form of the set $\mathcal{S}_{<t}$. This means that in the time-dependent context, $\mathcal{S}_{<t}$ contains tuples of values (s, t_s) that indicate a sampled location s together with the time $t_s \in [0, \mathcal{T}]$ the sample was obtained. This way, if a location is surveyed several times, it is represented by multiple entries with different sampling times in $\mathcal{S}_{<t}$.

Using this information, we formulate the time-dependent correlated team orienteering problem (GCorTOP-TD) as an extension of Equation (7.12) (Section 7.4). We account for the fact that information loses value over time by multiplying the distance-dependent weight w_{si} between a sampled location s and an interpolated location $i \in \mathcal{V}$ by a factor $(1 - \frac{t - t_s}{\mathcal{T}})$, where t indicates the point in time at which the informativeness measure is evaluated, with $t \in [0, \mathcal{T}]$. As the time since the last sampling time t_s increases, the

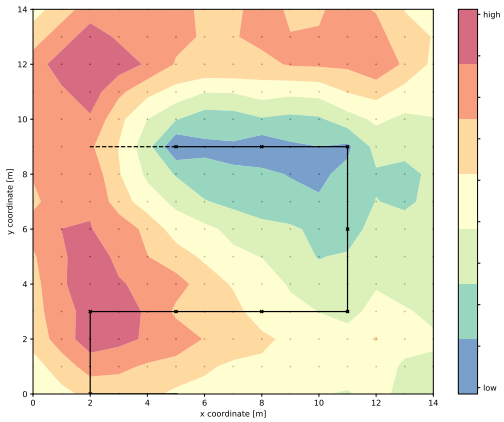

Figure 13.3: Informativeness of sampling locations using $\mathcal{I}^{LAT}$.

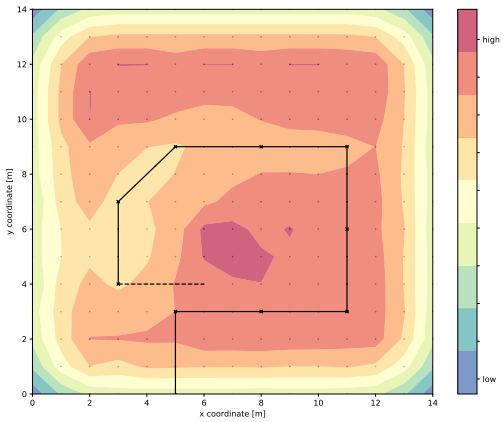

Figure 13.4: Informativeness of sampling locations using $\mathcal{I}^{GCorTOP\text{-}TD}$.

fraction $\frac{t-t_s}{\mathcal{T}}$ converges to 1, i.e., $(1 - \frac{t-t_s}{\mathcal{T}})$ converges to 0. The same assumption holds for the information that a location provides about itself at a later point in time. For this, we assume that $w_{ss} = 1$. This weight is similarly discounted as time progresses.

Based on this assumption, the GCorTOP-TD model estimates the informativeness of a set of completed samples $\mathcal{S}_{<t}$ relative to a point in time t as follows:

$$\mathcal{I}^{\text{GCorTOP-TD}}(\mathcal{S}_{<t}) = \sum_{i \in \mathcal{V}} \min\{u_i, \overbrace{\sum_{(s,t_s) \in \mathcal{S}_{<t}} \underbrace{\left(1 - \frac{t - t_s}{\mathcal{T}}\right) w_{si}}_{\text{reduced weight between } s \text{ and } i} u_i}^{\text{indirect information available about } i \text{ relative to time } t} \}. \tag{13.5}$$

Note that we assume that $t \geq t_s$, as $\mathcal{I}^{\text{GCorTOP-TD}}$ can only be reasonably evaluated with respect to samples obtained during the mission until t. We maintain our assumption that the sum of all indirect coverage values is equal or less to u_i to avoid redundant samples in close vicinity or within a short time (see Section 7.4).

Table 13.5 gives the definition of the modified model. Similar to the kernel-based approach given in Table 13.3, we define the informativeness $\mathcal{I}^{\text{KERNEL-TD}}$ of sampling candidates $\mathcal{S}_t$ as the incremental change in the objective value $\mathcal{I}^{\text{GCorTOP-TD}}$ with respect to the available samples $\mathcal{S}_{<t}$. For this purpose, all candidate locations $s \in \mathcal{S}_t$ are evaluated as if they were surveyed immediately at time t, i.e., their weights w_{si} are fully taken into account.

Parameters	Coverage radius $d^{cover} > 0$, baseline weight $\bar{w} \in (0, 1]$ and distance $d^{min} > 0$, priorities $u_i \geq 0$ for $i \in \mathcal{V}$, planning horizon length $\mathcal{T}$
Model data	Weights w_{ij} based on Equation (13.4), completed sampling locations $\mathcal{S}_{<t}$
Update	Add (s, t_s) to the set of completed samples $\mathcal{S}_{<t}$
Evaluate	$\mathcal{I}^{\text{KERNEL-TD}}(\mathcal{S}_t) = \mathcal{I}^{\text{GCorTOP-TD}}(\mathcal{S}_{<t} \cup \mathcal{S}_t) - \mathcal{I}^{\text{GCorTOP-TD}}(\mathcal{S}_{<t})$ $= \sum_{i \in \mathcal{V}} \min\{u_i, \left(\sum_{(s,t_s) \in \mathcal{S}_{<t}} \left(1 - \frac{t - t_s}{\mathcal{T}}\right) w_{si} + \sum_{s \in \mathcal{S}_t} w_{si}\right) \cdot u_i\}$ $- \sum_{i \in \mathcal{V}} \min\{u_i, \sum_{(s,t_s) \in \mathcal{S}_{<t}} \left(1 - \frac{t - t_s}{\mathcal{T}}\right) w_{si} u_i\}$

TABLE 13.5: Definition of the kernel-based coverage model with devaluation.

As indicated above, we measure the time $t - t_s$ that has passed since the last measurement relative to the length of the planning horizon $\mathcal{T}$. This means that information is devaluated more quickly the shorter the planning is, as little time remains to make additional samples. In contrast, in the case of longer planning horizons, UAVs are encouraged to travel further before returning for additional samples. Same as the weights

$w_{ji}, j, i \in \mathcal{V}$, this factor can be adjusted if information about the nature of the temporal correlation in the process is available.

We again represent this measure based on a scenario with unit priority, indicated in Figure 13.4. In this case, the UAV starts at position $(5, 0)$. Similar to the example for latency-based coverage, the area close to the starting position is surveyed first and becomes relevant again as the mission progresses, leading the UAV to return and take additional samples nearby. Note that overall, values for $\mathcal{I}^{GCorTOP\text{-}TD}$ are higher than for $\mathcal{I}^{GCorTOP}$ and $\mathcal{I}^{LAT}$, as the distance between samples both in the spatial as well as in the temporal dimension limit the value of the already obtained information.

13.4.3 Gaussian process models

Adaptive models make use of the observations received while the mission is in progress to adjust and improve the current environment representation. In this work, we use Gaussian process (GP) models that represent the surveyed distribution of gases. We compare different measures for modeling information gain based on these processes. These models can be applied to stationary processes as well as to phenomena changing over time by using appropriate kernel functions k (see also Section 3.2.2).

In this study, we consider three informativeness measures based on GP representations that follow the measures discussed in Section 3.2.4. The definition of these three models is given in Table 13.6. They make use of the same fundamental process representations and update strategies during the mission and differ in how $\mathcal{I}(\mathcal{S}, \mathcal{S}_{<t})$ is evaluated. In all cases, we parametrize the environment model using a kernel k_θ with hyperparameters θ that characterizes the initial belief about the nature of the correlation.

Process updates

While the mission progresses, UAVs take (potentially noisy) measurements of the observed phenomenon. We denote such a measurement taken at time t at a location $s \in \mathcal{V}$ as y_{st}. We refer to the set of collected measurements up to time t as $\mathcal{Y}_{<t}$. This set represents the entire information available at the time t a mission is updated. Note that the set $\mathcal{Y}_{<t}$ includes the *values* of the observations in contrast to the *locations* represented by $\mathcal{S}_{<t}$, which we have used in the models up to this point.

To make use of the information collected during the mission, the hyperparameters θ are updated such that they best fit the available data that is represented by the set $\mathcal{Y}_{<t}$. This yields a fitted kernel $k_{\hat{\theta}}$ with optimized hyperparameters $\hat{\theta}$. The fundamental idea behind this approach is to find a parametrization that matches the patterns exhibited by the collected samples as well as possible. We refer to Section 3.2.2 for an overview of how the parameters can be fitted to the obtained information. The fitted GP based on a kernel k_θ with hyperparameters θ conditioned on samples $\mathcal{S}_{<t}$ is denoted as $Z^{k_\theta}_{\mathcal{V}|\mathcal{S}_{<t}} \backsim \mathcal{GP}(\mu^{k_\theta}_{\mathcal{V}|\mathcal{S}_{<t}}, \Sigma^{k_\theta}_{\mathcal{V},\mathcal{V}|\mathcal{S}_{<t}})$ in the following.

Any predictions and variance estimates are based on such a fitted process model $Z^{k_{\hat{\theta}}}_{\mathcal{V}|\mathcal{S}_{<t}}$. This means that the estimated information value of locations can change over

time, and can increase and decrease depending on previous measurements. If, for example, the correlation between samples at larger distances is higher than initially assumed, the estimated value provided by locations close to previous sampling points decreases. If the model correctly identifies this property, it can increase the distance between sampling locations. This way, it uses the obtained knowledge about the correlation to cover a larger area. The computational effort for an update step in which newly obtained information is processed depends on how the parameters $\hat{\theta}$ are determined. `Scikit-learn` (Pedregosa et al., 2011), which is used in this work, uses a gradient-based multi-start approach to find locally optimal parameter sets.

Parameters	Kernel k_θ defining the prior covariance $\Sigma^{k_\theta}_{\mathcal{V},\mathcal{V}}$
Model data	Sampled locations $\mathcal{S}_{<t}$, corresponding measurements $\mathcal{Y}_{<t}$, posterior GP $Z^{k_{\hat{\theta}}}_{\mathcal{V}\mid\mathcal{S}_t}$ defined by the kernel function $k_{\hat{\theta}}$
Update	Add s to the set of sampled locations $\mathcal{S}_{<t}$ and add the corresponding measurement y_{st} to $\mathcal{Y}_{<t}$. Update kernel $k_{\hat{\theta}}$ such that the hyperparameters $\hat{\theta}$ are optimal with respect to the log-marginal likelihood function, i.e., $\hat{\theta} = \arg\max_\theta \log p(\mathcal{Y}_{<t}\mid\theta)$, and compute the posterior distribution $Z^{k_{\hat{\theta}}}_{\mathcal{V}\mid\mathcal{S}_{<t}}$ (see also Equation (3.12), Section 3.2.2)
Evaluate	Use one of three measures: $\mathcal{I}^{VAR}(\mathcal{S}_t,\mathcal{S}_{<t}) = \operatorname{tr}\left(\Sigma^{k_{\hat{\theta}}}_{\mathcal{S}_t,\mathcal{S}_t\mid\mathcal{S}_{<t}}\right)$ $\mathcal{I}^{ARV}(\mathcal{S}_t,\mathcal{S}_{<t}) = \frac{1}{\lvert\mathcal{V}\rvert}\left(\operatorname{tr}\left(\Sigma^{k_{\hat{\theta}}}_{\mathcal{V},\mathcal{V}\mid\mathcal{S}_{<t}}\right) - \operatorname{tr}\left(\Sigma^{k_{\hat{\theta}}}_{\mathcal{V},\mathcal{V}\mid(\mathcal{S}_{<t}\cup\mathcal{S}_t)}\right)\right)$ $\mathcal{I}^{MI}(\mathcal{S}_t,\mathcal{S}_{<t}) = \mathcal{H}\left(Z^{k_{\hat{\theta}}}_{\mathcal{V}\setminus(\mathcal{S}_{<t}\cup\mathcal{S}_t)\mid\mathcal{S}_{<t}}\right) - \mathcal{H}\left(Z^{k_{\hat{\theta}}}_{\mathcal{V}\setminus(\mathcal{S}_{<t}\cup\mathcal{S}_t)\mid(\mathcal{S}_{<t}\cup\mathcal{S}_t)}\right)$

Table 13.6: Definition of GP based models.

Informativeness measures

The core of the GP-based models is the way the information gain of a set $\mathcal{S}_t$ is determined. We study three variants that differ in the degree of complexity that is included in the models:

The measure $\mathcal{I}^{VAR}$, as the most basic model, estimates informativeness as the sum of the posterior variances associated with the candidate sampling locations in $\mathcal{S}_t$. That is, it computes the trace of the posterior covariance matrix $\Sigma^{k_{\hat{\theta}}}_{\mathcal{S}_t,\mathcal{S}_t\mid\mathcal{S}_{<t}}$ defined over the locations in $\mathcal{S}_t$. Hence, the model assigns the highest values to the most uncertain target locations. Usually, these tend to be those furthest from the surveyed locations.

The $\mathcal{I}^{VAR}$ model does not make full use of the information provided by the covariance matrix about the surveyed process. It selects the samples where the prediction is currently least certain, but these are not necessarily the sampling locations that can provide the most information about the remainder of the affected area. This can be achieved by applying more elaborate informativeness measures proposed in literature.

As the second measure, we use the average reduction in variance (ARV), see also Equation (3.21) (Section 3.2.4). The $\mathcal{I}^{ARV}$ measures the expected reduction of the variability in the GP posterior that can be achieved by the samples in $\mathcal{S}_t$. To this end, we compute the variability in the process over the entire target area $\mathcal{V}$ depending on the observed locations $\mathcal{S}_{<t}$ as the trace of the posterior covariance $\Sigma^{k_{\hat{\theta}}}_{\mathcal{V},\mathcal{V}|\mathcal{S}_{<t}}$. The expected variability assuming that the locations in $\mathcal{S}_t$ are surveyed as well is $\operatorname{tr}\left(\Sigma^{k_{\hat{\theta}}}_{\mathcal{V},\mathcal{V}|(\mathcal{S}_{<t}\cup\mathcal{S}_t)}\right)$, i.e., the trace of the posterior covariance matrix conditioned on sets $\mathcal{S}_{<t}$ and $\mathcal{S}_t$ together. The reduction in variance is the difference between these two values.

The ARV provides a more sophisticated measure of the improvement in prediction quality that can be expected by a set $\mathcal{S}_t$. However, this means that the posterior covariance $\Sigma^{k_{\hat{\theta}}}_{\mathcal{V},\mathcal{V}|(\mathcal{S}_{<t}\cup\mathcal{S}_t)}$ has to be computed for measuring the benefit obtained by a set of sampling locations. This has to be repeated for every evaluated set $\mathcal{S}_t$, introducing a significant overhead relative to $\mathcal{I}^{VAR}$.

As the third measure, we use the mutual information (MI) criterion, which measures the expected information gain provided by the samples about the remainder of the process, see also Equation (3.22) (Section 3.2.3). For sensor placement models, this measure has often yielded the best reported results. However, it is also the most expensive one (Krause et al., 2008). The MI criterion uses the entropy $\mathcal{H}(Z_{\mathcal{V}})$ to measure the uncertainty of a process $Z_{\mathcal{V}}$ (see Equation (3.17), Section 3.2.3). This measure is applied to the sampling locations $\mathcal{V} \setminus (\mathcal{S}_{<t} \cup \mathcal{S}_t)$ that are not included in either the previously selected samples nor in the candidate locations. The informativeness $\mathcal{I}^{MI}$ is then defined as the difference between the entropy of the process $Z^{k_{\hat{\theta}}}_{\mathcal{V}\setminus(\mathcal{S}_{<t}\cup\mathcal{S}_t)|\mathcal{S}_{<t}}$ that only considers previous measurements $\mathcal{S}_{<t}$ and the entropy of the process $Z^{k_{\hat{\theta}}}_{\mathcal{V}\setminus(\mathcal{S}_{<t}\cup\mathcal{S}_t)|(\mathcal{S}_{<t}\cup\mathcal{S}_t)}$ that has been conditioned on completed and candidate samples together.

The strength of these adaptive informativeness measures is that they assign the highest values to the targets that, accounting for all available information, will provide the most insights into the nature of the process. Furthermore, they consider these sets jointly–that is, they assess the information gain that is provided by all candidates $\mathcal{S}_t$ together depending on their locations and distances toward one another as well as to all remaining locations. However, the required computational effort is significantly higher than in case of nonadaptive models due to the need to fit the GP to the completed observations and to compute the GP posterior.

The three informativeness measures are illustrated in Figures 13.5 to 13.7. Again, black lines illustrate the executed missions and dashed lines the path toward the next sampling location that is greedily selected based on the underlying environment models. Figure 13.5 illustrates simple structure obtained when using $\mathcal{I}^{VAR}$ where each loca-

tion's informativeness corresponds to its posterior covariance. Consequently, the measure yields large values further away from the samples. Most importantly, this affects the locations that are near the border of the target region.

$\mathcal{I}^{ARV}$ prevents this effect, as is illustrated in Figure 13.6. This measure does not prioritizes the locations that are most uncertain themselves, but that reduce overall variance most. Again, this corresponds to locations further away from the obtained samples. However, the measure accounts for the fact that even though little information is known about the border locations, they also provide less insight into the phenomenon.

The last measure $\mathcal{I}^{MI}$ is represented in Figure 13.7. Based on the executed mission and obtained samples, this measure assigns high informativeness values to the area on the right hand side of the area. This is because previous measurements fall, roughly, on a vertical line, which means that little information has been obtained about the spatial correlation in the orthogonal direction.

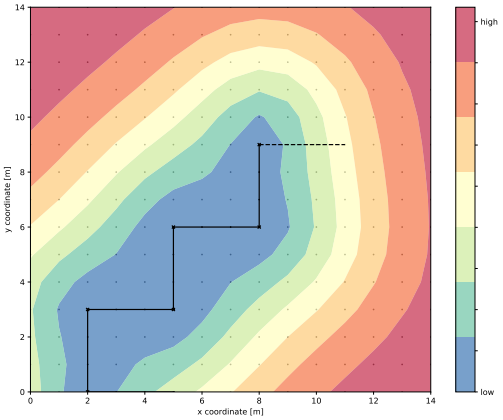

FIGURE 13.5: Informativeness of sampling locations using $\mathcal{I}^{VAR}$.

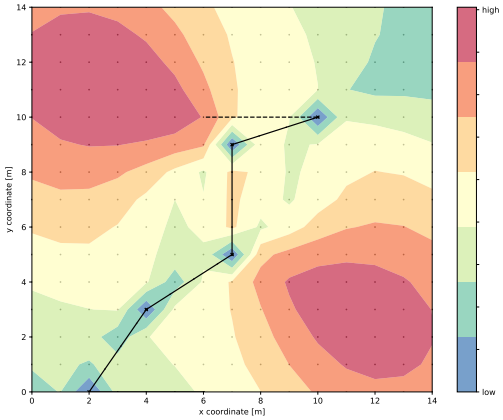

FIGURE 13.6: Informativeness of sampling locations using $\mathcal{I}^{ARV}$.

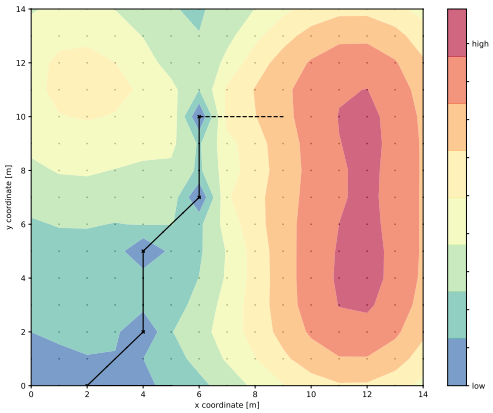

FIGURE 13.7: Informativeness of sampling locations using $\mathcal{I}^{MI}$.

13.5 Summary

In this chapter, we formally defined the AMPPES and introduced a range of planning approaches for solving this problem that is summarized in Table 13.7. These approaches include different strategies for the main components of solution concepts that we have identified in Chapter 11: The planning heuristic, the decision scope, and the environment representation.

We proposed approaches based on two planning concepts: First are model-based variants that use an informativeness measure to estimate the benefit of potential sampling locations. The second one is a geometric approach that seeks to distribute the available UAVs as well as possible across the target area. For model-based planning, we furthermore discussed possible decision scope variants. Besides a basic greedy heuristic, we introduced two concepts addressing longer planning horizons. In the first one, locations are selected heuristically and an optimal sequence through these locations is determined. In the other one, the problem of maximizing information gain within the considered planning horizon is solved exactly.

We furthermore proposed a wide range of environment models based on a unified modeling approach. The two core functionalities of such a model are an update step, which is used to adjust the model based on newly obtained measurements, and an evaluation method that returns some measure of informativeness for a set of candidate sampling locations, based on all previously obtained information. We sought to introduce environment models that are as varied as possible in terms of their expressiveness and computational effort. These approaches can be broadly grouped into three categories: First, we proposed models that represent the sensor range to assess where information has been collected and where it might not yet be sufficient. Second, we adapted these models such that they account for the time that has passed in between measurements as well. Third, we discussed probabilistic process models to study the effects of adaptive planning. Together with the proposed planning concepts, these models provide us with the means to study the trade-off between computational effort, adaptivity and prediction quality that is offered by these approaches, and to identify promising concepts for the AMPPES.

Environment models	Nonadaptive w.o. temporal aspects	Direct priority ($\mathcal{I}^{DIRECT}$) Disc-based ($\mathcal{I}^{DISC}$) Kernel-based ($\mathcal{I}^{KERNEL}$)
	Nonadaptive with temporal aspects	Latency-based ($\mathcal{I}^{LAT}$) Kernel devaluation ($\mathcal{I}^{KERNEL\text{-}TD}$)
	Adaptive	Maximum variance ($\mathcal{I}^{VAR}$) Avg. reduction in variance ($\mathcal{I}^{ARV}$) Mutual information ($\mathcal{I}^{MI}$)
Planning approaches	Geometric	LLoyd's algorithm
	Model-based	Greedy Sequential selection Integrated planning

Table 13.7: Summary of environment models and planning strategies.

14 Planning and simulation framework for online emergency surveillance

This chapter introduces the framework developed in this thesis for implementing and evaluating the approaches proposed in the previous chapter. Section 14.1 describes a modular planning architecture, using which we can which can combine the different planning strategies and environment models. To analyze and compare these strategies, we furthermore develop a discrete-event simulation that interacts with the planning strategies. This simulation is introduced in Section 14.2.

The essential entities that define this framework are illustrated in Figure 14.1. In practice, a human operator initializes the system, defines the planning problem and is informed about the results of the mission. The deployed UAVs communicate directly with the planning framework to receive sampling locations and to send status updates and obtained information. These entities are replaced by the simulation environment in this study, which represents, e.g., UAV movement and observations. This separation allows us to define and use common interfaces for communicating with the planning approach, which allows a direct transfer of the proposed mechanism to a real-world application.

14.1 Modular planning framework

We propose a solution architecture in which we can integrate the different concepts discussed in Chapter 13, i.e., the proposed environment models and the different planning heuristics. This allows an unbiased evaluation and comparison of solution approaches.

The basic scheme is depicted in Figure 14.2. In this figure, we distinguish between two groups of components that, together, fully define the planning strategy: The representation of the environment and the planning strategy. This framework furthermore defines interfaces that can be accessed by the UAVs for updating information and determining new target locations.

The environment representation indicated on top includes three core elements:

1. It stores the given input data, e.g., priorities that characterize the target area.

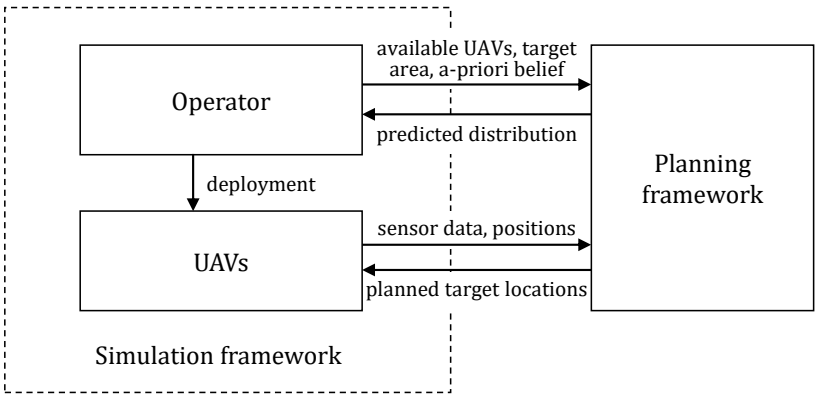

FIGURE 14.1: Central entities in the online planning and simulation framework.

2. It maintains information on the state of the deployed UAV system, e.g., the current positions of the available vehicles.

3. It includes one of the environment models discussed in Section 13.4.

The different environment models proposed in the previous chapter implement the interfaces defined in the framework, which allows us to freely exchange them against one another. They are initialized with the prior beliefs about the application scenario and target area. For example, in case of GP models, this belief describes the initial covariance matrix. Information on previously sampled locations and observed values is stored in the environment model and is updated throughout the mission.

The planning strategies evaluate these environment models in order to decide on UAV actions. To this end, we combine two elements:

1. A module that determines the currently permitted actions based on the input and system state, i.e., based on the given range restriction or length of the planning horizon.

2. The planning heuristic itself, which selects one or several targets from the set of permitted actions based on the implemented planning strategy and the environment knowledge encoded in the process model.

The result of the planning heuristics constitutes the next targets in the UAVs' missions.

Note that not all components are always necessary for implementing the planning concepts that we study here. For example, purely geometric approaches do not require an updated process model for making decisions.

14.2 Simulation framework

In order to evaluate the proposed approaches, the solution architecture introduced in Section 14.1 is embedded in a discrete-event simulation (DES). This simulation repre-

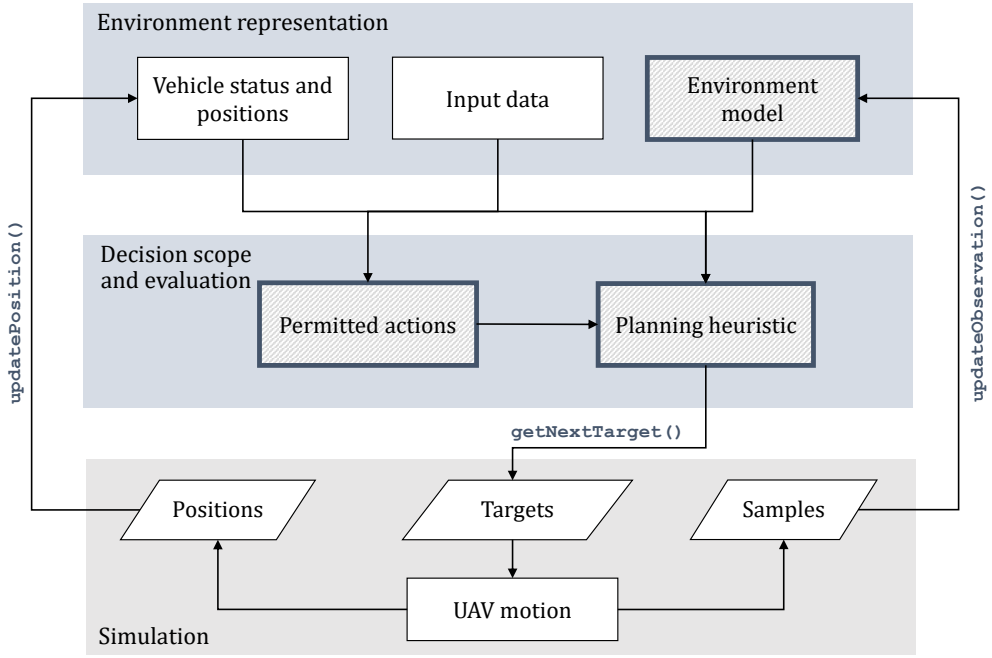

FIGURE 14.2: General online planning architecture. Elements that allow us to implement the various solution heuristics are hatched in grey. Interface functions for system status updates and target calculations are indicated in typewriter style.

sents the evolution of the system–including the UAVs and the surveyed phenomenon itself–through time. Each simulation run defines a scenario that is characterized by a population distribution, a distribution of contaminants and a number of UAVs. The situation then mirrors the movements of the UAVs through space and the information that is obtained while traveling, i.e., measurements including a possible measurement error.

In Figure 14.3, we give an activity diagram depicting the central loop of the simulation engine. A more detailed description in form of pseudocode can be found in Appendix C. The simulation represents the evolution of a planning scenario as a sequence of events, i.e., GPS position updates or plan updates, occurring at discrete points in time. It interacts with the planning architecture using the interface functions illustrated in Figure 14.2.

At the core of the simulation is a list of all pending events. Each event represents an instant in time at which the system state changes for any reason. The list of events is implemented as a priority queue sorted by event occurrence time in increasing order. The simulation progresses by removing the next event from the priority queue and resolving its effects until no events remain.

Each simulation run is initialized with complete input data on a specific scenario. A scenario is defined by the addressed target area, UAV configuration including take-off and failure times, and a generated distribution of contaminants. At the start of the simulation run, the event queue $\mathcal{Q}$ is initialized using the fixed UAV deployment and potential failure times along with the initial GPS position updates. Afterward, the simulation resolves events in increasing order of their respective occurrence time.

Event resolution depends on the type of event that occurs. In our simulation, we distinguish between five event types:

- Vehicle creation events
- GPS update events
- Vehicle failure events
- Observation events
- Events triggering an update of the planned mission.

Note that events that are concurrent in the simulation, i.e., events associated with the same point in time, are resolved in a strict hierarchy following the order indicated in the list above. The event resolution subroutines can also create new events for modeling the impact of a system state change or decision.

UAV creation events are resolved first. Otherwise, the planning strategy might receive updates from UAVs unknown to it. A new UAV is registered with the planner using the `addUAV()` function and is assumed to be ready immediately after registration. This means that a new event triggering an update of the plan of the corresponding UAV is created and inserted into the priority queue.

GPS updates and potential UAV failures are similarly announced to the planner by calling `updatePosition()` and `updateStatus()`, respectively. In case of a failure, we

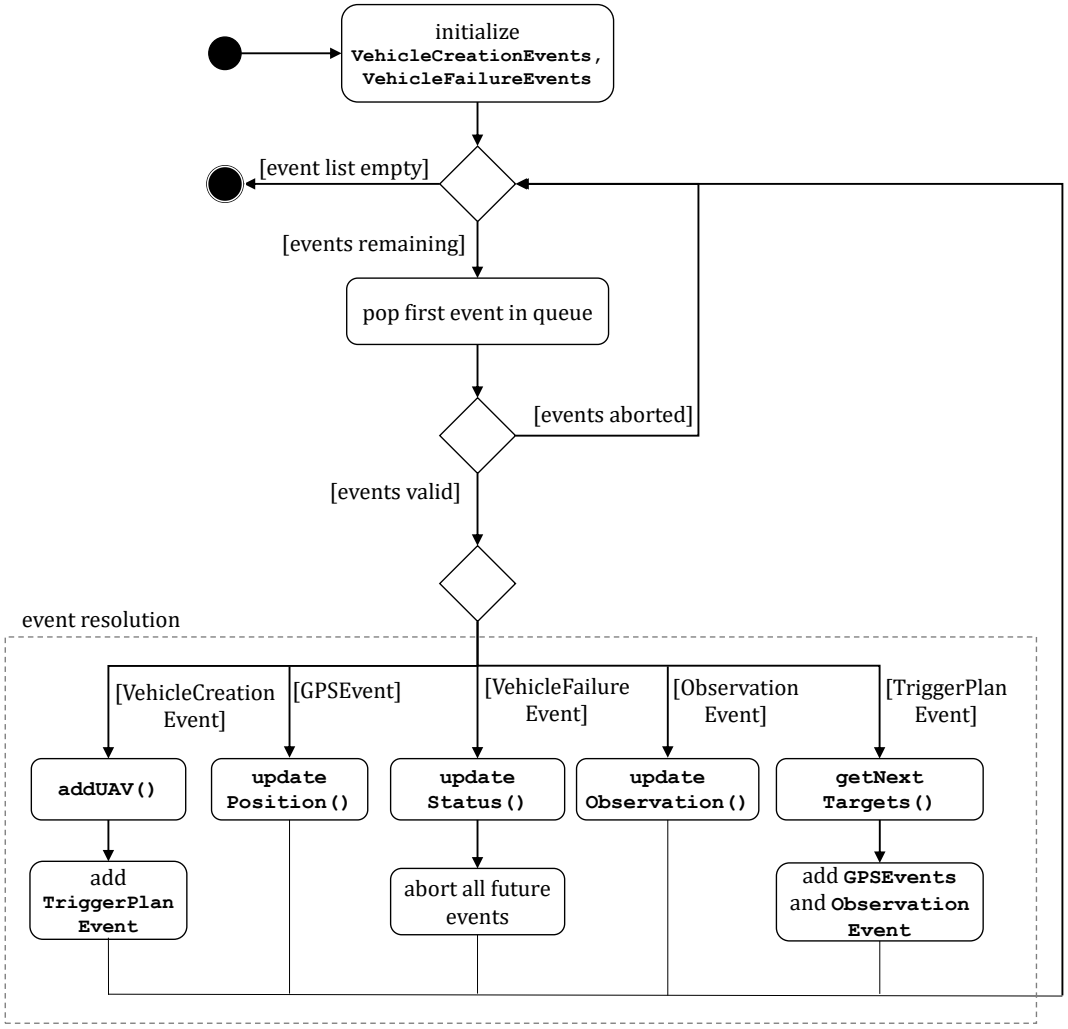

FIGURE 14.3: Activity diagram of the simulation engine. Calls to the interface functions offered by the planning strategy are indicated in typewriter style.

furthermore abort all future events associated with the UAV, as they are no longer valid. Note that aborted events are not removed from $\mathcal{Q}$, but are not resolved.

An observation event notifies the planning strategy that a new sample has been completed using the `updateObservation()` function. If desired, the simulation adds an observation error to the data.

Vehicle missions are updated last. This ensures that all information available at this point has been processed and can inform the next decision. Depending on the implemented planning strategy and decision scope, `getNextTargets()` returns one or more sampling locations for each UAV. In the simulation, we then model the UAV's trajectory as a straight line starting from its current position to the next destinations. GPS updates are created in fixed increments of time until the arrival at the target position. Furthermore, observation events are added for each completed sample. These events are inserted into $\mathcal{Q}$ together with the corresponding arrival times. Calls to the planning strategy are scheduled each time a new UAV is registered in the system and after each completed observation. This corresponds to an idle vehicle requesting new tasks.

Statistics on the performance of the planning strategy are maintained in the planning framework itself, e.g., computation time. After completing its run, the simulation engine assesses the quality of the obtained mission, e.g., by comparing predicted values with the true distribution that is unknown to the planner. For more details on these measures, we refer to the experimental study in Chapter 15.

14.3 Summary

In this chapter, we introduced the proposed planning and simulation framework that is the foundation of the study discussed in the next chapter. Together, these two components enable a comprehensive, unbiased evaluation of the planning strategies studied in this thesis: The solution architecture allows the implementation of different environment models and planning strategies within a shared planning framework. This framework offers interfaces that can be accessed by the deployed UAVs to send updates and receive new targets. These interfaces can also be accessed by the simulation engine, which imitates the evolution of the system and the surveyed process. By integrating different strategy-dependent components in the planning framework, we can compare the performance of different solution concepts in a controlled environment.

15 Simulation based assessment of online planning approaches

In this study, we explore the models and solution concepts proposed in Chapter 13 and implemented within the framework discussed in Chapter 14. Our main goal is to assess the different strategies with respect to the requirements that we have summarized at the beginning of Part III and to derive insights into the performance in terms of prediction quality and spatial coverage, the advantages, and the limitations of planning approaches.

We first introduce the instances on which this study is based in Section 15.1 and give an overview of the design and structure of the computational study in Section 15.2. In Section 15.3, we summarize the evaluation measures that are used throughout this chapter. We present the results of this study in Section 15.4. Section 15.5 concludes this chapter with a summary of insights and recommendations for the applications of the studied solution methods in practice.

15.1 Instance generation

The instances used in the evaluation of the AMPPES are based on the concepts discussed in detail in Part II of this thesis, see Section 9.1. Due to the large number of scenarios that are used in what follows, we do not generate instances for all possible combinations of input parameters—e.g., number of vehicles, instance size, or characteristics of the observed phenomenon. Instead, we use one base setting and assess the impact of changes in the instance characteristics in isolation. For each setting, we generate 20 instances to limit the effects of possible outliers.

The different parameters defining the base setting and the variations used in the following are summarized in Table 15.1. The base parameters are similar to those used for the largest instance size with an area of 2.5×2.5 km^2 and 2 UAVs in Part II. This is used as a baseline as the results for the offline case suggest that we can achieve a reliable prediction with high mission duration, while room for improvement remains (see, e.g.,

the summary of error measures in Figures 9.5 and 9.6).

The technical specifications characterizing the UAV and sensor types are the same as in Chapter 9. Variations address the target area, the available UAVs, and the surveyed distribution of gases. Similar to Section 9.1, we generate artificial distributions of contaminants by drawing samples from an unconditioned Gaussian process $Z_{\mathcal{V}}$ based on a covariance function k. To generate distributions of contaminants with different characteristics, we vary the parameters that define the covariance function. Again, we use a Matérn covariance function and vary the characteristic lengthscale l (see Rasmussen and Williams (2006) for a more detailed discussion of properties). Examples of distributions generated this way are represented in Figure 15.1. In the example depicted in Figure 15.1a, the phenomenon changes rapidly within short distances. In contrast, the distributions generated using the largest lengthscale parameter, depicted in Figure 15.1c, are much smoother. Instances in the base setting exhibit medium spatial variability with respect to the size of the target area, depicted in Figure 15.1b.

In practice, measurements are usually not completely accurate. This affects the quality of the final predictions, but can also impact the performance of the adaptive models. In order to account for these effects, we add noise to the sampled data such that we obtain noisy samples $y_i = z_i + e_i$, with $e_i \sim \mathcal{N}(0, \sigma)$. The standard deviation σ of the error term is determined as $\sigma = coe \cdot \mu$, where μ is the mean value of the true data and $coe \in \mathbb{R}^+$ indicates the error level relative to the mean value (i.e., the coefficient of variation).

15.2 Study design

In this study, we compare the planning heuristics and environment models proposed in Chapter 13. For an overview of all proposed approaches, we refer to Table 13.7. Our computational study is structured in four sections, which are summarized in Table 15.2.

In the first step, we assess the efficacy and efficiency of the proposed environment models. To this end, we combine them with the basic greedy heuristic and apply them to the instances in the base setting discussed in the previous section (see Table 15.1). This provides us with an indication of the computational effort and prediction quality that is

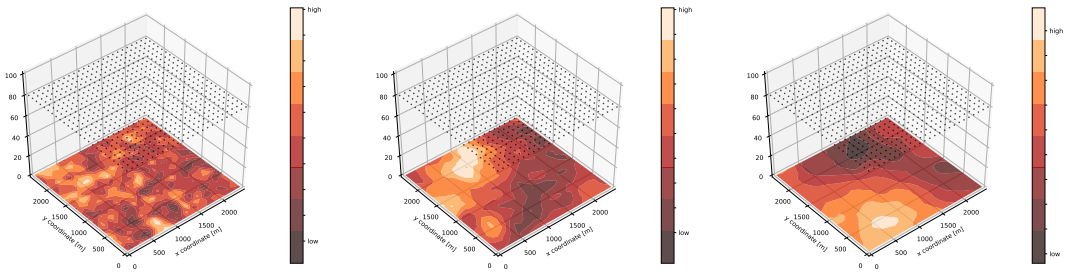

(A) Lengthscale l=100 m (B) Lengthscale l=400 m (C) Lengthscale l=800 m

Figure 15.1: Spatial distributions with varying variability.

Parameter		Base setting	Variants
Technical specification	Cruise speed v^{max} (in m/s)	7	-
	Acceleration a^{max} (in m/s^2)	2	-
	Sensing time τ_i (in s)	2	-
	Sensor resolution (pixel)	1,000 $\times$ 1,000	-
	Flight height h (in m)	100	-
	Ground sampling resolution *pxr* (in m^2/pixel)	0.01	-
Target area and resources	Available vehicles $\|\mathcal{M}\|$	2	1 to 5
	Starting locations	Random at border	-
	Mission duration T^{max} (in s)	1,200	600 to 1,800
	Target area sizes (km^2)	2.5 $\times$ 2.5	2.5 $\times$ 2.5 to 5 $\times$ 5
Process	Generator process kernel	Matérn covariance	-
	Spatial variability (lengthscale l)	Medium ($l = 400$ m)	100 m to 800 m
	Error level *coe*	0.05	0.0 - 0.2

TABLE 15.1: Parameter settings and instance characteristics for AMPPES study.

associated with these different approaches. We furthermore study the impact of the initialization of these models, i.e., the impact of the coverage radius and the weights that characterize nonadaptive models and the initial belief in the case of adaptive variants.

In the second step, we address the question of adaptive planning. Again, we compare the performance of all proposed environment models using a greedy approach. In contrast to the first step, we do not vary the configuration of the models but modify the characteristics of the surveyed phenomenon and the measurement error level. This way, we can assess the potential benefit of adaptive models over nonadaptive variants.

In the third step, we focus on the scalability of the proposed approaches. To this end, we use instances that are larger than those in the base setting, addressing a larger target area or a larger number of UAVs. We compare the greedy approach and Lloyd's algorithm to assess the scalability of model-based planning and geometric variants, respectively. Note that we do not consider model-based approaches with a larger decision scope. These are analyzed separately in the fourth part of this study. In this final step, we compare the greedy planning heuristic against the sequential selection of targets and the integrated approach that solves subproblems exactly. This way, we can assess the cost of increasing the decision scope and gain insights into the main reasons for potential improvements, i.e., either the improved sequencing of a heuristically selected set of locations or the optimal computation of routes.

The configurations of the planning approaches considered throughout this study are summarized in Table 15.3. Again, we define a base configuration for all models and planning heuristics. This initial configuration is based on the results of Part II: For nonadaptive planning, we use the same weights $w_{ij}, i, j \in \mathcal{V}$ for the GCorTOP model and a similar range for the covering radius d^{cover} of disc-based variants. These configurations have shown promising results for the offline case. Most importantly, they provide

Focus	Models	Planning concepts	Aim	Section
Efficacy, efficiency, robustness	All	Greedy, geometric approach	Comparison of environment models on a common basis, impact of model initializations and initial beliefs	Section 15.4.1
Adaptivity	All	Greedy	Comparison of adaptive and nonadaptive models, dependency on phenomenon characteristics and sampling error	Section 15.4.2
Scalability	All	Greedy, geometric approach	Performance on larger scenarios (target area size and number of UAVs), comparison of model-based and geometric planning	Section 15.4.3
Decision scope	All	Greedy, sequential selection, integrated planning	Assessment of costs and benefits of larger decision scopes, indication of the main contributions to potential solution improvements (selection or efficient sequencing)	Section 15.4.4

Table 15.2: Structure of the computational study.

a reasonable approximation of the distances at which the surveyed phenomenon exhibits positive spatial correlation in the base setting for the instances with a medium lengthscale of $l = 400$ m. To enable a fair comparison with the nonadaptive models, the adaptive models are similarly initialized with the correct information about the surveyed process, i.e., with the correct lengthscale.

We vary these parameters to assess the robustness of the approaches with respect to erroneous model configurations. To this end, both the adaptive and the nonadaptive variants can be initialized such that they either assume that the process either varies much more quickly or is much more smooth.

We furthermore study the impact of different moving planning horizons $\Delta\mathcal{T}$ and fixed decision horizons $\Delta\mathcal{T}^{fix}$. The moving planning horizon, i.e., the time window for which approaches with larger decision scopes determine a mission, is up to three minutes long. This is the limit of what is computationally feasible, based on preliminary experiments.

Parameter	Base configuration	Variants
Initial belief in adaptive models (lengthscale l in m)	400	100 - 800
Coverage radius d^{cover} (m) for coverage models	300	100 to 800
Baseline weight $\bar{w}$ for kernel-based models	0.5	-
Baseline distance d^{min} (m) for kernel-based model	100	-
Move restriction (m) for greedy variants	400	-
Moving planning horizon $\Delta\mathcal{T}$ (s)	n.a.	0 to 180
Fixed decision horizon $\Delta\mathcal{T}^{fix}$ (s)	n.a.	0 to 90

Table 15.3: Algorithm configurations for applied environment models and heuristics.

The fixed decision horizon, which indicates for how long this mission is executed until a new plan is determined, varies between 0 and 90 seconds, i.e., between an immediate replanning and an execution of large parts of the planned mission. Note that in all cases, we ensure that $\Delta\mathcal{T}^{fix} \leq \Delta\mathcal{T}$.

15.3 Performance measures for online planning

The first set of measures used in this chapter are those introduced in Section 9.2 for assessing the quality of solutions. First among them is the $PCov_d$ value (Equation (9.4))

$$PCov_d(\mathcal{S}) = \frac{\sum_{i\in\mathcal{V}:\exists j\in\mathcal{S} \text{ with } d_{ij}\leq d} u_i}{\sum_{i\in\mathcal{V}} u_i},$$

which determines the percentage of priorities within distance d to a sampling location. Here, $\mathcal{S}$ indicates the set of all sampling locations taken throughout the mission. We use the $PCov_d$ value as a first indication of how many highly prioritized targets are covered and how well the samples are spread across the target region. It formalizes the visual impression that an operator may have when looking at a planned mission.

Second, we use the weighted mean absolute error (WMAE) (Equation (9.7))

$$\text{WMAE} = \frac{1}{\sum_{i\in\mathcal{V}} u_i} \sum_{i\in\mathcal{V}} u_i \cdot |\widehat{Z}(i) - z_i|,$$

which evaluates the quality of the obtained prediction weighted with the priority of target locations. This directly represents the overall goal of the planning approaches of achieving high quality predictions.

To compare the performance of the solution approaches on a common basis, we evaluate the prediction quality using the same GP model for all strategies. This means that we use the same kernel function k_θ for all settings. The obtained samples $\mathcal{Y}_H$, which depend on the applied strategy H, are then used to compute the GP posterior based on the fitted kernel function $k_{\hat{\theta}}$. The quality of a prediction, i.e., the remaining uncertainty and the prediction error, is measured based on this posterior. To compare results over different instances and instance sizes, all results are based on normalized error measures, indicating the prediction error relative to the prediction quality at the beginning of the mission (indicated as rel. WMAE in what follows).

Following Dunke and Nickel (2016), we use counting distribution functions to compare the performance of different solution strategies in an online setting. Counting distribution functions allow the comparison of solution methods across a larger set of problem instances Λ. To this end, we denote the solution obtained using an algorithm configuration H as ω_{H}. The value for a solution measure f for a solution ω to instance $i \in \Lambda$ is denoted $f(\omega, i)$. The performance ratio of two configurations H_1 and H_2 depending

on the problem instance i is then defined as:

$$\zeta_{\mathrm{H}_1,\mathrm{H}_2}(i) = \frac{f(\omega_{\mathrm{H}_1}, i)}{f(\omega_{\mathrm{H}_2}, i)}. \tag{15.1}$$

In this chapter, we use WMAE as solution measure f, but other measures are equally possible.

To compare the two approaches H_1 and H_2 across the entire range of instances Λ, the counting distribution function $\mathcal{F}_{\mathrm{H}_1,\mathrm{H}_2}(\zeta)$ of the performance ratio ζ is then defined as:

$$\mathcal{F}_{\mathrm{H}_1,\mathrm{H}_2}(\zeta) = \frac{\sum_{i\in\Lambda} \mathbf{1}_{(-\infty,\zeta]}(\zeta_{\mathrm{H}_1,\mathrm{H}_2}(i))}{|\Lambda|}, \tag{15.2}$$

with indicator function

$$\mathbf{1}_{(-\infty,\zeta]}(\zeta') = \begin{cases} 1 \text{ if } \zeta' \leq \zeta \\ 0 \text{ otherwise.} \end{cases} \tag{15.3}$$

That is, $\mathcal{F}_{\mathrm{H}_1,\mathrm{H}_2}(\zeta)$ indicate the proportion of instances with a performance ratio of ζ or less. As discussed by Dunke and Nickel (2016), this provides additional information about the global performance of algorithms over all instances while indicating the variability of approaches and best-case and worst-case ratios.

An illustrative example is given in Figure 15.2. Note that, for the WMAE measure, a heuristic H_1 improves over an alternative approach H_2 if the WMAE is lower, i.e., if $\zeta_{\mathrm{H}_1,\mathrm{H}_2}(i) < 1$. If $\zeta_{\mathrm{H}_1,\mathrm{H}_2}(i) > 1$, H_2 yields the better result. In the example indicated in Figure 15.2, $\mathcal{F}_{\mathrm{H}_1,\mathrm{H}_2}(1)$ is approximately 0.75, i.e., in around 75 % of considered instances, H_1 has achieved a lower WMAE. Furthermore, we can see that while there are instances in which H_2 improved over H_1, this improvement is minor relative to the worst case outcome, i.e., the possible loss in solution quality in the worst case for H_2.

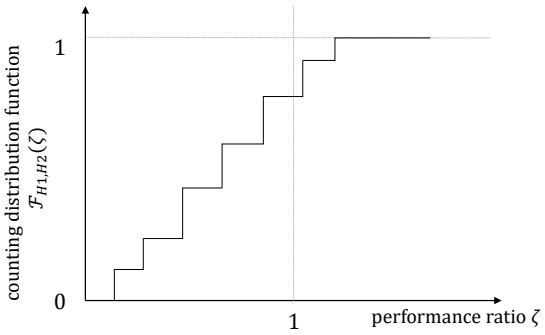

FIGURE 15.2: Illustrative example for a counting distribution function.

The computation times of the studied approaches are interpreted with respect to practical requirements. For offline planning, our project partners have estimated that a few

minutes are available before the UAVs can be deployed after arriving at the scene of an incident. In online planning, results should be available within a few seconds at the most to ensure that the mission progress is not delayed.

15.4 Computational results

The solution strategies and the simulation are implemented in Python 3.6 using `numpy` (Oliphant, 2006) and `scikit-learn` (Pedregosa et al., 2011). Note that these libraries support multi-threading, e.g., for matrix operations. To enable an unbiased comparison of computational effort, however, all experiments summarized here are run on a single core.

15.4.1 Comparison of environment models

In a first step, we compare the performance of the proposed environment models (see Table 13.7 for an overview of model variants proposed in this thesis). All results in this section are based on greedy maximization or our geometric approach. We compare the performance of approaches within the base setting characterized in Table 15.1, i.e., on target areas with a size of 2.5×2.5 km^2, a distribution with medium spatial variability and 2 UAVs. The models are initialized using the parameters given in Table 15.3 such that they fit the observed process relatively well. This section aims at providing a first comparison of the performances of adaptive and nonadaptive models, and to assess their robustness concerning the model parametrization.

Evaluation on base setting

In the following paragraphs, we compare the performance of the different models based on instances in our base setting using a greedy planning heuristic. We also include the geometric approach (Algorithm 7). Figure 15.3 summarizes the performance measures over all instances in this setting. The time per computation, i.e., the average computation time per call to the planning approach, is given on the left. Here, the geometric approach as well as $\mathcal{I}^{DISC}$ and $\mathcal{I}^{LAT}$ come with the lowest computation time, with 1-10 ms per call. The other models yield results in around 1 s or less, except for $\mathcal{I}^{MI}$ as the computationally most expensive one. Considering the practical limitations, the computation time of $\mathcal{I}^{MI}$ is still acceptable, as the impact relative to the travel time from one location to the next is low.

With respect to the $PCov_{300}$ measure, the nonadaptive models $\mathcal{I}^{DISC}$ and $\mathcal{I}^{KERNEL}$, which both directly address spatial coverage, perform best, with around 85 % to 90 % of priorities covered this way. The adaptive models $\mathcal{I}^{VAR}, \mathcal{I}^{ARV}, \mathcal{I}^{MI}$ result in a mean value for $PCov_{300}$ that is more than 10 % lower than the results obtained with nonadaptive models. Furthermore, these models lead to a higher variance in the results.

Finally, we look at the relative WMAE, indicated on the right-hand side of Figure 15.3. For all models, we observe a relatively high variance in the results. The general level of

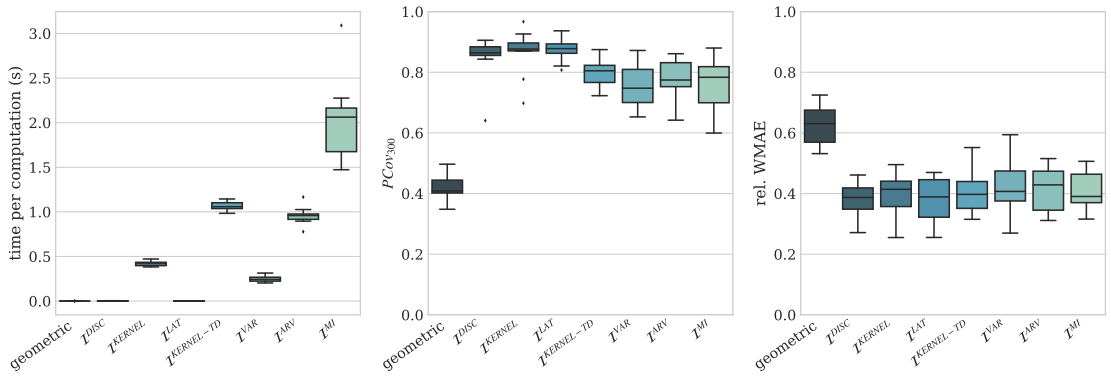

FIGURE 15.3: Comparison of modeling approaches on base settings: Computation time per iteration (left), $PCov_{300}$ (middle) and relative WMAE (right).

uncertainty in the prediction can be attributed to the nature of the created distributions, which are relatively coarse (see Figure 15.1b). Except for the geometric approach, the performances of the different models are similar. Minor improvements can be achieved by nonadaptive models over adaptive variants. We attribute the good performance of nonadaptive models to the excellent spatial coverage that they can achieve. Note that the three adaptive models yield similar results for WMAE even though they come with lower spatial coverage.

These effects are highlighted in three examples represented in Figure 15.4. The disc-based coverage model $\mathcal{I}^{DISC}$, in Figure 15.4a, leads to the survey of large parts of the target region. Even though the mission trajectories cross one another, sampling locations in general are well distributed. The mission illustrated in Figure 15.4b, based on $\mathcal{I}^{VAR}$, explores a smaller region. The figure furthermore illustrates a particular problem of this environment model: The objective is to visit locations where the current uncertainty is high. In this example, this leads to a focus on the border region, as the environment model does not account for the fact that these samples are less relevant for the remainder of the target region. This also explains why the prediction quality indicated in Figure 15.3 for $\mathcal{I}^{VAR}$ shows major outliers.

In contrast, $\mathcal{I}^{ARV}$, as an alternative adaptive model, focuses much more heavily on areas further away from the border of the target region. These locations provide more information about the area of interest than the ones selected in the previous example. This increases overall robustness relative to $\mathcal{I}^{VAR}$.

Finally, we asses the evolution of the information gain provided by the different approaches in Figure 15.5. The figure shows the relative WMAE as a function of the time that has passed since the beginning of the mission. The basic geometric approach achieves the best results within the first 200 s, as UAVs move from their takeoff location toward the inner regions of the target area. After this initial phase, the UAVs are close to weighted centers of their respective partitions and do not travel much further, which means that less information is obtained, leading to a major loss in prediction quality.

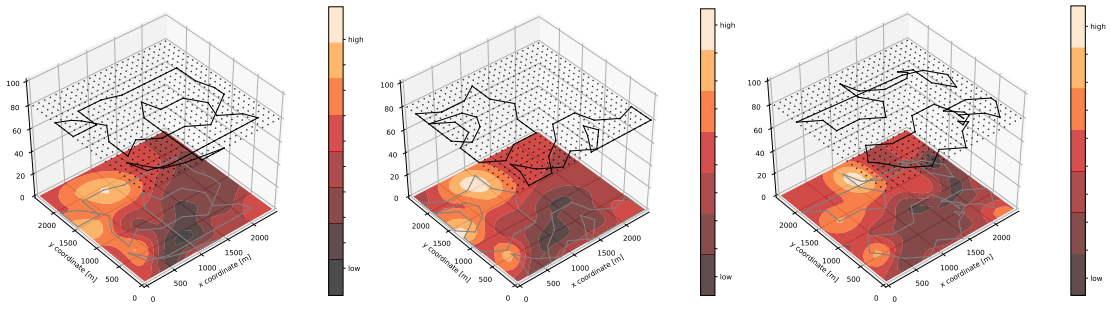

(A) $\mathcal{I}^{DISC}$ (B) $\mathcal{I}^{VAR}$ (C) $\mathcal{I}^{ARV}$

FIGURE 15.4: Missions obtained with different modeling strategies: $\mathcal{I}^{DISC}$ (nonadaptive), $\mathcal{I}^{VAR}$ and $\mathcal{I}^{ARV}$ (adaptive), superposed on the predicted distribution.

This problem is avoided in the different model-based variants. Again, we observe that the difference between the environment models is low. Overall, the best performance is achieved by the $\mathcal{I}^{DISC}$ and $\mathcal{I}^{KERNEL}$ models. We also see that the rate of improvement decreases as with increasing mission duration. At this point, large parts of the target region have been covered and the overall extent of the distribution can be mapped with reasonable accuracy. To improve the prediction further, samples would have to be taken in much closer density, which is not feasible within the restricted mission duration.

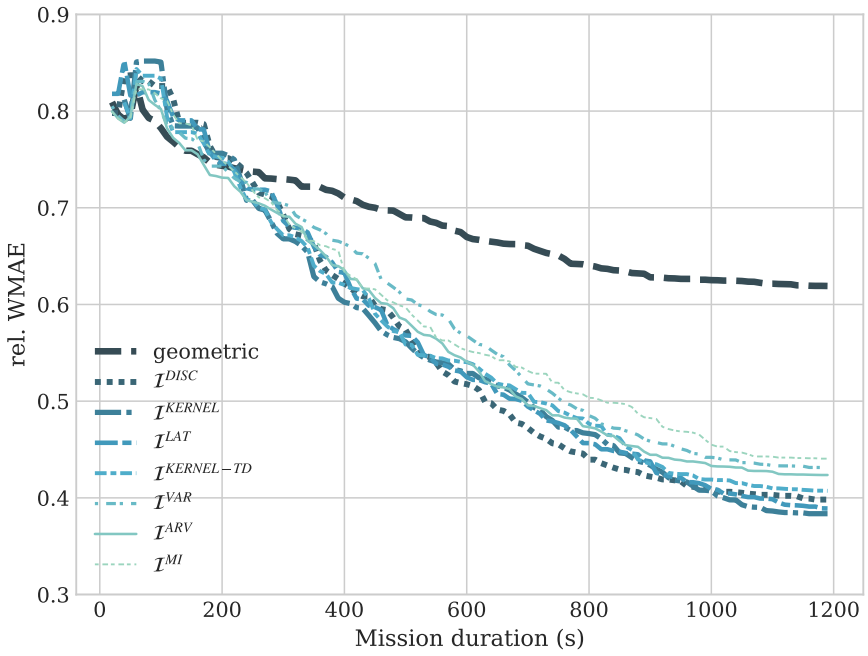

FIGURE 15.5: Evolution of WMAE relative to the mission duration, averaged over all instances. Note the scale on the vertical axis.

Impact of coverage radius in nonadaptive models

The results discussed up to this point have been achieved using a value for the coverage radius d^{cover} in nonadaptive models which fits the "true" spatial correlation in the data relatively well. In practice, we cannot always assume that this is the case. Therefore, we evaluate the impact of changes to the parametrization of the nonadaptive models, given a fixed spatial variability in the process.

The results for the four nonadaptive models, $\mathcal{I}^{DISC}$, $\mathcal{I}^{KERNEL}$, $\mathcal{I}^{LAT}$ and $\mathcal{I}^{KERNEL\text{-}TD}$, are summarized in Figures 15.6 and 15.7. Figure 15.6 indicates the relative WMAE that can be achieved for each model and value for d^{cover}. We can observe an impact of parameter variations on prediction quality for all models except $\mathcal{I}^{LAT}$. Errors are particularly high for low values for d^{cover}, where both prediction quality and the stability of the results are severely impacted by the changed parametrization. Overestimating d^{cover} also has a negative impact on results but to a lesser extent. $\mathcal{I}^{LAT}$ is affected by these parameter changes as well. However, the impact especially of underestimating d^{cover} is lower when applying this model. This is due to the consideration of the temporal dimension: As the time between subsequent samples is taken into account in addition to the spatial proximity, the influence that d^{cover} has on the results decreases. For a similar reason, the impact of d^{cover} on $\mathcal{I}^{KERNEL\text{-}TD}$ is less than the effects on its time-independent counterpart $\mathcal{I}^{KERNEL}$.

In Figure 15.7, we indicate the corresponding coverage $PCov_{300}$. Again, we can see a major impact of the parameter changes, with the best results obtained for values for d^{cover} that are close to the real range of the correlation in the observed data. Discrepancies are particularly high for $\mathcal{I}^{DISC}$ and $\mathcal{I}^{LAT}$. Here, increasing d^{cover} means that large parts of the target area are considered as "covered" after performing a sample. The models, therefore, have difficulty to distinguish between areas where a sufficient number of samples have been taken and other regions where there is a lack of information. The more flexible models $\mathcal{I}^{KERNEL}$ and $\mathcal{I}^{KERNEL\text{-}TD}$, in contrast, can account for the fact that coverage is low at high distances to the sampled location, which means that the impact of increasing d^{cover} diminishes. However, they still perform much worse for low values for d^{cover}. The results suggest that, even though these models can achieve very high performance, they are less robust to changes in the input parameters. Most importantly, their average performance is impaired by underestimating d^{cover}, as they do not achieve sufficient spatial coverage of the target area and do not have the means to compensate for this lack of exploration.

Impact of initial belief in adaptive variants

Similar to the nonadaptive models, the adaptive variants are initialized with a first belief about the nature of the spatial correlation in the process. This is done in the form of the GP prior, which depends on the initial parameterization of the applied kernel function. Again, we seek to assess the impact that an erroneous initial belief can have on the outcome of the mission. Therefore, we vary one parameter characterizing the prior belief, and assess the changes this variation has on the obtained results. In this study,

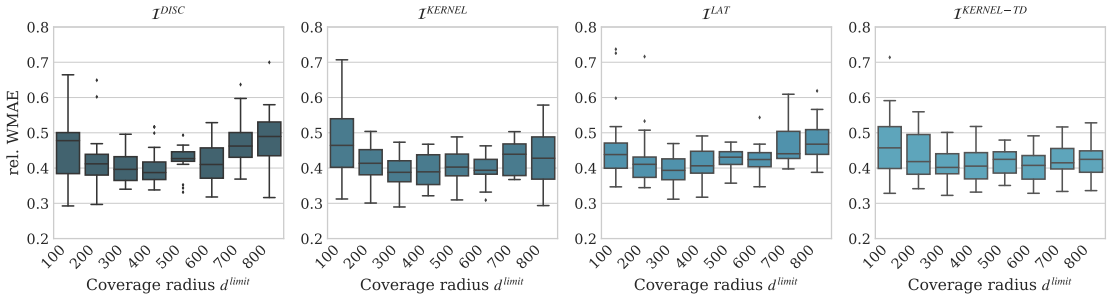

FIGURE 15.6: Impact of coverage radius d^{cover} on WMAE for nonadaptive environment models. Note the scale on the vertical axis.

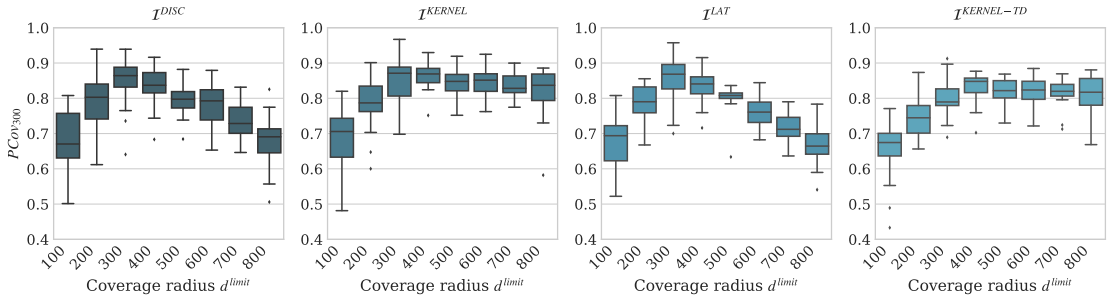

FIGURE 15.7: Impact of coverage radius d^{cover} on achieved coverage $PCov_{300}$ for nonadaptive environment models. Note the scale on the vertical axis.

we use the lengthscale l that, similar to d^{cover}, represents the extent of the correlation. The lengthscale used for creating the synthetic distributions of gases is $l = 400$ m in the configuration of the base instances which are considered in this section. We vary the lengthscale of the initial belief, i.e., the initial configuration of the GP represented in the environment models, with values ranging from $l = 100$ m to $l = 800$ m.

The impact of the initial belief on the prediction quality for the three adaptive models $\mathcal{I}^{VAR}$, $\mathcal{I}^{ARV}$ and $\mathcal{I}^{MI}$ is presented in Figure 15.8. We can see that, in contrast to nonadaptive models, the results do not depend much on the initial parametrization. Instead, the informativeness measure that is used has the largest effect, with $\mathcal{I}^{VAR}$ generally leading to higher prediction errors and $\mathcal{I}^{MI}$ yielding the best results. This pattern corresponds to our results in the base setting, depicted in Figure 15.3. It suggests that the adaptive models are successfully adjusting to the true correlation in the process.

Similar observations can be made when looking at the coverage measure $PCov_{300}$ in Figure 15.9. Again, the results do not achieve the same quality as the nonadaptive models in the base setting and show relatively high variability. However, they also do not depend

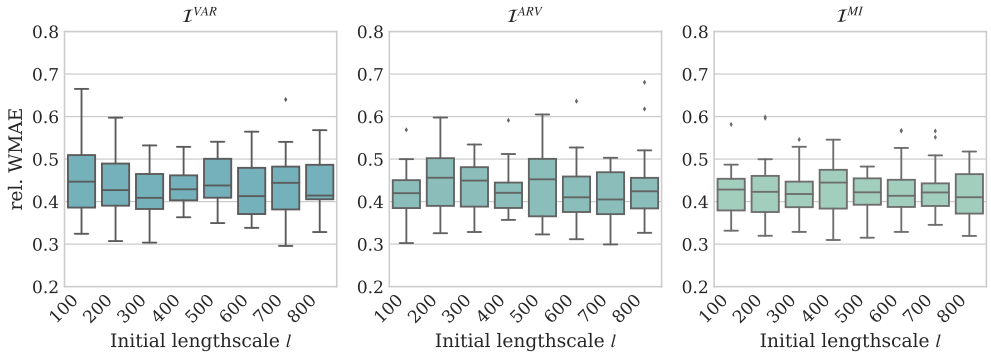

FIGURE 15.8: Impact of lengthscale l in the initial belief (Gaussian prior) on WMAE for adaptive environment models. Note the scale on the vertical axis.

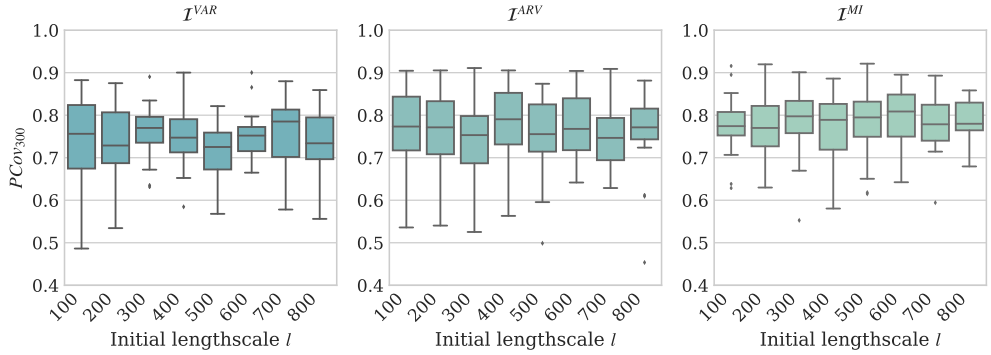

FIGURE 15.9: Impact of lengthscale l in the initial belief on achieved coverage $PCov_{300}$ for nonadaptive environment models. Note the scale on the vertical axis.

as much on the initial parameters. While it appears that larger lengthscales l lead to a slight decrease in the variability of the achieved coverage in the case of $\mathcal{I}^{ARV}$ and $\mathcal{I}^{MI}$, the overall level remains roughly the same. This allows avoiding the major loss in coverage and prediction quality that can be observed for nonadaptive variants.

15.4.2 Characteristics of the surveyed phenomenon

In this section, we assess to which degree the performance of the different models depends on the characteristics of the surveyed phenomenon. In particular, we are interested in studying in which situations adaptive models yield improvements over nonadaptive variants. To this end, we vary the characteristics of the surveyed phenomena and evaluate how well models perform, depending on these settings.

Spatial variability

First, we give the results on the achieved prediction quality, depending on the variability within the surveyed distribution of gases. Figure 15.10 represents the WMAE for three different settings: Left, we give results for instances with particularly high spatial variability. The figure in the middle summarizes results on instances with medium variability, while smooth phenomena are represented on the right. Note that the examples illustrated in Figure 15.1 share these characteristics.

The variability in the process has a major impact on the quality of a prediction that can be achieved within a limited amount of time. This effect is consistent over all modeling variants and is particularly obvious for processes with high spatial variability. Here, the available time is not sufficient for taking samples in the density that would be required for making more accurate predictions about these and similar processes.

However, the results on instances with high variability indicate that in these cases, adaptive planning approaches can improve prediction quality. While the predictions never achieve the same level of accuracy as is possible for processes with lower variability, $\mathcal{I}^{VAR}$, $\mathcal{I}^{ARV}$ and $\mathcal{I}^{MI}$ can decrease the WMAE by up to 10 % compared to nonadaptive variants. In contrast to the base instances discussed above, $\mathcal{I}^{VAR}$ performs best among these three models for these scenarios. This pattern changes as the spatial variability decreases. In these cases, nonadaptive planning concepts achieve a similar or even better performance, as can be seen on the right-hand side.

Figure 15.11 summarizes the impact that these changes have on the achieved coverage. As can be seen on the left-hand side, the improved prediction quality of the adaptive models come with a major decrease in the covered priorities. Again, the effect is particularly obvious for $\mathcal{I}^{VAR}$. This indicates that all three models correctly identify and adjust to the high variability, leading to samples that are located much more closely together. This enables a more accurate prediction, but the limited mission duration only allows surveying a relatively narrow region. However, the high variability in the process also means that little information can be inferred about the remainder of the area based on the sampled location. This might be the reason that $\mathcal{I}^{MI}$, which explicitly models the

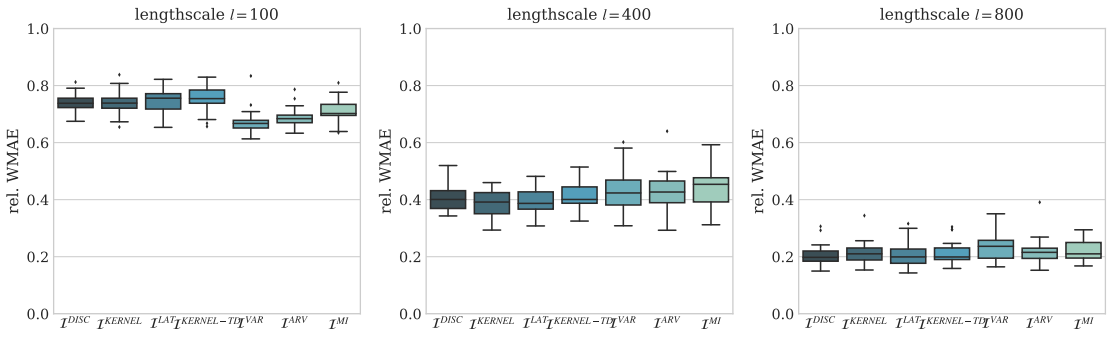

FIGURE 15.10: Impact of spatial variability on WMAE for different modeling approaches, depending on the spatial variability of the process.

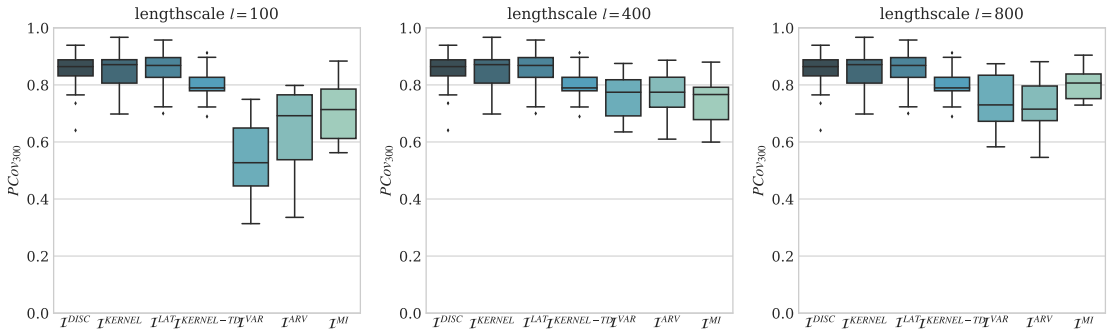

FIGURE 15.11: Impact of spatial variability on $PCov_{300}$ of different modeling approaches.

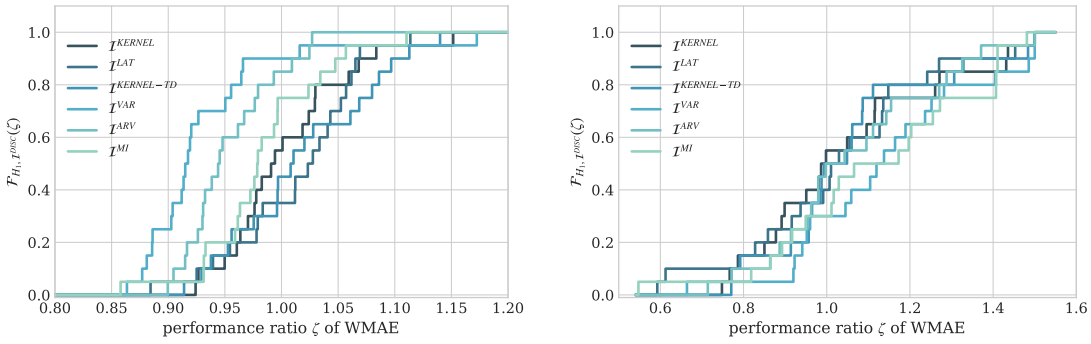

(A) High variability (lengthscale l=100 m)

(B) Low variability (lengthscale l=800 m)

FIGURE 15.12: Empirical counting distribution function $\mathcal{F}_{H_1,\mathcal{I}^{DISC}}(\zeta)$ of the performance ratio of WMAE for different models and levels of spatial variability. Note the different scales on the horizontal axis.

value of the inferred information, performs worse than the other adaptive variants. The coverage obtained by the nonadaptive variants is independent of the lengthscale in the model, as they do not provide the possibility to adjust to this information.

To obtain a better understanding of the relative performance of these approaches, depending on the variability in the process, we summarize the results using the counting distribution function for the WMAE measure. Figure 15.12a shows the results on instances with high variability for all environment models relative to $\mathcal{I}^{DISC}$, which we have chosen as a benchmark as it is the most basic model that considers spatial interdependencies. Note that $\zeta_{\mathrm{H},\mathcal{I}^{DISC}} < 1$ indicates that an approach H performs better than $\mathcal{I}^{DISC}$.

We can see that $\mathcal{I}^{VAR}$ is not only better on average, but performs better on 90 % of instances. $\mathcal{I}^{ARV}$ also improves over $\mathcal{I}^{DISC}$, but not to the same degree. Considering the performance of the other nonadaptive variants relative to $\mathcal{I}^{DISC}$, no conclusive result can be obtained. We observe the opposite behavior for instances with low variability, depicted in Figure 15.12b. Here, adaptive variants tend to decrease performance. This loss in prediction quality can be high, especially in contrast to the possible gains in Figure 15.12a.

An interesting observation is that the outcome in case of high process variability is more stable. In particular, the worst-case performance, i.e., the maximum deterioration relative to $\mathcal{I}^{DISC}$, is much more critical in case of a lengthscale $l = 800$ m. However, we believe that this is not due to the models themselves, but results from the nature of the studied scenarios: A high process variability means that the available flight time is not sufficient to obtain the number of samples that would be needed to provide an accurate prediction about the spatial distribution of gases. Hence, for all applied models, the achieved prediction quality is low. This limits the potential losses due to the application of a wrongly configured model.

Figure 15.13 illustrates these effects at the example of two missions planned using $\mathcal{I}^{VAR}$. The instance with high spatial variability, illustrated in Figure 15.13a, shows that samples are selected close to one another, which leaves large parts of the area unsurveyed. In contrast, the selected sampling locations for the scenario with low variability, which are shown in Figure 15.13b, are spread relatively wide apart. Furthermore, we can observe that $\mathcal{I}^{VAR}$ leads to samples at the border of the considered region. This not only limits the information that can be inferred about the phenomenon within this region but also means that the coverage of priorities diminishes, confirming the effects summarized in Figure 15.11.

Impact of observation noise

In addition to the nature of the phenomenon itself, information gain is also influenced by the observation error. In this study, this is modeled as white noise added to the obtained measurement, with a standard deviation that depends on an error level *coe* relative to the mean of the distribution. Values for *coe* have been determined in preliminary tests by increasing the value such that the effects are noticeable, but meaningful predictions

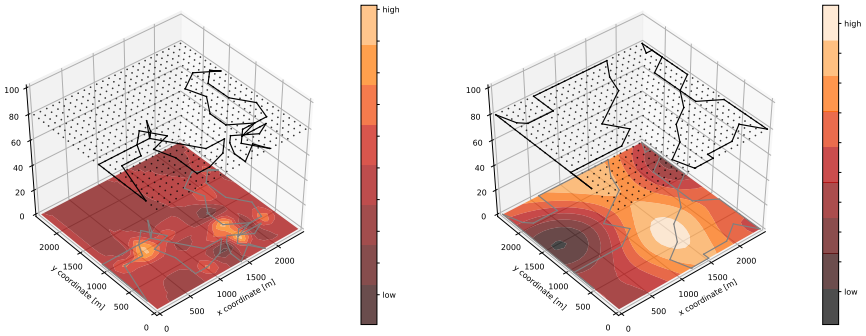

(A) High variability (lengthscale l= 100 m) (B) Low variability (lengthscale l = 800 m)

FIGURE 15.13: Missions based on $\mathcal{I}^{VAR}$ for instances with varying spatial variability.

are still possible on the noisy data. Note that, in all models, increasing the observation error impacts the final prediction quality. In the case of adaptive models, this factor can also influence the planned missions.

Figure 15.14 summarizes the prediction error for four models, depending on the error level. On the left-hand side, we give results for $\mathcal{I}^{DISC}$. This serves as a baseline for the other models, as the selected sampling locations with and without noise are identical. This means that all changes in the prediction quality can be directly attributed to the errors in the sampled data. The other three figures give results for the three adaptive models, where these errors can impact the missions.

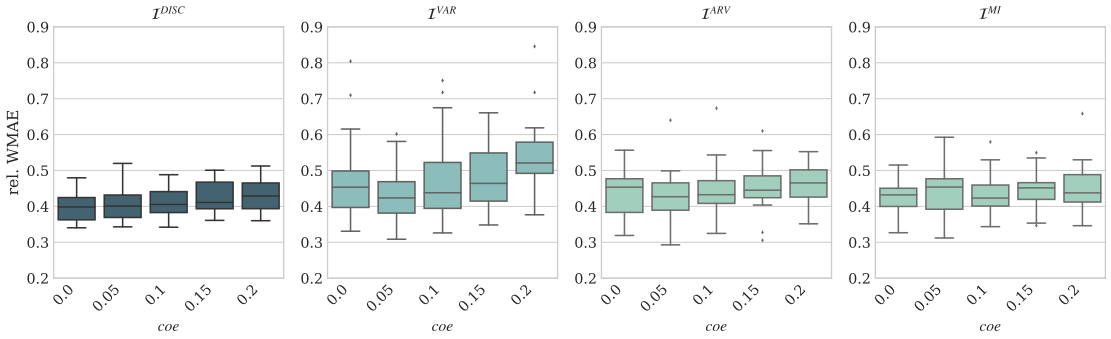

FIGURE 15.14: Impact of observation noise on prediction quality for different environment models. Note the scale on the vertical axis.

$\mathcal{I}^{VAR}$ again comes with the highest prediction error and is much less stable with respect to increases in the noise level than the other models. Most likely, increasing the observation noise means that the posterior covariance used for the prediction increases over all locations. This makes it more difficult for $\mathcal{I}^{VAR}$ to identify promising sampling candidates. The other two variants, in contrast, prove to be relatively stable with respect to the sampling error. Prediction quality is impacted in all cases, but the effect is not larger than in the case of $\mathcal{I}^{DISC}$.

15.4.3 Scaling up

In this section, we focus on the scalability of our modeling approaches. To this end, we increase the target area size from 2.5 × 2.5 km^2 to 5 × 5 km^2, i.e., to instances with up to 2,500 target locations. We furthermore include up to 5 UAVs.

Target area size

We summarize the results depicting the effect of increasing the target area in Figures 15.15 and 15.16 for scenarios with 3 UAVs. Figure 15.15 summarizes the prediction error for all considered models and target area sizes. The corresponding computation times are given in Figure 15.16. We again include the geometric approach as a benchmark for the model-based variants.

As expected, increasing the size of the considered area without adjusting the resources that are available for the survey negatively impacts the reliability of the prediction. Counterintuitively, the figure also shows that the prediction quality of the geometric approach can slightly improve as the target area size increases. We attribute this to the fact that in small-scale scenarios, UAVs quickly reach the centers of their respective partitioning of the target area and do not move much afterward. In larger scenarios, they have to travel further, thus spending a larger proportion of the available mission time actively surveilling the area. Additionally, as the target area size increases, achieving a balanced distribution of sampling locations over space becomes more important. This aspect is directly addressed by the geometric approach.

For all model-based approaches summarized in Figure 15.15, increasing the size of the target area comes with a consistent loss in prediction quality. The results do not vary significantly between the different scenarios and model variants. For the largest area, the WMAE of all models is close to the one achieved by the geometric approach.

The computation times show that the geometric approach is robust with respect to increasing target area sizes. Similarly, $\mathcal{I}^{DISC}$, $\mathcal{I}^{LAT}$ and $\mathcal{I}^{VAR}$ scale well due to their capability to very quickly evaluate the informativeness associated with candidate sampling locations. The computational effort of $\mathcal{I}^{KERNEL}$, $\mathcal{I}^{KERNEL\text{-}TD}$ and $\mathcal{I}^{ARV}$, in contrast, increases with the target areas size, reaching 5 to 10 s per computation for the largest scenarios. Considering the limits on flight and computation time, these computation times are no longer acceptable. $\mathcal{I}^{MI}$, which has been the slowest model in the base setting, becomes prohibitively expensive for larger areas. Here, computation times for determining the next sampling location can far exceed the travel time that would be necessary to move there.

Number of UAVs

One possibility to deal with the increasing size of the target area is to deploy more UAVs. For the largest scenarios with a target area size of 5 ×5 km^2, the impact of the number of UAVs on the prediction error is depicted in Figure 15.17. With one UAV, all approaches achieve a similar prediction quality. This includes the geometric approach. However,

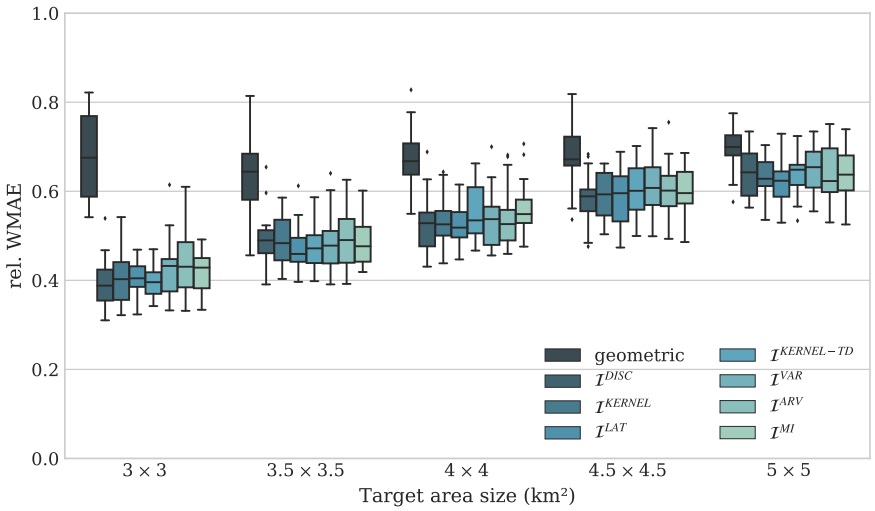

FIGURE 15.15: WMAE depending on target area size.

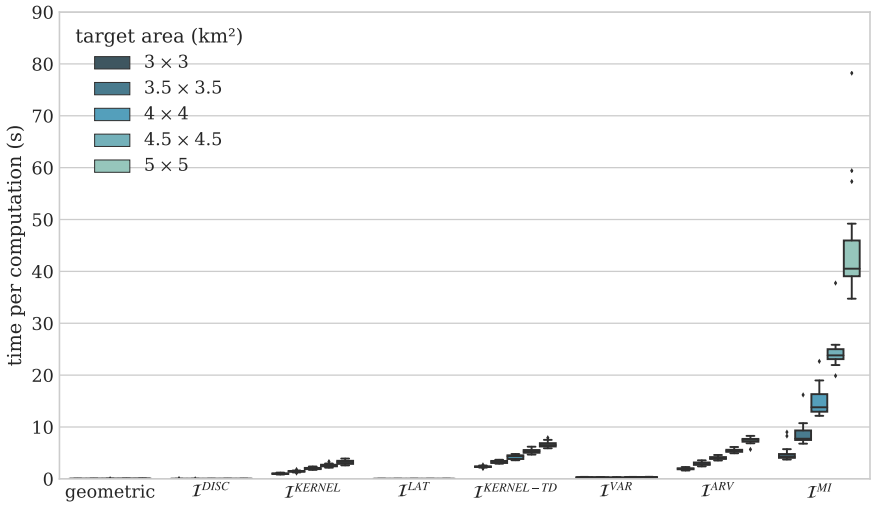

FIGURE 15.16: Computation time for different modeling strategies, depending on target area size.

rather than demonstrating that the geometric variant performs well, this indicates that all approaches fail at obtaining a reliable prediction for these scenarios. As we have expected, increasing the number of vehicles consistently helps to achieve lower prediction errors. The incremental benefit of additional vehicles is largest for the first few but diminishes when the number of UAVs is increased further.

This effect is also illustrated in the example scenario given in Figure 15.18. We show a large-scale instance with a simulated distribution of contaminants as indicated on the left and depict two missions, one using 2 UAVs and the other using 5, planned using the $\mathcal{I}^{MI}$ criterion. Here, 2 UAVs are not sufficient to cover the entire area, leaving room for further improvements. With 5 UAVs, however, the distribution of contaminants can be approximated reasonably well. The predicted distribution is too smooth compared to the ground truth indicated on the left, but otherwise, it closely mirrors the true data. To better represent the coarse structures would require samples to be taken in much closer density than is possible in this scenario.

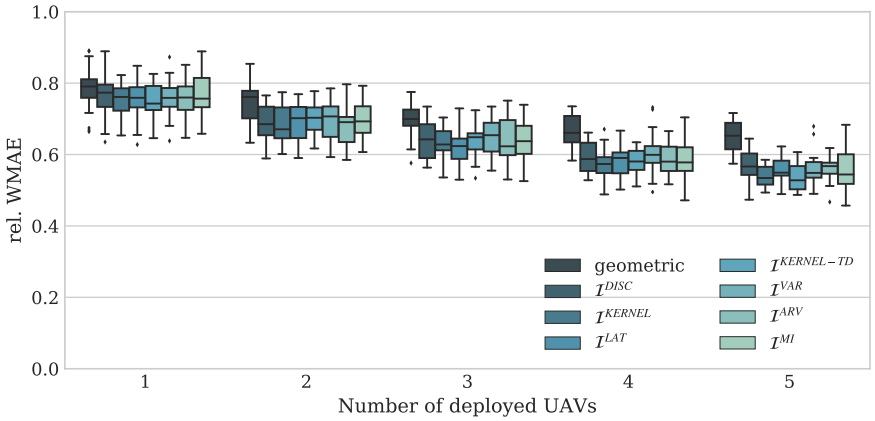

Figure 15.17: WMAE depending on the number of deployed UAVs.

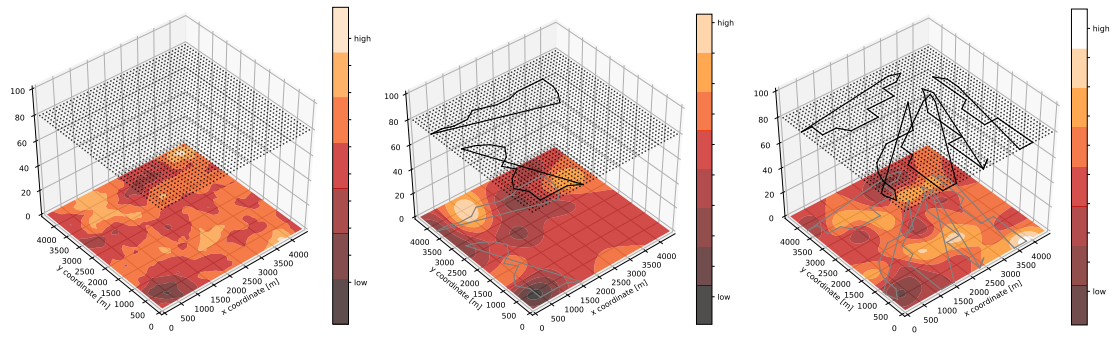

(A) True distribution (B) Prediction based on 2 UAVs (C) Prediction based on 5 UAVs

Figure 15.18: Missions and predictions obtained for a large-scale scenario ($4.5\times$ 4.5 km^2) for $\mathcal{I}^{MI}$ with 2 and 5 UAVs.

15.4.4 Decision scope analysis

After studying the environment representations in the previous sections in a greedy context, we now move on to larger decision scopes. We introduced these concepts in Section 13.3.1. Specifically, we proposed two different variants: First is a sequential selection heuristic where a set of sampling locations $\mathcal{S}_t$ is selected heuristically and a TSP is solved to determine the order in which these locations are visited. Second is an integrated approach where optimal solutions are determined for a moving planning horizon $\Delta\mathcal{T}$. Both variants are embedded in a receding horizon framework, where decisions are fixed for time $\Delta\mathcal{T}^{fix} \leq \Delta\mathcal{T}$. This means that the computed mission is not necessarily executed in its entirety within the moving horizon $\Delta\mathcal{T}$, but updated after a time $\Delta\mathcal{T}^{fix}$ has passed. This means that the larger planning horizon is only used to assess the impact of decisions, but the mission itself can adjust more quickly based on new information.

We tested both strategies on instances in our base setting, i.e., with two UAVs, a target area size of $2.5\times$ 2.5 km^2, and medium spatial variability. We combined them with all proposed environment models. For the sequential selection heuristic, we considered planning horizons with a length of up to 180 s, i.e., up to 10 % of the total flight time. The integrated variant, which is computationally more expensive, has been run with planning horizons of up to 90 s.

In Figures 15.19 and 15.20, we summarize the relative WMAE of all runs for the two planning approaches. Note that we set the fixed horizon $\Delta\mathcal{T}^{fix}$ to have the same length as the moving planning horizon for this evaluation. The effect of changing the $\Delta\mathcal{T}^{fix}$ is studied later in this section. The sequential selection heuristic is depicted in Figure 15.19. Contrary to our assumption that the selection and sequencing of multiple targets can lead to more efficient missions, thereby decreasing the prediction error, we observe that prediction quality nearly always stagnates or decreases with increasing length of the planning horizon. This is the case for both the nonadaptive models as well as the adaptive variants.

The integrated planning approach, summarized in Figure 15.20, demonstrates the opposite effect: Increasing the decision scope improves overall prediction quality with one exception ($\mathcal{I}^{LAT}$). This suggests that the most important aspect of larger decision scopes is the improved combination of sampling locations. We can observe this, for example, for $\mathcal{I}^{KERNEL}$ and $\mathcal{I}^{MI}$. Both strategies offer relatively complex representations of the interdependencies in the target area. The integrated approach can use this information to select sampling locations that work well together. In the sequential selection heuristic, the selection of targets is done in a heuristic fashion, which does not sufficiently account for the joint information gain of the sampling locations. As the decrease in solution quality demonstrates, the sequencing of targets alone does not provide additional benefits over a greedy approach.

The computation times associated with these results are given in Figures 15.21 and 15.22. In the first image, we give the results for the sequential selection heuristic. As expected, increasing the length of the planning horizon also increases the effort that is required for each computation. The most important factor is the time needed to

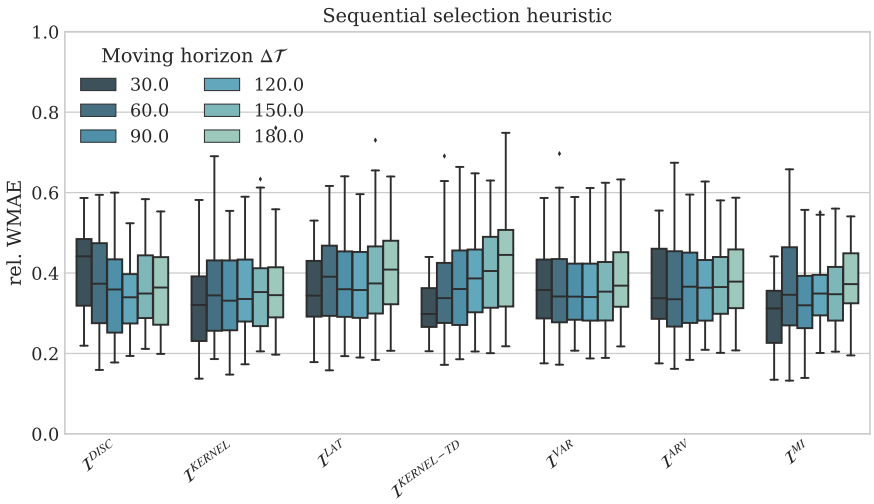

FIGURE 15.19: Relative WMAE for the sequential selection heuristic depending on the length of the planning horizon.

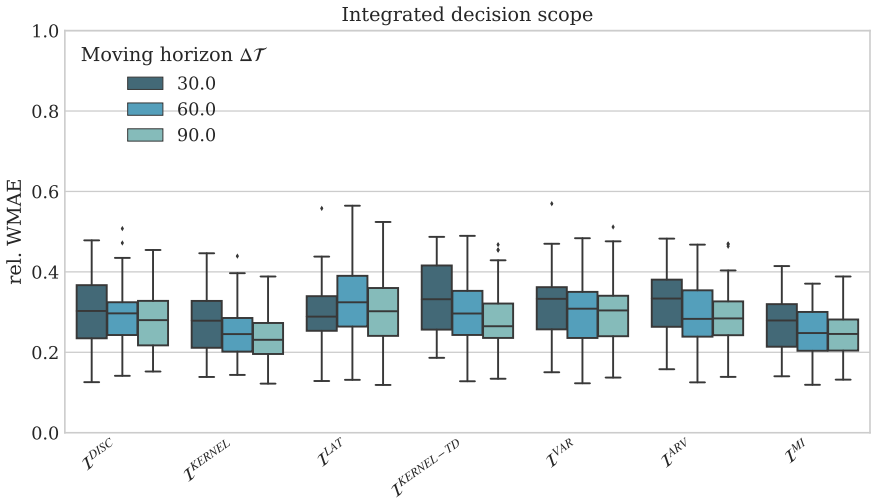

FIGURE 15.20: Relative WMAE for integrated planning depending on the length of the planning horizon.

evaluate the informativeness of a set of candidates: For models where this evaluation is inexpensive, e.g., for $\mathcal{I}^{VAR}$, which can be evaluated in $\mathcal{O}(|\mathcal{S}_t|)$, the increase in computation time is negligible. For more expensive models, however, even the sequential selection heuristic becomes prohibitively expensive for longer planning horizons, with computation times exceeding 10 s for every individual call.

The integrated solution approach, represented in Figure 15.22, shows an even more pronounced impact of the length of the planning horizon on the computation time. Here, computation times grow exponentially with the length of the planning horizon. Even for the less expensive models, e.g., $\mathcal{I}^{DISC}$ or $\mathcal{I}^{KERNEL}$, computation times quickly exceed the acceptable limits of a few seconds at most. Increasing the length of the planning horizon even more would mean that total computation time exceeds the available flight time.

To obtain a better understanding of the relative performance of different decision scopes and planning horizon lengths, Figure 15.23 provides more detailed results for two environment models, $\mathcal{I}^{KERNEL}$ and $\mathcal{I}^{MI}$, where the effects are particularly pronounced. The figure represents the counting distribution function of the prediction error relative to a greedy solution approach.

First, we look at the sequential selection heuristic, represented in the left column. For both environment models, we can see that the best results are obtained for shorter moving planning horizons $\Delta\mathcal{T}$. As the lengths of the horizon increases, the performance deteriorates further. For a planning horizon of more than 100 s, it performs worse than the greedy heuristic on virtually all instances. The fixed horizon $\Delta\mathcal{T}^{fix}$, represented in the form of different line styles, confirms this result: The longer this horizon is, the worse the performance of the heuristic approach becomes. The only improvements over a simple greedy approach that can be observed are obtained using the shortest possible horizon lengths of 30 s. Even then, the worst-case performance far outweighs these potential gains.

The right-hand column of Figure 15.23 summarizes the same information for the integrated decision scope. In accordance with our earlier observations, we can see the opposite effect compared to the sequential selection planning strategy: For integrated planning, the best results are obtained for longer planning horizons, especially when decisions are fixed. This confirms our assumption that the joint selection of several locations that together provide valuable information about the phenomenon is the main driver of improved prediction quality in these models.

15.5 Discussion and insights

To conclude this study, we summarize and discuss the results with particular attention to the requirements stated for AMPPES in Chapter 10. We then state the core insights and recommendations for the application of these approaches in practical settings.

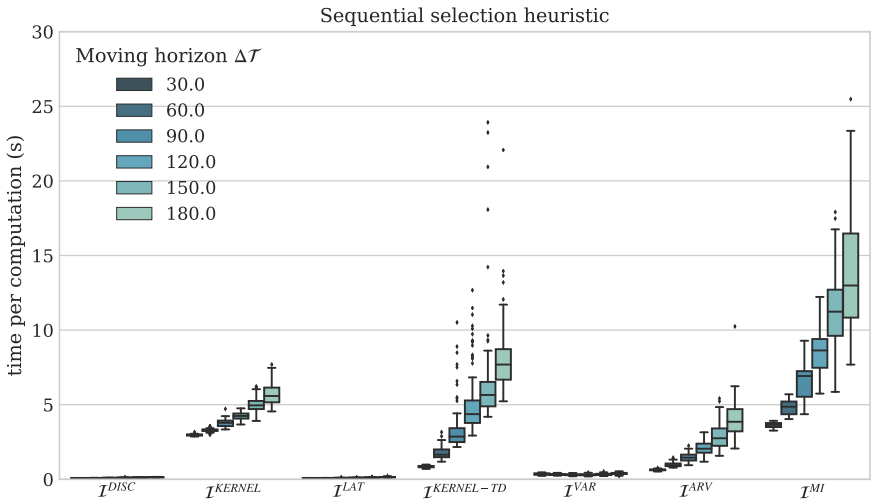

FIGURE 15.21: Computation time for sequential planning depending on the length of the planning horizon.

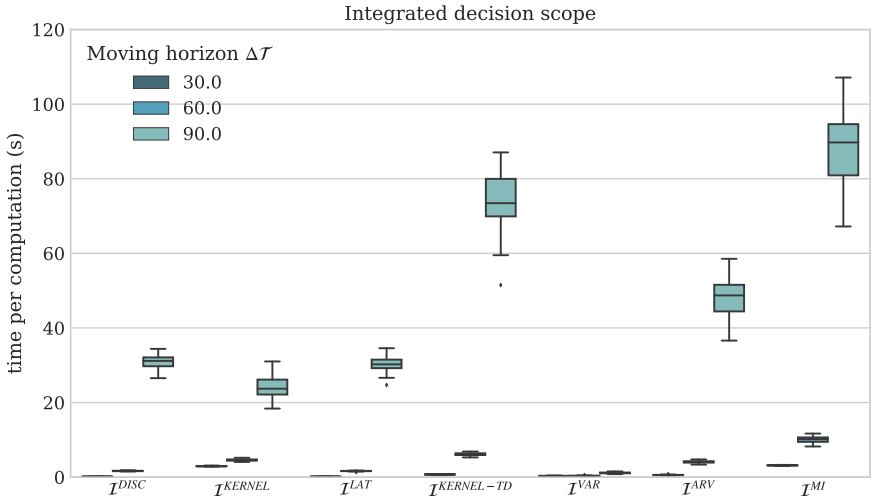

FIGURE 15.22: Computation time for integrated selection heuristic depending on the length of the planning horizon. Note that the scale of the vertical axis differs from the previous figure.

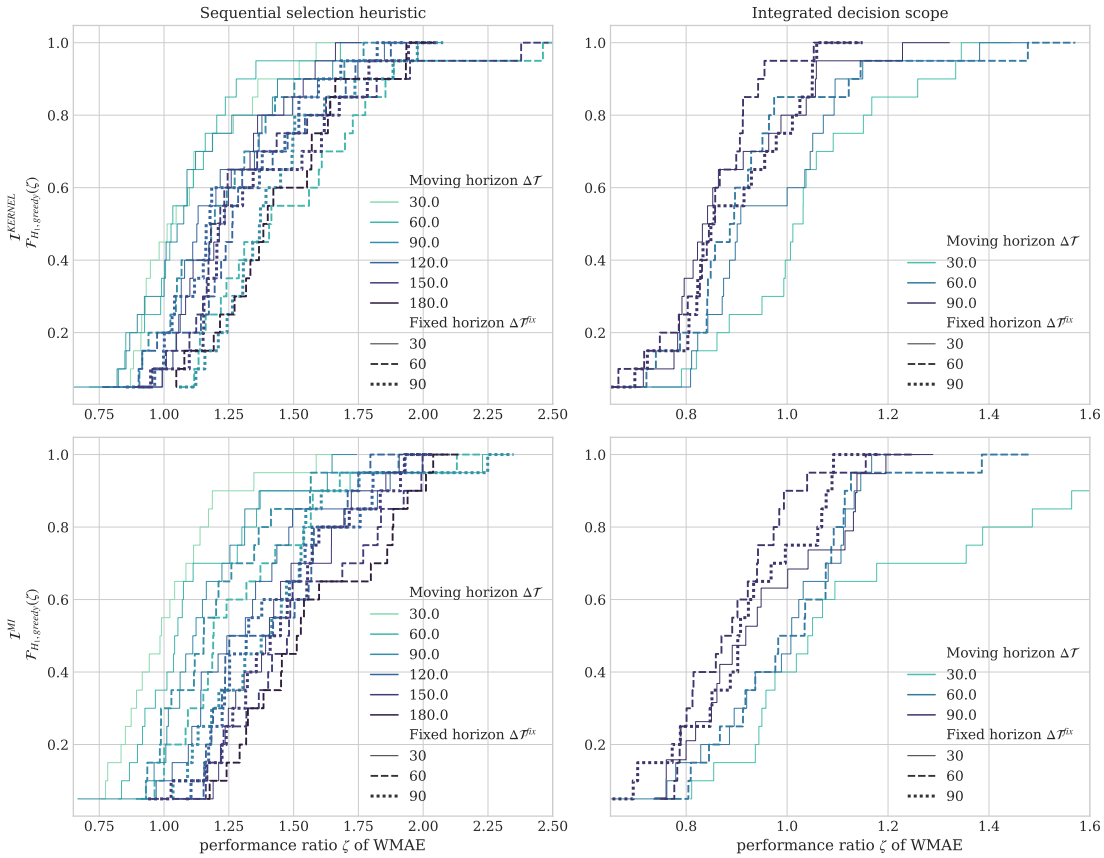

Figure 15.23: Empirical counting distribution functions of the performance ratio of WMAE relative to a greedy approach. Results for $\mathcal{I}^{KERNEL}$ are indicated in the first row, results for $\mathcal{I}^{MI}$ are given in the second one. Results for the sequential selection heuristic are represented on the left and for the integrated solution approach on the right in each row.

Environment models

To assess the proposed environment models, we applied them to a large range of instances using a greedy planning heuristic. We furthermore studied the impact of erroneous assumptions in the model configurations, i.e., assumptions about the range of the correlation that do not match the surveyed phenomenon.

The results demonstrated that the simple coverage-based approaches can, in many settings, achieve very good results at very low computation times, making them promising candidates for environmental surveillance considering the practical limitations on computation and flight times. In particular, the study confirmed that the main source for prediction quality is the selection of sampling locations that are well distributed within the considered area, an effect that we have already observed in the evaluation of the offline approaches in Section 9.3.

The main disadvantage of these nonadaptive models that we observed in our study is their inability to correct erroneous initial assumptions and parameter settings. Here, the basic coverage model $\mathcal{I}^{DISC}$ reached its limits as it does not provide the means to compensate for incorrect assumptions. The kernel-based variants $\mathcal{I}^{KERNEL}$ and $\mathcal{I}^{KERNEL\text{-}TD}$ were more stable with respect to wrong estimates of the correlation, as they provide a more natural representation of spatial coverage. However, all nonadaptive approaches were sensitive to an underestimation of the range of the correlation, which can come with a major loss in coverage and prediction quality.

Adaptivity

To assess the benefit of adaptive planning, we applied the different environment models to a set of instances with varying spatial characteristics, ranging from very coarse structures to smooth distributions of gases. This analysis showed that adaptive planning is beneficial in certain settings. Specifically, adaptive models provide the means to correct initial beliefs, which increases their stability over nonadaptive variants. However, they come at a much higher computational effort that limits their applicability to AMPPES. Moreover, we observed that the nonadaptive coverage-based yield reasonable results in many settings, despite their inability to adust to the obtained measurements.

The comparatively poor performance of adaptive modeling variants was especially surprising considering that the benefit of these variants has been shown for other applications. To give one example, the benefit of using $\mathcal{I}^{MI}$ as an informativeness criterion has been shown in sensor placement problems (e.g., Krause et al., 2008; Stranders, 2010). We attribute the different outcome in our study to two reasons: First, the additional routing constraints that we have to consider in the AMPPES limit the flexibility available for selecting sampling locations, which means that the full information of probabilistic models cannot be used. Second, a general challenge in the AMPPES is that we face tight constraints on the available resources for sampling, and the overall number of sampled locations is limited. This increases the importance of spatial coverage, which is more naturally represented in the nonadaptive variants. This also explains the good perfor-

mance of the geometric solution approach in the early stages of the mission.

Decision scopes

The decision scope variants–i.e., the considered planning horizon as well as the way decisions are made within this horizon–are studied by applying two approaches, the heuristic selection of targets and the optimal solution of subproblems. The results indicated that increasing the decision scope is not beneficial in all cases. In particular, we showed that the sequential selection heuristic, which involves a heuristic selection of sampling locations and a separation of the selection and the routing decisions, does not improve prediction quality. In contrast, the exact solution of subproblems generally resulted in lower prediction errors. This means that the main disadvantage of the greedy and the sequential selection heuristics are their inability to consider the joint benefit of several sampling locations together. The improved sequencing of selected targets and the resulting potential gains in the efficiency of the planned mission do not show improvements over simple greedy planning approaches.

Consequently, the largest improvements due to the integrated decision scope have been achieved in combination with the most complex environment representations–i.e., $\mathcal{I}^{KERNEL}$ among the set of nonadaptive variants, and $\mathcal{I}^{MI}$ as an adaptive model. However, these improvements came with computation times that prohibit their use in practical situations at this point.

Insights and recommendations

Over all instances, the difference between adaptive and nonadaptive approaches was surprisingly little during the course of this study. In other words, the advantage of adaptive planning has been mostly situational for the AMPPES. This is most likely due to two reasons: (1) If well parametrized, adaptive and nonadaptive models lead to similar sampling designs. These are characterized, for example, by the distances between sampled locations and the avoidance of areas near the border of the target region. (2) The practical limitations in emergency surveillance, especially the limited flight time, means that relatively few sampling locations can be taken. Hence, the routing constraints limit the flexibility in choosing sampling locations and, as a consequence, the potential of more sophisticated models. In practice, this often suggests using less complex approaches in many situations.

Based on this observation and the results of the computational study, we can offer the following insights and recommendations to practitioners:

- Nonadaptive modeling variants yield excellent results as long as they represent the "true" correlation reasonably well. As a consequence, if there is some prior knowledge of the correlation within the surveyed distribution of gases, e.g., due to surveillance flights at a higher altitude or the experience of the operator, these modeling variants are likely to suffice in many applications.

- Among the nonadaptive models, the gains in solution quality that can be attributed to the more detailed kernel-based representation ($\mathcal{I}^{KERNEL}$) outweigh the additional computational costs.
- The main disadvantage of nonadaptive models is their limited robustness with respect to erroneous assumptions. Including temporal aspects in the environment model helps to limit this effect to some degree, as the resulting missions are not only dependant on the assumptions about the spatial correlation.
- If there are doubts about the spatial processes, adaptive models offer a higher degree of robustness and the ability to correct wrong initial beliefs.
- Among the adaptive models, $\mathcal{I}^{VAR}$ often demonstrates reasonable results. However, the worst-case results and high variability in solution quality of this approach outweigh its advantages in terms of computation time and average-case performance.
- The $\mathcal{I}^{MI}$ measure is robust on a large case of scenarios. As long as its application is feasible for practical limitations on computation time and resources, it is the most promising model among the adaptive variants.
- If the mission time is very limited or if there is a large number of UAVs, geometric approaches offer a reliable starting heuristic to obtain information very quickly. This is similarly true if a simple rule-based approach is needed for planning missions immediately after the deployment of the UAVs before the models can be initialized. In case of longer planning horizons, it is likely beneficial to switch over to another planning strategy after the first samples have been obtained.
- If the available computation time allows planning missions within a larger planning horizons, the integrated planning variant, i.e., the optimal solution of the corresponding subproblem, should be preferred. Improving the sequence of sampling locations that are not optimal with respect to the considered planning horizon provides much less potential for the overall information gain.
- There does not seem to be an advantage in adapting the missions very quickly or very often. Instead, if there is sufficient time to plan missions for a longer moving planning horizon using the integrated approach, these missions should be executed in their entirety.

15.6 Summary of Part III

In *Part III* of this thesis, we focused on online mission planning approaches. We defined the adaptive mission planning problem for emergency surveillance (AMPPES) in Chapter 10 as the problem of combining online planning and online learning approaches so as to use the information that has been collected to improve the selection of sampling

locations while the mission is still in progress. In Chapter 11, we summarized solution concepts that have emerged in the context of environmental surveillance based on a classification scheme introduced in this thesis. In this classification, we explicitly distinguished between the applied environmental models, the planning strategies, and the decision scope that are used to determine missions. We completed this overview with a summary of online planning approaches in related domains.

In Chapter 12, we analyzed the solution concepts identified in the previous chapter with respect to the planning requirements stated for the AMPPES. We showed that the various planning concepts for environmental surveillance have not yet been compared and measured against one another. Specifically, it remained unclear whether expensive approaches–especially adaptive models, which react to obtained measurements, and heuristics with larger decision scopes–provide a benefit over easier strategies, considering the constraints on computation times in emergency surveillance.

In the remaining chapters in Part III, we sought to answer these questions, thereby identifying strategies that are applicable to the AMPPES. In particular, we focused on the trade-offs between computational effort and performance of different solution approaches, this way gaining a better understanding of when the use of more expensive adaptive models is justified over less complex strategies. In Chapter 13, we adapted existing environment representations and proposed new solution concepts to our use case. The proposed models range from simple disk-based coverage approximation to complex probabilistic models, which have been described using a unified modeling framework. We also adapted the GCorTOP introduced in Part II for online planning and introduced new strategies that take into account that information loses value over time.

To measure these models and approaches against one another on a broad and unbiased basis, we implemented all variants within a novel solution architecture. This architecture is combined with a discrete-event simulation. Both of these core components of our evaluation are discussed in Chapter 14. In contrast to existing approaches, which, with few exceptions, assess the performance of planning heuristics on a few selected data sets, this architecture allows the evaluation and comparison on a large number of scenarios. We summarized the results of an extensive simulative study in Chapter 15. We demonstrated the high performance of simple coverage-based approaches, which can achieve excellent spatial coverage and good predictions with a few milliseconds per computation. However, they are less robust with respect to wrong initial beliefs. The more detailed environment representation based on the GCorTOP introduced in Part II improves robustness and offers a well-balanced trade-off between solution time and computational effort. Adaptive models, which adjust based on the obtained measurements, offer advantages in situations with high variability in the surveyed random field and are particularly robust as they provide the means to correct wrong initial beliefs. Similarly, simple greedy planning heuristics demonstrate very good performance at low computation times. We demonstrated that increasing the decision scope can improve results as it is possible to select sets of sampling locations that, together, increase overall information gain. However, we also showed that currently, the computational effort for these variants is prohibitively high for practical applications in emergency surveillance.

Part IV

Synthesis

16 Conclusion and outlook

This thesis aimed at identifying efficient modeling and solution approaches for planning UAV missions for emergency surveillance, especially after large fires or chemical accidents. To this end, we have provided an overview of the technological and theoretical foundations for UAV-based emergency surveillance (Part I), introduced new models and heuristic solution strategies for the offline planning of informative missions (Part II), and obtained insights into the effective design of strategies for online mission planning (Part III). In this chapter, we offer a summary of our main findings and results throughout these three parts. We conclude this thesis with an outline of promising future research topics.

16.1 Summary and results

After an accident or emergency leading to the release of contaminants and hazardous substances, a quick first assessment of the extent of the contamination is crucial for the effective deployment of response units and the protection of the affected population. To this end, the distribution of contaminants can be interpolated based on a set of samples taken within the area of interest, using models that account for the spatial correlation between the contamination at different locations. In this thesis, we have developed solution approaches that enable the UAV-based rapid assessment of the distribution of hazardous substances across an area in time-sensitive situations.

The context in emergency surveillance poses specific challenges for solution approaches: The actual contribution of each surveyed sampling location in terms of prediction quality cannot be assessed in isolation. Instead, it depends on all other survey locations throughout the mission, as samples that are taken close to one another usually provide less information about the overall process compared to sampling locations further apart. This means that sophisticated models that can represent these interdependencies are needed. Moreover, the available computation time is limited due to the need to quickly provide relevant information to the emergency services. This prevents the application of detailed but expensive models for spatial interdependencies. Instead, the development of efficient approximative models and solution heuristics is necessary.

We approached these challenges from three different angles:

In *Part I*, we discussed the role of surveillance and situation assessment in emergency management and provided an overview of UAV and sensor systems as the technological basis for UAV-based surveillance. We summarized advances for interpolating and predicting spatial phenomena. Most important are probabilistic models that provide an estimation of the achieved prediction quality solely based on the locations selected for sampling, i.e., before actually starting the survey. Based on these models, several authors have proposed sampling strategies that minimize the expected prediction error. We then discussed how these models can be fitted to the collected data, which offers the possibility to correct wrong initial assumptions about the studied distribution. These overviews provided the foundation of the approaches introduced and analyzed in the remainder of this work.

In *Part II*, we introduced the mission planning problem for emergency surveillance (MPPES), which is defined as the problem of determining sampling locations for several UAVs such that the prediction accuracy at critical locations is maximal. We showed that literature in related fields has focused on either simple disc-based models for spatial coverage or on probabilistic models that are prohibitively expensive for our use case. We introduced the generalized correlated team orienteering problem (GCorTOP) as a way to model spatial interdependencies with reasonable accuracy. To solve even large scale instances with several hundred candidate target locations, we developed the two-phase multi-start adaptive large neighborhood search (2MLS). This heuristic combines a very efficient first phase targeted toward designing explorative vehicle tours with a more intensive second phase that seeks to improve these initial routes. We studied these approaches in an extensive computational study. On benchmark instances for the team orienteering problem, our approach yielded competitive results in terms of solution quality within computation times that are up to ten times faster than previously published approaches on the largest instances. We introduced a large set of benchmark instances for the MPPES based on real-world population data and simulated distributions of contaminants. Based on these instances, we highlighted the advantages of the GCorTOP in contrast to less detailed models for spatial interdependencies.

Part III has focused on the adaptive mission planning problem for emergency surveillance (AMPPES). The AMPPES is the problem of combining predictive models and mission planning approaches in an online framework that uses the trained models for selecting the next sampling locations. This promises to improve overall prediction quality, as the model can mirror the actual correlations in the surveyed distribution of gases more closely. However, it also introduces a significant computational overhead. In this work, we, therefore, aimed to determine in which situations adaptive planning approaches provide advantages over less complex approaches. To this end, we demonstrated that many existing approaches for online environmental mapping are derived from a few core concepts. Similar to the offline planning approaches in Part II, these range from simple coverage approximations to complex probabilistic variants. We adapted these concepts to our use case and introduced new concepts, notably the GCorTOP and variants that account for the time that has passed between samples in close proximity. These approaches were embedded in a simulation framework for comparing the approaches

against one another. In a study on a large number of scenarios, we demonstrated that non-probabilistic coverage models can compete with learning approaches, provided they are correctly parametrized. In contrast, adaptive variants based on trained process models outperform these variants when initial assumptions about the process are incorrect.

From a modeling perspective, we studied the trade-off between model complexity and efficiency in both offline and online planning. Overall, our results demonstrated the excellent performance of computationally inexpensive coverage-based representations, notably the GCorTOP model, relative to the more expensive probabilistic process models. However, this depends on whether or not these coverage models are correctly parametrized, i.e., whether they represent the true correlation in the process reasonably accurately. If this is not the case, the resulting missions can become detrimental as the fundamental assumptions based on which sampling locations are selected are wrong. In these situations, adaptive planning approaches based on trained process models yield superior results, as they can compensate for incorrect initial assumptions. In short, this shows that adaptive planning is not inherently better, but is more robust with respect to wrong model configurations.

Insights into what constitutes good models for spatial interdependencies also help in designing efficient solution approaches. We have exploited this in our 2MLS, which heavily promotes spatial coverage. This is also consistent with the results obtained using geometric approaches. These approaches achieve good results in the very early stages of a mission during which it is crucial to survey locations that are as widely spread across the affected area as possible. The lower computational effort associated with these approaches means that they are applicable in practice and scale to larger UAV fleets and target areas.

A similar balance between computation time and solution quality has to be found in the solution approaches that are based on these environment representations. In general, more complex solution approaches improve overall prediction quality. These approaches need to be capable of handling the interdependencies in the environment models, which means that the benefit of samples cannot be assessed in isolation, without considering the remainder of the mission. For this reason, the online approaches focusing on the greedy or sequential selection of samples can be detrimental to the overall prediction quality, as they cannot account for the joint contribution of a set of samples. We have shown that determining optimal samples within the planning horizon, in contrast, is beneficial. However, this comes at a high cost and is currently impractical for realistic scenarios. In this regard, efficient heuristics can offer a suitable compromise for time-sensitive applications.

Overall, the approaches and strategies proposed in this thesis provide a basis for the further development of decision support systems in emergency surveillance. Most importantly, the better understanding of the trade-offs between the computational cost and the obtained solution quality that are offered by the various approaches enables the development of more sophisticated methods in emergency surveillance.

16.2 Future work and outlook

We conclude this thesis by outlining promising lines of research build upon the results obtained in this thesis.

Regarding the development of planning methods for emergency surveillance, two important steps remain open for practical applications: First, the planning approaches need to be able to provide estimates about distributions that change over time, e.g., due to external factors such as wind. To this end, the proposed methods have to be extended such that they yield an estimate of the temporal correlations in the process. This also means that the duration of the missions has to be extended in order to provide a continuous assessment of the changing distribution of gases. In practice, it is necessary to assess how quickly a phenomenon can change until the provisioning of reliable information becomes impossible.

The second essential aspect is to consider the priorities within the target area as dynamic rather than treating them as static throughout the planning horizon. To give one example, it is possible to derive population maps based on mobile phone data. This way, emergency services would not depend on the quality of a-priori available information but can use up-to-date data. When the population is highly mobile, e.g., during public events or if an evacuation is in progress, this would ensure that information is always provided where it is most crucial. However, these dynamic priorities add a new dimension to an already challenging problem.

Another promising direction is the development of more sophisticated planning approaches that combine the strengths of individual concepts. A possibility for this is the integration of offline planning with online recourse actions: An initial mission can be determined quickly using simple models that incorporate the available information before the flight. During the mission, predictive models are trained using the incoming data. Only when a critical deviation between the initial assumptions and the survey results becomes revealed, is the mission updated using a corrected representation of the spatial correlations. This would increase the reliability of the outcome without the same computational costs as purely adaptive approaches. Another idea for addressing particularly large-scale scenarios is to combine a first phase based on highly efficient geometric planning approaches with later model-based strategies that exploit the information obtained during the first phase.

Given the plurality of solution approaches, a promising step is the development of automated decision support systems for configuring mission planning approaches based on the situation at hand. This would support emergency services in selecting appropriate mission planning approaches. In offline planning, such a system can then decide, e.g., between complete coverage and sampling-based approaches based on the available time and equipment. As the operation continues, it can be used to decide on the best models for subsequent surveillance flights, based on the insights obtained earlier.

Furthermore, our approaches have addressed centralized planning procedures for homogeneous UAV fleets with limited flight time. In practice, this can be extended in several ways: The rotary-wing UAVs considered here can be supplemented by fixed-wing

vehicles that allow the survey of larger areas but come with limited maneuverability. Additionally, UAVs might be recharged and deployed again, in which case later samples can improve early quick assessment. Integrating these aspects would allow emergency services to move from a rapid assessment to the monitoring and continuous surveillance of a larger affected area, thereby receiving updated information throughout the entire emergency response phase.

Finally, a centralized planning instance might not be available in emergency response. Also, communication is likely limited, and not all data and observations can be transmitted in real-time. Decentralized approaches promise to be more robust with respect to these limitations. However, this imposes additional restrictions on the computational effort, due to the restrictions of on-board processing units. Hence, it would be interesting to study the design of efficient decentralized models with lightweight communication effort based on our approaches.

A

Parameter tuning for 2MLS

This appendix gives detailed insights into the results of the parameter tuning phase for 2MLS. Specifically, this includes the starting solution concept used in Phase 1, the proposed insertion and removal strategies, neighborhood size parameters and the convergence criterion.

Aggregation and decomposition strategies

In Section 8.2.3, we have proposed two strategies for obtaining reduced problem representations. One is based on the aggregation of locations, using only selected representatives in Phase 1 of the 2MLS. The other is based on the decomposition of the target region, resulting in several single-vehicle problem representations. The iterations of the multi-start approach are limited to 4, i.e. we generate 4 different reduced problem instances for each approach. The maximum number of iterations of the ALNS to improve these solutions during Phase 1 is limited to 100.

In this section, we evaluate these two strategies based on their impact throughout the entire search. This means that we consider the quality of solutions found during Phase 2 of the search approach given the respective initial solution obtained in Phase 1. This provides a better indication of the impact these strategies have than would be possible by simply comparing the achieved initial solution.

For this purpose, we compare the convergence of the solution approach as search proceeds. Figure A.1 gives the average gap of the best found solution during search depending on the computation time. The vehicle decomposition strategy is represented as a green dashed line. The performance of the aggregation strategy is indicated in blue. For both strategies, the leftmost entry in the graph indicates the point in time at which the Phase 1 terminates as well as the corresponding solution quality. The following

points indicate average quality at 2 seconds, 5 seconds, and 10 seconds after the end of Phase 1. The second-to-last points for both strategies give the time at which the best found solution was achieved. Finally, the last point of each series indicates the termination of the search procedure.

We can see that the vehicle decomposition scheme quickly determines a very good starting solution. The average gap between the solution at the end of Phase 1 and the best-known solution is slightly higher than 4 %. Using the aggregation strategy, Phase 1 terminates more quickly, but the average gap is almost three times as high. However, this strategy is associated with a slightly higher average solution quality during the remainder of the search: We believe that this result indicates that the aggregation scheme obtains routes that are, in general, comparatively close to the best possible ones for each instance. This indicates that the strategy succeeds in anticipating the potential rewards for different regions and routes traverse the target area accordingly. Discrepancies between the initial routes and the best found ones can successfully be resolved during Phase 2.

In the case of the vehicle decomposition scheme, we can successfully achieve good decompositions and determine reasonable routes for each vehicle. However, if these routes differ significantly from the best possible routes, it is difficult for Phase 2 to correct this, as routes would have to be changed in their entirety.

Insertion strategies

We first assess the performance of the insertion and removal strategies proposed in Sections 8.2.5 and 8.2.6. As a first step, we apply the different insertion strategies for constructing starting solutions. For the sake of comparison, we also consider a random

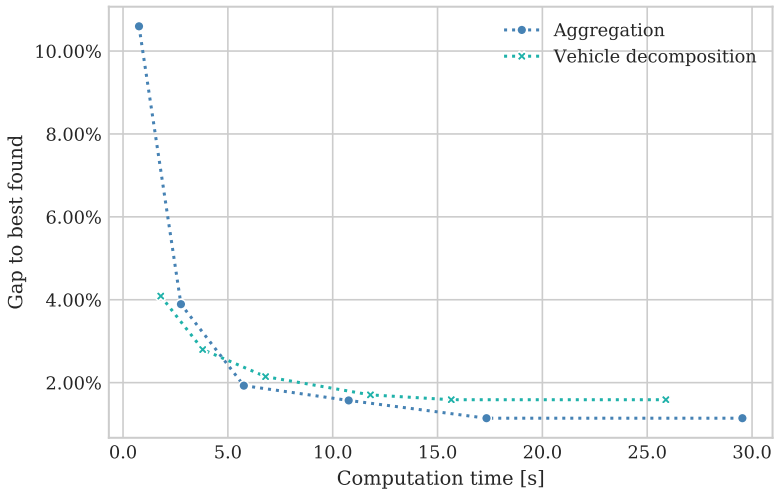

Figure A.1: Comparison of the convergence achieved with the proposed aggregation and decomposition schemes. The graph indicates the average gap during Phase 1 and Phase 2 of the 2MLS as an average of all solutions used for parameter tuning.

insertion order (Random). Vehicle routes are initialized with seeds and then completed using the respective strategies until no further improvement can be found. The lower the gap between these solutions and the best known one, the better we expect the corresponding strategy to perform during the ALNS. The results are depicted in Figure A.2 in form of a boxplot. The whiskers represent the 10^{th} and 90^{th} percentile.

For several instances, all proposed strategies yield the best known solution. Other than that, we can see that the two parametrizations of the priority-ratio strategies (indicated as BestRatio-1 and BestRatio-2 respectively) dominate in terms of average and worst-case solution quality. The cost-greedy strategy (CostGreedy) has the next best performance, even though it does not consider priority information for visit insertion. The maximum priority strategy (MaxPrio) and the region-based strategy (RegionTotal) also yield reasonable results.

Both the weighted center strategy (WeightedCenter) as well as the orienteering-regret variant (Regret-2) perform noticeably worse. In case of the WeightedCenter, this is possibly due to the fact that the weighted center is not a viable point of reference in case of an almost empty route. For Regret-2, it is possible that the preferred insertion of locations that are difficult to reach negatively impacts the routes.

To obtain a better understanding of the performance of these heuristics when embedded in the local search approach, we do not update the selection probabilities in the ALNS, thereby selecting each strategy with equal probability throughout the search. We can then compare the number of improving moves that have been achieved with each strategy. The results are summarized in Table A.1. We distinguish between two forms

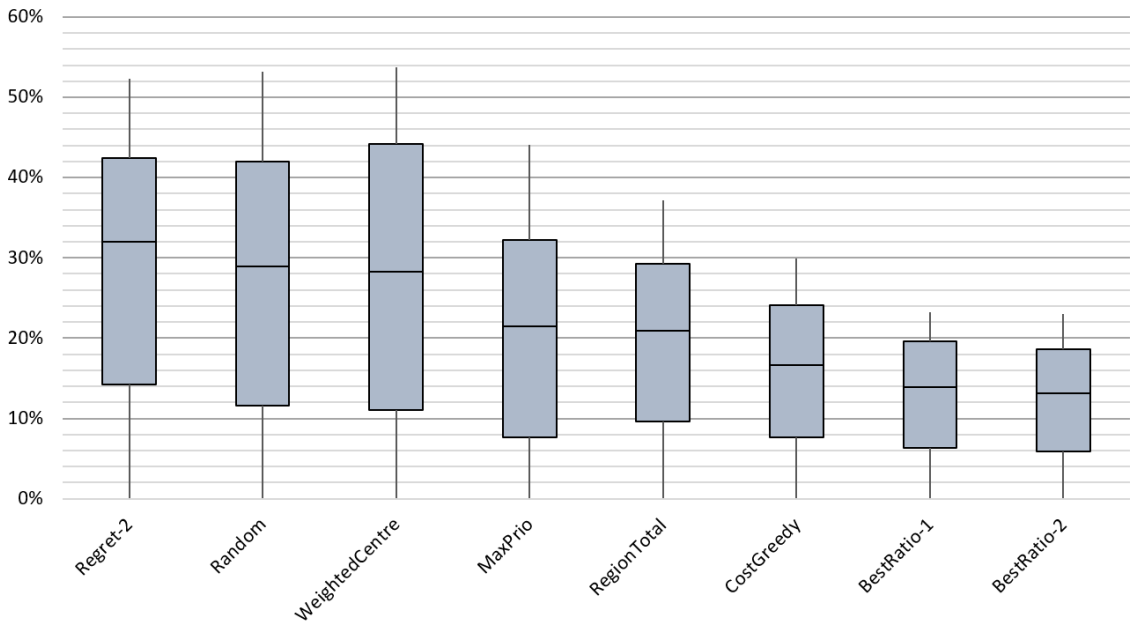

FIGURE A.2: Performance of insertion strategies: Average gap between starting solutions and best known solutions (BKS).

Insertion strategy	moves resulting in improvements ... of informativeness	... of route duration
BestRatio-2	**6.1 %**	**18.5 %**
BestRatio-1	**5.5 %**	**18.8 %**
MaxPrio	**5.2 %**	17.2 %
Regret-2	3.7 %	**21.4 %**
CostGreedy	3.5 %	18.2 %
RegionTotal	3.5 %	16.9 %
WeightedCenter	3.5 %	16.8 %
Random	3.3 %	16.4 %

Table A.1: Performance of insertion strategies: Improving moves during the ALNS relative to total number of executions per strategy. Best results are indicated in bold.

of improvement: First, we only consider the strict improvement of resulting informativeness (or the sum of priorities in case of a TOP model). The results are given in the middle column. Second, we take into account improvements in route duration with at least equal informativeness or total priority. This evaluation is indicated on the right. To facilitate the comparison, the three best-performing strategies in each category are indicated in bold. We can see that the BestRatio-1 and BestRatio-2 again perform particularly well, achieving good results in terms of both informativeness and route duration. The Regret-2 heuristic, which has performed worst in terms of average starting solution quality, yields improvements in terms of route duration. Based on these results, we select six strategies for our evaluation runs: Four of them, Ratio1, Ratio-2, CostGreedy, and MaxPrio, are selected due to their overall good performance. We furthermore choose to use the Regret-2 heuristic due to the potential for route improvement. Finally, we also select the RegionTotal strategy: While average performance during the ALNS is not outstanding, it performs reliably well with respect to the obtained starting solutions.

Removal strategies

Unlike the insertion heuristics, removal strategies cannot be evaluated independently of the ALNS. Instead, we have to assess whether they can provide room for improvements during the search.

Similar to the insertion strategies discussed above, we evaluate the number of moves that yield an improved solution relative to the total number of executions per strategy. The results can be found in Table A.2. With few exceptions, we can observe a clear distinction between strategies that perform well in terms of informativeness and those that yield improvements of total route duration. All in all, the priority-delta strategy (PriorityDelta) results in the largest number of improving moves. However, as we will see below, these improvements are minor compared to the results of the other strategies. The worst-angle removal strategy (WorstAngle) achieves good results in both aspects, informativeness and duration. In contrast, the route sparsification strategy (SparseRoute) yields similar results regarding informativeness but is associated with a particularly high

degradation of average route duration, possibly due to the major impact this strategy can have on routes. The remaining strategies perform well in terms of route duration rather than informativeness, with the region-based removal strategy (RegionTotal) as a rather extreme example.

In the next step, we evaluate the extent of the achieved improvements in order to get a better impression of the relative performance of these strategies. Figure A.3 indicates the average improvement in terms of informativeness of a newfound solution relative to the previous one. In this comparison, the RegionTotal strategy yields the best results. Hence, even though it does not obtain new solutions very often, we can conclude that if it *is* successful, it can achieve major improvements. The differences between the remaining strategies are not as large. Still, we can see that two of them perform worse than the others: In the case of the randomized-nearest neighbor removal strategy (RandNN), this corresponds to the average performance as indicated in Table A.2. The PriorityDelta strategy is more interesting, as this result seems to be at odds with the comparatively high ratio of improving moves in Table A.2. We believe that this is because many of the improvements made are rather incremental, with only a few locations in the routes exchanged with more promising targets. Based on these results, we decide against using the RandNN strategy in our evaluation, as it is outperformed by the others in terms of the number of improving moves, impact on route duration, and magnitude of the achieved improvements. We keep the remaining set of strategies to offer a broad range of possibilities to improve both target location selection and route durations.

Search neighborhood parameters

In the next step, we compare several parametrizations of the adaptive large neighborhood search. In Figure A.4, we give the performance depending on the average neighborhood size, i.e. the average number of visits removed by a removal strategy. The image indicates the runtime associated with a removal size as well as the achieved solution quality. The neighborhood size has a significant impact on overall computation time, as larger neighborhoods increase the effort necessary to rebuild vehicle routes. How-

Removal strategy	moves resulting in improvements	
	... of informativeness	... of route duration
PriorityDelta	**8.7** %	17.0 %
WorstAngle	**5.4** %	**20.8** %
SparseRoute	**5.4** %	12.4 %
SequNN	4.8 %	**20.4** %
WeightedCenter	3.3 %	18.0 %
RandNN	2.8 %	19.3 %
RegionTotal	1.7 %	**22.8** %

Table A.2: Performance of removing strategies: Improving moves during the ALNS relative to total number of executions per strategy. Best results are indicated in bold.

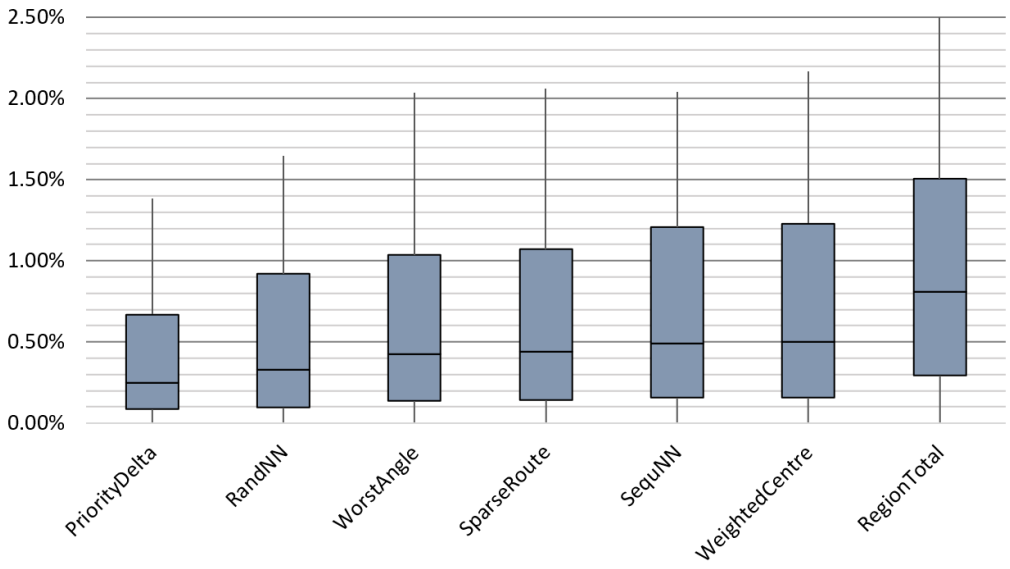

Figure A.3: Performance of selection strategies during the ALNS: Average improvement of accepted solutions with respect to the previous solution.

ever, they positively influence the achieved solution quality. This effect decreases with increasing neighborhood size, as larger neighborhoods more often lead to failures.

Figure A.5 reports the results obtained when varying the length of the segments removed from vehicle routes using the sequence-nearest neighbor removal strategy and the route sparsification strategy. The segment length depends on the average length of a vehicle route to adapt itself depending on the instance at hand: As vehicle capacities increase, tours become longer, which means that more locations have to be removed to change the route structures in a meaningful way. Variations of this parameter only have a minor influence on computation time, as the overall neighborhood size is fixed. Surprisingly, segment lengths that perform well are comparatively large, with the best results obtained for removing segments that include, on average, $25\ \%$ of the locations in a route. Further increasing the segment length has a detrimental impact on solution quality, as the initial route structure is no longer preserved, risking the loss of comparatively good segments.

Convergence criterion

Finally, we assess the influence of the convergence criterion, i.e. the number of non-improving moves after which search is stopped. The objective of this criterion is to terminate the search as early as possible if the current solution cannot be improved further. The results are reported in Figure A.6. Again, we can see a major impact on overall computation time as well as on solution quality. As we have already seen in Figure A.1, the first iterations during Phase 2 are associated with major gains in solution quality. We can, therefore, achieve reasonable results even when using a comparatively strict ter-

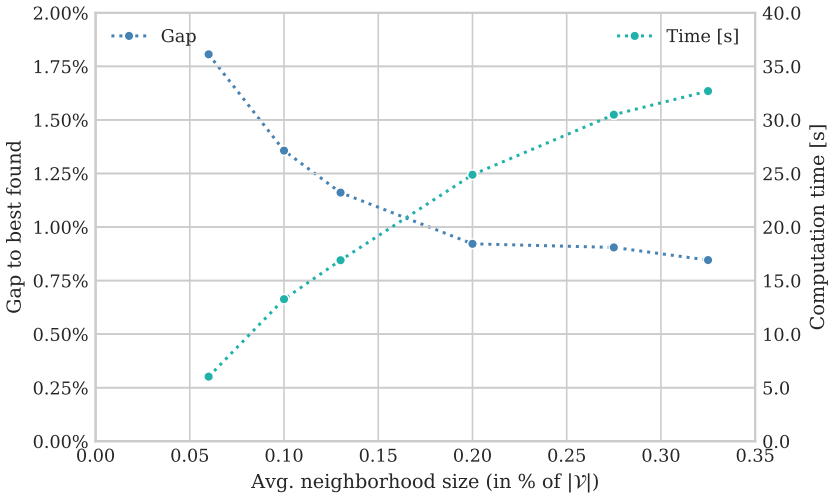

FIGURE A.4: Performance of the 2MLS depending on average neighborhood size.

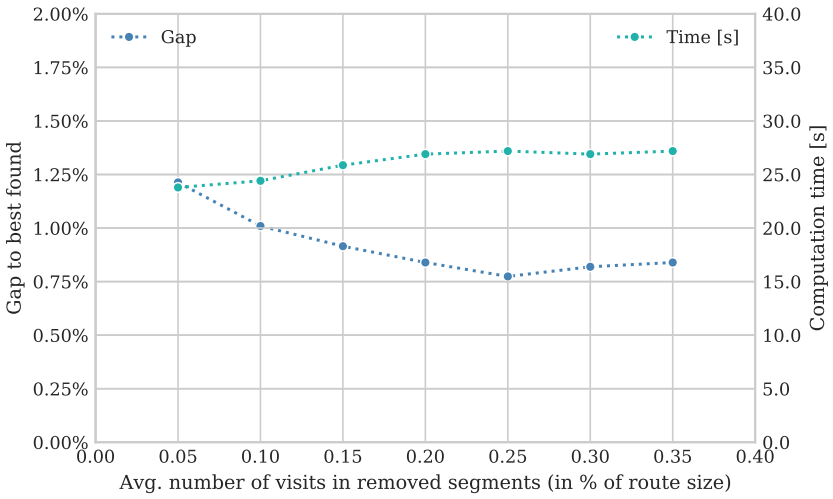

FIGURE A.5: Performance of the 2MLS depending on average segment size.

mination criterion of, e.g., 50 non-improving iterations. Finding further improvements later during the search becomes more difficult. As a consequence, increasing the convergence limit is associated with an increase in overall runtime, but not with further major improvements in solution quality.

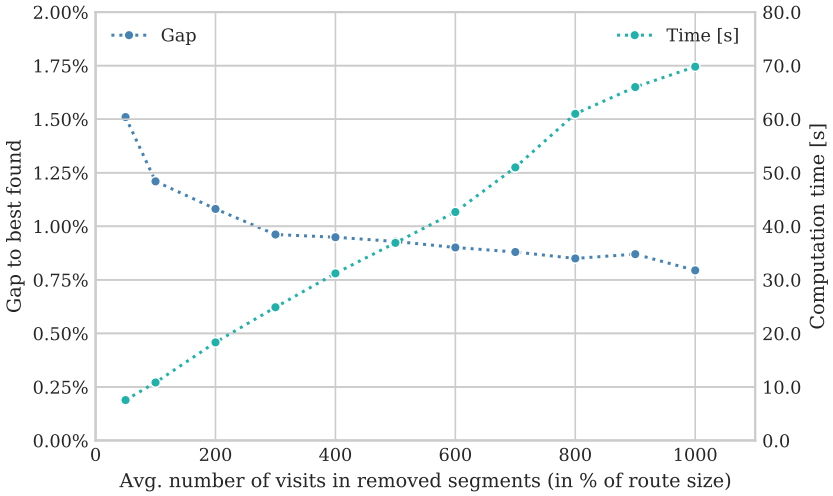

Figure A.6: Performance of the 2MLS depending on the convergence criterion.

B

Illustration of online planning strategies

The next pages serve as an illustration of the environment models and planning concepts introduced in Part III of this thesis. We refer to Chapter 13 for the definition of these models.

All planning approaches are demonstrated based on one scenario with priorities depicted in Figure B.1. For the sake of simplicity, we use a normalized target area with a size of 15×15 distance units. Vehicles move one distance unit per unit of time, total flight time is limited to 40 units of time. Performing an observation at a target location equally requires one unit of time. For greedy approaches, the maximum distance between samples is set to 3 distance units.

In the next sections, we illustrate the selection of the first six samples of one vehicle for each of the different strategies that we apply. To facilitate the comparison of approaches, we use normalized scales to indicate the relative importance of target locations, with the target that is currently seen as the most important one indicated in red. This allows demonstrating the evolution of the information represented in the models over time.

The range from within samples can be selected is indicated as a circle. The performed UAV trajectory is indicated in black, while the proposed next sampling locations are given as a dashed line extending from the current vehicle position.

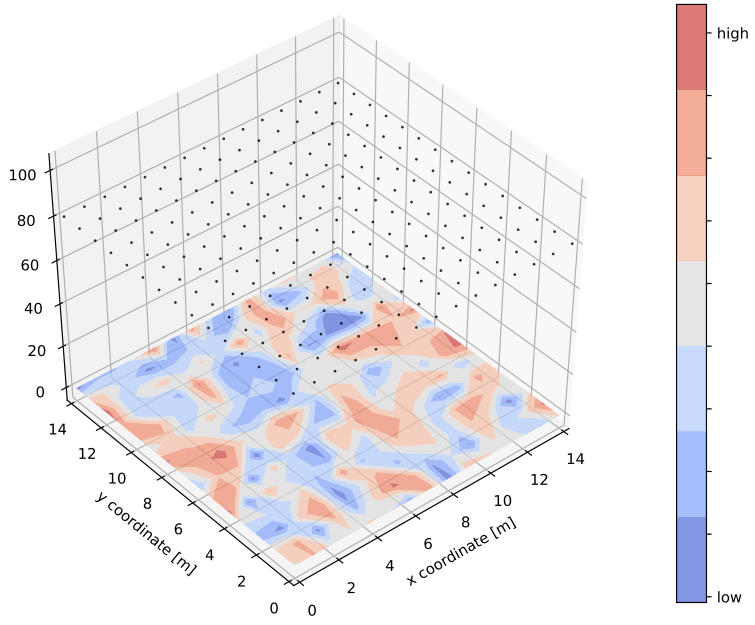

Figure B.1: Priority distribution used for illustrating planning strategies.

Direct coverage

Figure B.2 illustrates the selection of the fist samples using a direct coverage model. The algorithm iteratively selects the most highly valued target within the imposed range restriction for each step. Afterward, the target's priority is set to 0 and the algorithm proceeds with the next sample. The result is a narrow tour exploring highly prioritized samples in the immediate vicinity of the UAV takeoff location.

Disc-based coverage

The disc-based coverage model $\mathcal{I}^{DISC}$ is illustrated in Figure B.3. The initial distribution of the estimated informativeness resembles the model based on direct priority coverage discussed in the previous section. As the mission processes, the difference between the coverage count n_i at unsampled locations and the maximum coverage count $\max_{j \in \mathcal{V}} n_j$ increases. Priorities in the area surrounding the first sampling locations are decreased, while the estimated benefit of targets outside this region increases. This leads the UAV to explore previously unsurveyed areas with high priorities.

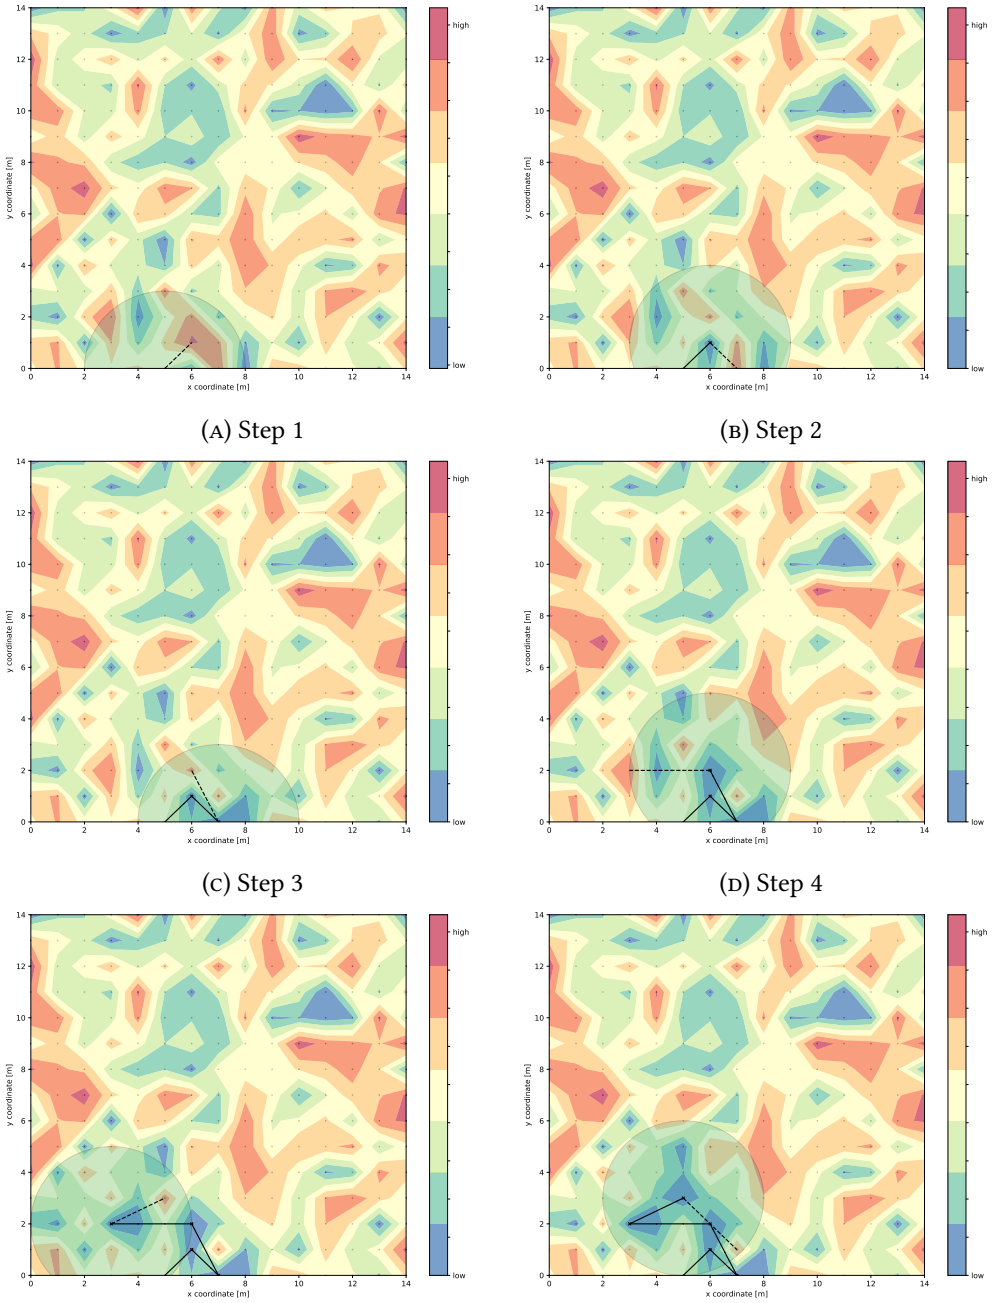

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.2: Environment model based on input priorities.

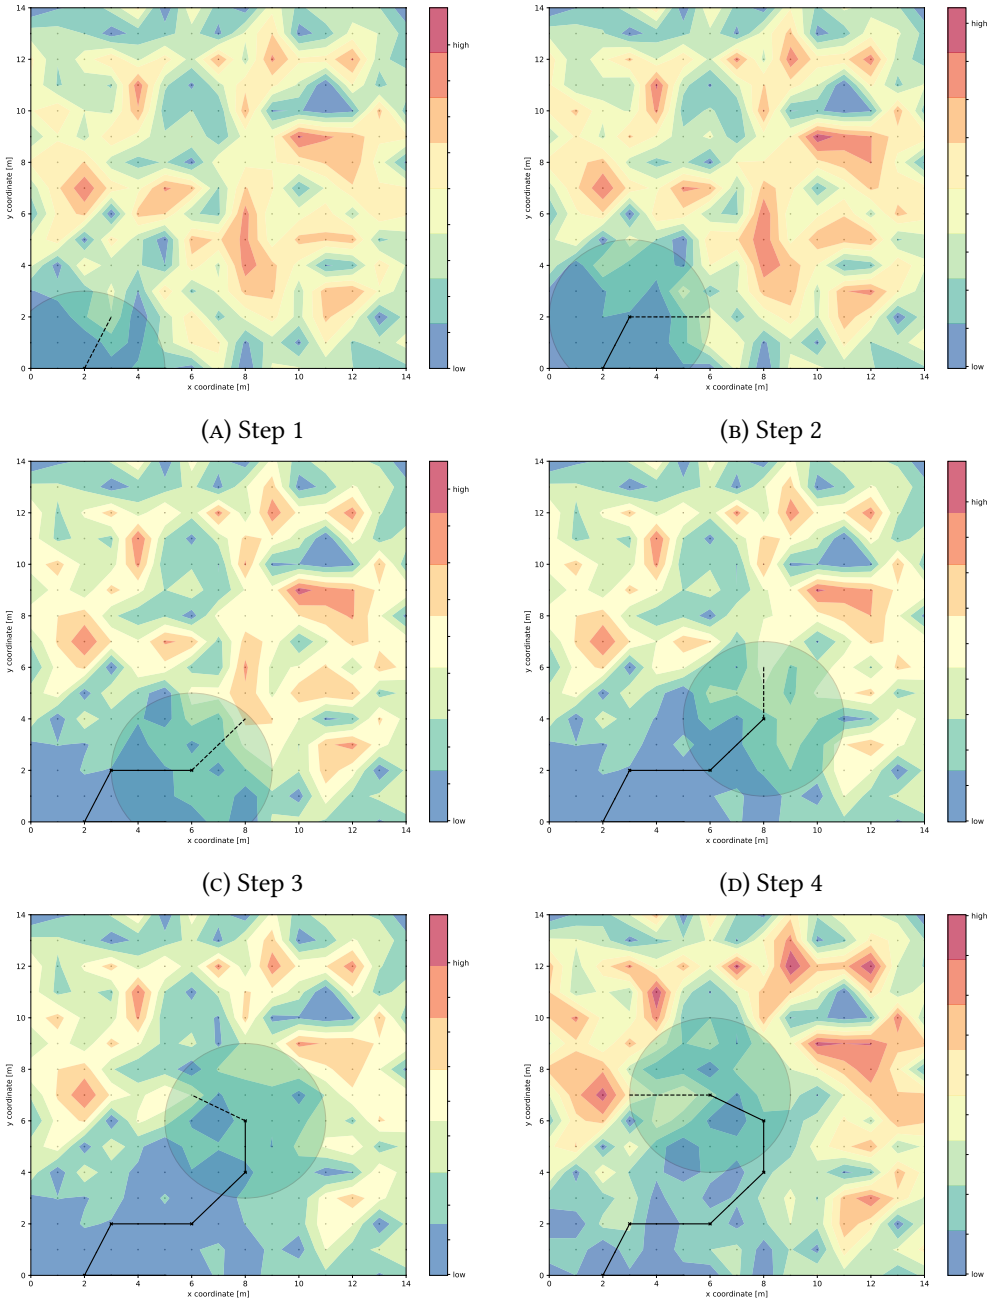

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.3: Environment model based on disc coverage $\mathcal{I}^{DISC}$.

Kernel-based coverage

The kernel-based measure $\mathcal{I}^{KERNEL}$ in Figure B.4 assigns high values to areas where several target come with high individual priority values. The UAV starts by moving toward the first one of these regions. The importance of sampling locations where observations have already been performed is set to 0, and the estimated benefit of samples in the vicinity of these nodes is decreased accordingly. Compared to a direct coverage model, this prevents redundant samples in close proximity. Instead, the UAV travels further, covering several locations with high priorities in the process.

Latency-based model

The latency-based information measure $\mathcal{I}^{LAT}$ is initialized similarly to the disc-based model such that the initial estimate of the target's benefit follows the distribution of input priorities (Figure B.5). The importance of areas that are not yet covered in a mission increases over time, while the estimated benefit in the vicinity of the last samples decreases. In contrast to the previous models, these areas become more relevant again as time progresses. This can be seen in the last two figures (Figures B.5e and B.5f), where the importance of those areas surveyed in the earlier iterations is increased relative to the area surrounding the current UAV position.

Kernel-based coverage with devaluation

Figure B.6 represents the first iterations based on the kernel-based model with devaluation ($\mathcal{I}^{KERNEL\text{-}TD}$). In the early stages of the mission, the estimated information gain corresponds to the kernel-based model discussed above (Appendix B), and the UAV moves in a similar fashion. Differences become more pronounced as the mission progresses. The relative importance of the areas sampled in the first steps increases again. Note furthermore that the estimated information gain of a location already surveyed never decreases as much as in the time-independent case: As this model assumes that a process varies in time, performing additional observations at an already surveyed location always provides additional information, specifically about the temporal correlation within the process.

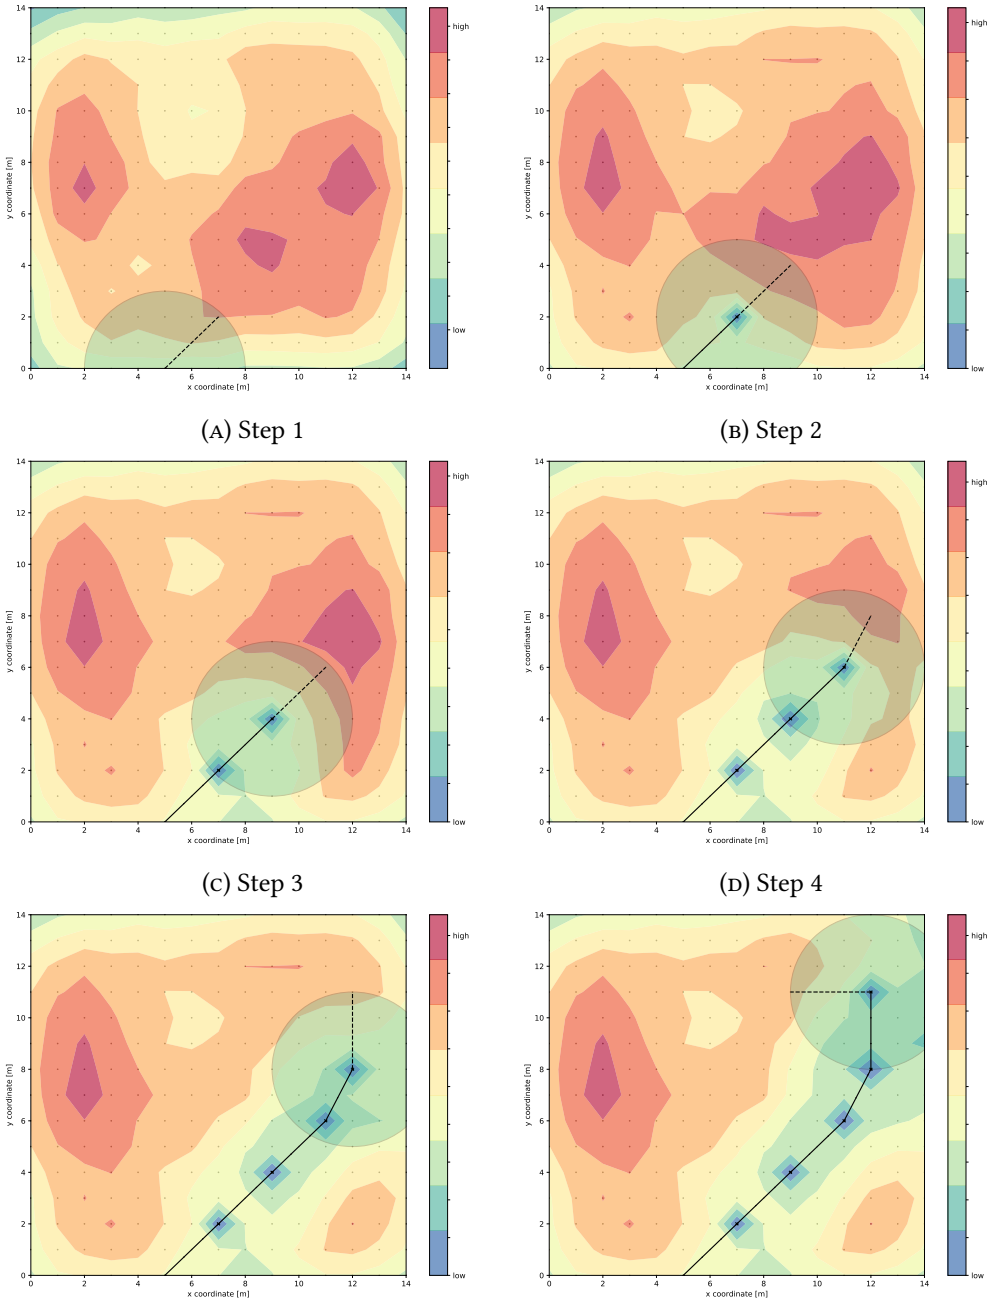

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.4: Environment model based on kernel-based coverage $\mathcal{I}^{KERNEL}$.

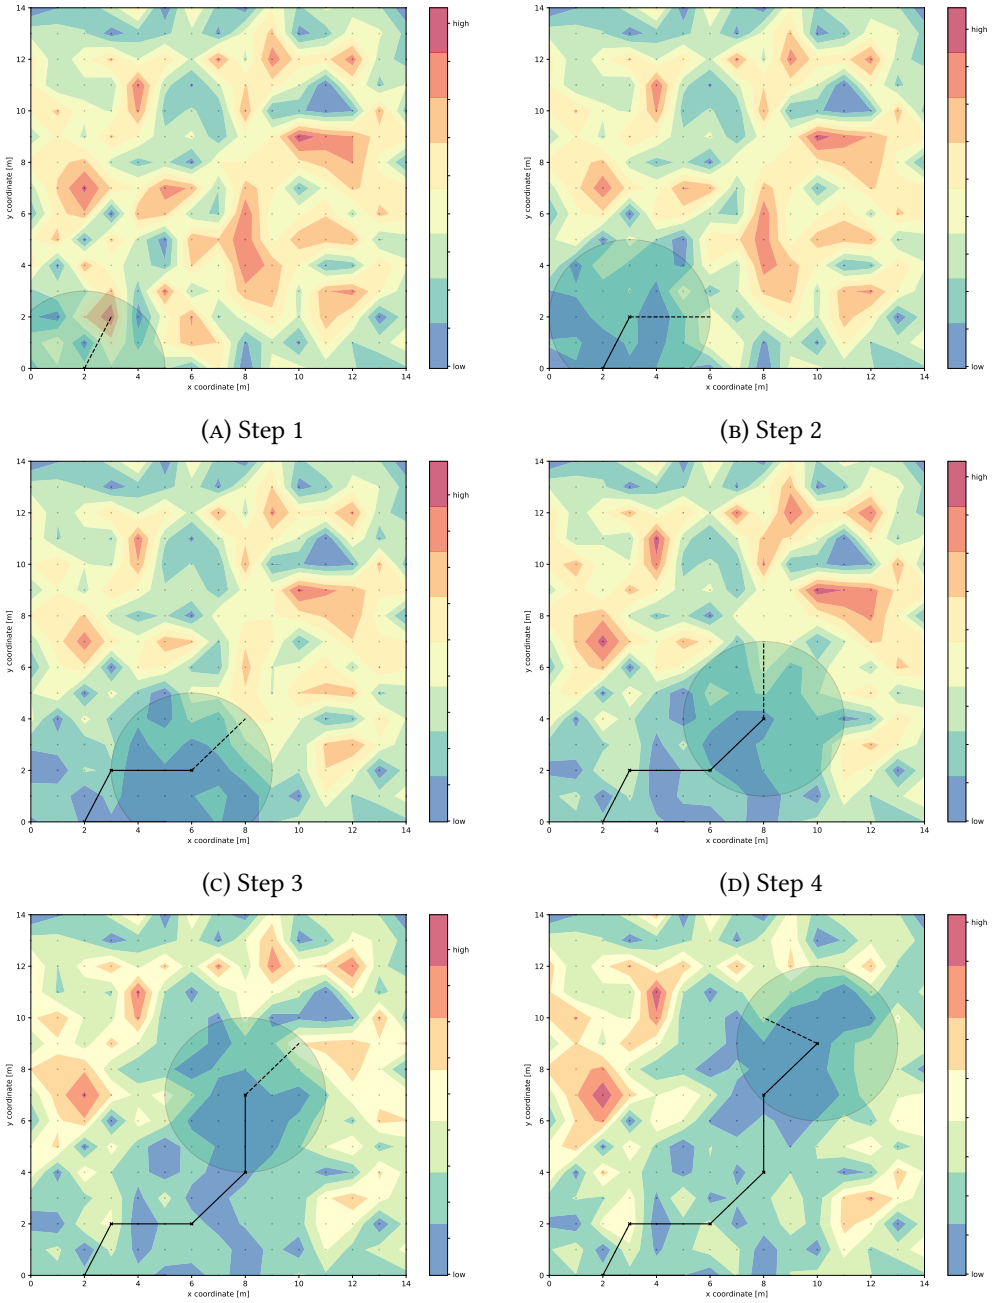

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.5: Environment model based on the reduction of sampling latency $\mathcal{I}^{LAT}$.

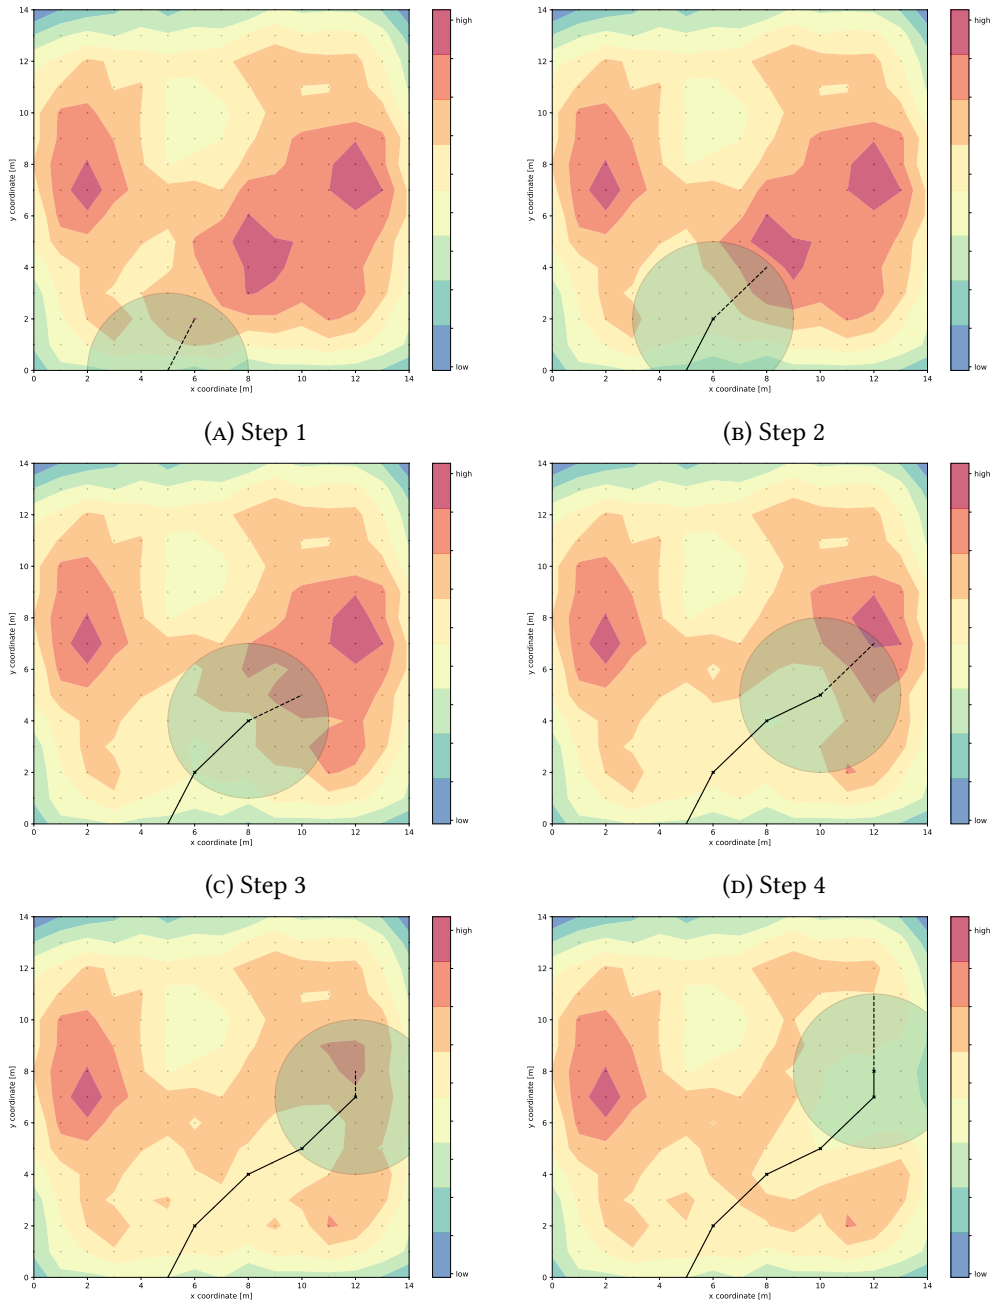

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.6: Environment model based on a kernel model with devaluation ($\mathcal{I}^{KERNEL\text{-}TD}$).

Variance maximization

As the first Gaussian process based model, the environment model based on the maximum process variance ($\mathcal{I}^{VAR}$) is represented in Figure B.7. Initially, i.e., before at least two samples are available and a model can be trained, this model uses the input priorities to evaluate targets. This changes in Figure B.7c and what follows, as the obtained samples allow an early estimate of the characteristics of the surveyed process. The posterior variance is low where samples have already been taken, as the prediction is likely to be accurate in this region. New samples are used to update the GP posterior, as can be seen in the changes between Figure B.7e and Figure B.7f. Here, the last sample taken leads to a change in the process model that decreases its certainty about the left-hand side of the region (e.g., coordinate (2,9)).

Average reduction in variance

Decision based on the average reduction in variance ($\mathcal{I}^{ARV}$) are illustrated in Figure B.8. As in the previous example, in the early stages of the missions, the lack of available samples does not allow an estimate of the variance within the process. Again, this changes in Figure B.8c. We emphasize that areas indicated as highly relevant here are not (necessarily) those with the highest current variance, but those that are likely to reduce prediction variance within the target area the most. This means that compared to the previous examples, targets close to the border of the target region are less valuable. Hence, the UAV first moves toward the center of the area. Then, it proceeds to move in the direction with high estimated benefits.

Mutual information

The last environment representation $\mathcal{I}^{MI}$ uses mutual information (MI) as an informativeness criterion. This is illustrated in Figure B.9. Again, early iterations do not have sufficient information available. This changes when two samples are processed. Similar to $\mathcal{I}^{ARV}$, we can see that $\mathcal{I}^{MI}$ decreases the importance of targets near the border of the considered area. Highly valued areas are far from the border and at larger distances from the previous samples. Therefore, the UAV again moves toward the center. As the first samples taken are nearly on a straight vertical line, they allow to estimate the covariance along this axis but provide little information about the interdependencies in the horizontal axis. For this reason, the model state represented in Figure B.9e suggests that samples that diverge from this first line of samples can provide more information about the process, which motivates the UAV movement indicated in Figure B.9f.

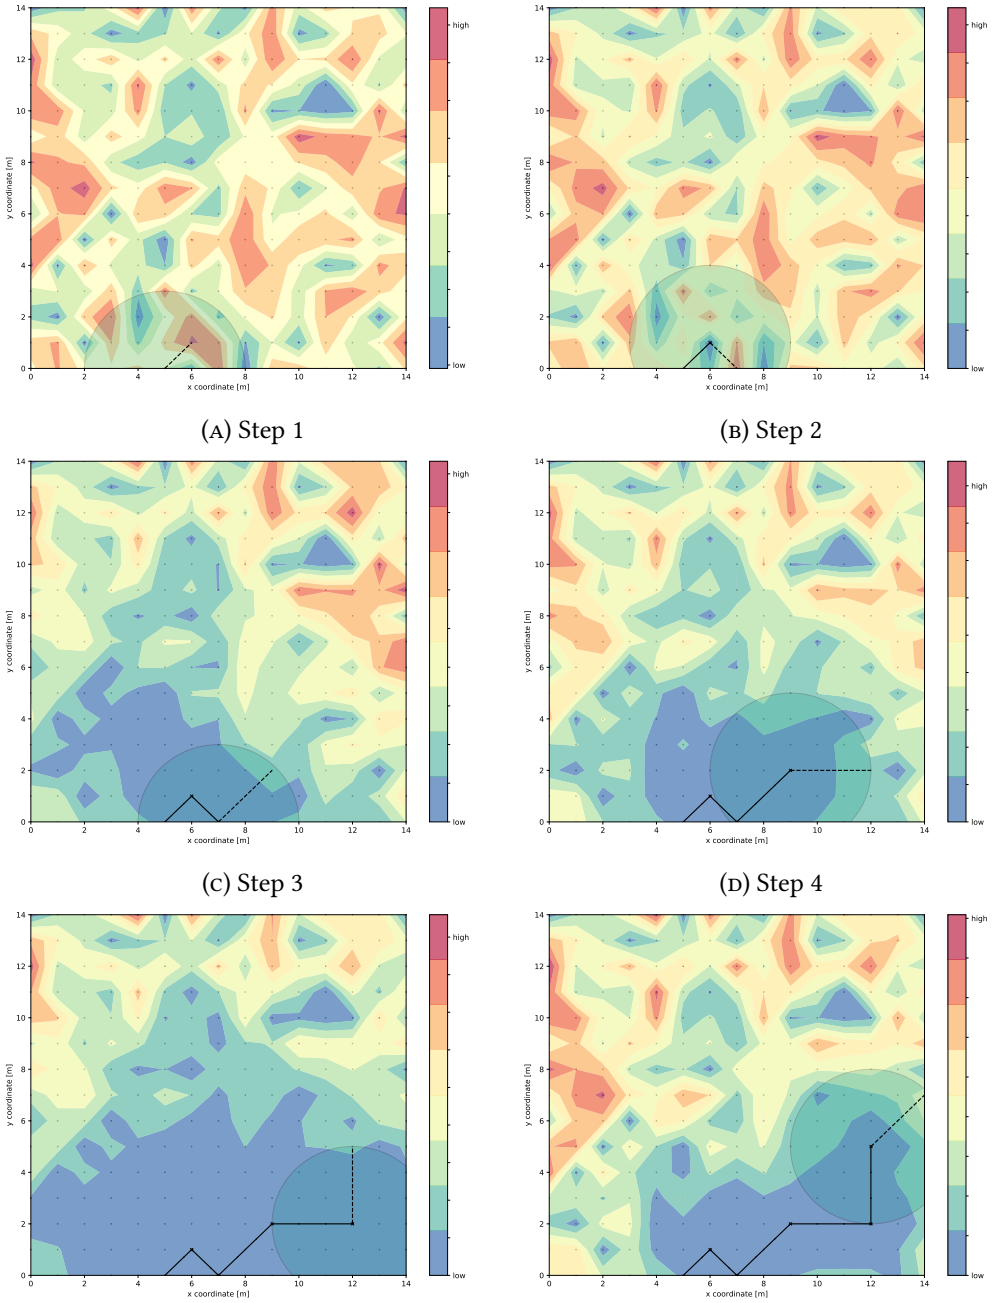

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.7: Environment model based on sampling location variance $\mathcal{I}^{VAR}$.

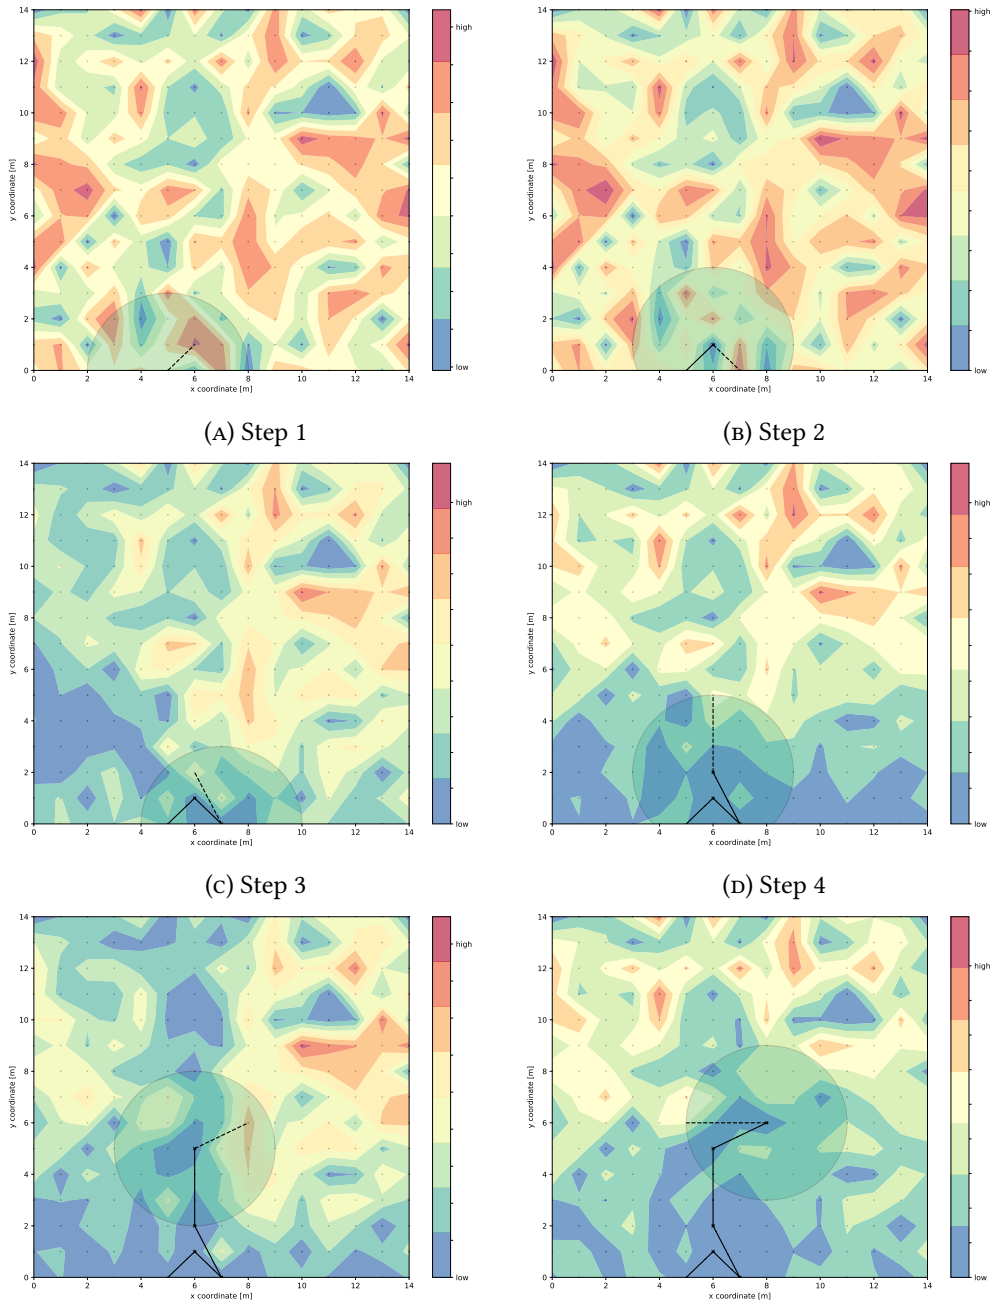

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.8: Environment model based on average reduction in variance $\mathcal{I}^{ARV}$.

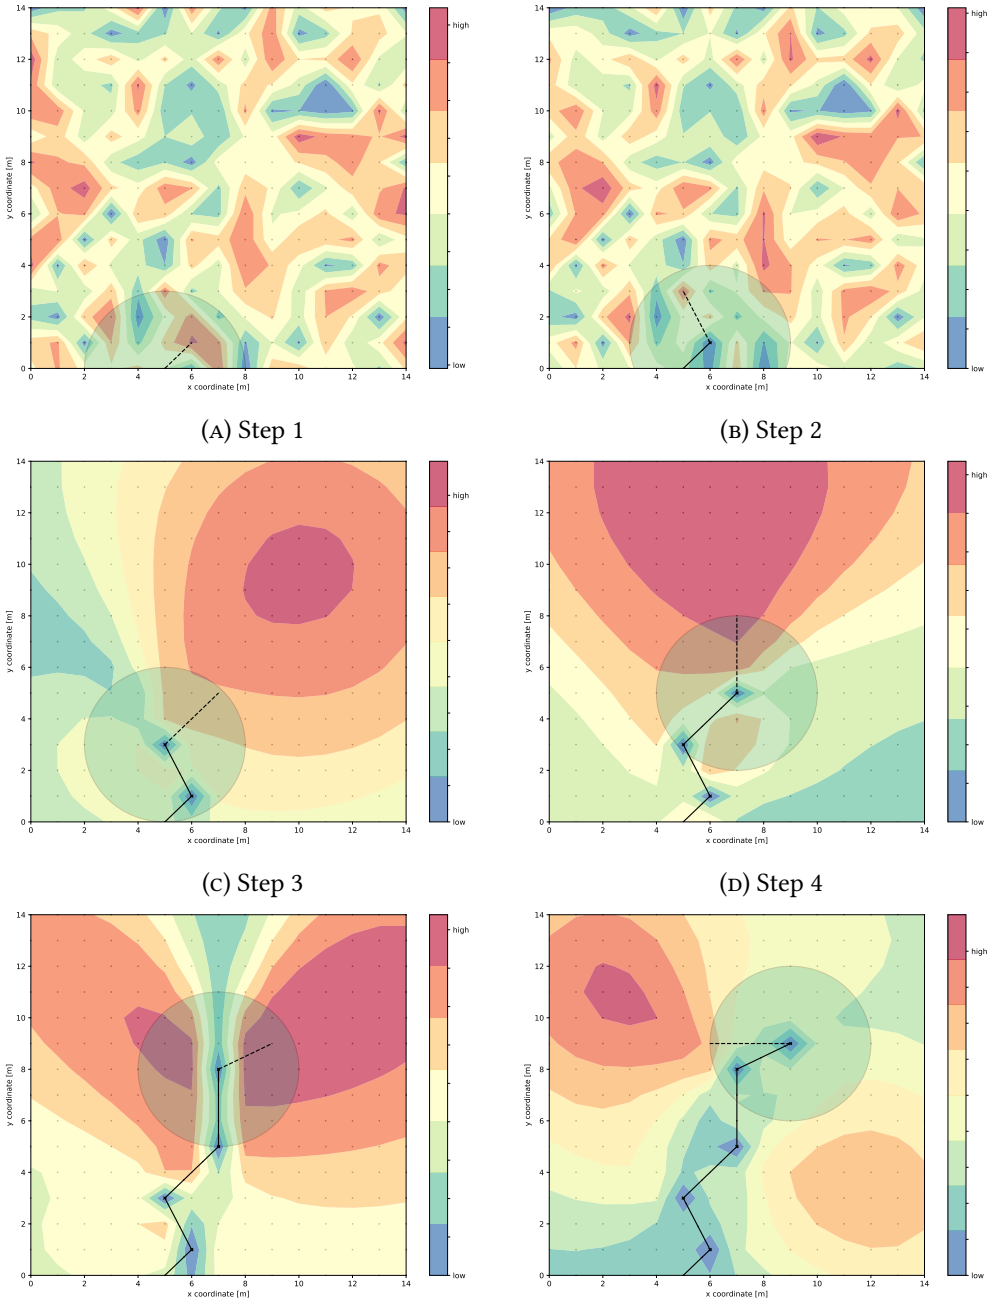

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.9: Environment model based on mutual information $\mathcal{I}^{MI}$.

Geometric approach

Figure B.10 illustrates several iterations based using LLoyd's algorithm, summarized in Algorithm 7). As this algorithm seeks to determine a partitioning of the target area between several UAVs, we adjust the base scenario considered in the other examples so that it includes two vehicles. The representation, however, focuses on one of these two UAVs. The figure demonstrates how the UAV iteratively moves toward the centroid of its partition, using the input priorities as a weight function. This centroid is indicated as a small red cross near position (9,7).

Sequential selection heuristic

Figure B.11 represents a sequential selection heuristic. The environment model used is the $\mathcal{I}^{ARV}$ model. The example furthermore uses a moving planning horizon of $\Delta\mathcal{T} = 10$ units of time and a fixed decision horizon of $\Delta\mathcal{T}^{fix} = 5$ units of time, i.e., planned targets can be changed, as only the first half is fixed before recomputing the sequence. Again, the algorithm first decides based on direct priorities. The resulting samples lead to a reduction of the ARV value toward the lower and left borders of the target region (Figures B.11c and B.11d), motivating the UAV to travel further in the subsequent steps. The structure of the solution differs significantly from the one obtained using the greedy algorithm which is illustrated in Appendix B. As the error variances tend to be higher for areas further away from the current sampling positions, increasing the length of the planning horizon means that the distances between subsequent sampling locations relatively large. The UAV travels further but obtains fewer samples during the mission.

Integrated decision scope

The integrated decision scope which determines the optimal sampling locations for a given planning horizon is illustrated in Figure B.12. We again use $\mathcal{I}^{ARV}$ for measuring the importance of target locations and set $\Delta\mathcal{T} = 10$ and $\Delta\mathcal{T}^{fix} = 5$. Especially compared to the sequential selection illustrate in the previous example, we can see a major impact of the used planning strategy: Rather than traveling far and taking fewer samples, the approach determines relatively short trajectories with several sampling locations. It moves toward the center of the area, where samples are promising, as we have seen using the examples given in Appendix B, rather than including samples near the border as in the previous example.

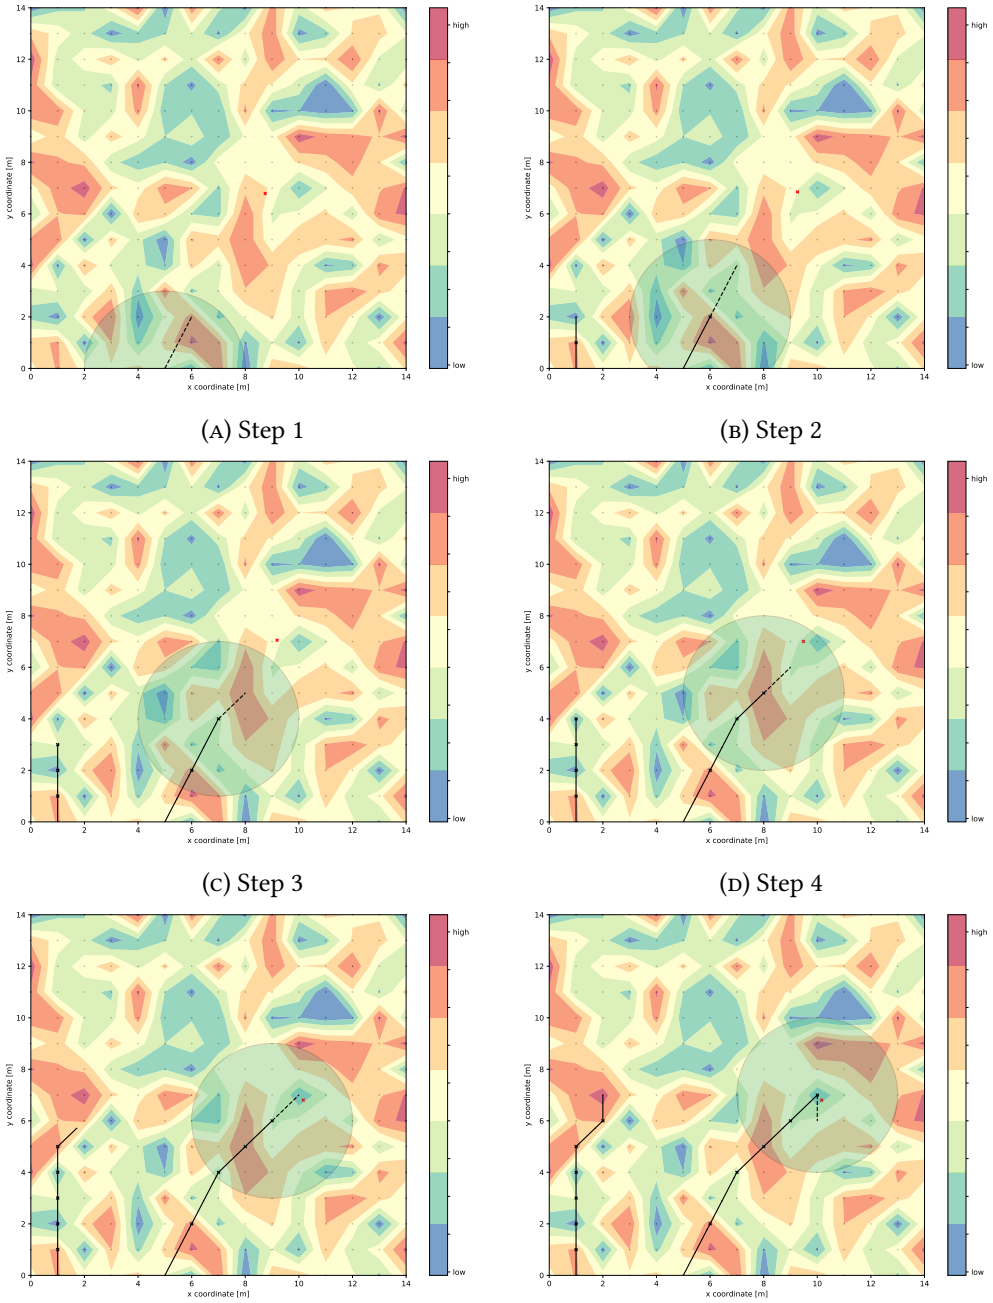

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.10: Geometric solution approach (LLoyd's algorithm).

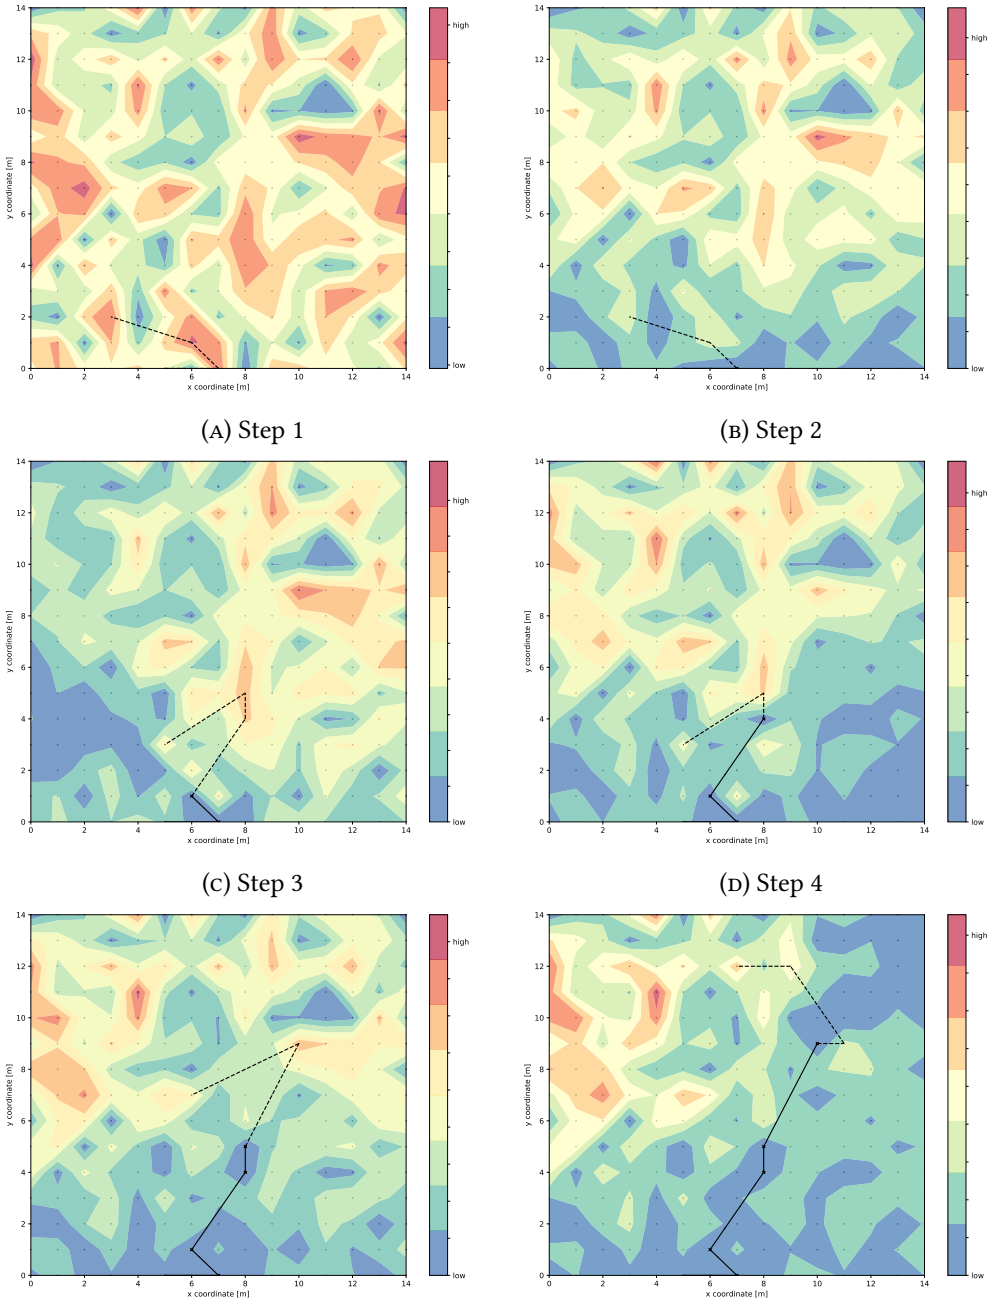

(A) Step 1

(B) Step 2

(C) Step 3

(D) Step 4

(E) Step 5

(F) Step 6

FIGURE B.11: Sequential selection heuristic based on $\mathcal{I}^{ARV}$.

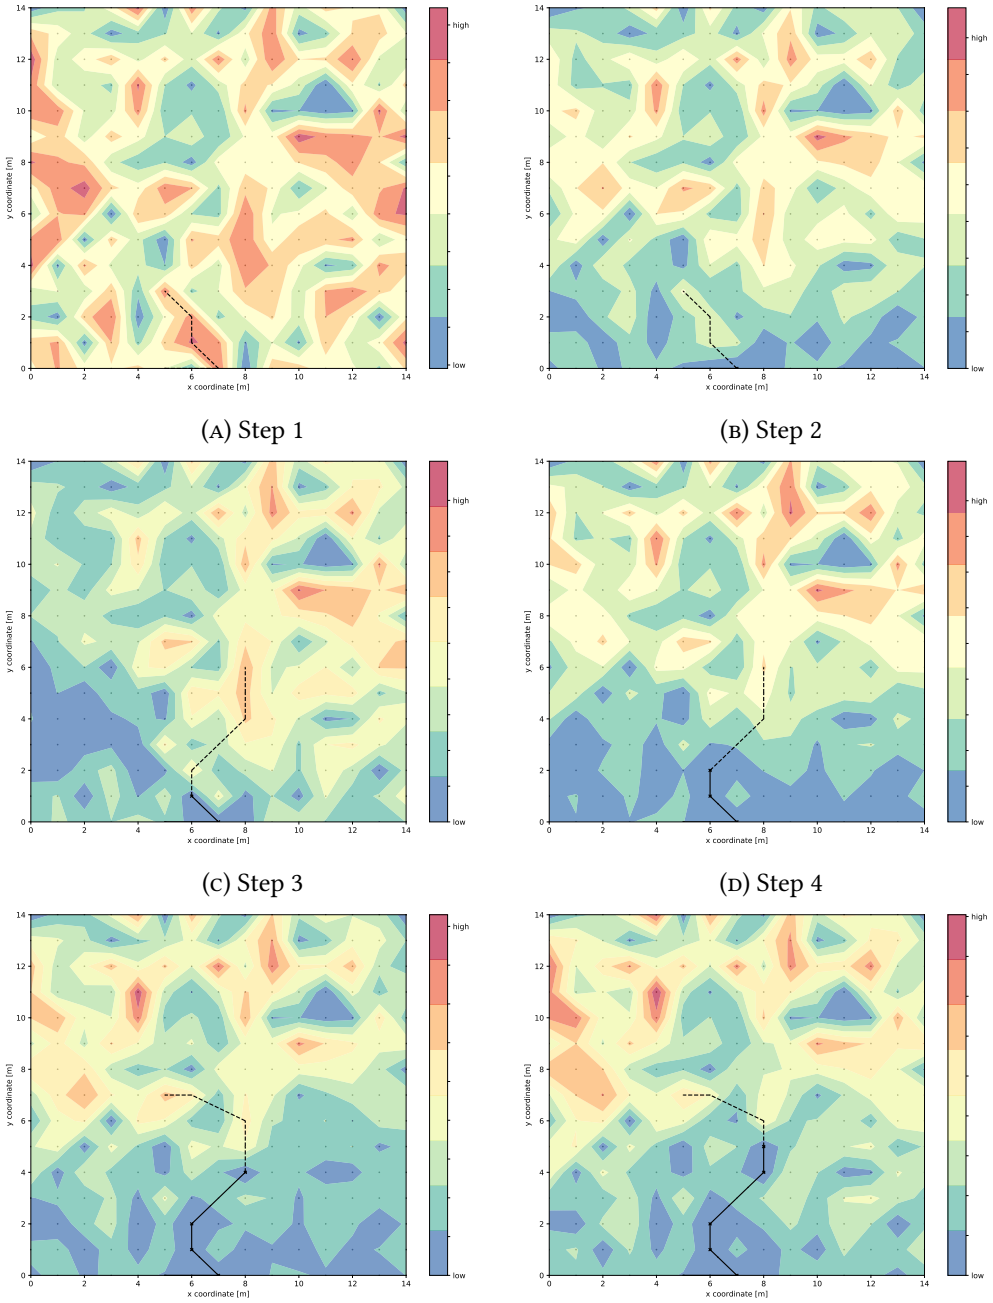

(A) Step 1 (B) Step 2

(C) Step 3 (D) Step 4

(E) Step 5 (F) Step 6

FIGURE B.12: Integrated decision scope for target selection based on $\mathcal{I}^{ARV}$.

C

Discrete event simulation

This overview describes the discrete event simulation implemented for evaluating planning strategies in more detail. Specifically, we give the pseudocode of the simulation engine in Algorithm 8 and specify the event resolution routines.

Algorithm 8: Simulation engine

Input: UAVs $\mathcal{M}$, target area $\mathcal{V}$ with priorities u_i, planning horizon $\mathcal{T}$, ground truth $z_{i,t}$ for $i \in \mathcal{V}, t \in \mathcal{T}$, planning heuristic `planner`

```
/* main simulation procedure */
Function run():
    initialize event queue Q = ∅, simulation clock t = 0;
    for all UAVs m in M do
        Q.insert(VehicleCreationEvent(entryTime, m));
        Q.insert(VehicleFailureEvent(failureTime, m));
        Q.insert(TriggerPlanEvent(entryTime, m));
    end
    while not Q.empty() and Q.peek().time ≤ T do
        event =Q.pop();
        set t to event.time;
        if event has not been aborted then
            event.resolve() ;
    end
    compute statistics;
end
```

Algorithms 9 to 13 depict the resolution of the update events that occur during the mission in the order in which they are resolved in each time step. The first four events only lead to an update of the data structures in the planning strategy, and, in case of a vehicle failure, the abortion of future events associated with the corresponding UAV sin the simulation.

Algorithm 13 combines two tasks: The algorithm triggers the planning strategy so that new targets are computed, and implements the mission resulting from these updates in the simulation. The planner is called each time a new UAV becomes available as well as each time an observation is completed. Note that it is possible that the mission does not change between subsequent calls, e.g., in case of longer fixed decision horizons. However, for the sake of simplicity, we do not distinguish between original and recomputed mission in this overview. Hence, each recomputation of a mission first leads to the abortion of all future events except, potentially, the failure of the UAV which is fixed for an instance. They are no longer valid as the mission might have changed. In a second step, the movements of the UAV through the newly determined set of samples is estimated, along with the completion times of samples, and events for completed observations and GPS updates along this route are inserted into the priority queue.

Algorithm 9: Vehicle creation

Input: Vehicle creation time t, vehicle ID m

```
1 Function VehicleCreationEvent.resolve():
2     planner.addUAV(t, m);
3     planner.updateStatus(t, m, status.READY);
4 end
```

Algorithm 10: GPS update

Input: Event time t, vehicle ID m, position $pos_{m,t}$

```
1 Function GPSEvent.resolve():
2     planner.updatePosition(t, m, pos_{m,t});
3 end
```

Algorithm 11: Vehicle failure

Input: Failure time t, vehicle ID m

Function `VehicleFailureEvent.resolve()`:
    `planner.updateStatus`$(t, m,$ `status.FAILED`$)$;
    for *all events* `event` *in* $\mathcal{Q}$ **do**
        if `event.time` $\geq t$ *and* `event.vehicleID` $= m$ **then**
            abort `event`;
    end
end

Algorithm 12: Observation event

Input: Sampling time t, vehicle ID m, ground truth z_i

Function `ObservationEvent.resolve()`:
    simulate noisy observation $y_i = z_i + e_i$ with $e_i \sim \mathcal{N}(0, \sigma)$;
    `planner.updateObservation`(t, y_i);
end

Algorithm 13: Trigger planning heuristic

Input: Event time t, vehicle ID m

Function `TriggerPlanEvent.resolve()`:
    $\mathcal{S} =$ `planner.getNextTargets`(m);
    for *all events* `event` *in* $\mathcal{Q}$ **do**
        if `event.time` $\geq t$ *and* `event.vehicleID` $= m$ *and* `event` *is not a* `VehicleFailureEvent` **then**
            abort `event`;
    end
    for *all sampling positions* $s \in \mathcal{S}$ **do**
        compute estimated observation completion time (`completionTime`);
        $\mathcal{Q}$.`insert`(`ObservationEvent`(`completionTime`, m, $z_{s,\texttt{completionTime}}$));
        $\mathcal{Q}$.`insert`(`TriggerPlanEvent`(`completionTime`, m));
        for *all positions* $pos_{m,t'}$ *en route to* s *in fixed increments of time* **do**
            $\mathcal{Q}$.`insert`(`GPSEvent`$(t', m, pos_{m,t'})$);
        end
    end
end

List of Figures

List of Tables

Acronyms

2MLS two-phase multi-start adaptive large neighborhood search.

ACO ant colony optimization.

ALNS adaptive large neighborhood search.

AMPPES adaptive mission planning problem for emergency surveillance.

APF artificial potential field.

ARV average reduction in variance.

BB branch-and-bound.

BBK Bundesamt für Bevölkerungsschutz und Katastrophenhilfe (Federal Office of Civil Protection and Disaster Assistance).

BC branch-and-cut.

BCP branch-and-cut-and-price.

BP branch-and-price.

CorTOP correlated team orienteering problem.

CSP covering salesman problem.

CTP covering tour problem.

DES discrete-event simulation.

FPR fast path relinking.

GCorTOP generalized correlated team orienteering problem.

GLS guided local search.

GP Gaussian process.

GRASP greedy randomized adaptive search procedure.

HGA hybrid genetic algorithm.

IDW inverse distance weighting.

ILP integer linear programming.

ILS iterated local search.

IPP informative path planning.

MAE mean absolute error.

ME mean error.

MI mutual information.

MIP mixed integer linear programming.

MIQP mixed integer quadratic programming.

MPC model predictive control.

MPPES mission planning problem for emergency surveillance.

MS-LS multi-start local search.

NN nearest neighbor interpolation.

OP orienteering problem.

PMA pareto mimic algorithm.

PSOiA particle swarm optimization-inspired algorithm.

RCSPP resource constraint shortest path problem.

RIG rapidly-exploring information gathering algorithm.

RMSE root mean squared error.

SA simulated annealing.

SOP set orienteering problem.

TOP team orienteering problem.

TS tabu search.

TSP traveling salesman problem.

UAV unmanned aerial vehicle.

UHGS unified hybrid genetic search.

VNS variable neighborhood search.

VRP vehicle routing problem.

WMAE weighted mean absolute error.

Symbols

Prediction and interpolation

$Z(s)$	Random variable
$\{Z(s)\}_{s\in\mathcal{V}}$	Random process defined over $\mathcal{V}$
z_s	Realization of the random variable $Z(s)$ at $s \in \mathcal{V}$
y_s	Noisy sample of the random variable $Z(s)$ with $y_s = z_s + e_s$
$\widehat{Z}(s)$	Interpolated value
Σ	Covariance matrix
$\mu(s)$	Mean function defined over $s \in \mathcal{V}$
$\mathcal{GP}$	Gaussian process with mean function μ and covariance matrix Σ
$\mathcal{N}(\mu, \sigma)$	Normal distribution with mean μ and standard deviation σ
$k(s, s')$	Kernel function defined over all $s, s' \in \mathcal{V}$
$\theta,\ \hat{\theta}$	Hyperparameters / optimal hyperparameters
$\gamma(d)$	Variogram
$N(d)$	Pairs of points at distance d
ν^2	Maximum process variance
l	Lengthscale $l > 0$
a	Variogram range $a > 0$
λ_s	Interpolation weight $\lambda_s \in [0, 1]$
$\mathcal{H}$	Entropy

e_s	Error or noise at $s \in \mathcal{V}$

MPPES and AMPPES models

$\mathcal{V}$	Locations comprising the target area with $\mathcal{V} \subset \mathbb{R}^2$
$\mathcal{S}$	Set of sensing locations with $\mathcal{S} \subseteq \mathcal{V}$
$\mathcal{M}$	Set of available vehicles $m \in \mathcal{M}$
s_m^{start}	Takeoff location of UAV $m \in \mathcal{M}$
s_m^{end}	Destination location of UAV $m \in \mathcal{M}$
$\mathcal{N}_S$	Set of UAV starting locations with $\mathcal{N}_S = \bigcup_{m \in \mathcal{M}} s_m^{start}$
$\mathcal{N}_E$	Set of UAV destination locations with $\mathcal{N}_E = \bigcup_{m \in \mathcal{M}} s_m^{end}$
$\mathcal{N}$	All locations, with $\mathcal{N} = \mathcal{V} \cup \mathcal{N}_S \cup \mathcal{N}_E$
$\mathcal{C}_i$	Covering neighborhood with $\mathcal{C}_i \subseteq \mathcal{V},\ i \in \mathcal{V}$
$\mathcal{I}(\mathcal{S})$	Informativeness value associated with samples $\mathcal{S}$
r	Route (sequence of planned visit locations)
Ω, Ω_m	Set of all feasible routes r / feasible routes for vehicle $m \in \mathcal{V}$
ω	Solution to the MPPES
y_r	Decision variable, 1 if $r \in \Omega_m$ is selected in a solution, 0 otherwise
x_i	Decision variable, 1 if $i \in \mathcal{V}$ is included in set $\mathcal{S}$, 0 otherwise
$pos = (p^x, p^y)$	Position and corresponding coordinates with $pos \in \mathbb{R}^2$
u_i	Priority associated with visiting $i \in \mathcal{V}$
τ_{ij}	Flight time between two locations or positions
τ_i	Sensing time at location $i \in \mathcal{V}$
d_{ij}	Airline distance between two locations or positions
T_m^{max}	Maximum flight time of UAV $m \in \mathcal{M}$
w_{ij}	Weight between nodes $i, j \in \mathcal{V}$ with $w \leq 1$
d^{min}	Baseline distance for inverse distance weighting, $d^{min} > 0$
$\bar{w}$	Baseline weight for inverse distance weighting relative to d^{min}, with $\bar{w} \leq 1$

Offline planning

$(\mathcal{S}, m, T_m, i)$	Label
$\mathcal{L}^{fw}, \mathcal{L}^{bw}$	Forward and backward label sets
P, P^{red}	Problem representation and reduced problem, resp.
ω^{red}	Solution to the reduced problem representation P^{red}
$\delta^{removal}$	Removal strategy
$\Delta^{removal}$	Set of removal strategies, with $\Delta^{removal} = \bigcup \delta^{removal}$
δ^{insert}	Insertion strategy
Δ^{insert}	Set of insertion strategies, with $\Delta^{insert} = \bigcup \delta^{insert}$
Γ^o	Open, i.e., unassigned locations, with $\Gamma^o \subseteq \mathcal{V}$
Γ^p	Planned (visited) locations in an emerging solution, with $\Gamma^p \subseteq \mathcal{V}$
Γ^u	Unvisited locations in an emerging solution, with $\Gamma^u \subseteq \mathcal{V}$
u_i^r	Remaining utility of a node $i \in \mathcal{V}$, relative to the emerging solution
u_i^m	Marginal utility of a node $i \in \mathcal{V}$, relative to the emerging solution
$\underline{\mathcal{I}}$	Lower bound on the objective value relative to the emerging solution
$\overline{\mathcal{I}}$	Upper bound on the objective value relative to the emerging solution
p^{rand}	Randomization factor for visit insertion
cost, $cost^{min}$	Cost (detour) associated with a visit insertion / minimum insertion cost
nh	Search neighborhood size, with absolute limits nh^-, nh^+
υ	Segment length factor relative to average route size, $\upsilon \in (0, 1)$
ϱ	Factor for distance in objective function
n^{reg}	Regret parameter, $n^{reg} \in \mathbb{N}^+$
det	Determinism parameter, $det > 0$
$\delta^{obj}_{i,n^{reg}}$	Difference between the objective value for the best and the n^{reg}-th best insertion position for i
δ^{obj}	Objective value delta between solutions in reheating scheme

Ψ, Ψ_0	Temperature in the reheating scheme initial temperature
ψ	Heating factor with $\psi > 1$
κ	Iteration limit for heating factor

Online approaches

$\mathcal{T}$	Planning horizon
$\Delta\mathcal{T}$	Moving planning horizon
$\Delta\mathcal{T}^{fix}$	Fixed decision horizon
$\mathcal{S}_{<t}$	Set of sampled locations until (excluding) time t
$\mathcal{Y}_{<t}$	Set of obtained measurements until (excluding) time t
$s_t, \mathcal{S}_t$	Sampling locations / set of sampling candidates selected at time t
$T_{\mathcal{S}_t}$	Travel time associated with path trough $\mathcal{S}_t$
V_m	Voronoi region of generator m
$\{V_m\}_{m\in\mathcal{M}}$	Voronoi tesselation determined by generators $m \in \mathcal{M}$
d^{cover}	Maximum coverage radius for disc-based environment models
d^{limit}	Maximum distance between subsequent sampling locations
τ_s^{lat}	Sampling latency at $s \in \mathcal{V}$: time since last sample in coverage radius
ζ	Performance ratio
$\mathcal{F}(\zeta)$	Counting distribution function for performance ratio ζ

Other

v_m^{max}	Maximum UAV speed $[\frac{\mathrm{m}}{\mathrm{s}}]$
a_m^{max}	Maximum UAV acceleration $[\frac{\mathrm{m}}{\mathrm{s}^2}]$
h	Flight height above ground [m]
f	Focal length [mm]
b	Horizontal sensor size [mm]
pix	Number of pixels per axis
pxr	Ground sampling resolution $[\frac{\mathrm{m}^2}{\mathrm{pixel}}]$

A	Covered ground area [m^2]
c	Center, representative
$\mathcal{Q}$	Priority queue
Λ	Set of problem instances
d	Distance
q	Exponent for ratio parametrization, $q > 0$
p	Probability, random number
i, j	Indices
$\mathcal{O}$	Time complexity (“big O notation”)
tr	Trace of a matrix (sum of eigenvalues)
$\|\cdot\|$	Euclidean norm
n, N	Count or number of elements

References

Aasen, H., A. Burkart, A. Bolten, and G. Bareth (2015). Generating 3D hyperspectral information with lightweight UAV snapshot cameras for vegetation monitoring: From camera calibration to quality assurance. *ISPRS Journal of Photogrammetry and Remote Sensing 108*, 245–259.

Allahyari, S., M. Salari, and D. Vigo (2015). A hybrid metaheuristic algorithm for the multi-depot covering tour vehicle routing problem. *European Journal of Operational Research 242*(3), 756–768.

Archetti, C., F. Carrabs, and R. Cerulli (2018). The set orienteering problem. *European Journal of Operational Research 267*(1), 264–272.

Archetti, C., A. Hertz, and M. G. Speranza (2007). Metaheuristics for the team orienteering problem. *Journal of Heuristics 13*(1), 49–76.

Arthur, D. and S. Vassilvitskii (2007). k-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1027–1035.

Bachmaier, M. and M. Backes (2011). Variogram or semivariogram? Variance or semivariance? Allan variance or introducing a new term? *Mathematical Geosciences 43*(6), 735–740.

Bachmann, U., N. Derakshani, M. Drobig, M. König, J. Mentfewitz, H. Prast, G. Uelpenich, M. Vidmayer, S. Wilbert, and M. Wolf (2015). Recommendations on sampling for hazard control in civil protection. `http://www.warnungs-inhalte.bund.de/SharedDocs/Downloads/BBK/DE/Publikationen/CBRN-Schutz/FiB5_2.Auflage_engl.html`. Online; accessed December 03, 2019.

BBK (2018a). CBRN-Schutz. `https://www.bbk.bund.de/DE/AufgabenundAusstattung/CBRNSchutz/cbrnschutz_node.html`. Online; accessed May 14, 2019.

BBK (2018b). Chemische Gefahren. `https://www.bbk.bund.de/DE/AufgabenundAusstattung/CBRNSchutz/Chemie/ChemGef/chemgef_node.html`. Online; accessed May 14, 2019.

BBK (2019). Mess- und Nachweistechnik. https://www.bbk.bund.de/DE/AufgabenundAusstattung/CBRNSchutz/Chemie/CMundNwtechn/cmundnwtechn_node.html. Online; accessed December 16, 2019.

Bektaş, T., P. P. Repoussis, and C. D. Tarantilis (2014). Dynamic vehicle routing problems. In *Vehicle Routing: Problems, Methods, and Applications, Second Edition*, 299–347. SIAM.

Binney, J., A. Krause, and G. S. Sukhatme (2010). Informative path planning for an autonomous underwater vehicle. In *IEEE International Conference on Robotics and Automation*.

Binney, J., A. Krause, and G. S. Sukhatme (2013). Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research 32*(8), 873–888.

Binney, J. and G. S. Sukhatme (2012). Branch and bound for informative path planning. In *IEEE International Conference on Robotics and Automation*, 2147–2154.

Boccardo, P., F. Chiabrando, F. Dutto, F. G. Tonolo, and A. Lingua (2015). UAV deployment exercise for mapping purposes: Evaluation of emergency response applications. *Sensors 15*(7), 15717–15737.

Bohling, G. (2005). Introduction to geostatistics and variogram analysis. C&PE 940, Kansas geological survey.

Brink, J. and E. Pebesma (2014). Plume tracking with a mobile sensor based on incomplete and imprecise information. *Transactions in GIS 18*(5), 740–766.

Bui, D.-M., T. Huynh-The, S. Lee, and Y. Yoon (2015). Complexity reduction for Gaussian process regression in spatio-temporal prediction. In *International Conference on Advanced Technologies for Communications (ATC)*, 326–331.

Cabreira, T., L. Brisolara, and P. R Ferreira (2019). Survey on coverage path planning with unmanned aerial vehicles. *Drones 3*(1), 4.

Casbeer, D. W., D. B. Kingston, R. W. Beard, and T. W. McLain (2006). Cooperative forest fire surveillance using a team of small unmanned air vehicles. *International Journal of Systems Science 37*(6), 351–360.

Caselton, W. F. and J. V. Zidek (1984). Optimal monitoring network designs. *Statistics & Probability Letters 2*(4), 223–227.

Central Office for Geotopography (2018). Service centre of the federal government for geo-information and geodesy. https://www.bbk.bund.de/DE/AufgabenundAusstattung/CBRNSchutz/cbrnschutz_node.html. Online; accessed August 15, 2018.

Chao, I.-M., B. Golden, and E. Wasil (1996a). A fast and effective heuristic for the orienteering problem. *European Journal of Operational Research 88*, 475–489.

Chao, I.-M., B. Golden, and E. Wasil (1996b). The team orienteering problem. *European Journal of Operational Research 88*, 464–474.

Chekuri, C. and M. Pal (2005). A recursive greedy algorithm for walks in directed graphs. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 245–253.

Copernicus EMS (2016a). Copernicus EMS is supporting national authorities in the rescue operations following the earthquake in central Italy occurred on 24 August 2016. `http://emergency.copernicus.eu/mapping/ems/copernicus-ems-supporting-national-authorities-rescue-operations-following-earthquake-central`. Online; accessed May 15, 2019.

Copernicus EMS (2016b). Upcoming call for tender related to EMS early warning system for forest fires (EFFIS). `http://emergency.copernicus.eu/mapping/ems/upcoming-call-tender-related-ems-early-warning-system-forest-fires-effis`. Online; accessed May 14, 2019.

Copernicus EMS (2018). Copernicus emergency management service. `http://emergency.copernicus.eu/`. Online; accessed May 15,2019.

Cortes, J., S. Martinez, T. Karatas, and F. Bullo (2004). Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation 20*(2), 243–255.

Cressie, N. A. and C. Wikle (2011). *Statistics for Spatio-Temporal Data*. John Wiley & Sons.

Curran, P. J. and P. M. Atkinson (1998). Geostatistics and remote sensing. *Progress in Physical Geography 22*(1), 61–78.

Current, J. R. and D. A. Schilling (1989). The covering salesman problem. *Transportation Science 23*(3), 208–213.

Dang, D.-C., R. El-Hajj, and A. Moukrim (2013). A branch-and-cut algorithm for solving the team orienteering problem. In *International Conference on AI and OR Techniques in Constriant Programming for Combinatorial Optimization Problems*, 332–339.

Dang, D. C., R. N. Guibadj, and A. Moukrim (2013). An effective PSO-inspired algorithm for the team orienteering problem. *European Journal of Operational Research 229*(2), 332–344.

Daniel, K., S. Rohde, N. Goddemeier, and C. Wietfeld (2011). Cognitive agent mobility for aerial sensor networks. *IEEE Sensors Journal 11*, 2671–2682.

Das, A. and D. Kempe (2008). Algorithms for subset selection in linear regression. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, 45–54.

Dunke, F. and S. Nickel (2016). A general modeling approach to online optimization with lookahead. *Omega 63*, 134–153.

Dunke, F. and S. Nickel (2019). Online optimization with gradual look-ahead. *Operational Research*, 1–35.

Euler, J. (2017). *Optimal cooperative control of UAVs for dynamic data-driven monitoring tasks*. Ph. D. thesis, Technische Universität Darmstadt.

European Commission, Joint Research Centre (JRC) (2015). GHS population grid, derived from GPW4, multitemporal (1975, 1990, 2000, 2015). `http://data.jrc.ec.europa.eu/dataset/jrc-ghsl-ghs_pop_gpw4_globe_r2015a`. Online, accessed April 14, 2019.

Fiorelli, E., N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni (2006). Multi-AUV control and adaptive sampling in Monterey Bay. *IEEE Journal of Oceanic Engineering 31*(4), 935–948.

Fischetti, M., J. J. S. Gonzalez, and P. Toth (1998). Solving the orienteering problem through branch-and-cut. *INFORMS Journal on Computing 10*(2), 133–148.

Flanigan, D. F. (1996). Short history of remote sensing of chemical agents. In *Electro-Optical Technology for Remote Chemical Detection and Identification*, Volume 2763, 2–18.

FwDV 100 (1999). Feuerwehr-Dienstvorschrift 100. Führung und Leitung im Einsatz. Führungssystem. `https://www.bbk.bund.de/SharedDocs/Downloads/BBK/DE/FIS/DownloadsRechtundVorschriften/Volltext_Fw_Dv/FwDV%20100.pdf?__blob=publicationFile`. Online; accessed May 14, 2019.

Gendreau, M., G. Laporte, and F. Semet (1997). The covering tour problem. *Operations Research 45*(4), 568 – 576.

Gendreau, M., G. Laporte, and F. Semet (2006). The maximal expected coverage relocation problem for emergency vehicles. *Journal of the Operational Research Society 57*(1), 22–28.

Glock, K. and A. Meyer (2020). Mission planning for emergency rapid mapping with drones. *Transportation Science 54*(2), 534–560.

Golden, B., Z. Naji-Azimi, S. Raghavan, M. Salari, and P. Toth (2012). The generalized covering salesman problem. *INFORMS Journal on Computing 24*(4), 534 – 553.

Gunawan, A., H. C. Lau, and P. Vansteenwegen (2016). Orienteering problem: A survey of recent variants, solution approaches and applications. *European Journal of Operational Research 255*(2), 315–332.

Hà, M. H., N. Bostel, A. Langevin, and L. M. Rousseau (2013). An exact algorithm and a metaheuristic for the multi-vehicle covering tour problem with a constraint on the number of vertices. *European Journal of Operational Research 226*(2), 211–220.

Hachicha, M., M. J. Hodgson, G. Laporte, and F. Semet (2000). Heuristics for the multi-vehicle covering tour problem. *Computers & Operations Research 27*(1), 29–42.

Haddow, G., J. Bullock, and D. P. Coppola (2017). *Introduction to emergency management*. Butterworth-Heinemann.

Harig, R. and P. Rusch (2011). Infrarot-Gefahrstoffkamera. *Forschung im Bevölkerungsschutz 14*.

Hensman, J., N. Fusi, and N. D. Lawrence (2013). Gaussian processes for big data. Technical report.

Hitz, G., E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart (2017). Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics 34*(8), 1427–1449.

Hollinger, G. A. and G. S. Sukhatme (2014). Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research 33*(9), 1271–1287.

Hutchinson, M., H. Oh, and W.-H. Chen (2017). A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors. *Information Fusion 36*, 130–148.

Ichoua, S., M. Gendreau, and J.-Y. Potvin (2006). Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science 40*(2), 211–225.

Indelman, V., L. Carlone, and F. Dellaert (2015). Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research 34*(7), 849–882.

Irnich, S., P. Toth, and D. Vigo (2014). The family of vehicle routing problems. In *Vehicle Routing: Problems, Methods, and Applications, Second Edition*, 1–33. SIAM.

Jawaid, S. T. and S. L. Smith (2015). Informative path planning as a maximum traveling salesman problem with submodular rewards. *Discrete Applied Mathematics 186*, 112–127.

Ke, L., C. Archetti, and Z. Feng (2008). Ants can solve the team orienteering problem. *Computers & Industrial Engineering 54*(3), 648–665.

Ke, L., L. Zhai, J. Li, and F. T. S. Chan (2016). Pareto mimic algorithm: An approach to the team orienteering problem. *Omega 61*, 155–166.

Kemna, S., J. G. Rogers, C. Nieto-Granda, S. Young, and G. S. Sukhatme (2017). Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2124–2130.

Keshtkaran, M., K. Ziarati, A. Bettinelli, and D. Vigo (2016). Enhanced exact solution methods for the team orienteering problem. *International Journal of Production Research 54*(2), 591–601.

Kilby, P. (2013). Tutorial: Constraint programming for the vehicle routing problem (slides). The 19th International Conference on Principles and Practice of Constraint Programming.

Krause, A. and D. Golovin (2014). Submodular function maximization. In L. Bordeaux, Y. Hamadi, and P. Kohli (Eds.), *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press.

Krause, A., H. B. McMahan, C. Guestrin, and A. Gupta (2008). Robust submodular observation selection. *Journal of Machine Learning Research 9*, 2761–2801.

Krause, A., A. Singh, and C. Guestrin (2008). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research 9*, 235–284.

Kuroki, Y., G. S. Young, and S. E. Haupt (2010). UAV navigation by an expert system for contaminant mapping with a genetic algorithm. *Expert Systems with Applications 37*(6), 4687–4697.

Lambrou, T. P. and C. G. Panayiotou (2013). Collaborative path planning for event search and exploration in mixed sensor networks. *The International Journal of Robotics Research 32*(12), 1424–1437.

Li, J. and A. D. Heap (2008). A review of spatial interpolation methods for environmental scientists. Technical report, Australian Government. Geoscience Australia Record 2008/23.

Lim, Z. W., D. Hsu, and W. S. Lee (2016). Adaptive informative path planning in metric spaces . *The International Journal of Robotics Research 35*, 585–598.

Low, K. H., J. M. Dolan, and P. Khosla (2008). Adaptive multi-robot wide-area exploration and mapping. In *International Joint Conference on Autonomous Agents and Multiagent Systems*, 23–30.

Ma, K.-C., L. Liu, and G. S. Sukhatme (2017). Informative planning and online learning with sparse Gaussian processes. In *IEEE International Conference on Robotics and Automation (ICRA)*, 4292–4298.

Manfreda, S., M. F. McCabe, P. E. Miller, R. Lucas, V. Pajuelo Madrigal, G. Mallinis, E. Ben Dor, D. Helman, L. Estes, G. Ciraolo, et al. (2018). On the use of unmanned aerial systems for environmental monitoring. *Remote Sensing 10*(4), 641.

Mayfield, H. T., D. Eastwood, and L. W. Burggraf (2000). Infrared spectral classification with artificial neural networks and classical pattern recognition. In *Chemical and Biological Sensing*, Volume 4036, 54–66.

Megerian, S., F. Koushanfar, M. Potkonjak, and M. B. Srivastava (2005). Worst and best-case coverage in sensor networks. *IEEE Transactions on Mobile Computing 4*(1), 84–92.

Melles, S., G. B. Heuvelink, C. J. Twenhöfel, A. van Dijk, P. H. Hiemstra, O. Baume, and U. Stöhlker (2011). Optimizing the spatial pattern of networks for monitoring radioactive releases. *Computers & Geosciences 37*(3), 280–288.

Mitas, L. and H. Mitasova (1999). Spatial interpolation. In P. Longley, M. Goodchild, D. Maguire, and D. Rhind (Eds.), *Geographical Information Systems: Principles, Techniques, Management and Applications*, Volume 1, Chapter 34.

Naji-Azimi, Z., J. Renaud, A. Ruiz, and M. Salari (2012). A covering tour approach to the location of satellite distribution centers to supply humanitarian aid. *European Journal of Operational Research 222*(3), 596–605.

Neumann, P. P. (2013). *Gas source localization and gas distribution mapping with a micro-drone*. Ph. D. thesis, Freie Universität Berlin.

Newaz, A. A. R., S. Jeong, H. Lee, H. Ryu, and N. Y. Chong (2016). UAV-based multiple source localization and contour mapping of radiation fields. *Robotics and Autonomous Systems 85*, 12–25.

Nguyen, L. V., S. Kodagoda, R. Ranasinghe, and G. Dissanayake (2015). Information-driven adaptive sampling strategy for mobile robotic wireless sensor network. *IEEE Transactions on Control Systems Technology 24*(1), 372–379.

Oliphant, T. (2006). NumPy: A guide to NumPy. `http://www.numpy.org/`. Online; accessed December 03, 2019.

Oliver, M. and R. Webster (2014). A tutorial guide to geostatistics: Computing and modelling variograms and kriging. *Catena 113*, 56–69.

Otto, A., N. Agatz, J. Campbell, B. Golden, and E. Pesch (2018). Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: A survey. *Networks 72*(4), 411–458.

Ozbaygin, G., H. Yaman, and O. E. Karasan (2016). Time constrained maximal covering salesman problem with weighted demands and partial coverage. *Computers & Operations Research 76*, 226–237.

Pavone, M., A. Arsie, E. Frazzoli, and F. Bullo (2011). Distributed algorithms for environment partitioning in mobile robotic networks. *IEEE Transactions on Automatic Control 56*(8), 1834–1848.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research 12*, 2825–2830.

Pěnička, R., J. Faigl, and M. Saska (2019). Variable neighborhood search for the set orienteering problem and its application to other orienteering problem variants. *European Journal of Operational Research 276*(3), 816–825.

Pillac, V., M. Gendreau, C. Guéret, and A. L. Medaglia (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research 225*(1), 1–11.

Pisinger, D. and S. Ropke (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research 34*(8), 2403–2435.

Ponda, S. S., L. B. Johnson, A. Geramifard, and J. P. How (2015). Cooperative mission planning for multi-UAV teams. In V. K. and G. Vachtsevanos (Eds.), *Handbook of unmanned aerial vehicles*, 1447–1490. Springer.

Popović, M., G. Hitz, J. Nieto, I. Sa, R. Siegwart, and E. Galceran (2017). Online informative path planning for active classification using UAVs. In *IEEE International Conference on Robotics and Automation (ICRA)*, 5753–5758.

Popovic, M., T. Vidal-Calleja, J. J. Chung, J. Nieto, and R. Siegwart (2019). Informative path planning and mapping for active sensing under localization uncertainty. Technical report.

Psaraftis, H. N., M. Wen, and C. A. Kontovas (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks 67*(1), 3–31.

QGIS Development Team (2017). Geographic Information System. Open Source Geospatial Foundation Project. `http://qgis.osgeo.org`. Online, accessed May 15, 2019.

Rasmussen, C. and C. Williams (2006). *Gaussian processes for machine learning*, Volume 1. MIT press, Cambridge.

Reggente, M. (2014). *Statistical gas distribution modelling for mobile robot applications*. Ph. D. thesis, Örebro university.

Righini, G. and M. Salani (2008). New dynamic programming algorithms for the resource constrained elementary shortest path problem. *Networks: An International Journal 51*(3), 155–170.

Ristic, B., D. Angley, B. Moran, and J. Palmer (2017). Autonomous multi-robot search for a hazardous source in a turbulent environment. *Sensors 17*(4), 918.

Ritter, T. (2017). *PDE-Based dynamic data-driven monitoring of atmospheric dispersion processes.* Ph. D. thesis, Technische Universität Darmstadt.

Rossi, R. E., J. L. Dungan, and L. R. Beck (1994). Kriging in the shadows: Geostatistical interpolation for remote sensing. *Remote Sensing of Environment 49*(1), 32–40.

Saatçi, Y. (2011). *Scalable inference for structured Gaussian process models.* Ph. D. thesis, University of Cambridge.

Saldaña, D., R. Assunção, and M. F. Campos (2015). A distributed multi-robot approach for the detection and tracking of multiple dynamic anomalies. In *IEEE International Conference on Robotics and Automation (ICRA)*, 1262–1267.

Schwager, M., D. Rus, and J.-J. Slotine (2011). Unifying geometric, probabilistic, and potential field approaches to multi-robot deployment. *The International Journal of Robotics Research 30*(3), 371–383.

Schwager, M., M. Vitus, S. Powers, D. Rus, and C. Tomlin (2017). Robust adaptive coverage control for robotic sensor networks. *IEEE Transactions on Control of Network Systems 4*(3), 462–476.

Seborg, D., D. Mellichamp, T. Edgar, and F. Doyle (2010). Model predictive control. In *Process Dynamics and Control*. John Wiley & Sons.

Settles, B. (2009). Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.

Singh, A., W. Kaiser, M. Batalin, A. Krause, and C. Guestrin (2007). Efficient planning of informative paths for multiple robots. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2204–2211.

Singh, A., A. Krause, C. Guestrin, and W. J. Kaiser (2009). Efficient informative sensing using multiple robots. *Journal of Artificial Intelligence Research 34*, 707–755.

Smith, R. N., M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme (2011). Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *Journal of Field Robotics 28*(5), 714–741.

Souffriau, W., P. Vansteenwegen, G. V. Berghe, and D. Van Oudheusden (2010). A path relinking approach for the team orienteering problem. *Computers & Operations Research 37*(11), 1853–1859.

Spruyt, P. (2017). Copernicus Emergency Management Service. https://www.b2match.eu/system/spaceweek2017-italy/files/1.Spruyt.pdf?1511873843. Presentation presented at Space Week, Rome, 21 - 23 November 2017. Online; accessed May 5, 2019.

Stachniss, C., C. Plagemann, and A. J. Lilienthal (2009). Learning gas distribution models using sparse Gaussian process mixtures. *Autonomous Robots 26*(2-3), 187–202.

Statistisches Bundesamt (Destatis) (2018). Ergebnisse des Zensus 2011. https://www.zensus2011.de/DE/Home/Aktuelles/DemografischeGrunddaten.html?nn=3066576. Online; accessed May 14, 2019.

Stranders, R. (2010). *Decentralised coordination of information gathering agents*. Ph. D. thesis, University of Southampton.

Stranders, R., E. M. De Cote, A. Rogers, and N. R. Jennings (2013). Near-optimal continuous patrolling with teams of mobile information gathering agents. *Artificial Intelligence 195*, 63–105.

Strobel, A. (2016). *Verteilte nichtlineare modellprädiktive Regelung von unbemannten Luftfahrzeug-Schwärmen*. Ph. D. thesis, Technische Universität Darmstadt.

Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography 46*, 234–240.

Todescato, M., A. Carron, R. Carli, G. Pillonetto, and L. Schenato (2017). Multi-robots Gaussian estimation and coverage control: From client–server to peer-to-peer architectures. *Automatica 80*, 284–294.

Truong, S. C., M.-I. Lee, G. Kim, D. Kim, J.-H. Park, S.-D. Choi, and G.-H. Cho (2016). Accidental benzene release risk assessment in an urban area using an atmospheric dispersion model. *Atmospheric Environment 144*, 146–159.

Tsiligirides, T. (1984). Heuristic methods applied to orienteering. *Journal of the Operational Research Society 35*(9), 797–809.

United States National Library of Medicine, National Institutes of Health (2018). Wireless information system for emergency responders (WISER). wiser.nlm.nih.gov. Online; accessed March 17, 2018.

Van Groenigen, J. W., W. Siderius, and A. Stein (1999). Constrained optimisation of soil sampling for minimisation of the kriging variance. *Geoderma 87*(3-4), 239–259.

Vansteenwegen, P., W. Souffriau, G. V. Berghe, and D. Van Oudheusden (2009a). A guided local search metaheuristic for the team orienteering problem. *European Journal of Operational Research 196*(1), 118–127.

Vansteenwegen, P., W. Souffriau, G. V. Berghe, and D. Van Oudheusden (2009b). Metaheuristics for tourist trip planning. In *Metaheuristics in the Service Industry*, 15–31. Springer.

Vansteenwegen, P., W. Souffriau, and D. Van Oudheusden (2011). The orienteering problem: A survey. *European Journal of Operational Research 209*(1), 1–10.

Vidal, T., N. Maculan, L. S. Ochi, and P. H. Vaz Penna (2015). Large neighborhoods with implicit customer selection for vehicle routing problems with profits. *Transportation Science 50*(2), 720–734.

Wackernagel, H. (2013). *Multivariate geostatistics: An introduction with applications*. Springer Science & Business Media.

Wang, M. and X.-S. Hua (2011). Active learning in multimedia annotation and retrieval: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST) 2*(2), 10.

Xu, Y., J. Choi, S. Dass, and T. Maiti (2013). Efficient Bayesian spatial prediction with mobile sensor networks using Gaussian Markov random fields. *Automatica 49*(12), 3520–3530.

Yilmaz, N., C. Evangelinos, P. Lermusiaux, and N. Patrikalakis (2008). Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering 33*(4), 522–537.

Yu, J., M. Schwager, and D. Rus (2014). Correlated orienteering problem and its application to informative path planning for persistent monitoring tasks. In *IEEE International Conference on Intelligent Robots and Systems*, 342–349.