Autonomous Coverage Path Planning For AUVs Considering Location Uncertainty

S.Y. Deen

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

Autonomous Coverage Path Planning For AUVs Considering Location Uncertainty

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

S.Y. Deen

September 14, 2017

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of Technology

The work in this thesis was supported by TNO. Their cooperation is hereby gratefully acknowledged.

The image on the cover page is owned by OceanScan-MST.

Delft University of Technology
Department of
Delft Center for Systems and Control (dcsc)

The undersigned hereby certify that they have read and recommend to the Faculty of Mechanical, Maritime and Materials Engineering (3mE) for acceptance a thesis entitled

Autonomous Coverage Path Planning For AUVs Considering Location Uncertainty

by

S.Y. Deen

in partial fulfillment of the requirements for the degree of

Master of Science Systems and Control

Dated: September 14, 2017

Supervisor(s):

dr.ir. J. Sijs

Reader(s):

prof.dr.ir. B. de Schutter

dr. J. Alonso-Mora

dr. A. Núñez

ir. J. Fransman

Abstract

Autonomous underwater vehicles (AUVs) are unmanned vehicles that operate underwater. These vehicles can be used for various operations, including scanning the ocean floor in order to search for objects. In such operations, the AUV navigates close to the ocean floor, using its sonar systems to map this ocean floor.

Before an AUV is launched on a search operation, usually very little is known about the underwater conditions at the search location. Currently, the AUV's coverage path is often defined before the start of the operation. However, the coverage path depends on the quality of the sonar images, which in turn depend on the unknown and varying underwater conditions. Therefore, it is unlikely that the predefined coverage path is optimal. In this final thesis project a coverage path planner is developed, which autonomously adapts the AUV's coverage path to the changing underwater conditions during an operation.

Since the radio waves carrying the global positioning system (GPS) signal do not travel far in seawater, the AUV needs to keep track of its location using an inertial navigation system (INS). Unfortunately, the error on the location measurements obtained from an INS grows over time, causing the AUV's believed trajectory to differ significantly from the AUV's true trajectory. As a result, parts of the search area might be believed to have been covered, while these parts have in reality not been visited. Taking the AUV's uncertain trajectory into account thus increases the complexity of the coverage path planning problem.

In order to keep track of which parts of the search area have been covered, and to which extent, a coverage map of the search area is maintained during an operation. This coverage map is constructed such that it takes the uncertainty on the AUV's location measurements into account. Three different coverage path planners have been developed that use this coverage map to autonomously plan the AUV's coverage path during an operation.

The three different coverage path planners have been implemented in MATLAB, and they have been evaluated in a simulation framework developed by TNO. The simulation results show that all three coverage path planners consistently achieve the required coverage of the

search area, in a scenario where no large changes in the underwater conditions occurred. Furthermore, one of the coverage path planners has been evaluated in three scenarios where large changes during the operation did occur. Halfway during these scenarios, either the strength of the underwater currents increased, the required coverage of the search area increased, or one of the AUV's sonar systems broke down. In all of these scenarios, the coverage path planner successfully adapted the coverage path such that the required coverage of the search area was still achieved.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to thank my supervisor dr.ir. J. Sijs, for his guidance, assistance and motivation during this final thesis project. You were always ready to help, and it has been a pleasure working together. I would also like to thank ir. J. Fransman, for his feedback on my presentation, and drs. A. Hensbergen, dr. P. Mohajerin Esfahani and dr. A. Cabo, for their useful advice.

Furthermore, I would like to thank ir. J. de Bruijckere, for always taking the time to think about my questions, and for always having the answers. Special thanks go to my brother, for his advice and support, and for making the long days spent at the university together a lot more fun. Finally, I would like to thank my parents, who I owe my graduation to in every possible way.

Thank you all!

Delft, University of Technology S.Y. Deen

September 14, 2017

Chapter 1

Introduction

1-1 Motivation

Autonomous underwater vehicles (AUVs) are unmanned vehicles that operate underwater. These vehicles are not directly nor remotely operated by a human operator. Instead, the AUV is allowed to take some operational decisions online during the operation. The fact that no human is situated in these vehicles makes them very suitable for long duration operations or operations in hazardous environments. Typical operations for which AUVs are used include the mapping of the ocean floor [1], inspection of underwater structures [2], mine countermeasures (MCM) [3] and marine research purposes [4]. An example of what an AUV can look like is shown in Figure 1-1.

Figure 1-1: OceanScan-MST's LAUV.

The light autonomous underwater vehicle (LAUV) [5] shown in Figure 1-1 (and on the cover page) is produced by OceanScan-MST. This is the type of AUV that has been bought by the Dutch organisation for applied scientific research TNO. The main purpose of this AUV is to scan the ocean floor, in order to search for objects.

While an AUV is scanning the ocean floor looking for objects, it is following a certain path. Since the usual objective of a search operation is to completely cover an area while searching, the path that an AUV follows during this search is called a coverage path.

Most of the research in path planning considers the problem of safely and efficiently getting from point A to point B. Although this type of path planning is also important in AUV search operations, the main objective in these search operations is to pass the AUV's sonar systems over the entire search area. The problem of finding a path that accomplishes this is called a coverage path planning (CPP) problem. It should be noted that a CPP problem is conceptually different from a conventional path planning problem. In conventional path planning problems the focus is on the destination, whereas in CPP problems the focus is on the journey [6].

The coverage path in AUV search operations usually has the shape of a lawnmower pattern. Two different lawnmower patterns to cover the same area are shown in Figure 1-2.

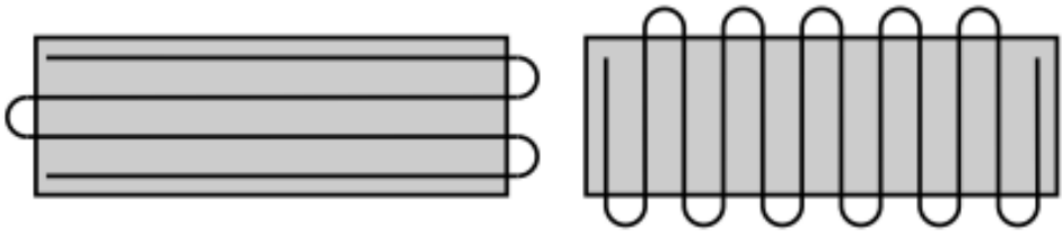

Figure 1-2: Two different lawnmower patterns to cover the same area.

There are two main reasons why the coverage path of an AUV on a search operation commonly has the shape of a lawnmower pattern. The first reason concerns the sonar characteristics, while the second reason concerns the uncertainty on the AUV's location measurements.

In order to scan the ocean floor, the AUV is equipped with two side-looking sonar systems, on both the port and starboard side of the AUV. In addition, the AUV is usually equipped with a forward-looking sonar (FLS) system to detect obstacles in front of the AUV. Although an FLS system could also be used in an attempt to cover the nadir gap [7], in practice only the side-looking sonar systems are used in order to scan the ocean floor. These side-looking sonar systems consist of either two side-scan sonar (SSS) systems or two synthetic aperture sonar (SAS) systems. For both of these sonar systems the image quality obtained while the AUV has been travelling in a straight line is far superior to the image quality obtained while the AUV has been following a curved trajectory. For this reason only straight tracks are used to cover an area, resulting in the lawnmower-shaped coverage paths shown in Figure 1-2.

The second reason for the lawnmower shape of the coverage paths is related to the way that the AUV obtains its location measurements. While the AUV is at the surface, it can make use of the global positioning system (GPS) to accurately obtain its location. However, due to the high electrical conductivity of seawater, the radio waves carrying this GPS signal do not travel far underwater. For this reason, the AUV has to rely on an inertial navigation

system (INS) to determine its location during an operation. An INS typically consists of three orthogonal accelerometers and three orthogonal gyroscopes. Using a dynamical model of the AUV and a state estimator, the estimated location and orientation of the AUV can be obtained from the measured rotational velocities and accelerations. When only the accelerometers and gyroscopes are used to estimate the location and orientation of the AUV, unknown biases on these sensors cause the estimated location and rotation to drift. Due to this integration drift, the error on the location and orientation estimates increases unbounded over time. How fast this error increases depends on a number of factors, but it is shown in [8] that the error increases relatively slowly when the AUV travels two parallel tracks in opposite ($\pm 180°$) directions. This is the second reason that the lawnmower pattern shown in Figure 1-2 is the standard coverage path used on AUV search operations.

Currently, the coverage path for a search operation is defined before the start of the operation. During the operation, the AUV follows this predefined coverage path until completion. In order to define a lawnmower-shaped coverage path, the distance between the straight tracks has to be known. This track-spacing distance depends mainly on the sonar range, the maximum distance where the resolution of the sonar images is still high enough to detect the objects of interest using automatic target recognition (ATR). However, this sonar range greatly depends on the environmental underwater conditions at the search location [9]. An example of how the underwater conditions can influence the sonar image quality is shown in Figure 1-3 [10].

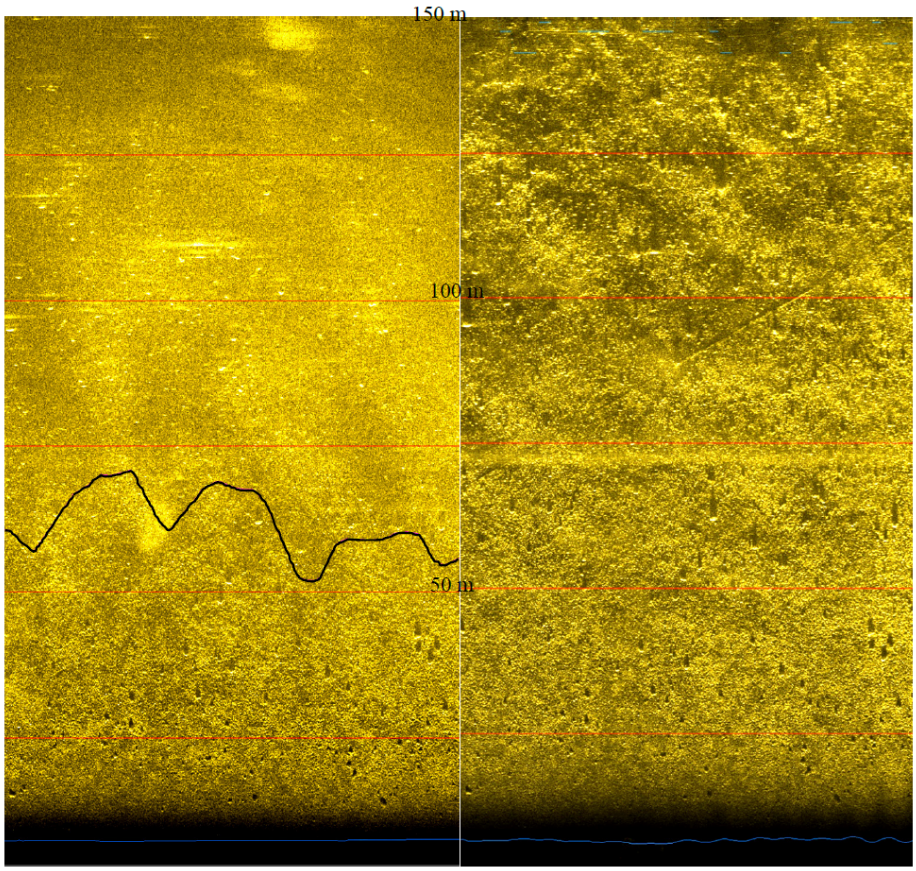

Figure 1-3: Two SAS images of the same scene, taken on different days.

Figure 1-3 shows two sonar images taken at the exact same location, one week apart. It can be seen that the sonar range in the left sonar image of Figure 1-3 is between 55 and 70 metres, while the sonar range in the right sonar image of Figure 1-3 is about 150 metres. This difference in sonar range is caused only by the environmental conditions, since the images were taken at the same location using the same AUV. This illustrates the influence that the underwater conditions can have on the sonar performance.

Since the underwater conditions at the search location are usually unknown at the start of the operation, it is unlikely that the track-spacing in the predefined coverage path is optimal. Furthermore, even if the predefined coverage path would be optimal at the start of the operation, changing underwater conditions during the operation could cause the track-spacing to become sub-optimal again for a complete coverage of the area. Both of these problems can be solved by adaptive coverage path planning. Instead of following the predefined coverage path, the AUV can in that case autonomously asses its performance during the operation and use this information to adjust its coverage path online, such that the coverage in the search area is optimized.

Special attention has to be paid to the fact that the uncertainty on the location measurements obtained through the INS grows over time. Because of the uncertainty on the location measurements, it becomes uncertain what trajectory the AUV has sailed exactly. This can cause the AUV to believe that it has covered the entire search area, while in reality some parts of the search area might have never been visited. For this reason, the growing uncertainty on the AUV's location measurements can not be ignored while planning a coverage path.

1-2 Problem Statement

The motivation in the previous section leads to the following problem statement of this final thesis project:

Design and implement an autonomous coverage path planning method for AUVs considering location uncertainty.

This problem statement can be divided into the three sub-problems listed below, which are all discussed in this final thesis project.

1. Obtain an autonomy framework.
2. Obtain a coverage map.
3. Obtain a coverage path planner.

The next section describes the related research on this topic, after which the outline of this final thesis project is presented.

1-3 Previous Work

The first sub-problem is to obtain an autonomy framework. In order to obtain persistent autonomy in an AUV, different system architectures can be used. Examples of currently used architectures in autonomous systems are the control architecture for robotic agent command and sensing (CARACaS) [11], the four-dimensional real-time control system (4D/RCS) [12] and the system architecture [13] used in the PANDORA project [14]. All of these architectures are more or less based on the observe, orient, decide & act (OODA) decision cycle [15]. The autonomy framework used in this final thesis project is inspired by these system architectures.

The second sub-problem is to obtain a coverage map. This coverage map has to keep track of the AUV's performance in the search area. A clever way of defining the AUV's performance on a search operation is described in the ESPRESSO model [16], developed by the North Atlantic treaty organization (NATO). In [17] a framework is presented that incorporates the uncertainty on the AUV's location measurements in the measure of performance described by the ESPRESSO model. In this final thesis project, this framework is implemented and adjusted to correct a mathematical error and to significantly reduce its computational cost, while considering the location uncertainty in two dimensions instead of in one dimension.

The third sub-problem is to obtain a coverage path planner. This planner makes use of the coverage map in order to adaptively plan a coverage path. A survey on path planning in AUVs can be found in [18]. In [19] a coverage path planning method is presented that attempts to minimize the uncertainty on the location measurements. The idea for a cost function based on the expected entropy reduction is presented in [20]. Inspired by these CPP methods, a new method to predict and optimize the coverage of the search area is presented, implemented and evaluated in this final thesis project.

1-4 Thesis Outline

The structure of this final thesis project report is based on the order of the three sub-problems. First, the system architecture is discussed in Chapter 2. This architecture consists of two parts, namely the simulation framework and the autonomy framework. In Chapter 3 it is discussed how the performance during a search operation is defined, how the uncertainty on the location measurements can be included into this performance and how to continuously update this performance. In Chapter 4 it is discussed how the resulting coverage map is used to adaptively plan a coverage path, after which the implemented coverage path planner is evaluated in Chapter 5. Finally, the conclusions of this final thesis project and suggestions for future work are presented in Chapter 6.

Chapter 2

System Architecture

2-1 Simulation Framework

Ideally, the autonomous coverage path planning method that considers location uncertainty should be tested on a real AUV. However, since there has been no access to a real AUV during this final thesis project, simulations have to be used in order to study the results of the developed autonomous CPP method. At TNO, a simulation framework that can simulate various AUV operations has been developed in MATLAB. This framework is used in this final thesis project in order to design, implement and evaluate the autonomous CPP method.

The used simulation framework provides an environment to simulate AUV operations. Discussing the entire framework is beyond the scope of this final thesis project, but to illustrate that the simulations take all the aspects that are relevant to this final thesis project into account, a selection of simulation components is mentioned below.

- The simulation contains a model of the AUV's kinematics in 3-D.
- The simulation contains a model of the INS that estimates the AUV's location and orientation, with increasing uncertainty over time.
- The simulation contains a GPS model, providing a location measurement when the AUV is near the surface.
- The simulation contains different navigation controllers in order to navigate the AUV to a specified location.
- The simulation contains a model for the two sonar systems mounted on the AUV, providing sonar image resolutions in the sonar range depending on the travelled trajectory.
- The simulation contains a complex underwater environment, with constant or varying currents and different types of ocean floor.

The simulation framework has been implemented in MATLAB using object-oriented programming techniques, by creating classes for the AUV's functions and components. These classes include for example a class for the kinematics of the AUV, a class for the INS, a class for the sonar systems, and an autonomy class. This final thesis project focusses on implementing an autonomy framework inside this autonomy class, enabling the AUV to autonomously plan a coverage path. In order to implement the autonomy framework, some small adjustments have been made to the other classes in the simulation framework as well.

The interactions between the simulation framework and the autonomy framework can be simplified by defining `Inputs` and `Outputs` as shown in Figure 2-1. Note that the inputs and outputs are defined relatively to the autonomy framework.

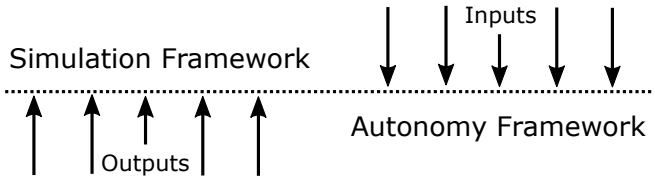

Figure 2-1: Simplified interactions inside the system architecture.

The inputs to the autonomy framework consist of all the information that the AUV can obtain from its sensors in the simulation framework, such as its sonar systems and INS or GPS sensor. The frequencies with which these inputs arrive can vary. The outputs that the autonomy framework provides to the simulation framework with a sampling frequency of 2 Hertz are listed below.

- Goal waypoint : $\in \mathbb{R}^3$
- Reference velocity [m/s] : $\in [0, 5]$
- Sonar systems : $\in \{\text{on}, \text{off}\}$
- Navigation controller : $\in \{\text{relaxed}, \text{aggressive}\}$

From the above outputs, changing the goal waypoint is the main way of controlling the coverage path of the AUV. This final thesis project introduces an extra parameter associated with the goal waypoint, named the mode of travel. The mode of travel can either be 'travel' or 'scan', and this mode specifies how the AUV should approach the goal waypoint. When the AUV is not on any of the straight tracks in the lawnmower pattern, it is approaching the next waypoint in the 'travel' mode. In this mode, the reference velocity is set to the AUV's maximum velocity of 5 m/s, the sonar systems are switched off, and the aggressive navigation controller is used. In this mode, the focus is on the destination. The second mode of travel, the 'scan' mode, is used when the AUV is on any of the straight tracks in the lawnmower pattern. In this mode, the reference velocity is set to 2 m/s, the sonar systems are switched on and the relaxed navigation controller is used. In this mode, the focus is on the journey. The reference velocity of 2 m/s has been chosen by studying the simulation results. The simulations show that increasing the reference velocity decreases the sonar image quality significantly, while decreasing the reference velocity does not significantly increase the sonar

image quality. The relaxed navigation controller is chosen because it controls the position of the AUV more smoothly than the aggressive navigation controller, resulting in a higher sonar image quality. The fact that the AUV controlled with the relaxed navigation controller has a significant smaller maximum turn rate (10°) than the AUV controlled with the aggressive navigation controller (45°) is not very important, since the scanning mode is only active on straight trajectories.

Although all of the four discussed outputs have to be provided to the simulation framework at a constant rate of 2 Hertz, introducing the mode of travel only leaves the two parameters that are shown below for the autonomy framework to decide on.

- Goal waypoint : $\in \mathbb{R}^3$
- Mode of travel : $\in \{\text{travel}, \text{scan}\}$

2-2 Autonomy Framework

The task of the autonomy framework inside the system architecture is to provide the simulation framework with the outputs discussed in the previous section. While doing so, the autonomy framework receives feedback from the simulation framework. However, the nature of the outputs makes this feedback loop a little different from the feedback loops known from classical control theory. This is illustrated in Figure 2-2 [21].

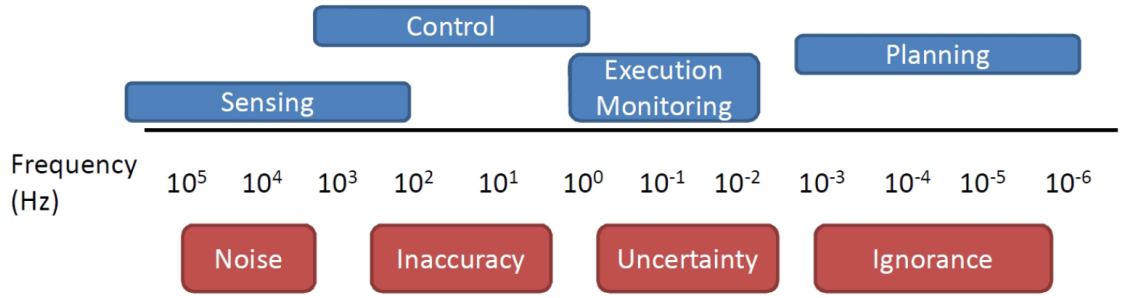

Figure 2-2: Different levels of reasoning activities.

In Figure 2-2 four different levels of reasoning activities can be seen, namely `Sensing`, `Control`, `Execution Monitoring` and `Planning`. These reasoning activities are ordered by the typical frequency that the events with which they are associated occur with. When the level of the reasoning activity decreases, this typical frequency increases. The mismatches between the coordinated reasoning activities and reality are shown at the bottom of Figure 2-2 in red. Normally the noise and inaccuracy are attempted to be minimised by signal processing and robust control respectively, while the uncertainty and ignorance are attempted to be minimised through planning and replanning.

Figure 2-2 shows that planning is a high-level reasoning activity, that coordinates the execution monitoring, control and sensing of the system. Since the problem statement of this final thesis project is to design and implement a coverage path planner, the focus of this thesis is on the planning and execution monitoring side of the reasoning activities. The sensing and control side of the reasoning activities is handled by the simulation framework, using the AUV sensor and actuator models and its controllers. How these lower-level reasoning activities are handled exactly is beyond the scope of this final thesis project, which focusses on the higher-level reasoning activity of deciding the waypoints and mode of travel of the AUV.

The autonomy framework has to decide on the waypoints and mode of travel using the AUV's sensor data obtained in the simulation framework. In autonomous systems, there are various ways to decide on the outputs given the inputs, but not all approaches are equally suitable in this final thesis project. Consider for example an autonomy framework consisting of a well trained artificial neural network (ANN), that maps inputs to desirable outputs. In this autonomy framework, the decision process is almost completely a black box. Although the weights and activation functions are known, there is no real insight in the decisions process. When it is for example autonomously decided to abort an operation and to surface the AUV, the operator needs to be able to check why the autonomy framework decided to do so. This illustrates that in this final thesis project a more transparent decision cycle for the autonomy framework is preferred.

The OODA loop [15] is such a transparent decision cycle. This decision cycle was originally developed for the military, but nowadays it is applied in a variety of fields. In fact, many system architectures for recent autonomous systems are based on this decision cycle. Examples are CARACaS [11], 4D/RCS [12] and the system architecture [13] used in the PANDORA project [22]. An abstract depiction of the OODA loop is shown in Figure 2-3.

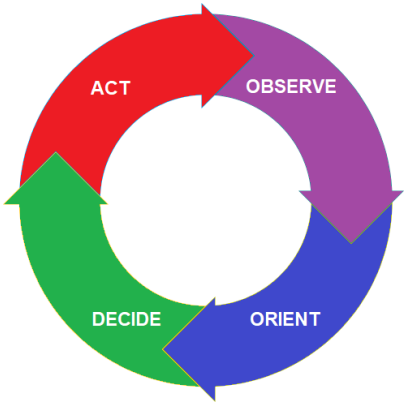

Figure 2-3: The abstract OODA loop.

In this final thesis project, an autonomy framework has been developed that is based on the abstract OODA loop shown in Figure 2-3. This autonomy framework is shown in Figure 2-4, and it is inside this autonomy framework that the autonomous coverage path planner that considers location uncertainty is implemented in this final thesis project. Figure 2-4 also shows how the developed autonomy framework interacts with the simulation framework inside the system architecture.

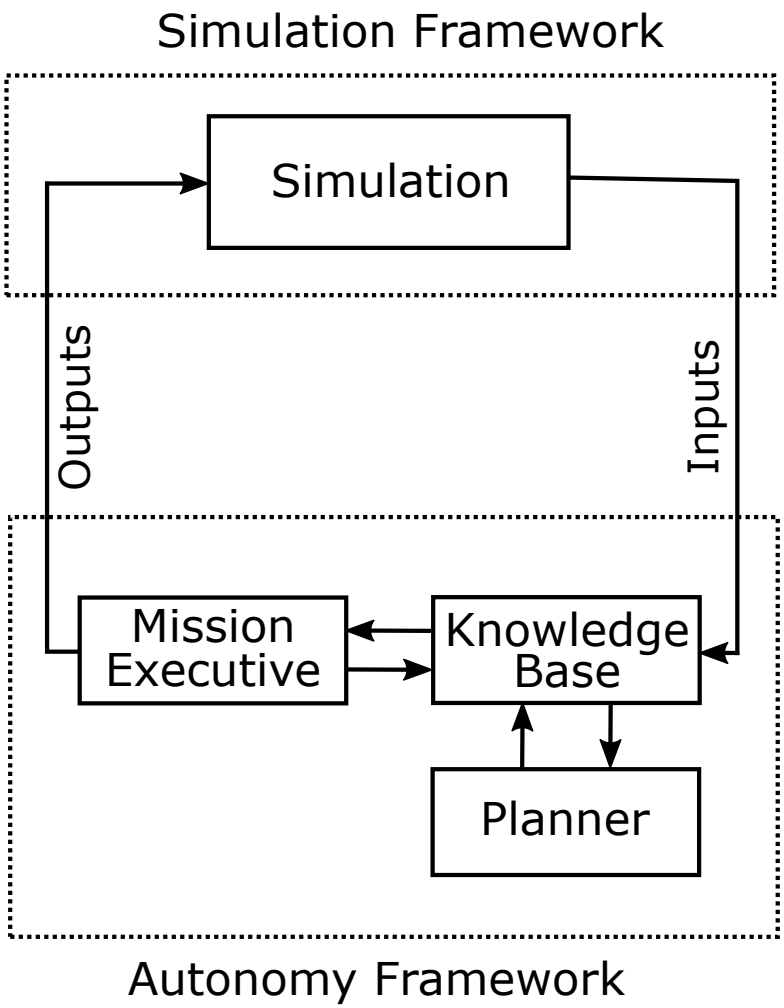

Figure 2-4: The simulation and autonomy frameworks inside the system architecture.

Figure 2-4 shows that the autonomy framework contains three agents, namely the `Knowledge Base` agent, the `Planner` agent and the `Mission Executive` agent. The inputs that the autonomy framework receives from the simulation framework are processed and stored inside the knowledge base agent. The knowledge base agent can also be initialised with additional expert knowledge by the operator. In the autonomy framework, all the relevant information about the current operation is centrally available for every agent inside this framework. The planner agent uses the information stored in the knowledge base agent to generate a plan, which in this final thesis project consists of a coverage path described by a set of waypoints and the corresponding modes of travel to these waypoints. This coverage path plan is stored inside the knowledge base agent, and translated by the mission executive agent into the required outputs that need to be provided to the simulation framework at a constant rate of 2 Hertz. The planner agent however is not required to run at this same frequency, since the planner agent can by bypassed as shown in Figure 2-4. In fact, the planner agent only generates a new plan when it finds that a condition that negatively affects the current plan has been changed inside the knowledge base agent, due to the simulation framework inputs being processed and stored.

As mentioned before, it can be seen in Figure 2-4 that the system architecture is based on the OODA loop. The observe phase of the OODA loop takes place in the simulation framework, where data from all the AUV's sensors is collected. In the orient phase of the OODA loop, this data is processed and stored into the knowledge base agent. During the decide phase of the OODA loop, the outputs of the autonomy framework are generated through interactions between the knowledge base, planner and mission executive agents. Finally, during the act phase of the OODA loop, the outputs of the autonomy framework result in the use of the AUV's actuators in the simulation framework. This means that the observe and act phases of the OODA loop are located inside the simulation framework, while the orient and decide phases of the OODA loop are located inside the autonomy framework.

In Chapter 3 the coverage map that considers location uncertainty is discussed, and it is now known that this coverage map is maintained during the orient phase of the OODA loop, inside the knowledge base agent. Subsequently, Chapter 4 discusses the optimization of the coverage path, which takes place during the decide phase of the OODA loop, mainly inside the planner agent.

Chapter 3

Coverage Under Location Uncertainty

3-1 Coverage Measurement

An AUV that is searching for objects in an area needs to keep track of what parts of this area it has covered already, as well as how well it has covered these parts. For this purpose, the search area is discretised into grid cells, such that a measure of coverage can be maintained for every grid cell in the area. How exactly the coverage is quantified is discussed later in this chapter, first it is discussed how the search area is discretised.

The smaller the size of the grid cells of the discretised area map gets, the closer this discretised map approaches the real (continuous) situation. Decreasing the grid size however increases the computational cost of maintaining the coverage map. In this final thesis project, the grid cells measure 2×2 metres. Since the AUV's location is discretised to the centre of the grid cell that it is located in, the maximum location error that can be caused by this discretisation is $\sqrt{2} \approx 1.41$ metres. The expected value of the location error that is caused by the discretisation of the AUV's location to the centre of a grid cell can also be computed, given that the location distribution of the AUV inside the grid cell is known. When the AUV's location is uniformly distributed inside a grid cell, the expected value of the error caused by the discretisation can be computed using the `Law of the unconscious statistician` [23], shown in Equation (3-1).

$$\mathrm{E}[g(\boldsymbol{X}, \boldsymbol{Y})] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) f_{\boldsymbol{X},\boldsymbol{Y}}(x, y)\, dx\, dy \tag{3-1}$$

When in Equation (3-1) a joint uniform distribution over the grid is taken for $f_{\boldsymbol{X},\boldsymbol{Y}}(x, y)$ while $g(x, y)$ describes the distance to the centre of the grid cell, Equation (3-1) simplifies to Equation (3-2), which has an exact solution.

$$\int_0^1 \int_0^1 \sqrt{x^2 + y^2}\, dx\, dy = \frac{1}{3}(\sqrt{2} + \sinh^{-1}(1)) \approx 0.765 \tag{3-2}$$

From Equation (3-2) it can be concluded that the expected value of the error introduced by discretising the position is less than 77 centimetres, which is very reasonable compared to the dimensions of the AUV.

In order to measure if a grid cell is covered, it does not suffice to only check if the sonar systems have passed over this grid cell. For example, the resolution of the sonar image that has been obtained from this grid cell can be too low to detect the object that is searched for. Due to characteristics of the sonar systems used in real operations, the resolution of the sonar images deteriorates when the AUV's trajectory becomes less smooth during the scanning, which can lead to a sonar image resolution below the threshold for detection of an object. In addition to the sonar image resolution, the ocean floor type can also prevent the detection of an object. When the ocean floor for example consists of a bed of seagrass, the AUV can never detect an object, regardless of the sonar image resolution.

In order to quantify the coverage in a grid cell, the ESPRESSO [16] model can be used. The ESPRESSO model, short for the extensible performance and evaluation suite for sonar, has been developed by the NATO by analysing a large number of trials with real AUVs. Many parameters have been found to affect the coverage, including the velocity and depth of the AUV, the frequency and range of the sonar systems, properties of the object of interest, the water depth, temperature, clarity and salinity and the type of ocean floor. The ESPRESSO model takes all these parameters into account, and maps them into a probability of detection (POD). This POD is used to quantify the coverage in a grid cell in this final thesis project.

When O is the event that there is an object located in a grid cell, then O^c is the event that there is no object located in that grid cell. When $\mathrm{P}(O)$ is defined as the probability of the event O, then $\mathrm{P}(O^c) = 1 - \mathrm{P}(O)$. In a similar manner the events D and D^c can be defined, which are the events that the AUV respectively detects or does not detect an object in a grid cell. The four possible intersections of the events O and D and their complements are shown in Figure 3-1.

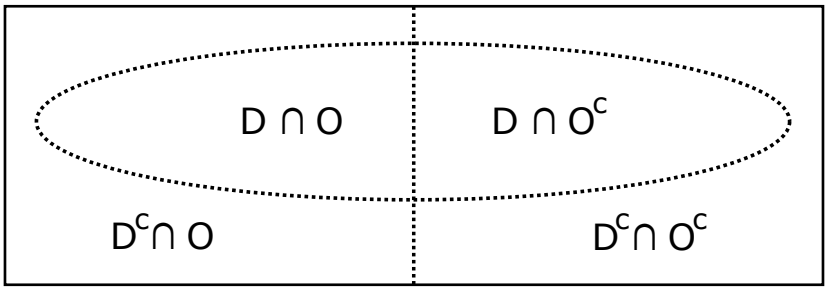

Figure 3-1: Intersections of two events and their complements.

Table 3-1 lists the four intersection shown in Figure 3-1 with a description, and classifies them as true or false positives or negatives.

Table 3-1: True and false observations.

$D \cap O$	True positive	There is a detection and an object
$D \cap O^c$	False positive	There is a detection but no object
$D^c \cap O^c$	True negative	There is no detection and no object
$D^c \cap O$	False negative	There is no detection but there is an object

Table 3-1 shows that there are two possible true observations of a grid cell, either there is a detection and an object, or there is no detection and no object. The POD is defined as the probability of a true observation of a grid cell, as shown in Equation (3-3).

$$\text{POD} = \text{P}(D \cap O) + \text{P}(D^c \cap O^c) \tag{3-3}$$

Consequently, the complement of the POD is defined as the probability of a false observation of a grid cell. There are two possible false observations of a grid cell, either there is a detection but no object, or there is no detection but there is an object. The complement of the POD is shown in Equation (3-4).

$$1 - \text{POD} = \text{P}(D \cap O^c) + \text{P}(D^c \cap O) \tag{3-4}$$

When all the parameters that affect the coverage are specified, the ESPRESSO model can generate a curve that defines the POD as a function of the lateral distance to the AUV. In Figure 3-2 [20] three of these curves are shown, generated for three different types of ocean floor. Figure 3-2 shows the impact of both the ocean floor type and the lateral distance to the AUV on the POD.

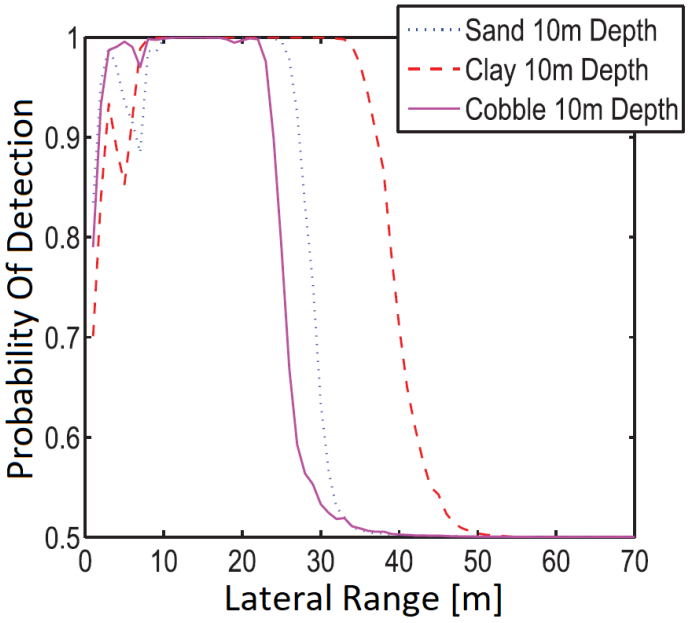

Figure 3-2: Three different curves generated by the ESPRESSO model.

In this final thesis project, a simplification of the ESPRESSO model is used. In this simplification, the POD is a function of the ocean floor type and of the maximum theoretical sonar image resolution only. This maximum theoretical sonar image resolution is estimated in the simulation framework, based on the sonar characteristics and on the estimated trajectory that the AUV has been following while acquiring the sonar images. This means that the maximum theoretical sonar image resolution is influenced by a number of parameters, including the AUV's velocity, the type of navigation controller that is currently active, the accuracy of the INS location and orientation measurements, the strength and direction of the underwater currents, and the distance to the AUV's sonar systems.

Once the maximum theoretical sonar image resolution and the ocean floor type in a grid cell are known, they can be mapped into a POD through the function $\mathcal{P}(r, o)$, where r [m] is the maximum theoretical sonar image resolution and o is the type of ocean floor. In this final thesis project a flat sandy ocean floor is considered, for which the PODs are given in Equation (3-5).

$$\mathcal{P}(r, \text{sand}) = \begin{cases} 1 & \text{if} \quad r \leq 0.5 \\ 0.95 & \text{if} \quad 0.5 < r \leq 0.6 \\ 0.9 & \text{if} \quad 0.6 < r \leq 0.7 \\ 0.8 & \text{if} \quad 0.7 < r \leq 0.9 \\ 0.5 & \text{if} \quad 0.9 < r \leq 1.4 \\ 0 & \text{if} \quad 1.4 < r \end{cases} \tag{3-5}$$

When the function shown in Equation (3-5) is used to determine the PODs during a simulation, the resulting POD surface plot after scanning a straight track can look like Figure 3-3. In Figure 3-3, the highest PODs are coloured yellow, while the lowest PODs are coloured blue.

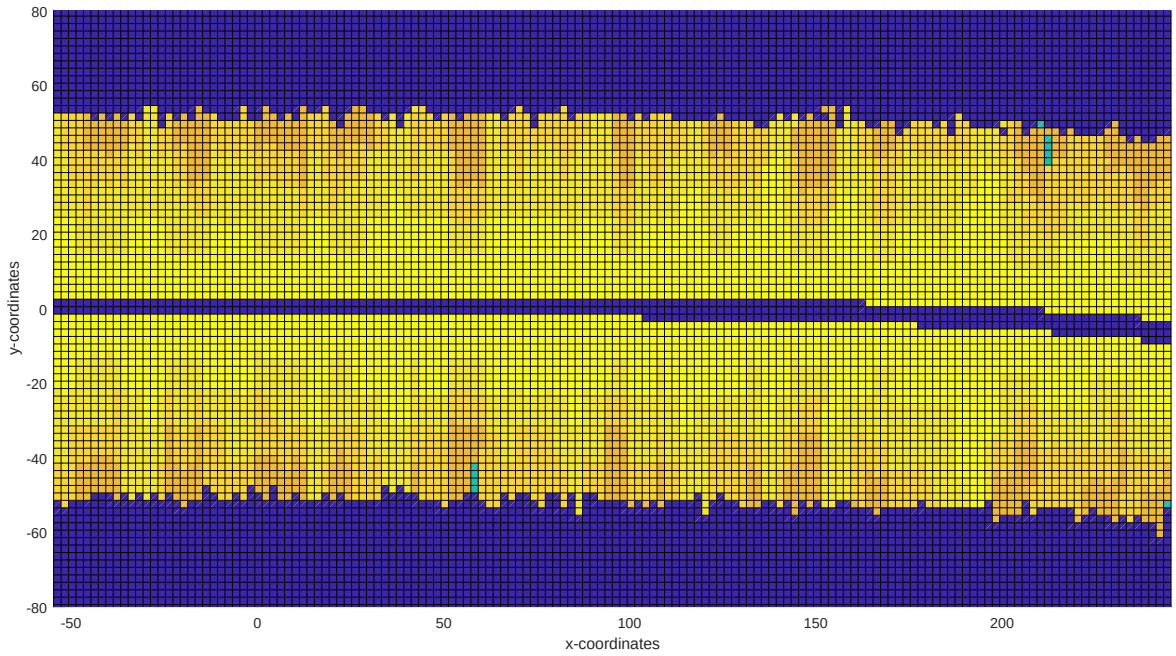

Figure 3-3: A surface plot of the PODs.

The PODs shown in Figure 3-3 have been obtained while the AUV was in scanning mode, as discussed in Chapter 2, and while there were virtually no disturbances in the form of underwater currents. The PODs have been obtained using the AUV's real location, which is not available through the AUV's sensors and thus also not available to the autonomy framework. Using the real location of the AUV enables the comparison of Figure 3-3 with reality, and the obtained PODs are deemed realistic by TNO. Note that the AUV has a deviation to the right when it travels in positive x-direction, while the underwater currents

were negligible during the simulation. This deviation is the result of the INS sensor drift, and it is discussed into more detail in the next section.

3-2 Location Uncertainty

In Chapter 1 it has briefly been discussed that an AUV has to use an INS to keep track of its location, since the radio waves that carry the GPS signal do not travel far in seawater due to its high electrical conductivity. This INS uses the kinematic equations to keep track of the AUV's location and orientation, by integrating the gyroscope measurements once and the accelerometer measurements twice. An INS usually consists of three of these orthogonal accelerometers and three orthogonal gyroscopes. Unfortunately, all inertial navigation systems suffer from integration drift. Integration drift occurs due to small errors in the accelerometer and gyroscope measurements. Since these measurement errors are integrated in the kinematic equations, they lead to an increasing error in the computed location and orientation over time. Since computing the location requires two integrations and computing the orientation only one, the error on the computed location grows faster than the error on the computed orientation. An example of how the location error in the horizontal plane of a modern INS grows over time is given in Figure 3-4 [19].

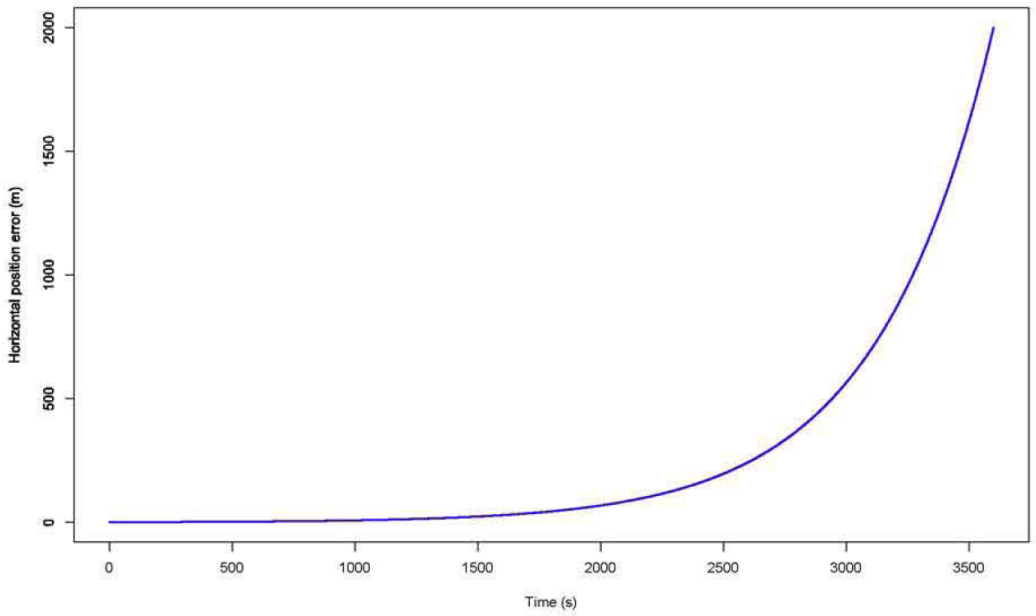

Figure 3-4: Horizontal plane error over time for a modern INS.

How fast the location error grows exactly depends on many factors, such as the AUV's velocity, the strength and direction of the underwater currents and the trajectory that has been followed by the AUV. The integration of additional sensors in the INS, such as a Doppler velocity log (DVL), a magnetic compass or a depth meter, can reduce the rate with which the location error increases. In fact, integrating the depth meter in the INS can bound the location error in the vertical direction. However, none of the additional sensors can prevent the location error from growing unbounded over time in the horizontal plane.

Given that the typical AUV search operation lasts between 3 and 8 hours, the horizontal plane location error will have grown to an unacceptable value before the end of the operation has been reached. In order to reset this location error during the operation, the AUV can interrupt the scanning of the ocean floor and surface, to obtain a GPS location measurement once it arrives at the ocean surface. This maneuver is called a GPS-fix. Performing a GPS-fix brings an additional cost regarding the duration of the operation, but this cost does not outweigh the benefit of practically eliminating the uncertainty on the AUV's location. In Figure 3-5, the simulated real and measured trajectory of an AUV on a search operation are shown.

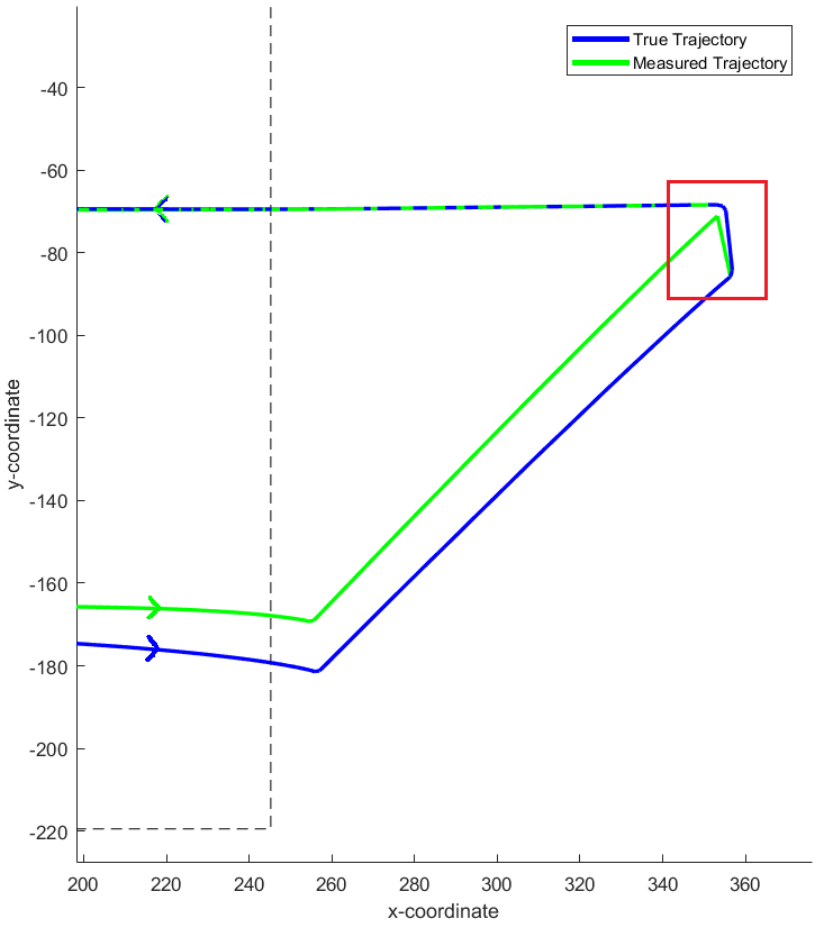

Figure 3-5: The real trajectory and the measured trajectory of an AUV during a GPS-fix.

The dotted black line in Figure 3-5 marks the boundary of the search area, and the AUV exits this search area after a 300 metre long straight track at the bottom left of Figure 3-5. At this point in time the difference between the real and the measured location of the AUV is roughly 12 metres, which is a realistic value after a 300 metre long straight track.

After the AUV has completed this straight track, it surfaces for a GPS-fix. Since Figure 3-5 is in 2-D, the z-coordinate of the AUV's trajectory can not be seen, but it can be seen that there is a jump in the measured location of the AUV within the right box at the top

right of Figure 3-5. This jump is due to the fact that at this moment the AUV has reached the surface, where it receives a GPS signal. Since this GPS location measurement is more accurate than the INS location measurement, the AUV's measured location becomes equal to the measured GPS location. This is what happens in Figure 3-5, where the measured location can be seen to jump to practically the real location.

Figure 3-5 shows how the location measurement uncertainty can periodically be reset by performing a GPS-fix. In this final thesis project, the AUV performs such a GPS-fix after every straight track of its lawnmower-shaped coverage path. The GPS-fixes are executed after every straight track in order to keep the location uncertainty at a minimum, and to start every straight track with approximately the same location uncertainty, which is discussed into more detail in Chapter 4.

The importance of keeping the location measurement uncertainty to a minimum regarding the coverage of a search area is illustrated in Figure 3-6 [17], which shows the results of a real-world experiment.

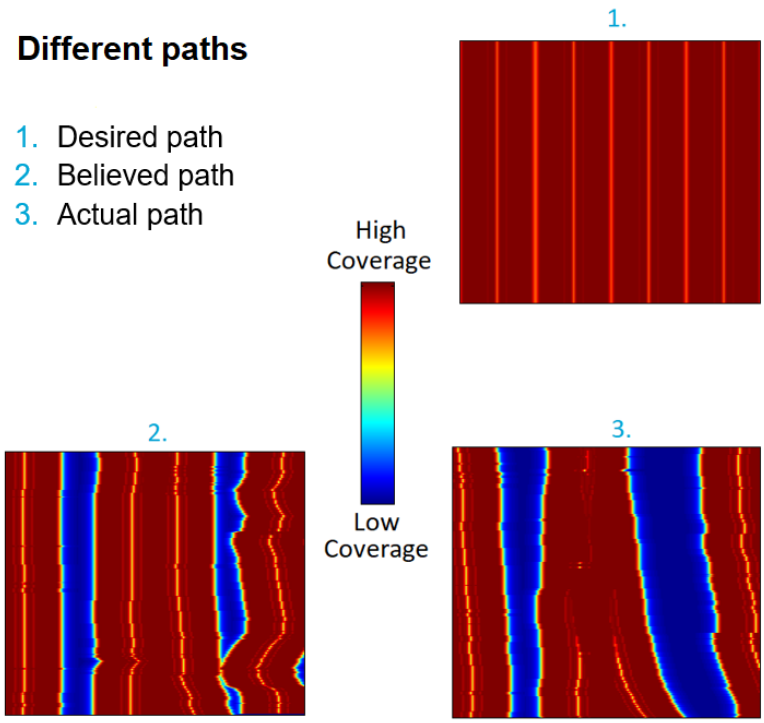

Figure 3-6: The coverage corresponding to three different types of paths.

In Figure 3-6 it is shown that there are three different types of paths. The first type of path is the `desired` path. If the AUV would be able to perfectly follow this path, the desired coverage of the area would be achieved. This desired coverage can be seen in the top right of Figure 3-6. The second type of path is the `believed` path. This is the path that is believed to have been followed by the AUV, based on the INS location measurements. The corresponding believed coverage, shown in the bottom left of Figure 3-6, can be seen to be worse than the desired coverage. The difference between the desired path and the believed path is caused

by external disturbances such as currents, and by constraints on the AUV's dynamics and controllers. The third type of path is the `actual` path, which is the path that the AUV has been following in reality. The coverage corresponding to this real path is path is shown in the bottom right of Figure 3-6. This real coverage can be seen to be worse than the believed coverage. The difference between the believed coverage and the actual coverage is caused by the error on the location measurements, obtained from the AUV's INS.

Figure 3-6 shows that when the uncertainty on the AUV's location is not included in the measure of coverage, the believed coverage can be too optimistic. This means that parts of the search area that are believed to have been covered by the AUV might not have been covered in reality. For this reason, the uncertainty on the AUV's location measurements should be included in the measure of coverage.

In order to include the uncertainty on the AUV's location measurements in the measure of coverage, an estimate of this location uncertainty has to be obtained. In the simulation framework, the INS provides an estimate of the AUV's location. In addition to the location measurement the INS provides an estimate of the uncertainty on this location measurement, by returning a vector named `set` and a vector named `cov` with every location measurement. The `set` vector provides the parameters that describe the distributions of three uniformly distributed random variables, one for each dimension, while the `cov` vector provides the parameters that describe the distributions of three normally distributed random variables, also one for each dimension. When the uniformly distributed and the normally distributed random variables are added in each dimension, a random variable that describes the location of the AUV is obtained for every dimension. Since the probability density functions of these random variables are known, the estimated uncertainty on the AUV's location measurements is also known.

Equation (3-6) [23] describes the probability density function of a uniform distribution. The a and b parameters are given for every dimension by respectively subtracting the `set` vector from the estimated location measurement vector and by adding the `set` vector to the estimated location measurement vector.

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{if } a \leq x \leq b \\ 0 & \text{if } x < a \text{ or } x > b \end{cases} \tag{3-6}$$

Equation (3-7) [23] describes the probability density function of a normal distribution. Here the μ parameter is equal to the location measurement and the σ parameter equals the corresponding `cov` value, for all three dimensions.

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \mathrm{e}^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \tag{3-7}$$

In the autonomy framework, the random variable describing the AUV's location can be obtained by adding the uniformly distributed and the normally distributed random variables. The probability density function of this resulting random variable is then given by the convolution of the probability density functions of the uniformly and normally distributed random variables. The convolution of two continuous functions in one dimension is shown in Equation (3-8) [24].

$$f(x) * g(x) = \int_{-\infty}^{\infty} f(\tau) g(x - \tau)\, d\tau \tag{3-8}$$

However, the autonomy framework works with discrete representations of the probability density functions of the uniformly and normally distributed random variables. This means that instead of the continuous convolution shown in Equation (3-8), the discrete convolution shown in Equation (3-9) [24] is required.

$$f[n] * g[n] = \sum_{i=-\infty}^{\infty} f[i]g[n-i] \tag{3-9}$$

The discrete convolution shown in Equation (3-9) is one-dimensional, since the convolutions are computed in the direction of each principal axis separately. The resulting three probability density functions can later be combined again to obtain a joint probability density function describing the AUV's location distribution in three dimensions, using Equation (3-10) [23]. Equation (3-10) holds because the uncertain locations in the direction of the principal axes are independent, and it holds for both the continuous and the discrete case.

$$\mathrm{P}(\boldsymbol{X} \leq a, \boldsymbol{Y} \leq b, \boldsymbol{Z} \leq c) = \mathrm{P}(\boldsymbol{X} \leq a)\mathrm{P}(\boldsymbol{Y} \leq b)\mathrm{P}(\boldsymbol{Z} \leq c) \tag{3-10}$$

An example of what the convolution of the probability density functions of a uniform distribution and a normal distribution can look like in one dimension is shown in Figure 3-7. This convolution corresponds to the INS parameters `set` $= 5$ and `cov` $= 1$, with the location measurement in the origin.

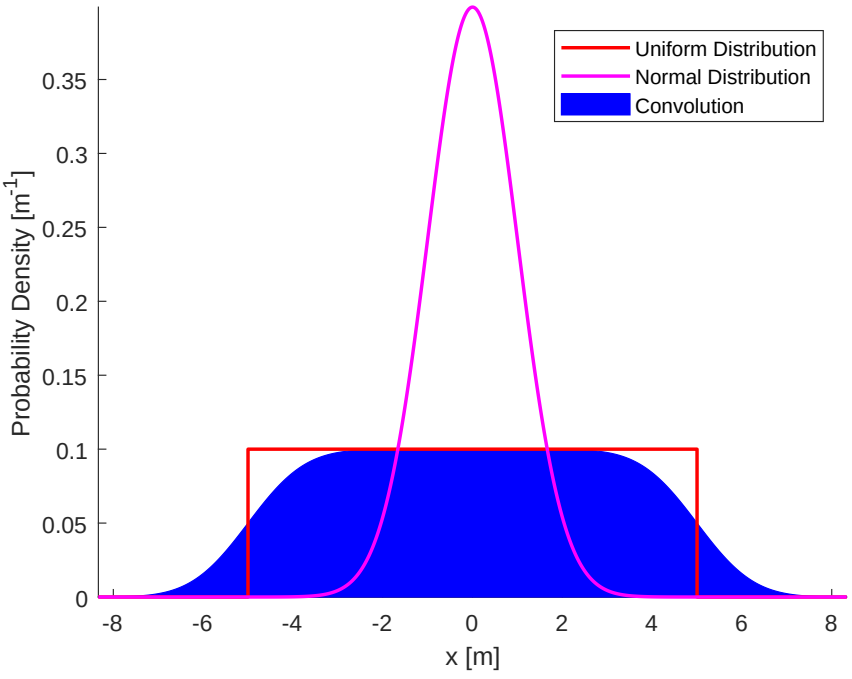

Figure 3-7: Convolution of a uniform distribution and a normal distribution.

In order to compute the probability that the AUV is located above a certain grid cell, the joint probability density function describing the AUV's location has to be integrated over this grid cell. Figure 3-7 shows in one dimension that the probability density of the convolution only

equals 0 when $x \to \pm\infty$, which means that the AUV can be located over the entire domain of the convolution. This implies that there is a non-zero probability for every grid cell that the AUV is located above that grid cell, although the majority of these probabilities will be very small. For this reason, the tails of the convolution are cut off in the autonomy framework. This is done symmetrically and such that the area under the truncated convolution equals 0.997, which is the same as the area under a normal distribution in the interval $(\mu-3\sigma, \mu+3\sigma)$.

In a real INS, the estimation of the uncertainty on the location measurement is provided by a single value, the horizontal plane error. This horizontal plane error gives the radius of a circle in the xy-plane centred at the location measurement, within which the real horizontal location of the AUV is estimated to be. The INS in the simulation framework however provides more information about the uncertainty on the location measurement, in the form of three location probability density functions. Since this amount of information is not available in reality, the truncated convolutions are in each dimension approximated by a uniform distribution. The approximated uniform distribution is defined such that the non-zero domain of this approximated uniform distribution equals the non-zero domain of the truncated convolution. An example of this is shown in Figure 3-8.

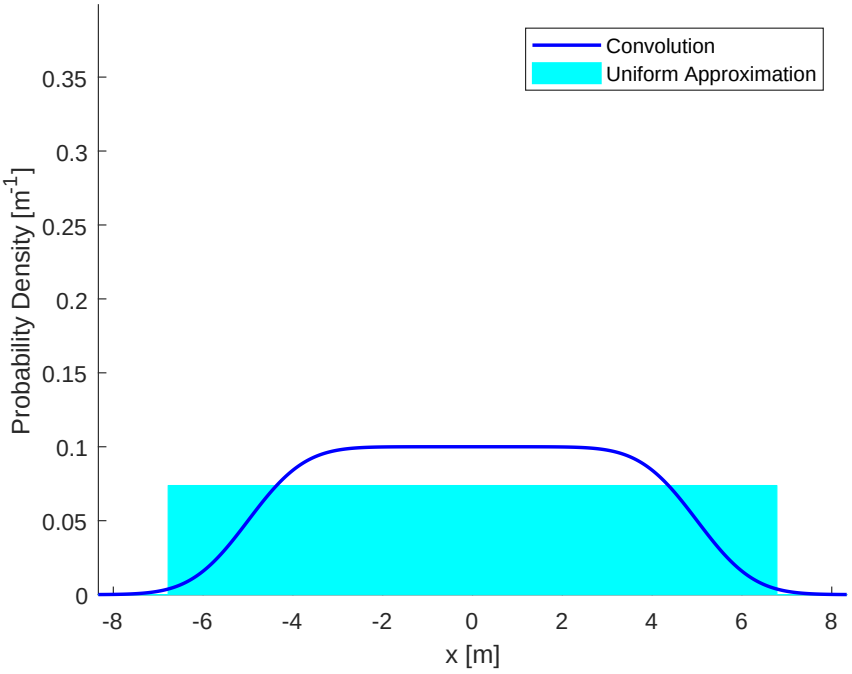

Figure 3-8: Uniform approximation of the convolution of a uniform and normal distribution.

The convolution shown in Figure 3-8 is equal to the convolution shown in Figure 3-7. Figure 3-8 shows that the uniform approximation underestimates the probability of the AUV being near the measured location, while it overestimates the probability of the AUV being located further away from the measured location. This approximated uniform distribution is however more realistic than the probability density function given by the convolution, and also computationally easier when the AUV's location has to be discretised.

In order to discretise the AUV's location to the centre of the grid cells in the search area, the approximated uniform distributions in the x-dimension and y-dimension need to be combined into a joint location probability density function first. Since both distributions are discrete in the autonomy framework, as well as independent, Equation (3-11) [23] can be used to compute the corresponding joint probability mass function. Equation (3-11) follows from Equation (3-10).

$$\mathrm{P}(\boldsymbol{X} = a, \boldsymbol{Y} = b) = \mathrm{P}(\boldsymbol{X} = a)\mathrm{P}(\boldsymbol{Y} = b) \tag{3-11}$$

Once the 2-D joint probability density function has been obtained, it can be integrated to obtain the probability that the AUV is located above a certain area. This can be done for every grid cell in the discretised search area. In order to reduce the computational cost, these integrations are only performed over the grid cells located within the non-zero domain of the joint probability density function. The boundaries of an approximated uniform joint probability density function during a simulation are shown in Figure 3-9.

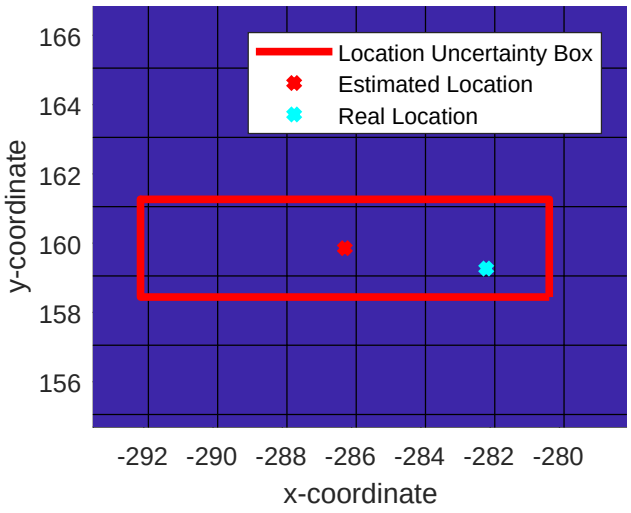

Figure 3-9: Estimation of the uncertainty on the INS location measurements.

In Figure 3-9 the boundaries of the 2-D approximated uniform probability density function are shown in red. The probability of the AUV being in a grid cell can be computed by multiplying the area of this grid cell that lies below the approximated uniform probability density function by the height of this function. This is the final step in using the INS measurements in the simulation framework to compute for every grid cell a probability that the AUV is located inside that grid cell. The location uncertainty box shown in Figure 3-9 has been computed at the end of a 300 metre long straight track, and it can be seen that the AUV can be located in 21 grid cells with non-zero probability at this moment. The real location of the AUV can be seen to lie within the computed location uncertainty box.

It is possible to use the truncated convolution probability density functions instead of the approximated uniform probability density functions in order to compute the probabilities of the AUV being in every grid cell. Computing an area under the joint distribution of two convolution probability density functions is however computationally more expensive, and corresponds less to reality than using the uniform approximations of the convolutions. That

is why in this final thesis project, the AUV's location is discretised by integrating the joint distribution of the uniform approximations of the convolutions of the uniform and normal location uncertainty probability density functions.

Since the AUV is not always located inside the search area during an operation, a grid larger than the search area needs to be considered in order to be able to discretise the AUV's location at any moment during the operation. Note that due to the way that the AUV's location is discretised, increasing the grid size does not affect the computational time of the AUV's location discretisation.

3-3 Coordinate Frames

The INS measurements of the AUV's location need to be described in a certain coordinate frame. In the scenario of an AUV on a search operation, there is more than one coordinate frame that is relevant to the operation. In fact, there are three coordinate frames that are important in this scenario. These three coordinate frames are listed below.

- Global coordinate frame
- Robot coordinate frame
- Sensor coordinate frame

The three coordinate frames are shown in Figure 3-10 [17].

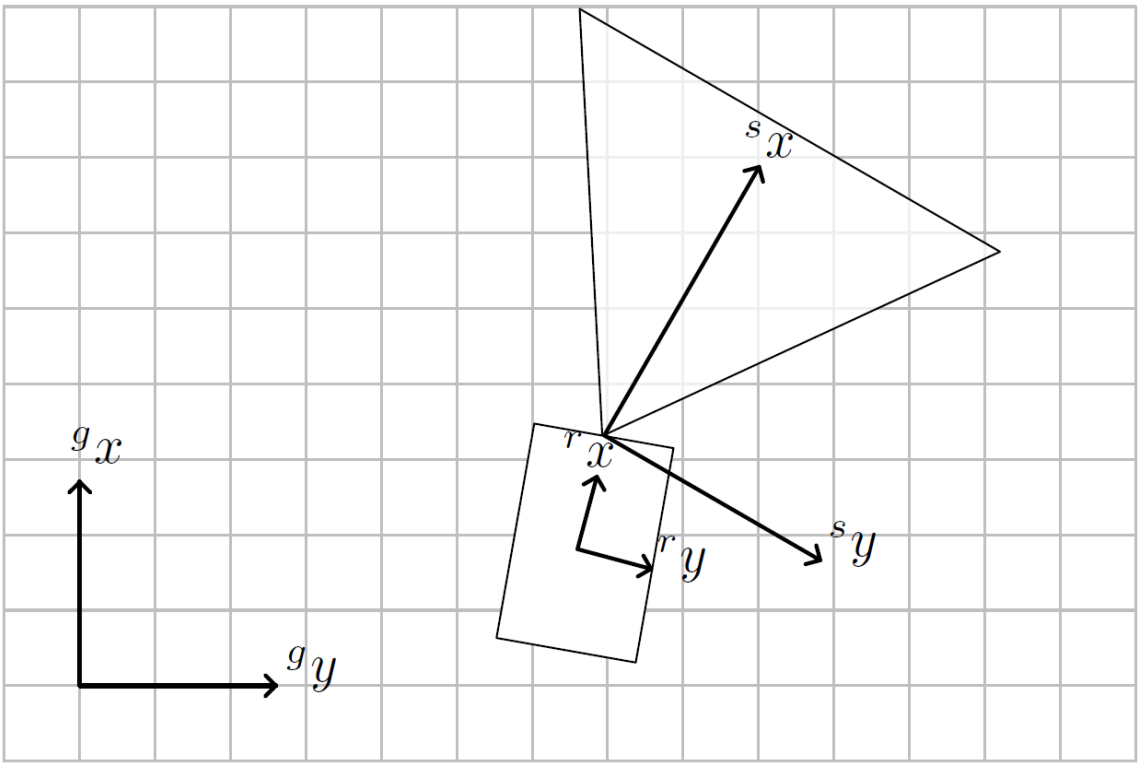

Figure 3-10: Three different coordinate frames.

In Figure 3-10, the axes corresponding to the global coordinate frame are labelled g_x and g_y, the axes corresponding to the robot coordinate frame are labelled r_x and r_y and the axes corresponding to the sensor coordinate frame are labelled s_x and s_y.

The global coordinate frame is a coordinate frame that has its axes fixed on the earth's surface. The global coordinate frame used by GPS has its origin located somewhere off the coast of Africa, but in the simulation framework the origin is chosen somewhat more convenient inside the search area. The location and orientation measurements made by the AUV's INS are described in this coordinate frame, and this is also the coordinate frame that the locations of the found objects should be described in. Contrary to the global coordinate frame, the robot coordinate frame has its axes fixed on the AUV. This means that when the AUV is moving, a location that is stationary in the robot coordinate frame will be moving in the global coordinate frame. Likewise, the axes of the sensor coordinate frame are fixed on a sonar system. When the sonar system moves relatively to the robot, locations that are stationary in the sensor coordinate frame will be moving in the global and robot coordinate frames. Locations in the sonar images are described in the sensor coordinate frame.

In this final thesis project, the sonar systems are not moving with respect to the AUV. For this reason, the robot coordinate frame can be discarded when the axes of the sensor coordinate frame are fixed on the AUV, since locations in sonar images obtained from both the sonar systems can be described in this coordinate frame now. This means that from the three coordinate frames shown in Figure 3-10, only the global and sensor coordinate frames are used. Another difference with Figure 3-10 is the way that the sensor coordinate frame is fixed on the AUV. In this final thesis project, the sensor coordinate frame is fixed on the AUV such that the AUV is always facing the positive y-axis of this coordinate frame.

Coordinates described in the sensor coordinate frame can be transformed into coordinates described in the global coordinate frame, and vice versa, provided that the current location and orientation of the sensor coordinate frame axes inside the global coordinate frame are known. Since the INS provides a measurement regarding the location and orientation of the AUV in the global coordinate frame, this measurement can be used to relate the two coordinate frames. These two coordinate frames are related through homogeneous transformation matrices, such that ${}^{b}_{a}T$ is the homogeneous transformation matrix that maps homogeneous coordinates from coordinate frame a to coordinate frame b.

The reason that homogeneous transformation matrices are used in this final thesis project, is because these transformation matrices can apply the translation of a vector through a matrix multiplication with this vector. In order to do so, they need to be multiplied with a vector describing homogeneous coordinates instead of Cartesian coordinates. Homogeneous coordinates add one dimension to the Cartesian coordinates. This extra dimension is a non-zero real number that scales the Cartesian coordinates. In this final thesis project this extra dimension is set to one, such that the Cartesian coordinates (x, y) in $\mathbb{R}^2$ are described by the homogeneous coordinates $(x, y, 1)$ in $\mathbb{R}^3$.

The sensor coordinate frame can be transformed into the global coordinate frame by rotating and translating the sensor coordinate frame such that it becomes equal to the global coordinate frame. The same can be done to transform the global coordinate frame into the

sensor coordinate frame. The required rotation is determined by the difference in orientation of the two coordinate frames. An estimated value for this difference in orientation can be obtained by the INS orientation measurement θ_m. This θ_m is the measured rotation [rad] of the AUV around the z-axis with respect to the positive x-axis, in the global coordinate frame. The required translation is determined by the difference in location of the origins of the two coordinate frames. An estimated value for this difference can be obtained by the INS location measurements x_m and y_m. These x_m and y_m are the measured x-coordinate and y-coordinate of the AUV's location, in the global coordinate frame.

Using the θ_m, x_m and y_m INS measurements, the homogeneous rotation and translation transformation matrices to transform coordinates from the sensor coordinate frame to the global coordinate frame are shown in Equation (3-12) [25] and Equation (3-13) [25] respectively.

$$ {}^{g}_{s}T_{rotation} = \begin{bmatrix} \cos\left(\theta_m - \frac{\pi}{2}\right) & -\sin\left(\theta_m - \frac{\pi}{2}\right) & 0 \\ \sin\left(\theta_m - \frac{\pi}{2}\right) & \cos\left(\theta_m - \frac{\pi}{2}\right) & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \sin(\theta_m) & \cos(\theta_m) & 0 \\ -\cos(\theta_m) & \sin(\theta_m) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3-12} $$

$$ {}^{g}_{s}T_{translation} = \begin{bmatrix} 1 & 0 & x_m \\ 0 & 1 & y_m \\ 0 & 0 & 1 \end{bmatrix} \tag{3-13} $$

The homogeneous transformation matrix that transforms coordinates from the sensor coordinate frame to the global coordinate frame can be computed by a matrix multiplication of ${}^{g}_{s}T_{translation}$ and ${}^{g}_{s}T_{rotation}$, as shown in Equation (3-14).

$$ {}^{g}_{s}T = {}^{g}_{s}T_{translation}\ {}^{g}_{s}T_{rotation} = \begin{bmatrix} \sin(\theta_m) & \cos(\theta_m) & x_m \\ -\cos(\theta_m) & \sin(\theta_m) & y_m \\ 0 & 0 & 1 \end{bmatrix} \tag{3-14} $$

The homogeneous transformation matrix that transforms coordinates from the global coordinate frame to the sensor coordinate frame can be constructed in a similar way. When transforming from the global coordinate frame to the sensor coordinate frame, the homogeneous rotation and translation transformation matrices are given by Equation (3-15) [25] and Equation (3-16) [25] respectively.

$$ {}^{s}_{g}T_{rotation} = \begin{bmatrix} \cos\left(\theta_m - \frac{\pi}{2}\right) & \sin\left(\theta_m - \frac{\pi}{2}\right) & 0 \\ -\sin\left(\theta_m - \frac{\pi}{2}\right) & \cos\left(\theta_m - \frac{\pi}{2}\right) & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \sin(\theta_m) & -\cos(\theta_m) & 0 \\ \cos(\theta_m) & \sin(\theta_m) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3-15} $$

$$ {}^{s}_{g}T_{translation} = \begin{bmatrix} 1 & 0 & -x_m \\ 0 & 1 & -y_m \\ 0 & 0 & 1 \end{bmatrix} \tag{3-16} $$

The homogeneous transformation matrix that transforms coordinates from the global coordinate frame to the sensor coordinate frame can again be computed by a matrix multiplication of the homogeneous translation and rotation transformation matrices, but this time in the reverse order. This is shown in Equation (3-17).

$$ {}^{s}_{g}T = {}^{s}_{g}T_{rotation}\ {}^{s}_{g}T_{translation} = \begin{bmatrix} \sin(\theta_m) & -\cos(\theta_m) & -x_m\sin(\theta_m) + y_m\cos(\theta_m) \\ \cos(\theta_m) & \sin(\theta_m) & -x_m\cos(\theta_m) - y_m\sin(\theta_m) \\ 0 & 0 & 1 \end{bmatrix} \tag{3-17} $$

It can be checked that the homogeneous transformation matrix ${}^{g}_{s}T$ shown in Equation (3-14) and the homogeneous transformation matrix ${}^{s}_{g}T$ shown in Equation (3-17) are each other's inverse, as is shown in Equation (3-18).

$$({}^{g}_{s}T)^{-1} = {}^{s}_{g}T \tag{3-18}$$

When the homogeneous transformation matrices ${}^{g}_{s}T$ and ${}^{s}_{g}T$ are matrix-multiplied with a vector containing a homogeneous coordinate, the resulting vector contains exactly one homogeneous coordinate as well. However, it has been discussed earlier that the INS measurements and estimated uncertainties are used to discretise the location of the AUV to the centre of a set of grid cells, with for every grid cell a probability that the AUV is located inside that grid cell. This means that instead of a deterministic location, the AUV's location is now a discrete random variable, and every possible outcome of this discrete random variable corresponds to a different set of coordinates describing the AUV's location. The AUV's orientation is not considered a random variable in this final thesis project, since the uncertainty on the orientation measurements is negligible compared to the uncertainty on the location measurements. This means that θ_m is equal to the measured orientation of the AUV, while there are as many sets of x_m and y_m as there are possible outcomes of the discrete random variable describing the AUV's location. For every possible outcome, the resulting homogeneous transformation matrices ${}^{g}_{s}T$ and ${}^{s}_{g}T$ can be computed, and they can all be matrix-multiplied with a vector containing homogeneous coordinates. The result is that a single location in one coordinate frame is transformed to a set of locations in the other frame, each with a probability equal to the probability that the AUV is located in the x_m and y_m locations that were used to compute the corresponding homogeneous transformation matrix.

Thus, a deterministic location in one coordinate frame is transformed into a discrete random variable describing this location in the other coordinate frame, due to the fact that the AUV's location is described by a discrete random variable.

3-4 Coverage Map

During a search operation, an AUV needs to keep track of the progress of this operation. It has been discussed that the coverage of an area can be quantified as a POD in this area, shown in Equation (3-3). Since the search area is divided into grid cells, the PODs can be maintained for every cell in this grid, creating a coverage map. This coverage map has to be maintained in the global coordinate frame, while the POD measurements are obtained in the sensor coordinate frame. It has also been discussed that since the exact location of the AUV is unknown, there is no deterministic relationship between these two coordinate frames. Therefore, it is unknown to what area in the global coordinate frame a measured POD exactly corresponds. The coverage map should include this location uncertainty when keeping track of the coverage in the search area.

In the simulation framework, the sonar systems provide sonar image resolutions in the area that has been inside the sonar range. When the ocean floor type is known to be flat and sandy, the PODs corresponding to these sonar image resolutions can be directly determined through Equation (3-5). How the sonar image resolutions are provided by the simulation framework is shown in Figure 3-11.

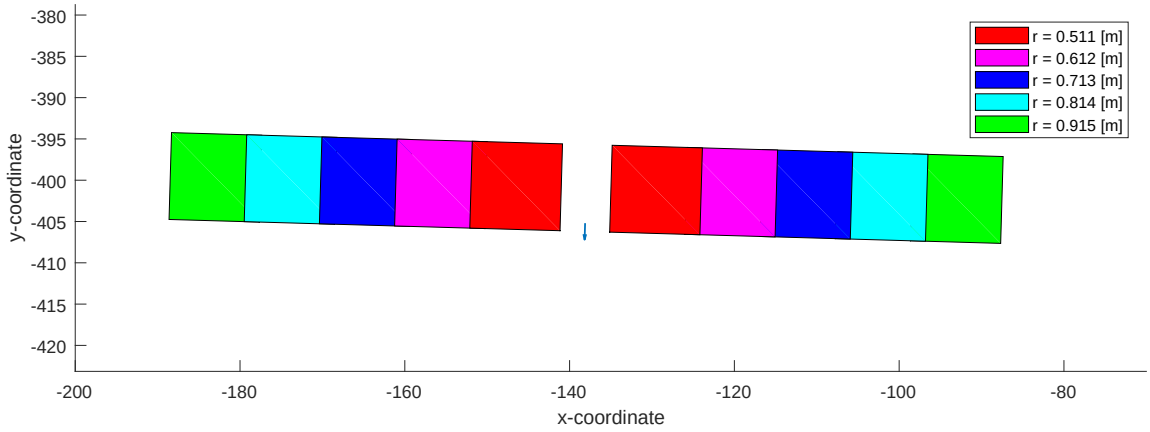

Figure 3-11: Sonar image resolutions in the global coordinate frame.

Figure 3-11 shows that the simulation framework provides the sonar image resolutions in the global coordinate frame. In Figure 3-11 the AUV is depicted by an arrow, which faces the same direction as the AUV. At this moment in time, the AUV is travelling in the negative y-direction of a straight track of a lawnmower-shaped coverage path. The AUV's two sonar systems each provide a number of polygons, with for each polygon a corresponding sonar image resolution. The shape and corresponding sonar image resolution of each polygon largely depends on the trajectory that has been followed by the AUV. The smoother this trajectory has been, the higher the sonar image quality will be.

The simulation framework constructs the sonar image resolution polygons in the sensor coordinate frame, and then transforms these polygons to the global coordinate frame using the current θ_m, x_m and y_m INS measurements. Since the location of the sonar image resolution polygons is deterministic inside the sensor coordinate frame while it is probabilistic inside the global coordinate frame, the autonomy framework transforms the sonar image resolution polygons shown in Figure 3-11 back to the sensor coordinate frame. For this transformation, the current θ_m, x_m and y_m INS measurements are used in the homogeneous transformation matrix shown in Equation (3-17). The resulting sonar image resolution polygons in the sensor coordinate frame are shown in Figure 3-12.

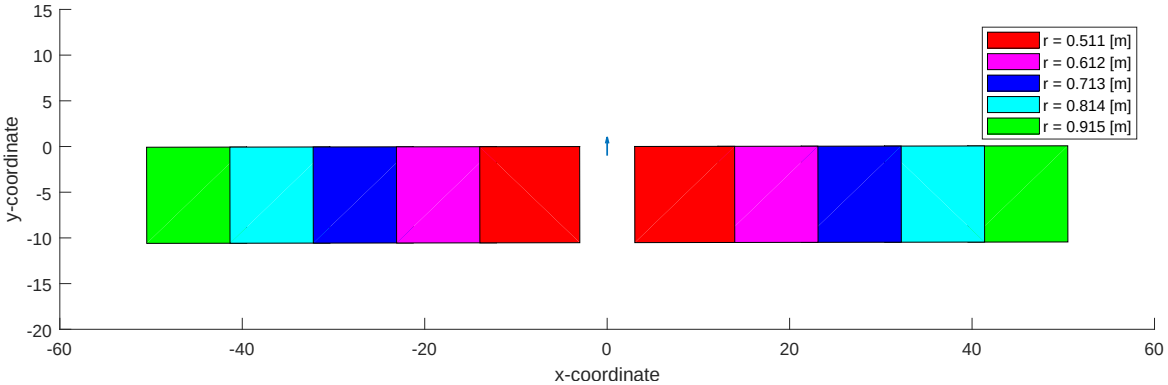

Figure 3-12: Sonar image resolutions in the sensor coordinate frame.

It can be seen that Figure 3-12 describes the sonar image resolution polygons in the sensor coordinate frame, since the AUV is at the origin of the coordinate frame facing the positive y-axis. Describing the sonar image resolution polygons in the sensor coordinate frame is more realistic than describing the sonar image resolution polygons in the global coordinate frame. The reason for this is that the locations of the sonar image resolution polygons are known exactly relatively to the AUV, so these locations can be described deterministically in the sensor coordinate frame. By transforming the locations of the sonar image resolution polygons to the global coordinate frame, which happens inside the simulation framework, the assumption is made that the AUV is located in the exact location that is measured by the INS, which is an invalid assumption.

In order to maintain a coverage map, the sonar image resolutions and the resulting PODs need to be known in the global coordinate frame. This means that the sonar image resolutions that are now available in the sensor coordinate frame need to be transformed to the global coordinate frame. This can be done by constructing a homogeneous transformation matrix for every possible outcome of the discrete random variable describing the AUV's location.

In [17], every grid cell in the search area is transformed from the global coordinate frame to the sensor coordinate frame using these homogeneous transformation matrices, resulting in a discrete random variable describing the location of this grid cell in the sensor coordinate frame. For every possible outcome of this discrete random variable, it is checked what the sonar resolution at this location in the sensor coordinate frame is. Since the probability of every possible outcome is known, the probabilities of the corresponding sonar image resolutions are also known. These sonar image resolutions can be used to determine the PODs, resulting in a discrete random variable describing the POD in a grid cell in the global coordinate frame. By doing this for every grid cell in the search area, a coverage map is obtained that contains a discrete random variable describing the POD for every grid cell in the search area.

In this final thesis project, the approach taken in [17] is reversed. Instead of transforming every grid cell in the search area from the global coordinate frame to a discrete random variable describing the location of this cell in the sensor coordinate frame, the sonar image resolution polygons are transformed from the sensor coordinate frame to the global coordinate frame. Since the AUV's location is uncertain, the vertices of every polygon in the sensor coordinate frame transform into a discrete random variable describing the location of this vertex in the global coordinate frame. This way, a single sonar image resolution polygon in the sensor coordinate frame transforms into as many sonar image resolution polygons in the global coordinate frame as there are possible outcomes of the discrete random variable describing the AUV's location. The probability of every polygon being at the correct location in the global coordinate frame is equal to the probability that the AUV is in the location that has been used to compute the corresponding homogeneous transformation matrix. Since the probabilities of the polygon locations and the corresponding sonar image resolution are known, it can be checked for every grid cell of the search area in the global coordinate frame what the probability of every sonar image resolution is. These sonar image resolutions can be used to determine the PODs, resulting in a discrete random variable describing the POD in the global coordinate frame for every grid cell in the search area. The coverage map is the collection of these discrete random variables describing the POD for every grid cell in the search area in the global coordinate frame.

The reason that in this final thesis project the approach taken in [17] is reversed, is because it reduces the computational cost significantly. If the total number of possible AUV locations is denoted by N_l, the total number of sonar image resolution polygons in the sensor coordinate frame by N_p and the total number of grid cells in the search area by N_c, then the autonomy framework in this final thesis project has to check what points are located inside a polygon $N_l \cdot N_p$ times. Using the approach in [17], it has to be checked what points are located inside a polygon $N_c \cdot N_p$ times. Since $N_l \ll N_c$, the approach taken in this final thesis project has to check what points are located inside a polygon significantly fewer times. Furthermore, the approach taken in this final thesis project requires $4 \cdot N_p \cdot N_l$ transformation matrix multiplications, while the approach taken in [17] requires $N_c \cdot N_l$ transformation matrix multiplications. Since $4 \cdot N_p \ll N_c$, the approach taken in this final thesis project also requires significantly fewer transformation matrix multiplications. When in the approach taken in [17] not every grid cell in the search area is considered, but only the grid cells within the sonar ranges of the AUV, the approach taken in this final thesis project still has proven to be significantly faster. The comparison between the approach taken in [17] and the approach taken in this final thesis project is also shown in Table 3-2, where the number of times that it has to be checked what points are inside a polygon and the number of transformation matrix multiplications that are required are compared.

Table 3-2: A comparison of the required computations for two different approaches.

	Inside polygon	Transformations
Approach taken in [17]	$N_c \cdot N_p$	$N_c \cdot N_l$
Final thesis project	$N_l \cdot N_p$	$4 \cdot N_p \cdot N_l$

Another difference between the approach taken in this final thesis project and the approach taken in [17], and also with the approach taken in [19], is that in the last two approaches only the AUV's location uncertainty in lateral direction is taken into account, while in this final thesis project the AUV's location uncertainty in both longitudinal and lateral direction is taken into account.

When a coverage map is constructed using the approach taken in this final thesis project, the computational cost is affected by the way that this approach is implemented in the autonomy framework. One way of implementation is to transform all the N_p sonar image resolution polygons from the sensor coordinate frame to the global coordinate frame, for every possible outcome of the discrete random variable describing the AUV's location separately. So for any possible AUV location, the locations of the N_p sonar image resolution polygons in the global coordinate frame can be obtained. These sonar image resolution polygons in the global coordinate frame can be used to determine the sonar image resolution in every grid cell in the search area, with the probability of these resolutions equal to the probability that the AUV is in the location that has been used to construct the transformation matrix. This has to be done for every possible location of the AUV, resulting in N_l maps of size N_c describing the sonar image resolution for every grid cell in the search area. Each of these resolution maps then corresponds to a single probability. The N_l maps and their corresponding probabilities can then be used to determine the discrete random variable describing the POD in every grid

cell in the search area. Note that the number of required resolution maps grows when the uncertainty on the AUV's location grows.

Another way that the approach taken in this final thesis project can be implemented, is by transforming a single sonar image resolution polygon from the sensor coordinate frame to the global coordinate frame, for all the N_l possible outcomes of the discrete random variable describing the AUV's location at once. This way, N_l different sonar image resolution polygons in the sensor coordinate frame are obtained, all corresponding to the same sonar image resolution, but all with a different probability. For every grid cell in the search area, the probability for this resolution can then be determined. This has to be done for every sonar image resolution polygon in the sensor coordinate frame, resulting in N_p maps of size N_c describing a probability for every grid cell in the search area. Each of these probability maps then corresponds to a single sonar image resolution. The N_p maps and their corresponding sonar image resolutions can then be used to determine the discrete random variable describing the POD in every grid cell in the search area. Note that the number of required probability maps is equal to the constant number of sonar image resolution polygons. Table 3-3 summarises the comparison of the two implementations.

Table 3-3: A comparison between resolution maps and probability maps.

	Number of maps	Size of maps
Resolution maps	N_l	N_c
Probability maps	N_p	N_c

Table 3-3 shows that since a resolution map and a probability map are of the same size N_c, maintaining resolution maps is computationally more favourable when $N_l < N_p$. When $N_l > N_p$, it becomes computationally more favourable to maintain probability maps instead.

For both implementations, the resulting coverage map contains a discrete random variable describing the POD in every grid cell in the search area. Figure 3-13 shows what the probability mass function of such a discrete random variable describing the POD in a grid cell can look like.

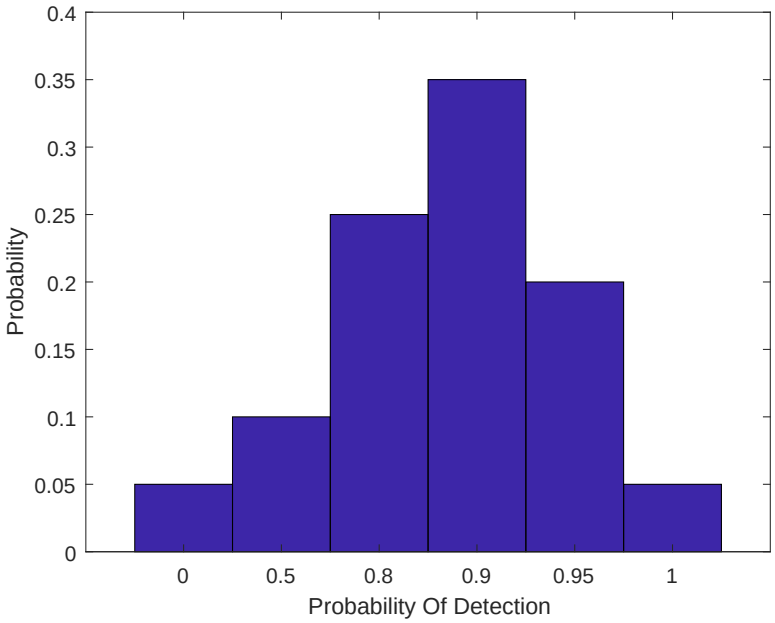

Figure 3-13: The POD as a discrete random variable.

If the location uncertainty had not been taken into account while creating the coverage map, the POD in the grid cell shown in Figure 3-13 would be a single deterministic value. When the location uncertainty is taken into account, the POD is described by a discrete random variable instead. This means that the POD in a grid cell can have different possible values, where the probabilities of each of these different values are known. These probabilities are obtained from the probability mass function of the discrete random variable describing the POD in the grid cell. The probability mass function of the POD of Figure 3-13 is shown again in Table 3-4.

Table 3-4: Probability mass function of the POD in a grid cell.

a	0	0.5	0.8	0.9	0.95	1
$\mathrm{P}(\mathrm{POD} = a)$	0.05	0.1	0.25	0.35	0.2	0.05

When the probability mass function of a discrete random variable is known, the probabilities of its n possible outcomes $a_1, a_2, \ldots, a_n$ are known. The expected value of this discrete random variable can then be computed using Equation (3-19) [23].

$$\mathrm{E}[\boldsymbol{X}] = \sum_{i=1}^{n} a_i \, \mathrm{P}(\boldsymbol{X} = a_i) \tag{3-19}$$

Using Equation (3-19), the expected value of the POD shown in Figure 3-13 can be calculated to equal 0.805. Note that the expected value depends on the possible outcomes of the discrete random variable. The higher the number of possible outcomes of the discrete random variable, the closer the probability mass function will start approaching the probability density

function of a continuous random variable. This will allow the calculation of a more accurate expected value, but it will also increase the computational cost of maintaining the coverage map. When the total number of outcomes of the discrete random variable describing the POD in a grid cell is denoted by N_o, then the coverage map contains $N_o \cdot N_c$ elements. Increasing N_o not only increases the size of the coverage map in (large) steps of N_c, it also increases the computational cost of combining two discrete random variables, discussed later in this section. In this final thesis project, the possible outcomes of the discrete random variables describing the PODs in the grid cells are given in Equation (3-5). These possible outcomes have been found to provide realistic POD values in the simulation framework, as shown in Figure 3-3, while having a reasonable computational time.

The coverage maps that have been discussed so far, have been constructed from a single sonar image. When the simulation framework provides a new sonar image, the autonomy framework can use the exact same procedure to obtain a coverage map based on this new sonar image. These two coverage maps now have to be combined into a single map. In the remainder of this thesis, a coverage map corresponding to a single sonar image is called a `coverage measurement`, and the map combining all these coverage measurements is referred to as the `coverage map`.

In order to combine the coverage measurements into a coverage map, a conservative method is preferred. The reason for this is that an overestimation of the coverage can have significantly more serious consequences than an underestimation of the coverage. Underestimating the coverage can increase the duration of an operation, since the underestimation could require certain parts of the area to be scanned again. Overestimating the coverage, on the other hand, can lead to the AUV not finding the object that it is searching for. A reasonably conservative way to combine two discrete random variables describing the PODs in a grid cell can be obtained by taking the maximum of these two discrete random variables. The reasoning behind this is that by taking the maximum of two coverage measurements, the resulting coverage can not be larger than any of the coverage measurements already obtained. Since taking the maximum of two discrete random variables is different from taking the maximum of two deterministic variables, it is now first explained how the maximum of two discrete random variables can be obtained.

Since a discrete random variable is described by its probability mass function, the goal is to obtain the probability mass function of the maximum of the two discrete random variables, which are each also described by their probability mass function. These two discrete random variables describe the POD in the same grid cell, but have been obtained from two different sonar images. In order to be able to compute the maximum of these two discrete random variables, they have to be independent. The reason for this is that the probability mass functions of both the discrete random variables are known, but their joint probability mass function is not. When the two random variables are independent however, this joint probability mass function can be computed using Equation (3-11). Two random variables are independent when a realisation of the first random variable does not affect the probability distribution of the second random variable. In an AUV search operation, this seems to be true for two random variables resulting from two different sonar images, describing the POD in the same cell. Since the PODs depend mostly on the sonar image resolution, which in turn depends mostly on the distance to the AUV and on the trajectory that the AUV has been

following, a POD distribution resulting from one sonar image does not seem to be related to a POD distribution resulting from another sonar image obtained later during the operation. Although it seems likely, independence of two random variables resulting from two different sonar images describing the POD in the same grid cell can not not be guaranteed. However, since independence is both likely and required to combine two coverage measurements by taking their maximum, random variables describing the POD in the same grid cell resulting from different sonar images are assumed to be independent in this final thesis project. Note that this independence is only assumed for random variables describing the POD in a grid cell that have been obtained from two different sonar images. Due to the way that the sonar model is implemented inside the simulation framework, multiple random variables describing the POD in the same grid cell are obtained during a single straight track in the coverage path of the AUV. Since the random variable describing the POD in this grid cell is in reality obtained from a single sonar image, the multiple random variables describing the POD in the same grid cell obtained from a single straight track can not be assumed to be independent. This means that these random variables can not be combined using the maximum function, so instead the random variable with the lowest expected value is discarded. This means that random variables describing the POD in a grid cell that have been obtained from the same straight track are assumed to be dependent, while random variables describing the POD in a grid cell obtained from different straight tracks are assumed to be independent.

In [17], two different coverage measurements of the same grid cell are also combined by taking their maximum. However, the equation that is used to compute the distribution of the maximum of two random variables is incorrect. In this final thesis project, the equation to compute the probability mass function of the maximum of two independent random variables has been corrected in the following way. When $\boldsymbol{X}$ and $\boldsymbol{Y}$ are two independent discrete random variables with the same possible outcomes $a_1, a_2, \ldots, a_n$, the probability mass function of the maximum of these two independent discrete random variables can be computed through Equation (3-20).

$$\begin{aligned} \mathrm{P}(\max(\boldsymbol{X},\boldsymbol{Y}) = a_k) \;\; = \;\; & \mathrm{P}(\boldsymbol{X} = a_k) \sum_{i=1}^{k} \mathrm{P}(\boldsymbol{Y} = a_i) \;\; + \\ & \mathrm{P}(\boldsymbol{Y} = a_k) \sum_{i=1}^{k} \mathrm{P}(\boldsymbol{X} = a_i) \;\; - \;\; \mathrm{P}(\boldsymbol{X} = a_k)\mathrm{P}(\boldsymbol{Y} = a_k) \end{aligned} \tag{3-20}$$

When Equation (3-20) is used to compute $\mathrm{P}(\max(\boldsymbol{X},\boldsymbol{Y}) = a_k)$ for all possible a_k, the probability mass function of the maximum of the two independent discrete random variables is obtained. Equation (3-20) can however only be used to compute the probability mass function of the maximum of two independent discrete random variables. A more general version of Equation (3-20) is shown in Equation (3-21) [23], which can be used to compute the probability mass function of the maximum of n independent discrete random variables $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n$, as long as all these independent discrete random variables have the same possible outcomes.

$$\mathrm{F}_{\boldsymbol{Z}}(a) = \prod_{i=1}^{n} \mathrm{F}_{\boldsymbol{X}_i}(a) \tag{3-21}$$

Equation (3-21) makes use of the cumulative distribution functions of the independent discrete random variables, where the cumulative distribution function of a random variable $\boldsymbol{X}$ is given

by Equation (3-22) [23].

$$\mathrm{F}_{\boldsymbol{X}}(a) = \mathrm{P}(\boldsymbol{X} \leq a) \tag{3-22}$$

In Equation (3-21), $\mathrm{F}_{\boldsymbol{Z}}$ is the cumulative distribution function of the maximum of the independent discrete random variables $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n$. When in Equation (3-21) $\mathrm{F}_{\boldsymbol{Z}}(a)$ is computed for every possible outcome a, the cumulative distribution function of the maximum of $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n$ has been obtained.

Equation (3-23) [26] shows how the probability mass function of the maximum can be obtained from this cumulative distribution function of the maximum.

$$\mathrm{P}(\boldsymbol{X} = a) = \mathrm{F}_{\boldsymbol{X}}(a) - \lim_{0<\epsilon\to 0} \mathrm{F}_{\boldsymbol{X}}(a - \epsilon) \tag{3-23}$$

When in Equation (3-23) $\mathrm{P}(\boldsymbol{X} = a)$ is computed for every possible outcome a, the probability mass function is obtained.

Equation (3-21) can be used to prove that Equation (3-24) holds, when $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n$ are n independent discrete random variables with the same possible outcomes.

$$\max(\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n) = \max(\max(\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_{n-1}), \boldsymbol{X}_n) \tag{3-24}$$

The result shown in Equation (3-24) is very useful in this final thesis project. During an AUV search operation, multiple sonar images of a grid cell can be obtained. This results in multiple independent discrete random variables describing the POD in this grid cell. The discrete random variable describing the POD in this grid cell in the coverage map is obtained by taking the maximum of all those independent discrete random variables. Equation (3-24) states that it is not necessary to store all these discrete random variables until the end of the operation, in order to obtain the maximum of all these discrete random variables at once. Instead, the coverage map can be updated every time a new coverage measurement is obtained, by taking their maximum. After the coverage map has been updated, the coverage measurement can be discarded. When a new coverage measurement is obtained this process can be repeated, such that only one discrete random variable describing the POD is required in every grid cell in the coverage map, regardless of how many coverage measurements have been obtained. This means that the coverage map has the Markov property, since all future states of the coverage map can be computed with this coverage map and future coverage measurements only.

Note that it is required to maintain the probability mass function of a discrete random variable describing the POD in every grid cell in the coverage map in order to correctly update the coverage map, even when only the expected value of this discrete random variable is of interest. This is due to the fact that taking the maximum of the expected values of two discrete random variables is not the same as taking the expected value of the maximum of two discrete random variables. This is shown in Equation (3-25).

$$\mathrm{E}[\max(\boldsymbol{X}, \boldsymbol{Y})] \neq \max(\mathrm{E}[\boldsymbol{X}], \mathrm{E}[\boldsymbol{Y}]) \tag{3-25}$$

The inequality shown in Equation (3-25) also becomes clear when the coverage map that considers location uncertainty is compared with the coverage map that does not consider location uncertainty. When the mean of the expected values of the PODs in every grid cell in the coverage map that considers location uncertainty is computed, this mean expected value

is at least as high as the mean expected value of the PODs in every grid cell in the coverage map that does not consider location uncertainty. It might seem as if the mean expected value changes due to the location uncertainty dispersing the POD in every grid cell over the surrounding grid cells, but this is not the case. The dispersion does change the expected value in individual grid cells, but since the probabilities that the AUV is in a grid cell add up to one when they are summed over the search area, the mean expected value over the search area does not change through this dispersion. However, the mean expected value in the search area does change due to the way that different coverage measurements of a grid cell are combined. When the probability mass function of the maximum of two independent discrete random variables $\boldsymbol{X}$ and $\boldsymbol{Y}$ with the same possible outcomes is obtained, the expected value of this probability mass function is at least as high as the maximum of the expected values of the two discrete random variables that have been combined. Indeed it can be proven that Equation (3-26) holds, which is a stronger version of Equation (3-25).

$$\mathrm{E}[\max(\boldsymbol{X}, \boldsymbol{Y})] \geq \max(\mathrm{E}[\boldsymbol{X}], \mathrm{E}[\boldsymbol{Y}]) \tag{3-26}$$

Equation (3-26) shows that combining two coverage measurements by taking their maximum can increase the expected value of the POD in a grid cell, even when both coverage measurements are identical. When the uncertainty on the location measurements is not considered, the POD in every grid cell is deterministic, in which case taking the maximum of identical coverage measurements can never increase the POD in that grid cell.

Every time a new coverage measurement is obtained, N_c new discrete random variables are obtained, one for every grid cell in the search area. However, most of these discrete random variables in the coverage measurement describe the POD in grid cells that have not been inside the sonar ranges of the AUV. When the discrete random variables describing the POD in such an uncovered grid cell are called $\boldsymbol{C}_{\text{uncovered}}$, the probability mass function of $\boldsymbol{C}_{\text{uncovered}}$ is shown in Table 3-5.

Table 3-5: Probability mass function of the POD in an uncovered grid cell.

a	0	0.5	0.8	0.9	0.95	1
$\mathrm{P}(\boldsymbol{C}_{\text{uncovered}} = a)$	1	0	0	0	0	0

In order to update the coverage map, the maximum of the coverage map and the coverage measurement has to be computed for every grid cell in the search area. This means that the maximum of two independent random variables has to be computed N_c times. However, when the maximum of $\boldsymbol{C}_{\text{uncovered}}$ and any other probability mass function of a discrete random variable with the same possible outcomes $\boldsymbol{C}$ is taken, the probability mass function of this maximum equals the probability mass function of $\boldsymbol{C}$. This is shown in Equation (3-27).

$$\max(\boldsymbol{C}, \boldsymbol{C}_{\text{uncovered}}) = \boldsymbol{C} \tag{3-27}$$

The result of Equation (3-27) is that when a new coverage measurement is obtained, only the grid cells for which the probability mass function of the POD is unequal to $\boldsymbol{C}_{\text{uncovered}}$ have to be combined with the coverage map, by taking their maximum. Since this way only a fraction of the N_c grid cells have to be combined, the result shown in Equation (3-27) greatly reduces the computational cost of updating the coverage map.

When the AUV is scanning on a straight track of the lawnmower pattern, the autonomy framework obtains many sonar images from the simulation framework. When the coverage map is updated for every new sonar image, as described in this chapter, the expected values of the PODs in a part of the search area are shown in Figure 3-14.

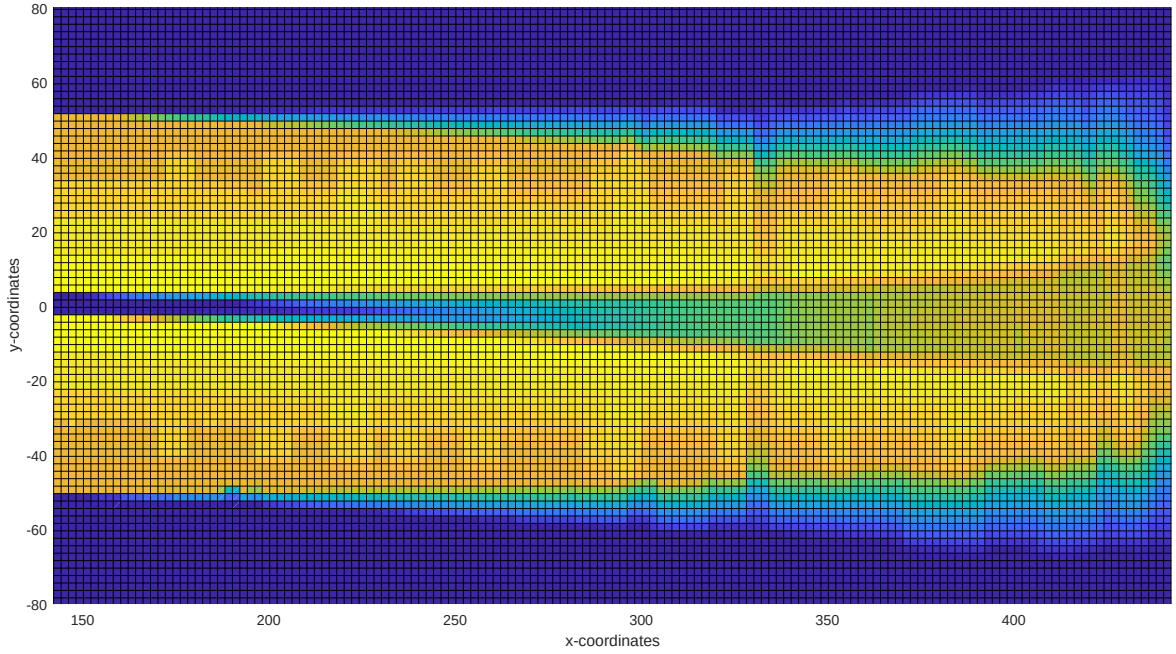

Figure 3-14: A surface plot of the expected values of the PODs under location uncertainty.

In Figure 3-14, the highest expected values of the PODs are coloured yellow, while the lowest expected values of the PODs are coloured blue. While the AUV was travelling this straight track, there were virtually no external disturbances in the form of underwater currents. In Figure 3-14, the AUV has been travelling in the positive x-direction. Surface plots showing the probability of every POD in the grid cells shown in Figure 3-14 separately can be seen in Appendix A.

It is interesting to compare the expected values of the coverage map shown in Figure 3-14, which is based on the AUV's uncertain trajectory, with the coverage map based on the measured trajectory and with the coverage map based on the real trajectory. These three maps are obtained using the same sonar model in the simulation framework, but they are transformed into the global framework according to the three different trajectories.

The coverage map showing the POD in every grid cell based on the measured trajectory is shown in Figure 3-15, which has been obtained under the exact same circumstances as the coverage map shown in Figure 3-14. The coverage map shown later in Figure 3-16 has also been obtained under these same circumstances.

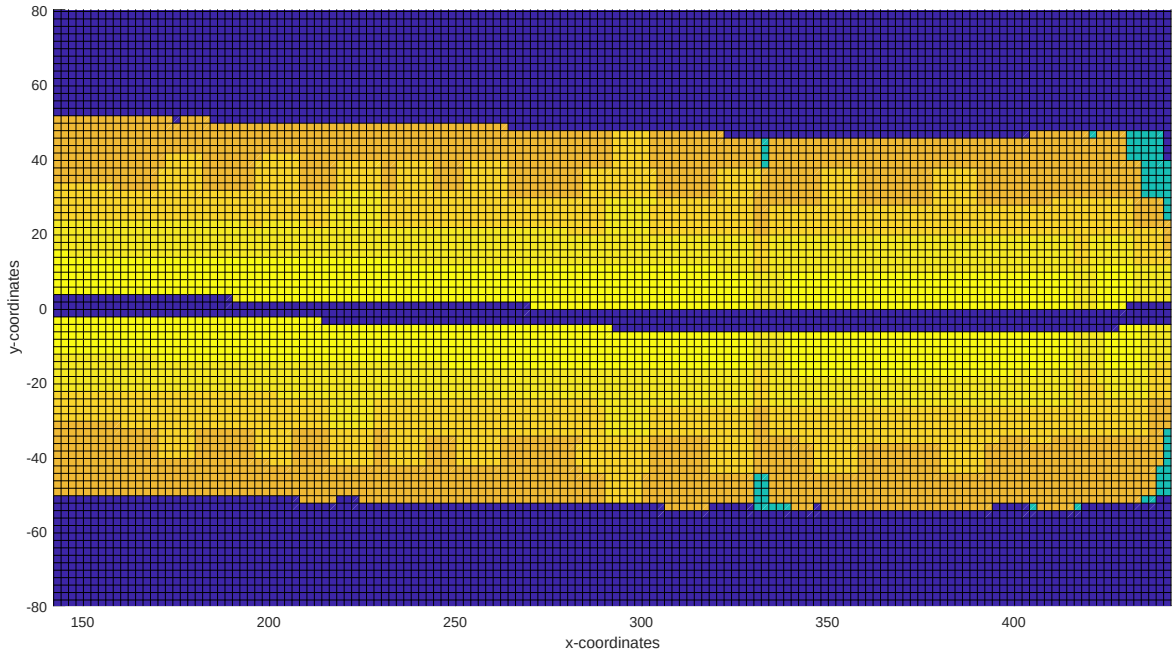

Figure 3-15: A surface plot of the PODs obtained without considering location uncertainty.

The coverage map showing the POD in every grid cell based on the real trajectory is shown in Figure 3-16. Note that since the coverage map shown in Figure 3-16 has been obtained using the AUV's real trajectory, this coverage map can only be used for comparison, since the real trajectory of the AUV is not available to the autonomy framework.

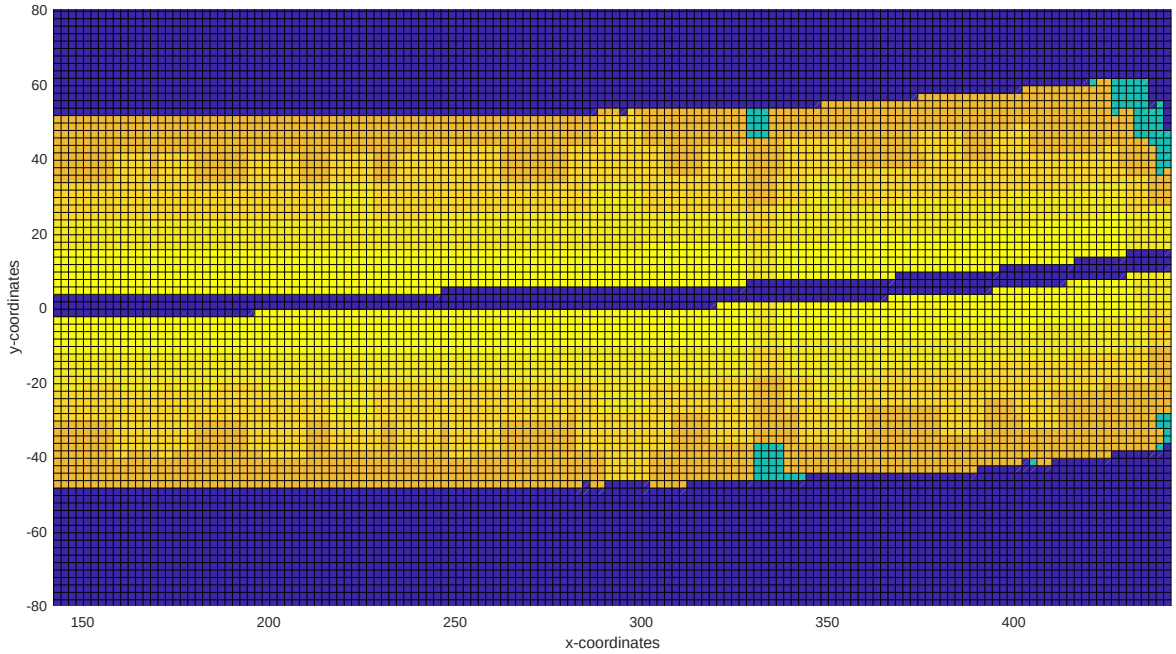

Figure 3-16: A surface plot of the PODs based on the real trajectory.

When the coverage map is maintained using the measured locations without considering the uncertainty on these measurements, a comparison between Figure 3-15 and Figure 3-16 shows that this coverage maps is very similar to the coverage map corresponding to the real trajectory at the start of the straight track. This is due to the fact that the uncertainty on the location measurements is initially small. However, since the uncertainty on the location measurements grows over time, Figure 3-15 and Figure 3-16 look different near the end of the straight track. By looking at the nadir gap, located directly below the AUV, Figure 3-15 shows that the AUV believes it has travelled in an almost perfectly horizontal straight line, while Figure 3-16 shows that this has not been the case in reality.

When the uncertainty on the AUV's location measurements is considered, the difference between the believed path and the actual path is taken into account, which can be seen by comparing Figure 3-14 with Figure 3-16. Again, both maps are very similar at the start of the straight track, since the uncertainty on the location measurements is still small at this moment. Near the end of the straight track however, Figure 3-14 shows that the AUV is not very certain any longer where in the global coordinate frame the sonar images have been obtained exactly. As a result, the PODs in every grid cell are dispersed over the surrounding grid cells, changing the PODs from deterministic values into discrete random variables. Since the measured trajectory differs significantly from the real trajectory at the end of the straight track, it is an unrealistic assumption to assume that the sonar images have been obtained from the AUV's measured location exactly. For this reason, the coverage map constructed inside the autonomy framework that takes the uncertainty on the location measurements into account, is the coverage map that is used by the coverage path planner discussed in Chapter 4.

Chapter 4

Autonomous Coverage Path Planning

4-1 Coverage Path Planning Problem

In Chapter 3 it has been discussed how an AUV can keep track of how well it has covered a search area, while taking the uncertainty on the INS location measurements into account. The AUV achieves this by maintaining a coverage map. This coverage map is used by the path planner, which enables the AUV to autonomously plan its coverage path during an operation. The location of this planner inside the autonomy framework can be seen in Figure 2-4.

When the coverage path has the shape of a lawnmower pattern, only the straight tracks of this lawnmower pattern influence the coverage in the search area, since the connections between adjacent straight tracks are located outside the search area. In this final thesis project, the coverage path is planned by optimizing the locations of these straight tracks. In order to determine the location of the coverage path through an optimization, it needs to be decided with which parameters the coverage path is described, since these parameters are the parameters that are going to be optimized. One way to describe a lawnmower-shaped coverage path is by defining a parameter for every straight track, where this parameter describes the location of this straight track in the search area. An advantage of describing a coverage path this way is that the straight tracks can be placed on the search area in any possible way. A disadvantage is that the number of parameters that has to be optimized is equal to the number of straight tracks, which means that the optimization algorithm has to optimize a large number of parameters for an average coverage path. A second disadvantage is that the number of required straight tracks is not known beforehand, so the optimization algorithm does not know how many parameters it has to optimize. Therefore, a coverage path is not described with one parameter for the location of every straight track in this final thesis project, but with the following three parameters instead.

- A_n: The number of straight tracks in the coverage path.
- A_d: The distance between adjacent straight tracks in the coverage path.
- A_l: The location of the first straight track of the coverage path in the search area.

The set of the optimization parameters A_n, A_d and A_l is referred to as A. With the three parameters in A every lawnmower-shaped coverage path can be described, provided that the straight tracks are spaced equidistantly. The optimization parameters are depicted in Figure 4-1.

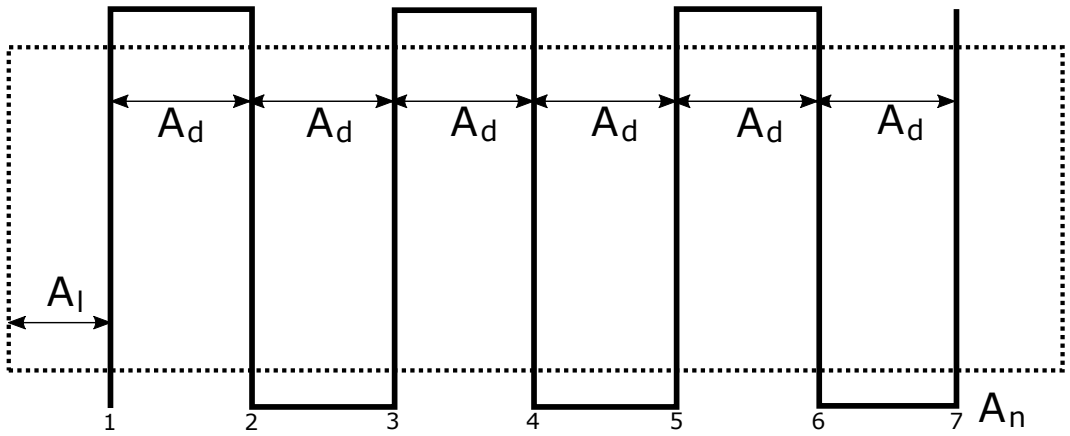

Figure 4-1: Optimization parameters.

Before A can be determined through an optimization, the AUV needs to obtain an estimate of the coverage that it can achieve with a straight track. Since the underwater conditions are usually unknown before the start of an operation, the AUV has to travel an initial straight track in order to estimate this resulting coverage. For every straight track in the coverage path, as well as for the initial track, four waypoints are generated. These waypoints are shown in Figure 4-2, in both the xy-plane and in the xz-plane. In Figure 4-2 it can be seen that the first waypoint of every straight scanning track is located at the surface, such that the AUV can reach this waypoint using the more accurate GPS location measurements instead of the less accurate INS location measurements. It can also be seen that the AUV surfaces after every straight track to execute a GPS-fix, as discussed in Chapter 3. Since the waypoints at the surface are located at the same relative location for every straight track, the AUV reaches the search area with approximately the same INS location measurement uncertainty for every straight track. This is important when the coverage that results from one straight track is used to predict the coverage that results from another straight track, which is discussed into more detail later in this chapter. The middle two waypoints of every straight scanning track are located at the scanning height of 6 metres above the ocean floor, each located 10 metres outside of the search area. The reason that the AUV starts scanning the ocean floor just outside the area, is because this increases the sonar image quality of the sonar images obtained at the boundaries of the search area. Between the two middle waypoints the AUV is in the scan mode, as discussed in Chapter 2, otherwise the AUV is in the travel mode.

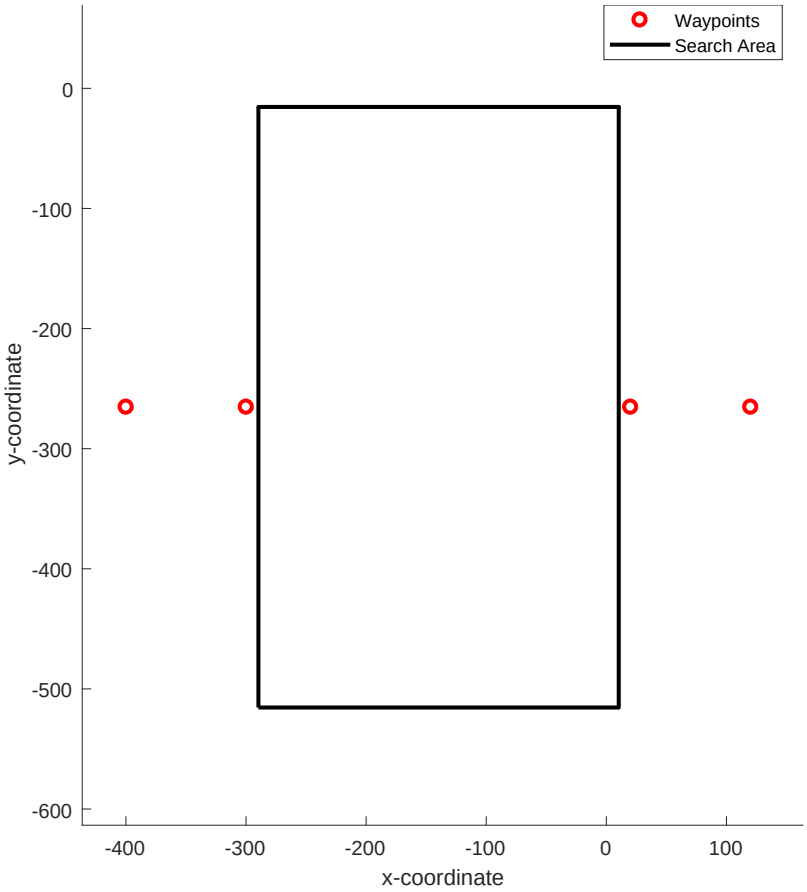

(a) The waypoints in the xy-plane.

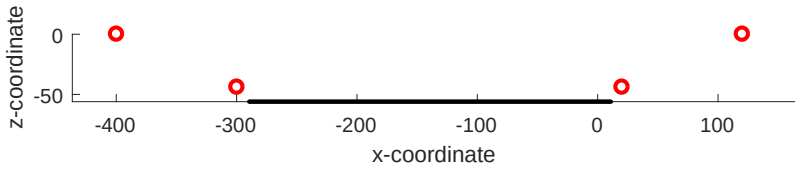

(b) The waypoints in the xz-plane.

Figure 4-2: The waypoints for a straight scanning track.

Once the AUV has completed the initial straight track, the resulting coverage of the search area can be used to determine the optimization parameters describing the coverage path through an optimization. The details of this optimization are discussed later in this chapter, for now it is important to note that after the initial track has been completed, the entire coverage path for the remainder of the operation is determined through an optimization. The reason that the entire coverage path is planned instead of only the next straight track, is that this way the total number of required straight tracks can be estimated, which can be placed on the search area in an optimal way. When instead only the location of the next straight track is optimized, it can happen that the final straight track needs to be placed close to the boundary of the search area, thus not fully utilising the potential coverage that can be achieved by following this final straight track.

Although the planner in the autonomy framework plans the entire coverage path after the initial straight track, only the next straight track of this planned coverage path is executed. Once the AUV has completed this next straight track, the entire remainder of the coverage path is planned again. The reason that the AUV replans the coverage path after every straight track, instead of planning a coverage path once and then following this coverage path for the entire operation, is that after a new straight track a new estimate of the current coverage performance is obtained. Since this latest estimate is assumed to resemble the current coverage performance more accurately than the previous estimate, replanning after every straight track enables the AUV to adapt its coverage path to changing underwater conditions affecting the AUV's coverage performance.

After a straight track has been completed, the coverage map maintained in the knowledge base contains the updated probability density functions describing the PODs in every grid cell in the search area. When these probability density functions have been obtained by combining multiple coverage measurements from different straight tracks, it is impossible to determine what coverage has been achieved exactly by the latest straight track only. For this reason, a second coverage map is maintained in the knowledge base, which is reset every time the AUV starts a new straight track. By making use of this partial coverage map, the coverage performance corresponding to the most recent straight track can easily be computed. When a straight track has been close to the boundary of the search area however, this coverage performance can be incomplete, since the coverage is only maintained on the search area. When this happens, the incomplete coverage performance estimation is rejected, and instead the previously estimated coverage performance is used to replan the coverage path. Incomplete estimations of the coverage performance near the boundaries of the search area can be avoided by maintaining the coverage in an area larger than the search area, but this increases the computational cost of both maintaining the coverage map and of the path planning.

In the next section, entropy reduction is introduced as the main objective to maximise by the optimization algorithm. The subsequent section discusses how the current estimate of the coverage performance can be used to predict the coverage of the search area in the future. The final section of this chapter discusses how this prediction is used to determine the parameters describing the coverage path through an optimization.

4-2 Entropy Reduction

The planner agent located inside the autonomy framework plans the AUV's coverage path by attempting to achieve a certain objective. For an AUV search operation, an apparent choice for this objective is to maximise the coverage in the search area while penalising the time needed to complete the operation. Typically the constraints are on the duration of the operation, and on the minimum level of coverage in the search area.

In the coverage map, every grid cell in the search area contains a discrete random variable describing the POD in that grid cell. In Chapter 3 it has been discussed that the POD has been defined as the probability of a true observation of a grid cell, so detecting an object when it is there, or not detecting an object when it is not there. This has been shown in Equation (3-3). Consequently the complement of the POD has been defined as the probability of a

false observation of a grid cell, so detecting an object when it is not there, or not detecting an object while it is there. This has been shown in Equation (3-4).

In order to define if a grid cell has been covered correctly or incorrectly, a new binary random variable $\boldsymbol{D}$ is introduced. This binary random variable $\boldsymbol{D}$ describes the probabilities of correct and incorrect coverage of a grid cell. Since a grid cell has been covered correctly when a true observation of this grid cell has been made, the probability of correctly covering a grid cell is equal to the POD in that grid cell. Likewise, a grid cell is covered incorrectly when a false observation of this grid cell has been made, so the probability of incorrectly covering a grid cell is equal to the complement of the POD. Since the uncertainty on the location measurements has caused the POD to become a discrete random variable $\boldsymbol{C}$ in the coverage map, the probability of correctly covering a grid cell can be described by the expected value of $\boldsymbol{C}$ in this grid cell. Since $\boldsymbol{D}$ is a binary random variable, the probability of incorrectly covering a grid cell is then equal to the complement of correctly covering a grid cell. The resulting probability mass function of the binary random variable $\boldsymbol{D}$, which describes the binary coverage of a grid cell, is shown in Table 4-1.

Table 4-1: Probability mass function of the binary coverage of a grid cell.

a	correct	incorrect
$\mathrm{P}(\boldsymbol{D} = a)$	$\mathrm{E}[\boldsymbol{C}]$	$1 - \mathrm{E}[\boldsymbol{C}]$

When planning a coverage path, it seems an obvious choice to plan this coverage path such that the increase in $\mathrm{P}(\boldsymbol{D} = \text{correct})$ over the entire search area is maximised. When this approach is taken however, every increase in $\mathrm{P}(\boldsymbol{D} = \text{correct})$ is considered to have the same benefit. In reality, this is not the case. When the probability of correct coverage of a grid cell is for example increased from 0 to 0.3, there is little benefit in knowing this, since either way there is still too little information about this grid cell. When the probability of correct coverage of a grid cell is increased from 0.7 to 1 however, it is clear that this is very useful to know. When the objective is to increase $\mathrm{P}(\boldsymbol{D} = \text{correct})$, both increases will yield the same result regarding this objective. A way to improve this is to maximise the entropy reduction [20], instead of maximising the increase in $\mathrm{P}(\boldsymbol{D} = \text{correct})$.

Information entropy, also called Shannon entropy, is a concept from information theory. This entropy tells something about the amount of uncertainty associated with the realization of a discrete random variable. When all the outcomes of a discrete random variable are equally probable, the entropy has its maximum value. On the other hand, when the probability of one of the outcomes of the discrete random variable approaches 1, the entropy approaches 0. The entropy $\mathrm{H}(\boldsymbol{X})$ of a discrete random variable $\boldsymbol{X}$ with the possible outcomes $a_1, a_2, \ldots, a_n$ can be computed through Equation (4-1) [27].

$$\mathrm{H}(\boldsymbol{X}) \;=\; -\sum_{i=1}^{n} \mathrm{P}(\boldsymbol{X} = a_i) \log_b(\mathrm{P}(\boldsymbol{X} = a_i)) \tag{4-1}$$

When the discrete random variable is a binary random variable, the probabilities of its both possible outcomes are p and $1 - p$. When the base of the logarithm in Equation (4-1) is set

to $b = 2$, such that the entropy is normalised, Equation (4-1) simplifies to Equation (4-2) for a binary random variable $\boldsymbol{X}_{\text{binary}}$.

$$\mathrm{H}(\boldsymbol{X}_{\text{binary}}) \;=\; -p\log_2(p) - (1-p)\log_2(1-p) \tag{4-2}$$

In Equation (4-2), the entropy is defined to be 0 when $p = 0$ or $p = 1$. This is consistent with the limit shown in Equation (4-3), which can be proved using `l'Hospital's Rule` [28].

$$\lim_{0<x\to 0} x\log(x) = 0 \tag{4-3}$$

In Figure 4-3 the entropy of a binary random variable is shown as a function of p in the interval [0,1], according to Equation (4-2).

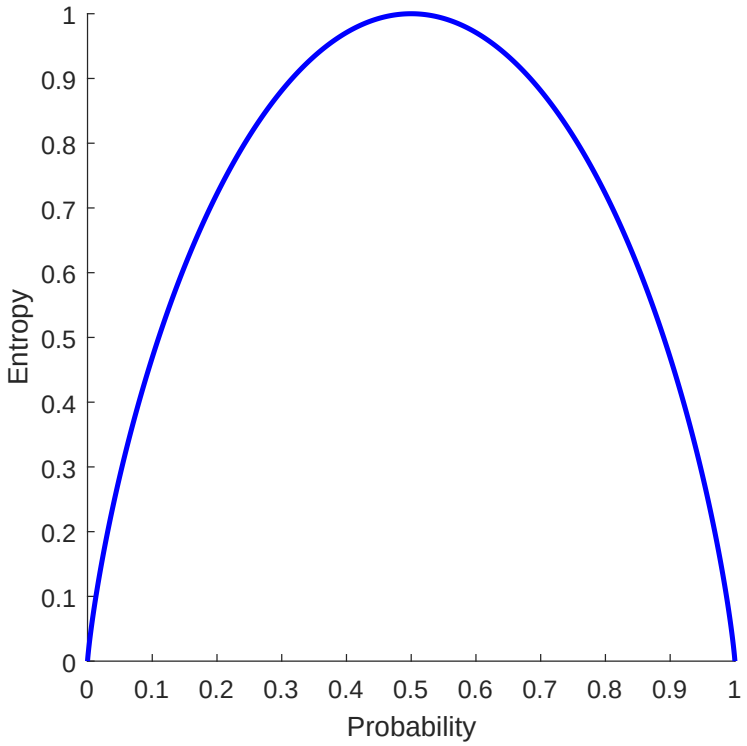

Figure 4-3: Entropy as a function of the probability for a binary random variable.

Figure 4-3 shows that the entropy is symmetrical around its maximum value. The location of the maximum is as expected, since at $p = 0.5$ the probabilities of both outcomes of the binary random variable are equal. When the binary random variable under consideration is $\boldsymbol{D}$, as shown in Table 4-1, it can be seen that it does not matter if $\mathrm{P}(\boldsymbol{D} = \text{correct})$ or $\mathrm{P}(\boldsymbol{D} = \text{incorrect})$ is taken as the probability in Figure 4-3, since the resulting entropy is the same due to its symmetry.

The reason that the entropy has been introduced, is to be able to maximise the entropy reduction. For this purpose, a symmetric entropy as shown in Figure 4-3 is not suitable. Instead, the entropy should be at its maximum when $\mathrm{P}(\boldsymbol{D} = \text{correct}) = 0$, and at its minimum

when $\mathrm{P}(\boldsymbol{D} = \text{correct}) = 1$, such that the entropy is always reduced when $\mathrm{E}[\boldsymbol{C}]$ increases. This desired entropy function can be obtained by scaling and translating the entropy function shown in Figure 4-3. When the probabilities on the x-axis of Figure 4-3 are p_{original}, the shifted entropy as a function of p_{shifted} instead of p_{original} is obtained by substituting Equation (4-4) into Equation (4-2).

$$p_{\text{shifted}} = 2 \cdot p_{\text{original}} - 1 \tag{4-4}$$

The relation shown in Equation (4-4) is only meaningful when p_{original} is in the interval [0.5,1], and when p_{shifted} is in the interval [0,1]. Furthermore, p_{original} can be equal to both $\mathrm{P}(\boldsymbol{D} = \text{correct})$ and $\mathrm{P}(\boldsymbol{D} = \text{incorrect})$, while p_{shifted} is only equal to $\mathrm{P}(\boldsymbol{D} = \text{correct})$, or the equivalent $1 - \mathrm{P}(\boldsymbol{D} = \text{incorrect})$. The resulting shifted entropy as a function of p_{shifted} in the interval [0,1] is shown in Figure 4-4.

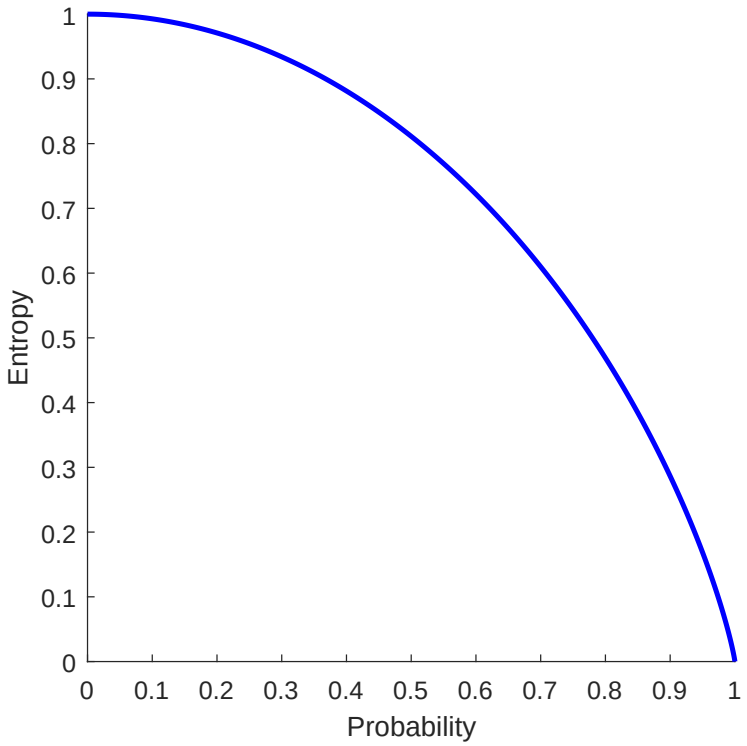

Figure 4-4: Shifted entropy as a function of the probability for a binary random variable.

Figure 4-4 clearly shows that the shifted entropy reduces for every increase in p_{shifted}. Since p_{shifted} is equal to $\mathrm{P}(\boldsymbol{D} = \text{correct})$, Equation (4-5) describes the shifted entropy as a function of the binary random variable $\boldsymbol{D}$.

$$\begin{aligned} \mathrm{H}(\boldsymbol{D}) \;=\; & -\frac{1}{2}(1 + \mathrm{P}(\boldsymbol{D} = \text{correct}))\log_2\left(\frac{1}{2}(1 + \mathrm{P}(\boldsymbol{D} = \text{correct}))\right) \\ & -\frac{1}{2}(1 - \mathrm{P}(\boldsymbol{D} = \text{correct}))\log_2\left(\frac{1}{2}(1 - \mathrm{P}(\boldsymbol{D} = \text{correct}))\right) \end{aligned} \tag{4-5}$$

Equation (4-5) has been obtained by substituting Equation (4-4) into Equation (4-2). Using a logarithmic identity and the fact that $\mathrm{P}(\boldsymbol{D} = \text{correct}) = 1 - \mathrm{P}(\boldsymbol{D} = \text{incorrect})$, Equation

(4-5) can be further simplified to Equation (4-6).

$$\begin{aligned} \mathrm{H}(\boldsymbol{D}) \;=\; & -\frac{1}{2}(1+\mathrm{P}(\boldsymbol{D}=\text{correct}))\log_2\left(1+\mathrm{P}(\boldsymbol{D}=\text{correct})\right) \\ & -\frac{1}{2}\mathrm{P}(\boldsymbol{D}=\text{incorrect})\log_2\left(\mathrm{P}(\boldsymbol{D}=\text{incorrect})\right) \;+\; 1 \end{aligned} \tag{4-6}$$

Since the probability mass function of $\boldsymbol{D}$, shown in Table 4-1, is a function of the probability mass function of $\boldsymbol{C}$, the shifted entropy shown in Equation (4-6) can also be written as a function of the discrete random variable $\boldsymbol{C}$. The resulting function is shown in Equation (4-7).

$$\begin{aligned} \mathrm{H}(\boldsymbol{C}) \;=\; & -\frac{1}{2}(1+\mathrm{E}[\boldsymbol{C}])\log_2\left(1+\mathrm{E}[\boldsymbol{C}]\right) \\ & -\frac{1}{2}(1-\mathrm{E}[\boldsymbol{C}])\log_2\left(1-\mathrm{E}[\boldsymbol{C}]\right) \;+\; 1 \end{aligned} \tag{4-7}$$

The shifted entropy as a function of $\boldsymbol{C}$, as shown in Equation (4-7), is very useful in this final thesis project. Since the coverage map contains a probability mass function of the discrete random variable $\boldsymbol{C}$ describing the POD for every grid cell in the search area, Equation (4-7) can be used to compute the shifted entropy in this grid cell directly. In the remainder of this final thesis project, the shifted entropy shown in Figure 4-4 will simply be referred to as the entropy.

The usefulness of maximising the entropy reduction instead of maximising the increase in $\mathrm{P}(\boldsymbol{D}=\text{correct})$ becomes clear by returning to the example given earlier in this section. This example stated that when $\mathrm{P}(\boldsymbol{D}=\text{correct})$ is increased from 0 to 0.3, this yields the same benefit as an increase from 0.7 to 1, while the latter increase is obviously more useful. When the entropies for this example are compared, it is found that an increase of $\mathrm{P}(\boldsymbol{D}=\text{correct})$ from 0.7 to 1 leads to an entropy reduction over 9 times larger than the entropy reduction for an increase from 0 to 0.3. Increasing $\mathrm{P}(\boldsymbol{D}=\text{correct})$ from 0 to 0.78 amounts to the same entropy reduction as increasing $\mathrm{P}(\boldsymbol{D}=\text{correct})$ from 0.78 to 1, which corresponds well with the benefits of this increase in reality.

4-3 Coverage Prediction Models

In order to predict the entropy reduction that will occur if the AUV follows a certain coverage path, a model of how a coverage path influences the entropy reduction in the search area is required. In this final thesis project, the entropy reduction is predicted using a projection of the partial coverage map, which contains the coverage corresponding to the latest straight track only. This method has not been encountered in literature before.

This section is divided into two subsections. In the first subsection, three different models to estimate the latest coverage performance are discussed. In the second subsection, it is discussed how the coverage performance in the future is predicted using the three different models and their coverage estimations.

Estimation Of Latest Coverage Performance

In the first section of this chapter it has been discussed that when the search operation starts, the AUV first plans an initial straight scanning track in the search area. The objective of this initial straight track is to get an estimate of the resulting coverage after a straight track, since this is unknown before the start of the operation. Whenever the AUV is in the scanning mode, the coverage map discussed in Chapter 3 is constantly updated. After the initial scanning track has been completed, a coverage map such as shown in Figure 3-14 has been obtained. For every grid cell in the search area, this coverage map contains a discrete random variable $\boldsymbol{C}$ that describes the POD in that grid cell.

The coverage map can be used to estimate the coverage corresponding to the latest straight track, and to predict the future coverage of the search area. Three different models that can estimate and predict the coverage from the coverage map are briefly introduced now, and discussed into more detail later in this subsection. `Coverage Performance Model 1` uses the coverage map containing a discrete random variable $\boldsymbol{C}$ in every grid cell directly, in order to estimate and predict the coverage performance in the search area. Instead of using the coverage map with a discrete random variable $\boldsymbol{C}$ in every grid cell, a map containing the expected value of $\boldsymbol{C}$ in every grid cell can be used instead. The expected value in every grid cell is computed through Equation (3-19), and `Coverage Performance Model 2` uses this map containing the expected value of $\boldsymbol{C}$ in every grid cell to estimate and predict the coverage performance in the search area. A third option is to use Equation (4-7) to construct a map that contains an entropy in every grid cell instead of an expected value of $\boldsymbol{C}$. This map is used by `Coverage Performance Model 3` to estimate and predict the coverage performance in the search area.

In the autonomy framework, the partial coverage map resulting from a straight track is used to estimate the coverage performance resulting from this straight track. This coverage performance is later used to predict the coverage in the search area after completing the coverage path. In order to estimate the coverage performance achieved during the latest straight track, the partial coverage map corresponding to this straight track is averaged in the direction of the straight track. This results in a projection of the partial coverage map onto the AUV's lateral plane, giving the average coverage as a function of the lateral distance to the AUV. In the remainder of this section, it is discussed into more detail how these projections estimating the coverage performance are obtained for each of the three different coverage performance models.

Coverage Performance Model 1

The first coverage performance model uses the partial coverage map that contains a discrete random variable $\boldsymbol{C}$ in every grid cell. This coverage performance model preserves the most information about the coverage map out of the three models. In order to obtain the average discrete random variable $\boldsymbol{C}$ as a function of the lateral distance to the AUV, the average of the discrete random variables $\boldsymbol{C}$ has to be computed in the direction of the straight tracks of the coverage path, over the entire length of the search area. Since a discrete random variable is described by its probability mass function, the average of a number of discrete random variables can be computed by computing the average of their probability mass functions.

This average probability mass function can be obtained by computing the average probability for every discrete possible outcome of the probability mass functions separately. When the average probability mass functions in the direction of the straight track are computed over the entire length of the search area, the average probability of every outcome of $\boldsymbol{C}$ as a function of the lateral distance to the AUV can look like the projections shown in Figure 4-5.

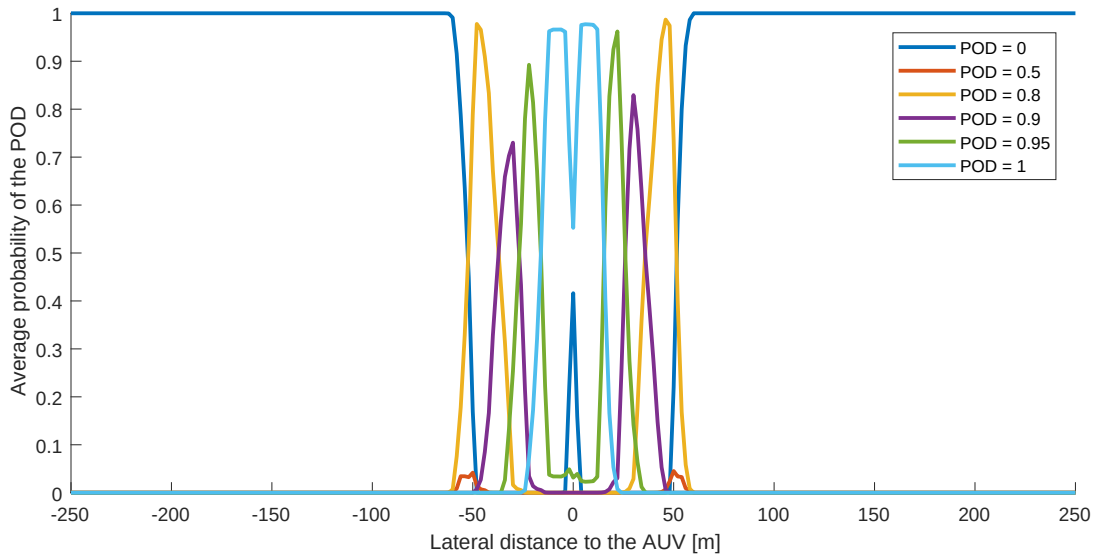

Figure 4-5: Average POD probability as a function of the lateral distance to the AUV.

In Figure 4-5, lateral distances to the left of the AUV are shown as negative, while lateral distances to the right of the AUV are shown as positive. Figure 4-5 shows that the sonar range in this case was approximately 60 metres, since $\mathrm{P}(\boldsymbol{C} = 0) = 1$ everywhere outside of this range. The presence of the nadir gap also shows in Figure 4-5, since $\mathrm{P}(\boldsymbol{C} = 1)$ significantly decreases right below the AUV.

Coverage Performance Model 2

The second coverage performance model takes the partial coverage map that has been used by the first coverage performance model, and computes the expected value of $\boldsymbol{C}$ in every grid cell. Subsequently, the averages of these expected values are computed in the direction of the straight track over the entire length of the search area, such that the average expected value of the POD as a function of the lateral distance to the AUV is obtained. This is however exactly the same as computing the average probability mass function projection shown in Figure 4-5, and then computing the expected values corresponding to this projection, since it can be proven that Equation (4-8) holds.

$$\mathrm{E}[\mathrm{mean}(\boldsymbol{C}_1, \boldsymbol{C}_2, \ldots, \boldsymbol{C}_n)] \;=\; \mathrm{mean}(\mathrm{E}[\boldsymbol{C}_1], \mathrm{E}[\boldsymbol{C}_2], \ldots, \mathrm{E}[\boldsymbol{C}_n]) \tag{4-8}$$

The resulting average expected value of $\boldsymbol{C}$ as a function of the lateral distance to the AUV, obtained under the same conditions as the projections shown in Figure 4-5, is shown in Figure 4-6. In Figure 4-6, the sonar range and the nadir gap are clearly visible.

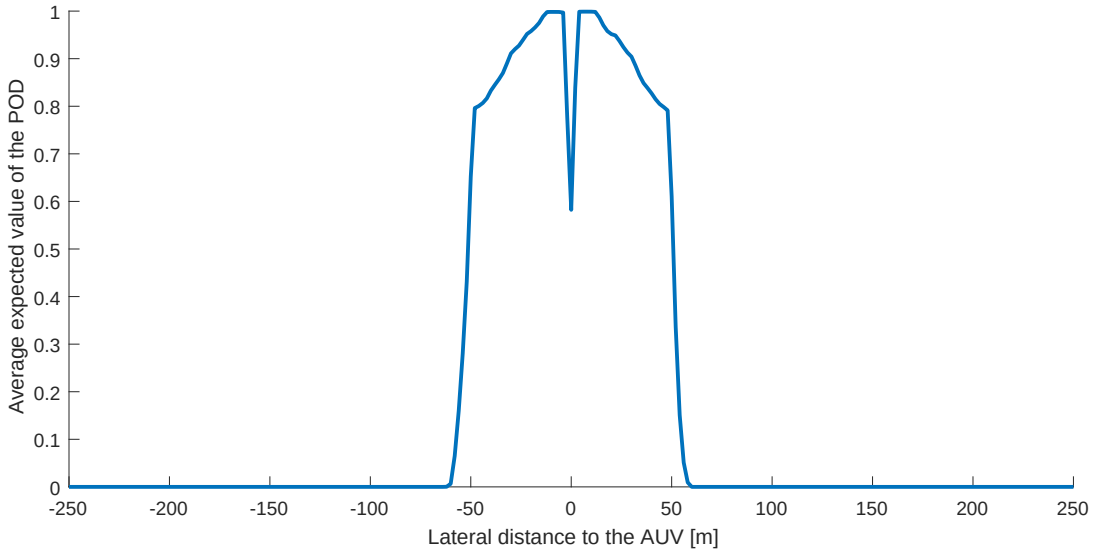

Figure 4-6: Average POD expectation as a function of the lateral distance to the AUV.

Coverage Performance Model 3

The third coverage performance model starts by taking the partial coverage map containing a discrete random variable $\boldsymbol{C}$ in every grid cell to compute the expected value of $\boldsymbol{C}$ in every grid cell, equal to the approach taken by the second coverage performance model. The map containing the expected values of $\boldsymbol{C}$ is then used to compute the entropy in every grid cell. From the map with an entropy in every grid cell, the averages of the entropies in the direction of the straight track are computed, over the entire length of the search area. Figure 4-7 has been obtained under the same conditions as Figure 4-5 and Figure 4-6, and it shows the average entropy of the expected value of $\boldsymbol{C}$ as a function of the lateral distance to the AUV.

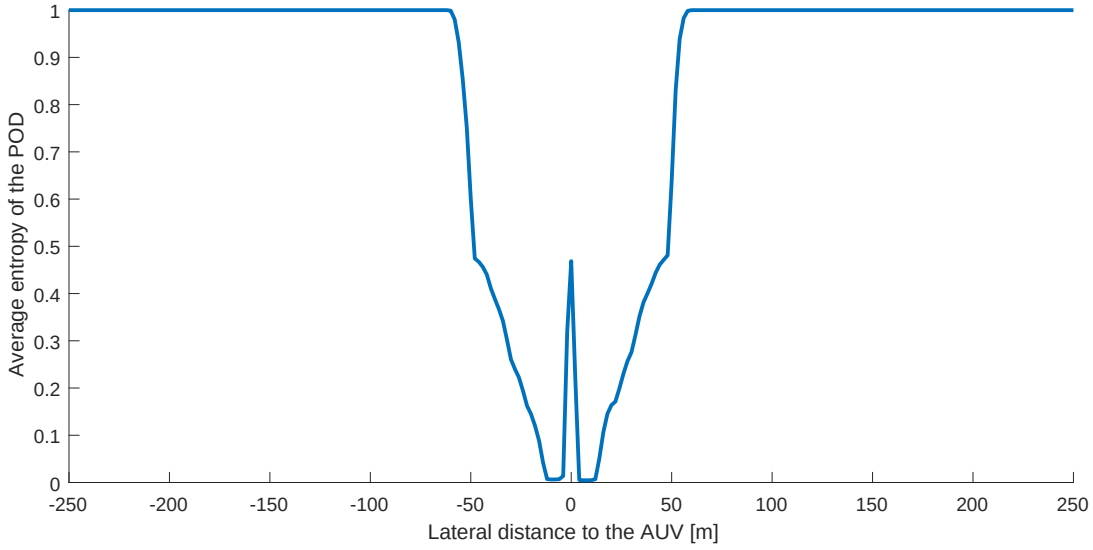

Figure 4-7: Average POD entropy as a function of the lateral distance to the AUV.

Just as in Figure 4-5 and Figure 4-6, the sonar range and nadir gap are also clearly visible in Figure 4-7. Note that although it may seem as if Figure 4-7 can be obtained by flipping Figure 4-6 upside down, this is not the case.

When instead of first constructing a map with an entropy in every grid cell and then computing the average entropy in the direction of the straight track for every grid cell, the entropies are computed from the average expected values of $\boldsymbol{C}$ shown in Figure 4-6, a different projection than the projection shown in Figure 4-7 is obtained. Since the domain of the entropy function is given by the range of the expected value of $\boldsymbol{C}$, the entropy function is only defined in the interval [0,1], as shown in Figure 4-4. In this interval the entropy function is a strictly concave function, which means that `Jensen's inequality` [29] can be used to prove that Equation (4-9) holds.

$$\mathrm{H}(\mathrm{mean}(\mathrm{E}[\boldsymbol{C}_1], \mathrm{E}[\boldsymbol{C}_2], \ldots, \mathrm{E}[\boldsymbol{C}_n])) \;\geq\; \mathrm{mean}(\mathrm{H}(\mathrm{E}[\boldsymbol{C}_1]), \mathrm{H}(\mathrm{E}[\boldsymbol{C}_2]), \ldots, \mathrm{H}(\mathrm{E}[\boldsymbol{C}_n])) \tag{4-9}$$

From Equation (4-9) it follows that when the entropies are computed from the average expected value projection shown in Figure 4-6, these entropies can be too pessimistic. Therefore, the third coverage performance model first computes the entropies in every grid cell, after which the average entropy is taken in the direction of the straight track, in order to obtain a projection such as shown in Figure 4-7.

Prediction Of Future Coverage Performance

The three projections corresponding to the three different coverage performance models can each be used to predict the coverage resulting from following the next straight track in the coverage path. The projections describe the average $\boldsymbol{C}$, $\mathrm{E}[\boldsymbol{C}]$ or $\mathrm{H}(\mathrm{E}[\boldsymbol{C}])$ that resulted from the latest straight track, as a function of the lateral distance to the AUV. Since a GPS-fix is obtained in the same location before each straight track, as shown in Figure 4-2, the location measurement uncertainty at the start of the next straight track can be assumed to be similar to the location measurement uncertainty at the start of the straight track from which the projections have been obtained. When the underwater conditions during the next straight track are also assumed to be similar to the underwater conditions during the latest straight track, the resulting coverage after a new straight track can be predicted with the projection obtained from the latest straight track.

Since the next straight track starts on the opposite side of the search area compared to the previous straight track, the projection has to be mirrored, since the AUV's deviation from the coverage path tends to be to the same side. This mirrored projection is then assumed to describe the average coverage performance of the AUV on the next straight track, as a function of the lateral distance to the AUV. When the location of the next straight track is known, the mirrored projection can be centred around this location, such that the average coverage of the search area in the direction of the straight track is described by the initial projection and the mirrored projection together, provided that these projections do not overlap. If the projections do overlap, the two projections have to be combined. How two projections are combined is discussed for each of the coverage performance models later in this subsection.

Instead of predicting the coverage after the next straight track by combining the initial and mirrored projection, the coverage after the next straight track can also be predicted by using the complete partial coverage map. When the partial coverage map of the latest straight track is mirrored and centred around the location of the new straight track, the combination of these partial coverage maps predicts the coverage in the search area after the next straight track. When this resulting partial coverage map is averaged in the direction of the straight tracks, a projection of the coverage in the search area has been obtained again, but this projection is different from the one obtained by combining projections instead of partial coverage maps. The difference between using projections and partial coverage maps is discussed for each of the three different coverage performance models separately later in this subsection. Since combining partial coverage maps instead of projections is computationally significantly more expensive, all three coverage performance models predict the coverage of the search area after a new straight track using projections of the partial coverage maps instead of the complete partial coverage maps.

The coverage resulting from following every subsequent straight track in the coverage path can be obtained by repeating the method used to predict the coverage after the second straight track, so by mirroring and combining the initial projection for every straight track in the coverage path. When this is done for every straight track in the coverage path, a projection of the resulting estimated average $\boldsymbol{C}$, E[$\boldsymbol{C}$] or H(E[$\boldsymbol{C}$]) in the search area has been obtained. Thus, a method has been obtained to estimate the coverage of the search area after following a certain coverage path, based on the latest coverage performance of the AUV.

Coverage Performance Model 1

The first coverage performance model uses the projection shown in Figure 4-5, describing the average discrete random variable $\boldsymbol{C}$ as a function of the distance to the AUV. When two or more of these projections overlap, they are combined by computing the probability mass function of the maximum of the average probability mass functions of the overlapping projections. The probability mass function of the maximum of discrete random variables is computed through Equation (3-21) and Equation (3-23). When instead of the projections the partial coverage maps are combined, Equation (4-10) shows that the new projection obtained by combining two partial coverage maps is different from the new projection that has been obtained by combining the projections of the two partial coverage maps.

$$\begin{aligned} \max(\text{mean}(\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n), \text{mean}(\boldsymbol{Y}_1, \boldsymbol{Y}_2, \ldots, \boldsymbol{Y}_n)) \neq \\ \text{mean}(\max(\boldsymbol{X}_1, \boldsymbol{Y}_1), \max(\boldsymbol{X}_2, \boldsymbol{Y}_2), \ldots, \max(\boldsymbol{X}_n, \boldsymbol{Y}_n)) \end{aligned} \tag{4-10}$$

In Equation (4-10) $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_n$ are n discrete random variables describing the POD in a line of grid cells in the direction of the straight tracks in the first partial coverage map, $\boldsymbol{Y}_1, \boldsymbol{Y}_2, \ldots, \boldsymbol{Y}_n$ are the n discrete random variables in the same grid cells in the second partial coverage map, and the two partial coverage maps are independent.

Coverage Performance Model 2

The second coverage performance model uses the projection shown in Figure 4-6, describing the average expected value of $\boldsymbol{C}$ as a function of the distance to the AUV. When two or more of these projections overlap, they are combined by computing the maximum of the overlapping

expected values of the two projections. When instead of the projections the partial coverage maps are combined, the maximum of the expected values in the overlapping grid cells is taken.

When $x_1, x_2, \ldots, x_n$ are the expected values in a line of grid cells in the direction of the straight tracks in the first partial coverage map, and $y_1, y_2, \ldots, y_n$ the expected values in the same grid cells in the second partial coverage map, it can be proven that Equation (4-11) holds.

$$\begin{aligned} \max(\text{mean}(x_1, x_2, \ldots, x_n), \text{mean}(y_1, y_2, \ldots, y_n)) \ \leq \\ \text{mean}(\max(x_1, y_1), \max(x_2, y_2), \ldots, \max(x_n, y_n)) \end{aligned} \tag{4-11}$$

Equation (4-11) shows that when two projections are combined, the resulting expected values can be lower than the expected values resulting from the combination of two partial coverage maps. This lower expected value is preferred over an expected value that is higher than the expected value resulting from the combination of two partial coverage maps, due to the serious possible consequences of overestimating the coverage discussed in Chapter 3.

Coverage Performance Model 3

The third coverage performance model uses the projection shown in Figure 4-7, describing the average entropy as a function of the distance to the AUV. When two or more of these projections overlap, they are combined by computing the minimum of the overlapping entropies of the two projections. When instead of the projections the partial coverage maps are combined, the minimum of the entropies in the overlapping grid cells is taken.

When $x_1, x_2, \ldots, x_n$ are the entropies in a line of grid cells in the direction of the straight tracks in the first partial coverage map, and $y_1, y_2, \ldots, y_n$ the entropies in the same grid cells in the second partial coverage map, it can be proven that Equation (4-12) holds.

$$\begin{aligned} \min(\text{mean}(x_1, x_2, \ldots, x_n), \text{mean}(y_1, y_2, \ldots, y_n)) \ \geq \\ \text{mean}(\min(x_1, y_1), \min(x_2, y_2), \ldots, \min(x_n, y_n)) \end{aligned} \tag{4-12}$$

Equation (4-12) shows that when two projections are combined, the resulting entropies can be higher than the entropies resulting from the combination of two partial coverage maps. This higher entropy is preferred over an entropy that is lower than the entropy resulting from the combination of two partial coverage maps, again due to the serious possible consequences of overestimating the coverage.

4-4 Optimization

In the previous section, a method to predict the coverage of the search area after following a certain coverage path, based on the latest coverage performance, has been discussed. This method can now be used to determine the location of the coverage path through an optimization, which maximises the entropy reduction in the search area.

In this section, the optimization parameters that have been discussed in the first section of this chapter are discussed into more detail first. Then, the optimization problem is presented and discussed for all three different coverage performance models separately, and the section

concludes with discussing the optimization algorithm that is used to solve the optimization problems.

Optimization Parameters

In the first section of this chapter it has been discussed that the optimization parameters consist of the number of straight tracks A_n, the distance between adjacent straight tracks A_d and the location of the first straight track in the search area A_l, all shown in Figure 4-1. Furthermore, in Chapter 3 it has been discussed that the AUV's location is described by a discrete random variable, which possible outcomes are always at the centre of a grid cell. So when the AUV attempts to follow a straight track that is not located through the centre of grid cells, the AUV's location is discretised away from this straight track when the coverage map is constructed. For this reason, the planner is restricted to plan the straight tracks through the centre of grid cells in this final thesis project.

The restriction that the planner can only locate the straight tracks through the centre of grid cells results in A_l being an integer, which describes the index of these grid cells in the search area. This implies that A_d is an integer as well. Since A_n is clearly also an integer, the coverage path in this final thesis project is described by three integer parameters. To give an example of the lower and upper bounds on these optimization parameters, the lower and upper bounds that have been used to optimize a coverage path for a 300×500 metre search area are shown in Table 4-2 for all three optimization parameters.

Table 4-2: Optimization parameters and their domains.

Parameter	Domain
A_n	$\mathbb{Z} \in [1, 20]$
A_d	$\mathbb{Z} \in [1, 75]$
A_l	$\mathbb{Z} \in [1, 250]$

In order to scan a 300×500 metre search area, the straight tracks of the coverage path are placed such that their length is minimised, in order to keep the uncertainty on the AUV's location measurements as low as possible during the operation. This means that the location of a straight track can be described by a single location on the 500 metre long axis. Since it has been discussed in Chapter 3 that the size of the grid cells in the search area is 2×2 metres, Table 4-2 shows that the optimization parameter A_l is restricted to be an integer in the interval [1,250]. This interval can not be decreased in size, since the optimization algorithm should be free to place a straight track anywhere inside the search area. Table 4-2 also shows that the maximum number of straight tracks is 20, and that the maximum distance between two planned straight tracks is 150 metres. In theory it is possible that the distance between two straight tracks is 500 metres, but in order to decrease the search space of the optimization algorithm this maximum distance is decreased to 150 metres. Note that it is still possible for the AUV to travel two tracks that are 500 metres apart, by for example planning 5 straight tracks with a distance of 100 metres in between them. Since only the first straight track of the planned coverage path is executed, the next straight track can be planned 500 metres

away from the first straight track, such that the AUV travels two consecutive tracks that are 500 metres apart. However, a scenario where this would be required is highly unlikely.

Optimization Problem

Now that the optimization parameters have been defined, the optimization problem can be stated. Since the prediction of the coverage of the search area after following a coverage path is different for each of the coverage performance models, the optimization problem is also different for each of these models. However, the general idea behind the optimization problem is the same for all models. For this reason, the general optimization problem is introduced first, after which the optimization problems used by each of the coverage performance models are discussed separately. The general optimization problem is shown in Equation (4-13).

$$\begin{aligned} \arg\min_{A} \quad & -H_R(A) + w \cdot A_n \\ \text{s.t.} \quad & c - \bar{C}(A) \leq 0 \end{aligned} \tag{4-13}$$

A brief description of the terms shown in Equation (4-13) is presented in Table 4-3, after which the optimization problems are discussed separately for each of the coverage performance models.

Table 4-3: A brief description of the terms in the general optimization problem.

A	Set of optimization parameters describing the coverage path
A_n	Number of straight tracks in the lawnmower pattern
H_R	Function calculating the entropy reduction
$\bar{C}$	Function calculating the average coverage of the search area
w	Weight factor associated with the cost of a straight track
c	Constant describing the required average coverage of the search area

Coverage Performance Model 1

When the first coverage performance model is used to predict the coverage resulting from following a certain coverage path, the optimization problem is as shown in Equation (4-14).

$$\begin{aligned} \arg\min_{A} \quad & -H_{R,p}(A) + w \cdot A_n \\ \text{s.t.} \quad & c_1 - \bar{E}_p(A) \leq 0 \end{aligned} \tag{4-14}$$

In the optimization problem shown in Equation (4-14), the non-linear function $H_{R,p}$ calculates the entropy reduction in the search area when the coverage path is described by A. In order to compute the entropy reduction, the current average entropy in the search area needs to be known, as well as the predicted average entropy in the search area after following the coverage path. The difference between these entropies gives the average entropy reduction. First the current average entropy in the search area can be computed from the coverage map, by computing the projection from the partial coverage map maintained in the knowledge base, and by then computing the average expected values and subsequently the average entropies of this projection, after which the current average entropies in the direction of the straight

tracks have been obtained. Secondly, the predicted entropy after following the coverage path has to be computed. Given the coverage path parameters A, the projection can be used to predict the average probability mass functions after following the coverage path, as discussed in the previous section. From these average probability mass functions, the predicted average expected value of the POD in the search area $\bar{E}_p$ can be computed, as well as the predicted average entropies after completing the coverage path A. When the difference of these predicted average entropies and the current average entropies is taken and summed over the area, the entropy reduction in the search area $H_{R,p}$ is obtained. Furthermore, the parameter c_1 describes the required average expected value of the POD in the search area, providing a lower bound on $\bar{E}_p$. Finally, the problem statement contains the product $w \cdot A_n$. When this product would be excluded from the problem statement, the optimization algorithm would find an optimal solution where the number of tracks is equal to its upper bound. In this final thesis project, the value of w is chosen such that it is higher than the maximum possible value of the entropy reduction, such that a lower number of straight tracks is always preferred over a higher entropy reduction, provided that the minimum coverage constraint is satisfied. This value of w is used in every optimization problem in this final thesis project, regardless of the coverage performance model.

Coverage Performance Model 2

When the second coverage performance model is used to predict the coverage resulting from following a certain coverage path, the optimization problem is as shown in Equation (4-15).

$$\begin{aligned} \arg\min_{A} \quad & -H_{R,e}(A) + w \cdot A_n \\ \text{s.t.} \quad & c_1 - \bar{E}_e(A) \leq 0 \end{aligned} \tag{4-15}$$

In the optimization problem shown in Equation (4-15), the non-linear function $H_{R,e}$ calculates the entropy reduction in the search area when the coverage path is described by A. This is done by first computing the projection from the partial coverage map maintained in the knowledge base, and by then computing the average entropies of this projection, such that the current average entropies in the direction of the straight tracks have been obtained. Next, the predicted average entropy after following the coverage path has to be computed. In order to do so, the projection is used to predict the average expected value $\bar{E}_e$ in the search area after completing the coverage path described by A. From the average expected values, the predicted average entropies after completing the coverage path can be computed. When the difference of the predicted average entropies and the current average entropies is taken and summed over the area, the entropy reduction in the search area $H_{R,e}$ is obtained. Again, the parameter c_1 describes the required average expected value of the POD in the search area, providing a lower bound on $\bar{E}_e$.

Coverage Performance Model 3

When the third coverage performance model is used to predict the coverage resulting from following a certain coverage path, the optimization problem is as shown in Equation (4-16).

$$\begin{aligned} \arg\min_{A} \quad & -H_{R,h}(A) + w \cdot A_n \\ \text{s.t.} \quad & \bar{H}_h(A) - c_2 \leq 0 \end{aligned} \tag{4-16}$$

In the optimization problem shown in Equation (4-16), the non-linear function $H_{R,h}$ calculates the entropy reduction in the search area when the coverage path is described by A. This is done by first computing the projection from the partial coverage map maintained in the knowledge base, such that the current average entropies in the direction of the straight tracks have been obtained. Next, the predicted average entropy after following the coverage path has to be computed. In order to do so, the projection is used to predict the average entropies in the search area $\bar{H}_h$, after completing the coverage path described by A. When the difference of the predicted average entropies and the current average entropies is taken and summed over the area, the entropy reduction in the search area $H_{R,h}$ is obtained. The parameter c_2 describes the required average entropy in the search area, providing an upper bound on $\bar{H}_h$.

Optimization Algorithm

In this final thesis project, a genetic algorithm (GA) is used to optimize the coverage path. Since all of the optimization problems are highly non-linear, and virtually impossible to write as an explicit function of the optimization parameters A, which are all integers, an evolutionary algorithm is a suitable choice. Out of the different evolutionary algorithms, the GA is chosen since it is readily available in MATLAB and since it can explicitly handle integer optimization parameters. It is not claimed that this is the best suited optimization algorithm for these optimization problems, but MATLAB's default GA consistently finds usable solutions to the optimization problems presented in this section, within an acceptable time.

Since the optimization algorithm can consider any combination of the optimization parameters within the bounds shown in Table 4-2, the coverage path described by these parameters can be partly outside of the search area. When this happens, the value of the optimization function is not computed according to Equation (4-14), Equation (4-15), or Equation (4-16), but it is immediately awarded a fixed large penalty value instead, such that the amount of time spent on these invalid coverage paths is kept to a minimum. Since the planner only accepts solutions of the optimization algorithm for values of the optimization function well below this penalty value, coverage paths located partly outside the search area are never accepted.

Something similar happens when the constraint of the optimization problem is not satisfied. In that case, the computed value of the optimization function is increased with half of the value of the fixed penalty awarded when the coverage path is located partly outside of the search area. This way the optimization algorithm always favours valid coverage paths over invalid coverage paths, but since the threshold of accepting a solution also lies below half of the penalty value, a valid coverage path that violates the constraint is never accepted either.

Any coverage path that does not violate the constraint and lies completely within the search area is accepted as a solution to the optimization problem. Due to the randomness involved in the GA, and due to the presence of local minima, the GA does not always find the same solution to an optimization problem. For this reason, the planner lets the optimization algorithm find an acceptable solution a number of times, after which the coverage path corresponding to the lowest value of the optimization function is taken as the planned coverage path. In this final thesis project the best out of 3 acceptable solutions is taken, thus increasing the

probability of finding the global minimum.

In this chapter it has been discussed how the autonomy framework can plan a coverage path using the coverage map considering location uncertainty discussed in Chapter 3. It has also been discussed how this coverage path can autonomously adapt to changing underwater conditions. The three different models used to autonomously and adaptively plan the AUV's coverage path have not been encountered in literature before. In Chapter 5, the results of the different discussed path planning methods are evaluated in different scenarios.

Chapter 5

Evaluation

5-1 Simulation Quantity

In the previous chapters it has been discussed how an AUV can autonomously plan its coverage path, while adapting to changing underwater conditions and while taking the uncertainty on the location measurements into account. It has also been discussed that there are three different coverage performance models that can be used by the autonomous coverage path planner. In order to evaluate the performance of these three different path planners in different scenarios, various AUV operations have been simulated. The AUV operations can be simulated by connecting the autonomy framework with the simulation framework, as discussed in Chapter 2 and shown in Figure 2-4. This means that the AUV operations are simulated in MATLAB, since both the autonomy framework and the simulation framework have been implemented in MATLAB.

The simulation of a single AUV operation depends on many different pseudorandomly generated numbers. The numbers generated by a pseudorandom number generator (PRNG) appear to be realisations of a random variable, while they are actually generated through a deterministic algorithm. This makes pseudorandomly generated numbers very suitable for simulating an AUV operation, since they can model the randomness involved in an AUV operation, while the simulation results remain reproducible. Pseudorandomly generated numbers are for example used inside the simulation framework to determine the starting location of the AUV with respect to the search area, to model the external disturbances acting on the AUV and to add noise and biases to the INS location measurements, while inside the autonomy framework they are extensively used by the GA.

Since there are so many pseudorandomly generated numbers involved in the simulation of an AUV operation, in general not much can be concluded from the outcome of a single simulation. When one coverage path planner for example achieves a higher average coverage of the search area than another coverage path planner, this could be due to the fact that the first coverage path planner simply got lucky. When the pseudorandomly generated numbers in a

simulation lead to very small external disturbances and INS noise and biases for example, this will greatly benefit the average coverage of the search area. When the same coverage path planners are compared again while using a different seed for the PRNG, it could be possible that the second coverage path planner outperforms the first coverage path planner instead.

Every simulation that uses a different seed for the PRNG can be seen as a realisation of the random variable describing the average coverage of the search area. From the `Central Limit Theorem` [23] it follows that the mean of n realisations of a random variable is equal to the mean of this random variable, when $n \to \infty$. This means that when an infinite amount of simulations would be run, each with a different seed for the PRNG, the mean average coverage of the search area would be known. In fact, the mean of every variable of interest would be known after an infinite amount of simulations. Unfortunately an infinite number of simulations can clearly not be run, in fact running 500 simulations already took over 300 hours. Since this amount of time to evaluate a single AUV operation scenario was not available in this final thesis project, every scenario has been evaluated with 50 simulations instead.

In order to get an idea of how much the mean of a variable after 50 simulations differs from the mean of this variable after 500 simulations, a single scenario has been simulated 550 times, using a different seed for the PRNG for every simulation. For each of the variables used in the evaluation of this final thesis project, it has been checked how much the mean of these variables after 50 simulations differs from the mean of these variables after 500 simulations. The result are shown in Table 5-1, and the evaluated variables are discussed into more detail later in this chapter. The relative differences shown in Table 5-1 have been computed according to Equation (5-1).

$$\text{Relative difference} \;=\; \left(\frac{\text{mean(50 simulations)}}{\text{mean(500 simulations)}} - 1\right) \cdot 100\% \tag{5-1}$$

The absolute differences that are shown in Table 5-1 have been computed according to Equation (5-2).

$$\text{Absolute difference} \;=\; \text{mean(50 simulations)} - \text{mean(500 simulations)} \tag{5-2}$$

Table 5-1 shows that the largest relative difference between two means is 1.0%, and that the corresponding absolute difference is 3.5 seconds. All other absolute differences of the evaluated variables are so small that the mean values obtained after 50 simulations can be considered accurate representations of the mean values after 500 simulations, for this particular simulated AUV operation. Although the absolute difference of 3.5 seconds in the mean total time spent in the GA is of a different order of magnitude than the other absolute differences, it is not a very significant difference given that the mean total time spent in the GA is roughly 5.5 minutes.

Table 5-1: Differences between 50 and 500 simulations regarding the means of variables.

Variable	Relative	Absolute
Average E[$\boldsymbol{C}$] in the search area (uncertain trajectory)	$+0.0\%$	$+2.0 \cdot 10^{-4}$
Average E[$\boldsymbol{C}$] in the search area (real trajectory)	-0.1%	$-8.1 \cdot 10^{-4}$
Average H(E[$\boldsymbol{C}$]) in the search area (uncertain trajectory)	-0.2%	$-2.4 \cdot 10^{-4}$
Average H(E[$\boldsymbol{C}$]) in the search area (real trajectory)	$+0.8\%$	$+1.3 \cdot 10^{-3}$
Number of straight tracks in the coverage path	-0.4%	$-2.2 \cdot 10^{-2}$
Total time spent on the GA [s]	-1.0%	-3.5

Note that although the results presented in Table 5-1 might seem to suggest that 50 simulations are sufficient to obtain an accurate estimate of the true means of the evaluated variables, there is not much that can be concluded from these results. Firstly, the 500 simulations have been used as a substitute for running an infinite amount of simulations, and they do thus not represent the true mean of the evaluated variables. Secondly, the results presented in Table 5-1 have been obtained for one particular AUV operation scenario, and might thus not be representative for other simulated AUV operation scenarios. However, it is time-wise simply not possible to run 500 or more simulations for every simulated scenario, and the results shown in Table 5-1 at least prove that for one of the AUV operations of interest, the means of the evaluated variables after 50 simulations are an accurate approximation of the means of the evaluated variables after 500 simulations. Therefore, every AUV operation that is evaluated in this final thesis project is simulated 50 times, and the PRNG uses a different seed for each of these simulations.

5-2 Simulation Results

In Chapter 4 it has been discussed that the autonomous coverage path planner developed in this final thesis project can make use of three different coverage performance models. In this section, these three coverage performance models are first compared, and after that one of the coverage performance models is evaluated in different dynamic scenarios.

Comparison Of Coverage Performance Models

In order to evaluate the three coverage performance models, the same AUV operation scenario is simulated for each of these coverage performance models. This means that the only difference between the simulations is the coverage performance model used by the autonomous coverage path planner, while all other scenario parameters are identical across the simulations. The most important scenario parameters are the dimensions of the search area, which measures 300×500 metres, the strength of the currents, which are very weak, and the required average coverage of the search area. For the first two coverage performance models this required average coverage is represented by c_1 in the optimization problems shown in Equation (4-14) and Equation (4-15), and this required average expected value of the POD in the search area is set to $c_1 = 0.9$. The required average coverage of the search area for the third coverage performance model is represented by c_2 in the optimization problem shown in

Equation (4-16), and Equation (4-7) is used to compute that the required average entropy in the search area is set to $c_2 = \mathrm{H}(0.9) \approx 0.286$. As discussed in the previous section, 50 simulations are run for each of the different coverage performance models, using a different seed for the PRNG for every simulation. However, the simulations of the three different coverage performance models use the same 50 seeds for the PRNG, such that the three different coverage performance models are evaluated under the same simulation conditions.

In Figure 5-1, the average expected values of the POD in the search area at the end of the operation are shown for the three different coverage performance models. The average expected values of the POD in the search area shown in green are computed using the coverage map that considers location uncertainty, which has been discussed in Chapter 3. This is the coverage map that is used by the AUV to plan its coverage path. As a comparison, the average expected values of the POD are shown in blue for the coverage map that is constructed using the AUV's real trajectory, which is not available to the autonomy framework. This coverage map is used to check how the coverage that the AUV believes that it has achieved compares to the coverage that the AUV has actually achieved. The height of the green and blue bars shows the mean of the 50 simulations, while the error bars in red show the range of these 50 simulations. The exact values of the means and ranges shown in Figure 5-1 are listed in Appendix B.

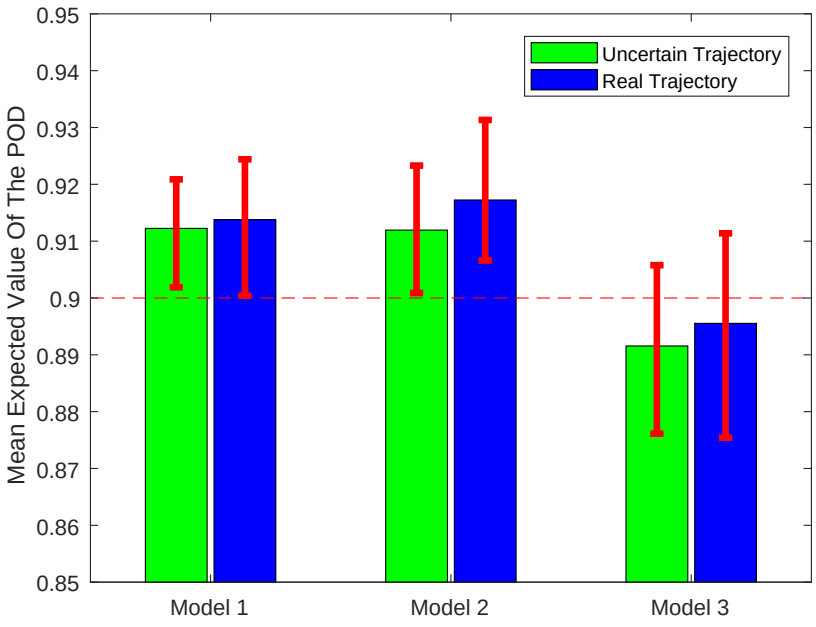

Figure 5-1: Mean POD expectations for the different models and trajectories ($c_2 = 0.286$).

Since the AUV plans its coverage path using the coverage map that considers the uncertainty on the location measurements, the ranges of the green bars should be entirely above the required average expected value of the POD of 0.9. Figure 5-1 shows that this is the case for the first two coverage performance models, meaning that the required average expected value of the POD was achieved in all 50 simulations for both of these coverage performance models.

Figure 5-1 also shows that the third coverage performance model does not achieve this average expected value of the POD in the search area. The reason for this is that the third coverage performance model required an average entropy in the search area of $\mathrm{H}(0.9) \approx 0.286$, instead of an average expected value of the POD of 0.9. Figure 5-2 shows that the third coverage performance model thus also achieved its required average coverage of the search area, since the range of its green bar is entirely below 0.286. The exact values of the means and ranges shown in Figure 5-2 are listed in Appendix B.

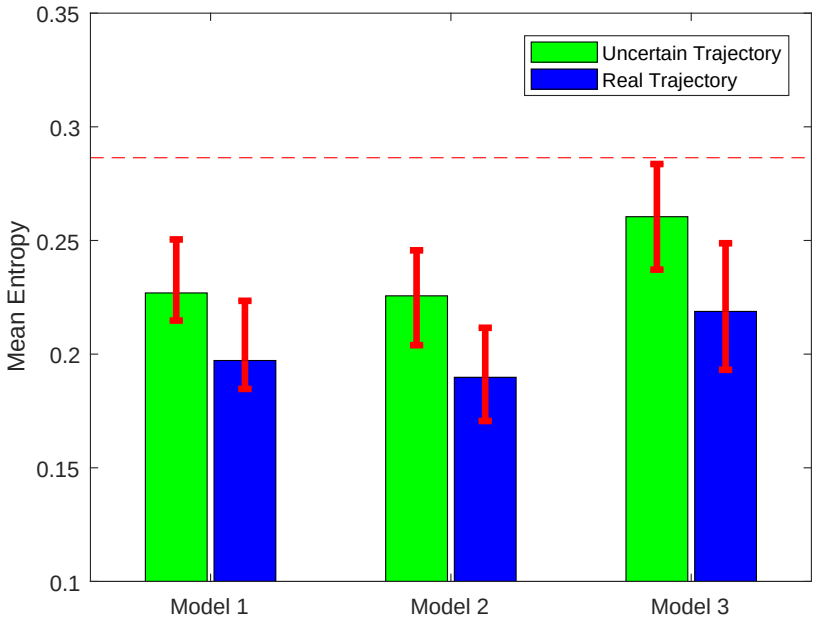

Figure 5-2: Mean entropies for the different models and trajectories ($c_2 = 0.286$).

Both Figure 5-1 and Figure 5-2 show that the first two coverage performance models outperform the third coverage performance model, regarding both the average expected value of the POD and the average entropy in the search area at the end of the operation. This can be understood by looking at the relation between the expected value of the POD and the entropy, as shown in Figure 4-4. The shape of the entropy as a function of the expected value of the POD shows that when the entropies of a set of grid cells are averaged, the corresponding expected value of the POD will always be equal to or higher than the average of the expected values of the POD of these grid cells. Equation (4-9) can be used to prove that this always holds. The consequence is that the third coverage performance model is too optimistic about the average expected value of the POD in the search area that it has achieved. In order to achieve a similar coverage performance of all three coverage performance models, such that a fairer comparison of these three coverage performance models can be made, the required average entropy in the search area for the third coverage performance model is decreased to $c_2 = \mathrm{H}(0.92) \approx 0.242$. In the remainder of the evaluation of the three different coverage performance models, the average coverage requirements are thus set to $c_1 = 0.9$ and $c_2 = 0.242$.

In Figure 5-3 the mean expected value of the POD is shown again for the different coverage performance models and trajectories, this time with the new coverage requirement for the third coverage performance model. The exact values of the means and ranges shown in Figure 5-3 are listed again in Appendix B. Figure 5-3 shows that the the means of the 50 simulations for each of the three coverage performance models are now similar, although the range of the simulations is larger for the third coverage performance model than for the first two coverage performance models.

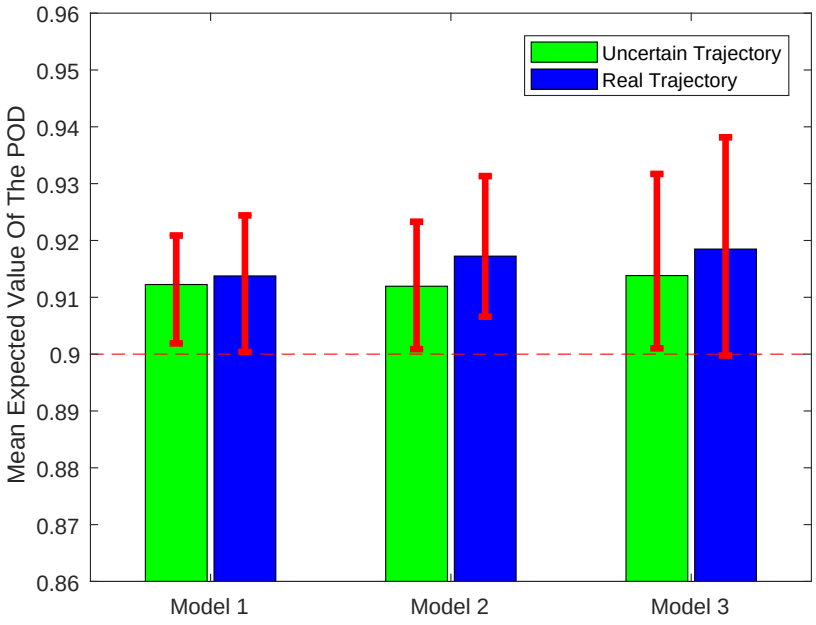

Figure 5-3: Mean POD expectations for the different models and trajectories ($c_2 = 0.242$).

The mean entropies in the search area are shown again for the different coverage performance models and trajectories in Figure 5-4, according to the new coverage requirement for the third coverage performance model. The exact values of the means and ranges shown in Figure 5-4 are listed in Appendix B again. Figure 5-4 shows that all 50 simulations of the third coverage performance model achieve the required mean entropy in the search area at the end of the operation. Figure 5-4 also shows that the first two coverage performance models do not achieve this desired average entropy in the search area for all simulations, which can be explained by the fact that the first two coverage performance models have to achieve a certain average expected value of the POD in the search area instead of an average entropy, which Figure 5-3 shows they do achieve. Figure 5-4 shows that the mean entropies of the 50 simulations are now similar as well for each of the three coverage performance models, due to the new average coverage requirement for the third coverage performance model.

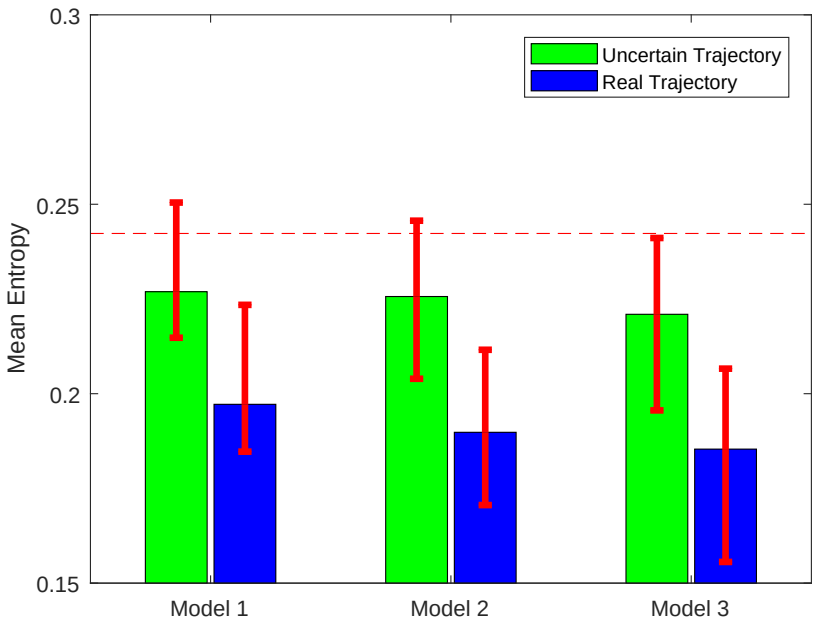

Figure 5-4: Mean entropies for the different models and trajectories ($c_2 = 0.242$).

Figure 5-1, Figure 5-2, Figure 5-3 and Figure 5-4 show that the mean coverage according to the coverage map constructed using the AUV's real trajectory always outperforms the mean coverage according to the coverage map constructed using the AUV's uncertain trajectory, for all of the coverage performance models. This could be caused by the AUV's INS, since when it estimates the AUV's trajectory to be less smooth than the trajectory that the AUV has actually been sailing, the coverage corresponding to the estimated trajectory will be lower than the coverage corresponding to the real trajectory. The result is that the AUV is actually performing better than it believes it is doing, which is desirable, since it has been discussed in Chapter 3 that underestimation is preferred over overestimation.

In Figure 5-3 and Figure 5-4, some differences between the three coverage performance models can be seen, but these differences are relatively small. A reason for this is the weight factor w, shown in the general optimization problem in Equation (4-13). This weight factor is chosen such that the cost of a straight track is always larger than the maximum entropy reduction in the search area. This means that all three coverage performance models plan a coverage path with the minimal number of straight tracks that can still satisfy the constraint, where the straight tracks are placed such that the entropy reduction is maximised. Since all three models usually place the straight tracks at similar locations, and they also require roughly the same amount of straight tracks, the resulting coverage maps are consequently similar as well.

Since the mean expected value of the POD and the mean entropy of the search area at the end of the operation can be seen to be similar for all three coverage performance models in Figure 5-3 and Figure 5-4, it is interesting to compare the amount of time it took the

coverage performance models to achieve this coverage. It is both checked how long it took the AUV to complete the operation, as well as how much time the planner has spent on planning the coverage path.

The time it takes the AUV to complete the evaluated scenario mainly depends on the number of straight tracks in the coverage path. Table 5-2 shows how many straight tracks the coverage path contained in each of the 50 simulations, for all three coverage performance models.

Table 5-2: Frequency table of the number of straight tracks in the coverage path.

Tracks	Model 1	Model 2	Model 3
8	8	6	0
9	42	35	33
10	0	8	16
11	0	1	1

From Table 5-2 it can be computed that the the average number of straight tracks in the coverage path is 8.84 for the first coverage performance model, 9.08 for the second coverage performance model and 9.36 for the third coverage performance model. Table 5-2 also shows that coverage performance model 1 never required more than 9 straight tracks in its coverage path, while coverage performance model 3 never required less than 9 straight tracks in its coverage path. The fact that coverage performance model 1 on average requires the least amount of straight tracks can be explained by the fact that this coverage performance model uses the complete distributions of the POD in every grid cell, which the other two coverage performance models do not. Therefore the first coverage performance model can predict the coverage resulting from a straight track with the highest accuracy of the three coverage performance models, such that the least amount of extra straight tracks is necessary to compensate for the mismatch between the predicted and realised coverage after a straight track. Table 5-2 suggests that the mismatch between the predicted and the realised coverage is slightly larger for the third coverage performance model than for the second coverage performance model. Note that although the absolute differences between the average required number of straight tracks might seems small, it could be that these absolute differences increase when the dimensions of the search area increase.

For every simulation, it has been recorded how much time was spent inside the GA. Since the AUV plans a new coverage path after every straight track that it has completed, the average time spent on planning a coverage path can be computed by dividing the total time spent in the GA by the number of straight tracks that have been planned. Note that the time spent to plan a coverage path at the start of the operation will be larger than the time spent to plan a coverage path near the end of the operation, but since the three coverage performance models required roughly the same amount of straight tracks, the average times spent on planning a coverage path are still comparable for the three coverage performance models. This comparison is shown in Table 5-3, which lists the minimum, mean and maximum average time spent in the GA out of 50 simulations, for all three coverage performance models.

Table 5-3: Average time [s] spent on planning a coverage path by the GA.

	Model 1	Model 2	Model 3
Min	87.3	9.6	8.9
Mean	115.1	14.0	14.5
Max	159.6	21.2	23.4

Note that unlike all other results presented in this chapter, the results in Table 5-3 are not exactly reproducible, since they depend on the system that the simulations are run on. A comparison between two different laptops however showed that the ratios of the mean average times spent in the GA for the three different coverage performance models remained roughly the same.

Table 5-3 shows that the average time spent on planning a coverage path is roughly the same for the second and third coverage performance models. This can be explained by the fact that the required computations for both methods are very similar. Table 5-3 also shows that the first coverage performance model spends significantly more time on planning a coverage path than the other two coverage performance models, which can be explained by the fact that this coverage performance model requires significantly more computations since it works with probability mass functions describing the POD in a grid cell, instead of working with scalar expected values of the POD or scalar entropies in a grid cell. From Table 5-2 and Table 5-3 it can be seen that the first coverage performance model requires on average the least amount of straight tracks in its coverage path, but it requires the most time to plan this coverage path.

When the ocean floor is located at a depth of 50 metres, as in the evaluated scenario, it takes the AUV roughly 30 seconds to obtain its GPS-fix, after which the location of the next straight track needs to be known in order to continue the operation. So when the planning of a coverage path takes more than 30 seconds, the total duration of the operation increases due to this planning. Since it takes the AUV about 160 seconds to travel a 320 metre long straight track, the lower average number of straight tracks required by the first coverage performance model does not outweigh the extra time it spends on planning this coverage path, at least in this particular scenario.

From the presented results it can be concluded that coverage performance model 1 and coverage performance model 2 achieve the required average expected value of the POD for every simulation that has been run, in the coverage map using the uncertain trajectory as well as in the coverage map using the real trajectory. It can also be concluded that coverage performance model 3 achieves the required average entropy for every simulation that has been run, also in both the the coverage map using the uncertain trajectory and in the coverage map using the real trajectory. Furthermore, the required number of straight tracks in the coverage path is roughly the same for the three coverage performance models, while coverage performance model 1 spends significantly more time on planning these coverage paths then the other two coverage performance models.

Dynamic Scenarios

In the previous subsection, the three coverage performance models have been compared. In the AUV operation scenario that they have been compared in, no large changes in the environment, the operation objective or the AUV itself occurred. It has been discussed that in this AUV operation scenario, all three coverage performance models achieved their required coverage performance. In order to evaluate if the autonomous coverage path planner developed in this final thesis project can autonomously adapt to changing conditions during an AUV operation as well, the autonomous coverage path planner is evaluated in three different dynamic scenarios. These scenarios are dynamic since significant changes in the environment, the operation objective or the AUV itself occur during the operation. The changes that occur during the three dynamic scenarios are listed below.

- The strength of the underwater currents increases during the operation.
- One of the sonar systems breaks down during the operation.
- The required average coverage of the search area changes during the operation.

These three dynamic scenarios are simulated 50 times each, while the planner uses coverage performance model 2. In the previous subsection it has been discussed that the average coverage performance of the three different coverage performance models is very similar, so it is expected that the other two coverage performance models will yield similar results. In the remainder of this chapter, the autonomous coverage path planner using the second coverage performance model is evaluated in the three dynamic scenarios.

Increasing Currents

In this dynamic scenario, the strength of the underwater currents is increased halfway during the operation. At the start of the operation, the strength of the underwater currents is equal to the strength of the underwater currents in the scenario simulated in the previous section. Halfway during the operation however, the strength of these underwater currents is increased by a factor 4. The required average expected value of the POD in the search area was set to $c_1 = 0.8$.

As the underwater currents increase, the trajectory of the AUV becomes less smooth, which results in a decreasing sonar performance. The real trajectory of the AUV during this simulation scenario is shown in Figure 5-5. In Figure 5-5 the AUV starts the operation at the marker at the beginning of the blue trajectory. At the marker where the blue trajectory changes into the red trajectory, the currents increased by a factor 4. The AUV ends the operation at the marker at the end of the red trajectory. Figure 5-5 shows that initially, the straight tracks are spaced widely apart, since the underwater currents are weak so the sonar performance is high. When the currents increase however, this sonar performance decreases, such that the straight tracks in the coverage path need to be spaced closer together. Figure 5-5 shows that the autonomous coverage path planner indeed decreases the spacing between the straight tracks, after the currents have been increased.

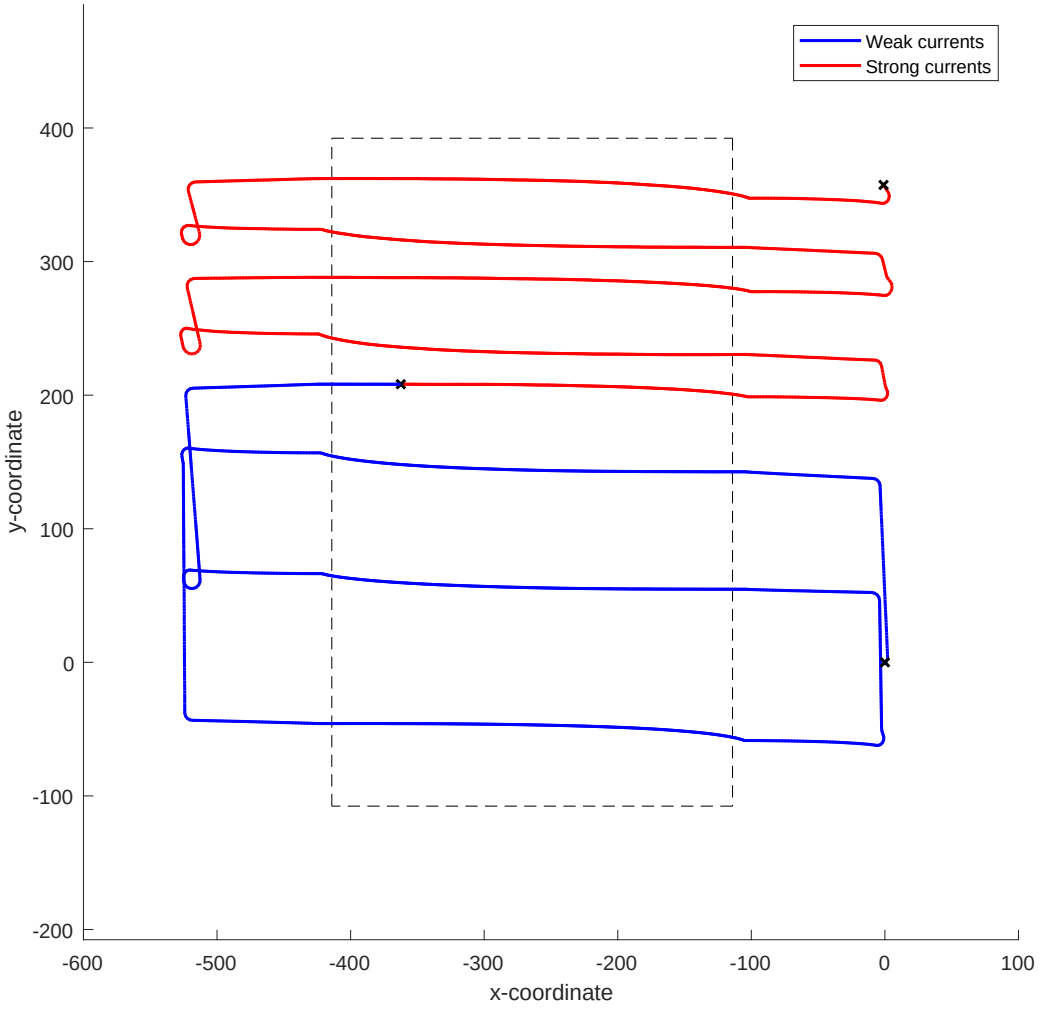

Figure 5-5: Real trajectory of the AUV when the currents increase.

This scenario has been simulated 50 times with 50 different seeds for the PRNG, and the AUV achieves the required average expected value of the POD in the search area in all but one of these simulations, both in the coverage map considering the AUV's uncertain trajectory as well as in the coverage map considering the AUV's real trajectory. The reason that in one simulation the required coverage performance is not achieved is not caused by the autonomous coverage path planner, but by the way that the stopping criterion of the simulations has been implemented. As soon as the planned coverage path only contains one straight track, the simulation terminates after the completion of this straight track. Due to the mismatch between the predicted coverage and the achieved coverage after a straight track however, it is possible that the AUV requires an additional straight track to achieve its required coverage performance in the search area, but it is not allowed to execute this straight track since the simulation is terminated. Unfortunately there was no time to run the

50 simulations again with the corrected stopping criterion. However, it can be assumed that when this stopping criterion of the simulations is corrected, the AUV achieves its coverage requirement in all simulations. Figure 5-5 shows that the autonomous coverage path planner adapts its coverage path to the increasing underwater currents very well.

Broken Sonar System

In this dynamic scenario, one of the sonar systems breaks down halfway during the operation. At the start of the operation, both the sonar system on the left side of the AUV and the sonar system on the right side of the AUV are working correctly. Halfway during the operation however, the sonar system on the left side of the AUV breaks down, such that only the sonar system on the right side of the AUV can be used to cover the search area. The required average expected value of the POD in the search area was set to $c_1 = 0.9$, and the underwater currents are weak.

The real trajectory of the AUV during this simulated scenario is shown in Figure 5-6. In Figure 5-6 the AUV starts the operation at the marker at the beginning of the blue trajectory, the sonar system on the left side of the AUV breaks at the marker where the blue trajectory changes into the red trajectory, and the AUV ends the operation at the marker at the end of the red trajectory. Figure 5-6 shows that initially, the AUV travels roughly equally spaced straight tracks to achieve its required average coverage of the search area. When the left sonar system breaks down however, the AUV travels two consecutive straight tracks in almost the same location, but in opposite direction, in order to compensate for the fact that it can only use its right sonar system.

For this scenario, 50 simulations with 50 different seeds for the PRNG have been run again, and the AUV achieved its required average expected value of the POD in the coverage map based on the real trajectory in all simulations, and in the coverage map based on the uncertain trajectory in all but one simulations. The fact that the AUV did not receive its coverage performance in one simulation was caused by the incorrect implementation of the stopping criterion again.

Figure 5-6 shows that even when one of the AUV's sonar systems breaks down, the autonomous coverage path planner developed in this final thesis project is still able to plan a coverage path that achieves the required average expected value of the POD in the search area.

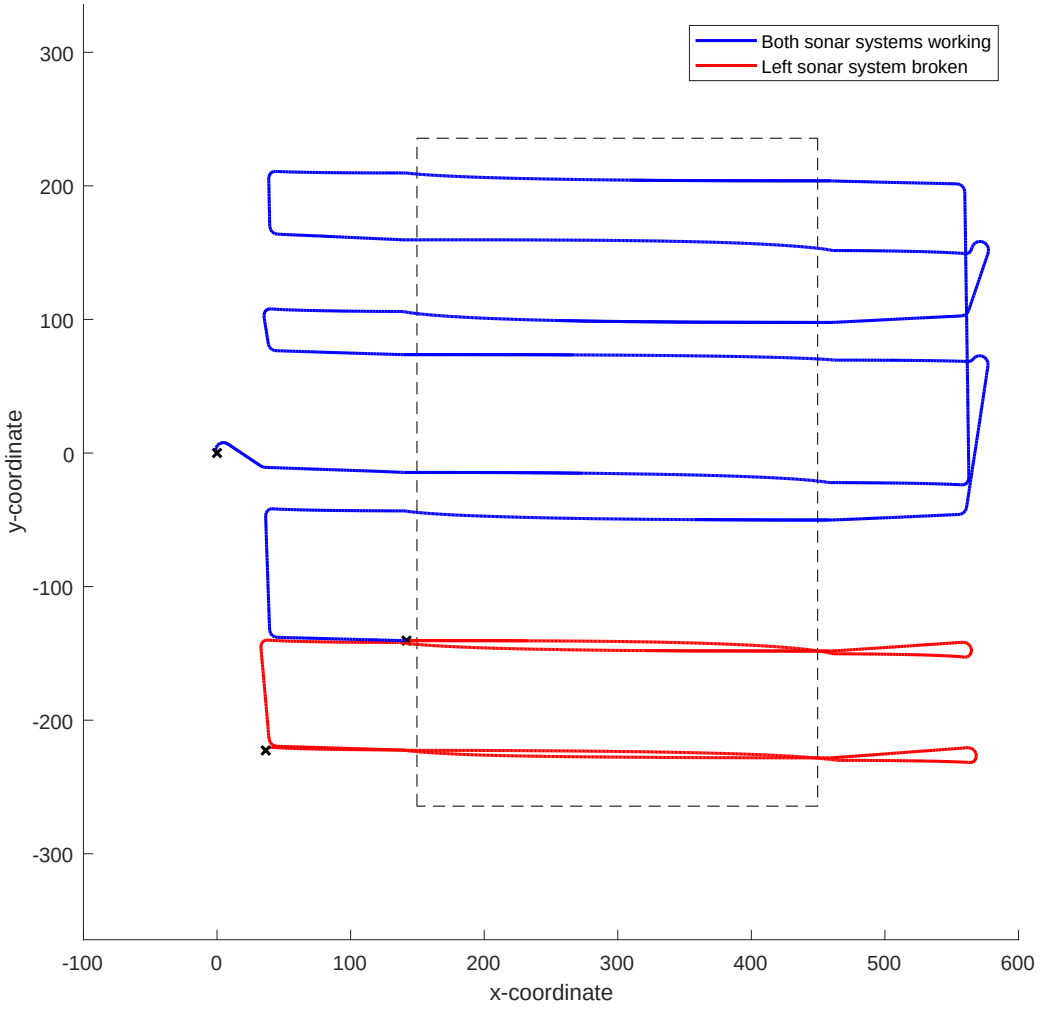

Figure 5-6: Real trajectory of the AUV when the left sonar system breaks down.

Changing Operation Objective

In this dynamic scenario, the required average coverage of the search area changes halfway during the operation. Although this is not a very likely scenario, the required average coverage of the search area could change when for example new information about the search area arrives during the operation. This scenario is however mainly evaluated to see if the autonomous coverage path planner is able to plan a coverage path that achieves the newly required average coverage of the search area, when it has already started covering the search area with a different required average coverage. Initially, the required average expected value of the POD in the search area is set to $c_1 = 0.7$, but halfway during the operation this required average expected value of the POD in the search area is increased to $c_1 = 0.9$. In this scenario, the underwater currents are weak.

Figure 5-7 shows the real trajectory of the AUV during this simulated scenario. Again, the AUV starts the operation at the marker at the start of the blue trajectory, the operation objective is increased at the marker where the blue trajectory changes into the red trajectory, and the operation ends at the marker at the end of the red trajectory. Figure 5-7 shows that the AUV starts covering the search area with widely spaced straight tracks, since the required coverage performance is initially relatively low. When the operation objective is increased however, the AUV has to return to the areas that it has covered with the lower average coverage requirement already, in order to increase the average coverage in these areas.

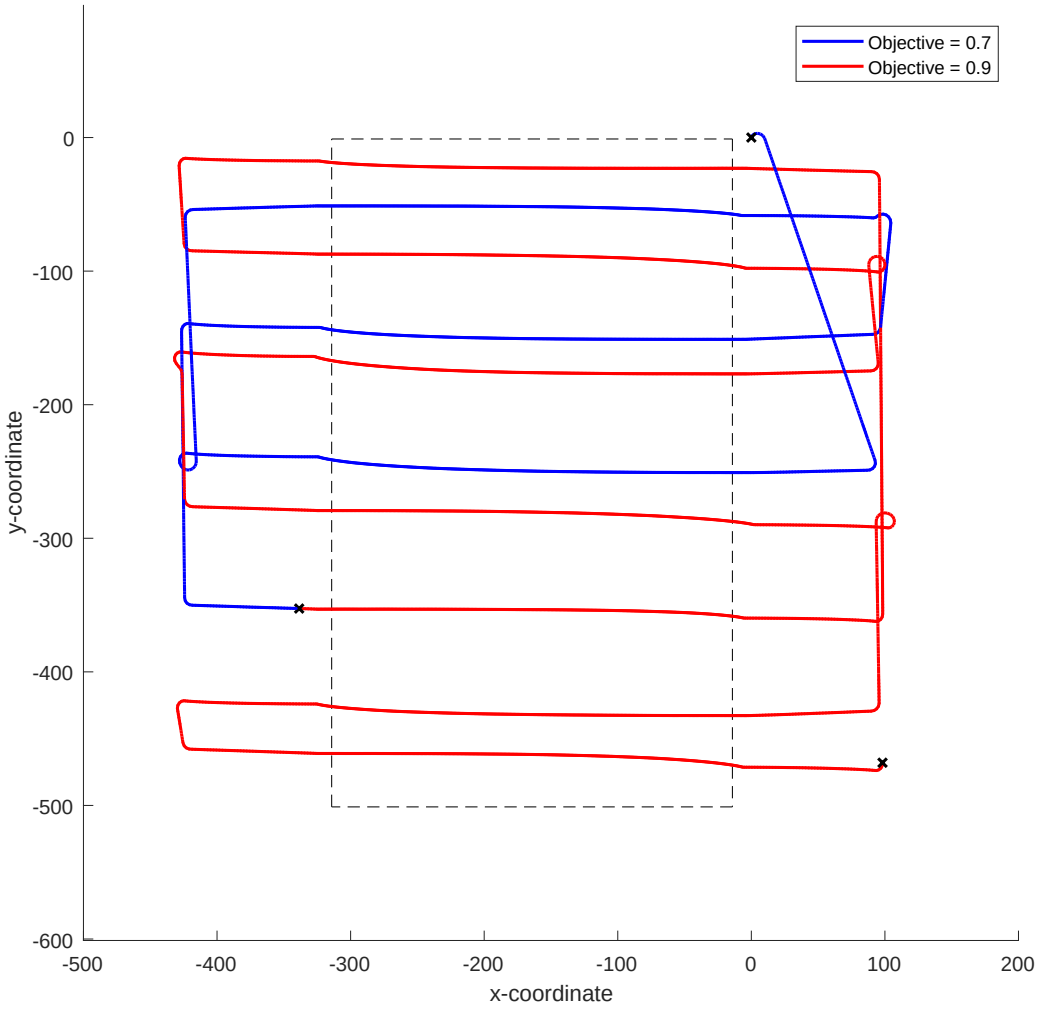

Figure 5-7: Real trajectory of the AUV when the operation objective changes.

This scenario has been simulated 50 times as well, with 50 different seeds for the PRNG, and the new required average expected value of the POD in the search area in the coverage map based on the real trajectory of the AUV is achieved in all simulations. The new required

average expected value of the POD in the search area in the coverage map based on the AUV's uncertain trajectory is achieved in all but one simulations, again due to the incorrect implementation of the stopping criterion.

Figure 5-7 shows that when the required coverage performance is increased during an operation, the autonomous coverage path planner developed in this final thesis project is able to adapt to this changed operation objective, by planning a coverage path that achieves the new required coverage performance.

Chapter 6

Conclusions

6-1 Problem Statement

The problem statement of this final thesis project has been to design and implement an autonomous coverage path planner for AUVs that considers location uncertainty. In order to do so, a constantly updating coverage map has been developed, which keeps track of the current coverage in the search area while taking the uncertainty on the AUV's location measurements into account. This coverage map has been used to develop an autonomous coverage path planner, which adaptively plans a coverage path that results in the required coverage of the search area. Finally an autonomy framework has been developed, in which the autonomous coverage path planner that considers location uncertainty has been implemented, such that its performance could be evaluated in the simulation framework provided by TNO.

In order to keep track of the coverage in the search area, the search area has been divided into grid cells. The coverage of these grid cells has been defined as the probability that a true observation in these grid cells has been obtained. A simplified model to relate a sonar image to a probability of a true observation has been developed. Since it is uncertain to which grid cells an obtained sonar image exactly belongs, due to the uncertainty on the AUV's location measurements, a sonar image obtained in the sensor coordinate frame is mapped to an uncertain location in the global coordinate frame. The result is that the coverage of a grid cell is described by a random variable, instead of by a deterministic variable. By including the location uncertainty in the coverage map, the fact that the AUV's believed path differs from the AUV's actual path is taken into account. The idea for this coverage map has been encountered in literature before, but the implementation in this final thesis project is novel in the sense that it is computationally significantly faster, mathematical errors have been corrected, and the AUV's location uncertainty is considered in two dimensions instead of in one.

The coverage map that considers the location uncertainty is used by the autonomous coverage path planner, which plans the AUV's coverage path by turning it into an optimization problem. In order to do so, three different coverage performance models have been developed that each estimate the most recent coverage performance of the AUV, and use this to predict the coverage of the search area after the AUV has completed the coverage path. Since these coverage performance models use the latest estimated coverage performance to predict the future coverage performance, the autonomous coverage path planner is able to adapt to changing conditions during the operation. The optimization problem has been defined such that the AUV achieves a desired average coverage of the search within the shortest possible duration to complete the operation, while the entropy reduction in the search area is maximised. The optimization problem is solved with a GA. The autonomous coverage path planner developed in this final thesis project has not been encountered in literature before.

In MATLAB an autonomy framework has been developed, in which the autonomous coverage path planner that considers the location uncertainty has successfully been implemented. This autonomy framework stores and processes the relevant information obtained from the AUV's sensors, and uses this information to dynamically plan and replan the AUV's coverage path. The autonomy framework has been implemented such that it can be integrated in the simulation framework provided by TNO.

TNO's simulation framework provides an environment to simulate AUV operations. Since the results of a simulation depend on the realizations of many pseudorandom variables, it has been checked how many simulations are required in order to obtain results that can be considered accurate. For every AUV operation scenario that has been evaluated, 50 simulations have been run. The 50 simulations for each of the three different coverage performance models in the same AUV operation scenario have been compared, and all coverage performance models achieve the required average coverage of the search area in all simulations.

Besides comparing the three coverage performance models, the performance of the autonomous coverage path planner has also been evaluated in different dynamic scenarios. The simulations show that the AUV still achieves the required average coverage of the search area when the strength of the underwater currents increases during the operation. The increasing underwater currents decrease the coverage performance of the AUV, and when this happens the autonomous coverage path planner successfully decreases the distance between the straight tracks of the coverage path. The simulations also show that the AUV can successfully adapt to a change in the required average coverage of the search area during an operation, by returning to parts of the search area that had been covered according to the lower coverage requirement before, to then cover these areas according to the new coverage requirement. Finally, the simulations show that even when one of the two AUV's sonar systems breaks down, the autonomous coverage path planner is still able to plan a coverage path that achieves the required average coverage of the search area. Thus, it can be concluded that an autonomous coverage path planner for AUVs that considers location uncertainty has successfully been developed and implemented in this final thesis project.

6-2 Future Work

Although an autonomous coverage path planner for AUVs that considers location uncertainty has successfully been developed, implemented and evaluated, this coverage path planner is not perfect. This section discusses how the research presented in this final thesis project could be continued. The suggested future work is summarized below.

- Tune or change the optimization algorithm.
- Change the optimization parameters.
- Evaluate using more simulations.
- Include additional ocean floor types.

One component of the coverage path planner that can most certainly be improved, is the optimization algorithm that is used. In this final thesis project MATLAB's default GA has been used to solve the optimization problems, since this algorithm was readily available and it consistently solved the optimization problems within an acceptable time. However, it has not been investigated if the default optimization options of MATLAB's GA are the most suited options for the optimization problems presented in this final thesis project. Changing some of the optimization options might lead to faster solutions, or reduce the probability of finding a local minimum, or do both. Besides using a GA, other optimization algorithms should also be considered. Given the nature of the optimization problem, other evolutionary algorithms seem the most obvious choice.

Besides changing the optimization algorithm itself, alternative ways to describe a coverage path in the search area should be investigated, resulting in different optimization parameters. Currently the coverage path can only contain equidistant straight tracks which can not be placed under an angle, meaning that the performance of the coverage path planner can almost certainly be improved in scenarios that require a non-equidistant track spacing, or straight tracks under a different angle. Allowing the coverage path these extra degrees of freedom will however increase the number of parameters needed to describe the coverage path, consequently increasing the search space for the optimization algorithm. When the optimization parameters are changed, it should therefore also be checked if the optimization algorithm has to be tuned or changed again.

In this final thesis project, every scenario has been evaluated with the simulation results obtained from 50 simulations. For one particular scenario, it has been shown that the results obtained after 50 simulations are representative of the results obtained after 500 simulations. The first problem with this is that 500 is not an extremely large number of simulations to compare with, and the second problem is that the obtained result only holds for one particular scenario. Both problems can be solved by evaluating the simulation results of a larger number of simulations for every future scenario.

Finally, additional types of ocean floor could be added to the currently flat sandy bottom. When the ocean floor is expanded with areas containing sand dunes or seagrass, the coverage

path planner is presented with new problems. The coverage of areas with sand dunes for example depends on the angle under which the sonar images of these areas have been obtained, while the coverage of areas containing seagrass will always be low. Therefore, including a more complex ocean floor promises to be an interesting direction for future work on the research presented in this final thesis project.

Appendix A

A Comparison Of Coverage Maps

In this appendix, three coverage maps are compared for every possible POD. The highest probabilities of the POD are shown in yellow, the lowest probabilities of the POD in blue.

The same sonar performance model can be used with three different trajectories to compute the PODs in a search area. On the next page, the PODs based on the real trajectory, the measured trajectory and the uncertain trajectory of the AUV are shown. The real trajectory is only for comparison, since this trajectory is not available to the autonomy framework. The measured trajectory is the trajectory as measured by the AUV's INS, and this trajectory is available from the simulation framework. The uncertain trajectory takes the uncertainty on the measured trajectory into account, by processing the measured trajectory inside the autonomy framework.

On the subsequent pages, the probability of every discrete possible value of the POD separately is shown for all three trajectories.

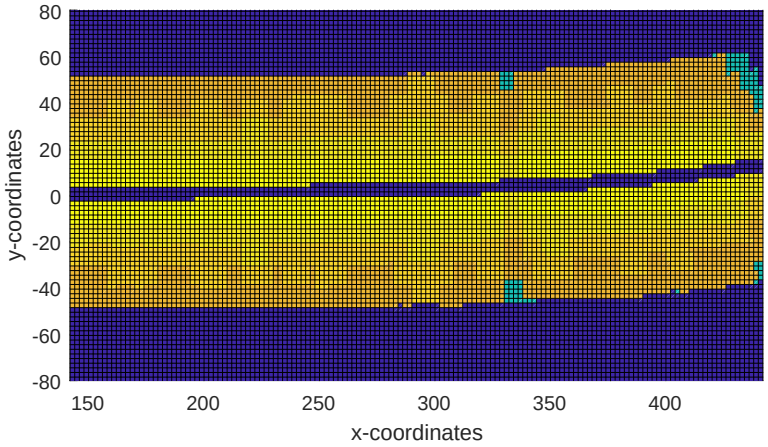

Figure A-1: The PODs based on the real trajectory.

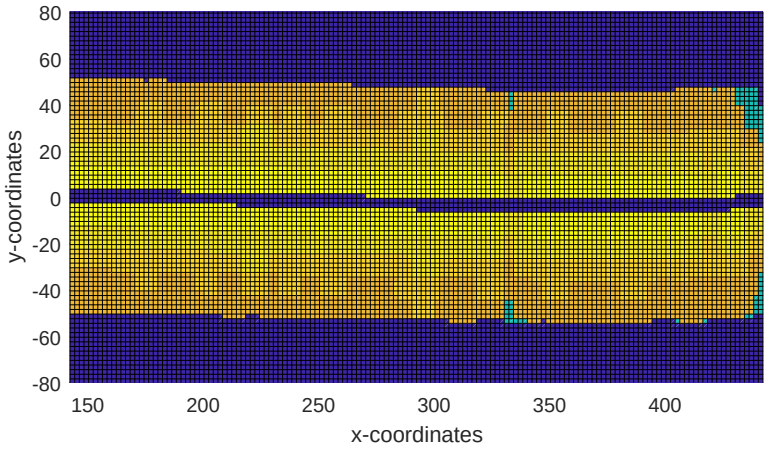

Figure A-2: The PODs based the measured trajectory.

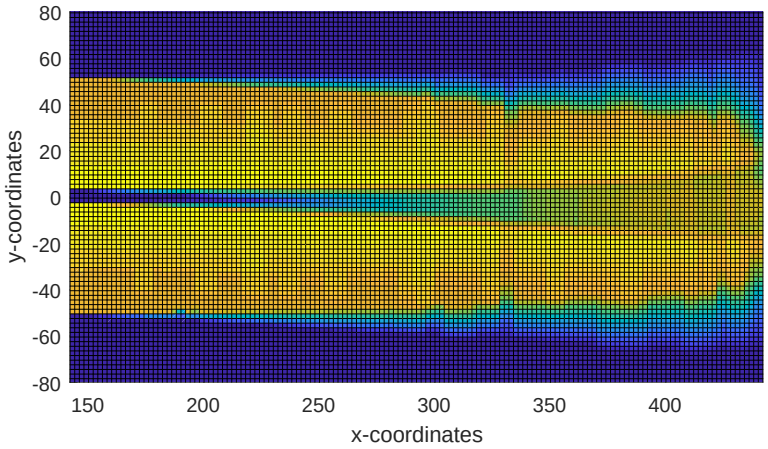

Figure A-3: The expected value of the PODs based on the uncertain trajectory.

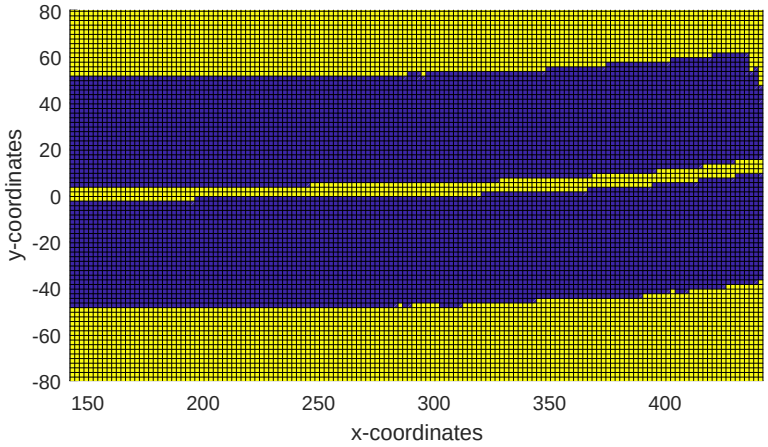

Figure A-4: Probability that the POD $= 0$ based on the real trajectory.

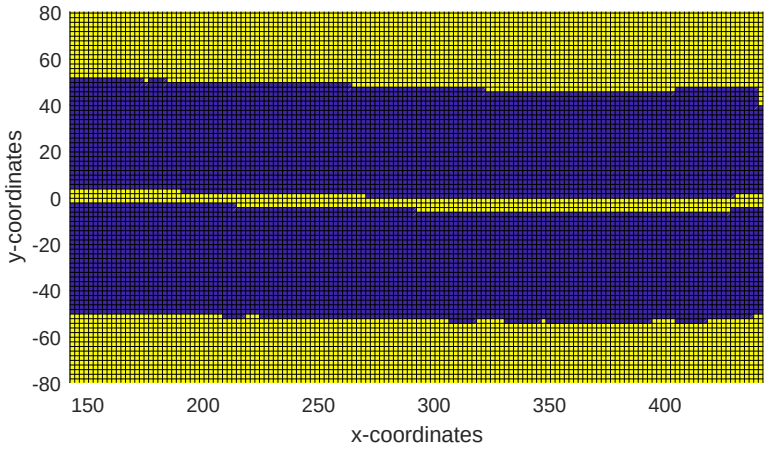

Figure A-5: Probability that the POD $= 0$ based the measured trajectory.

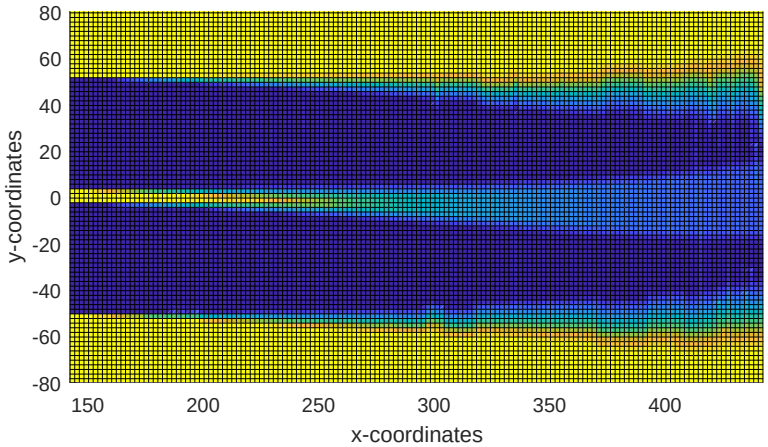

Figure A-6: Probability that the POD $= 0$ based on the uncertain trajectory.

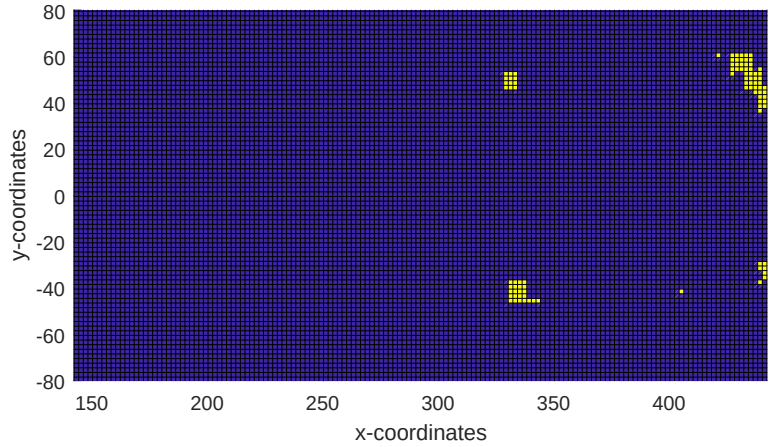

Figure A-7: Probability that the POD $= 0.5$ based on the real trajectory.

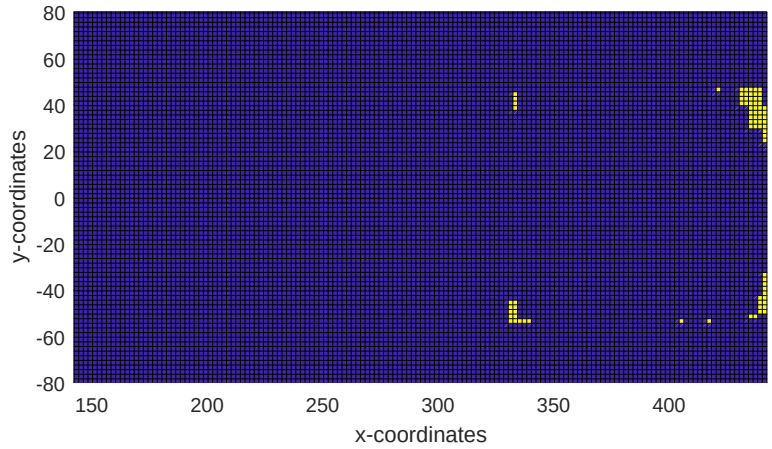

Figure A-8: Probability that the POD $= 0.5$ based the measured trajectory.

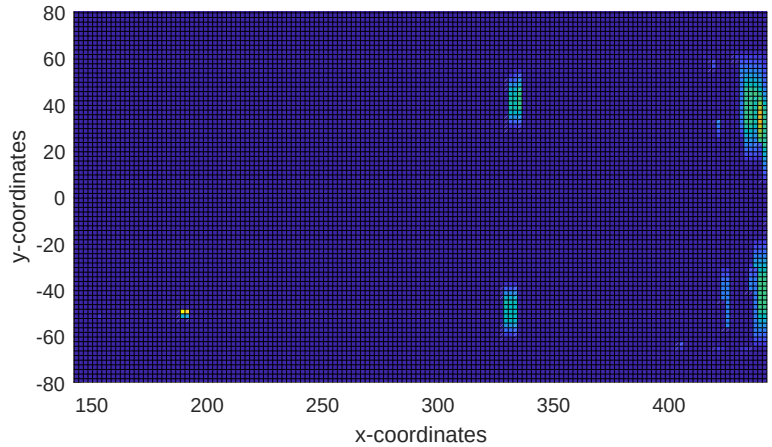

Figure A-9: Probability that the POD $= 0.5$ based on the uncertain trajectory.

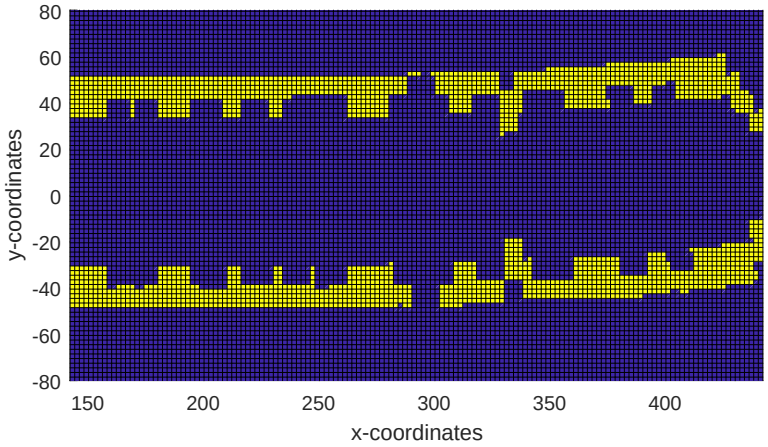

Figure A-10: Probability that the POD $= 0.8$ based on the real trajectory.

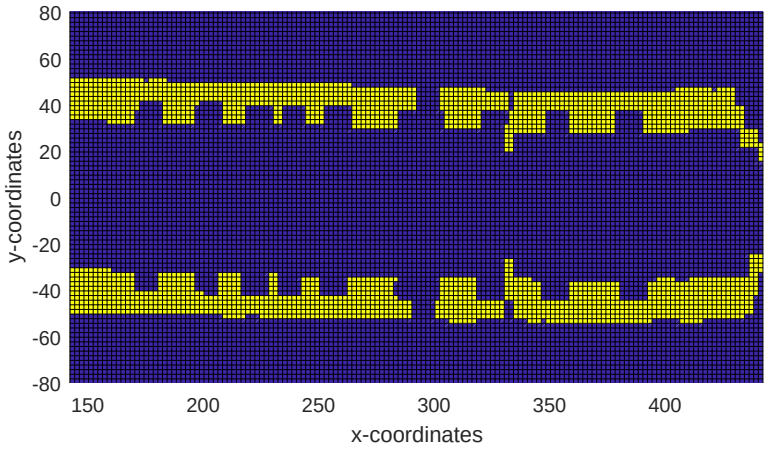

Figure A-11: Probability that the POD $= 0.8$ based the measured trajectory.

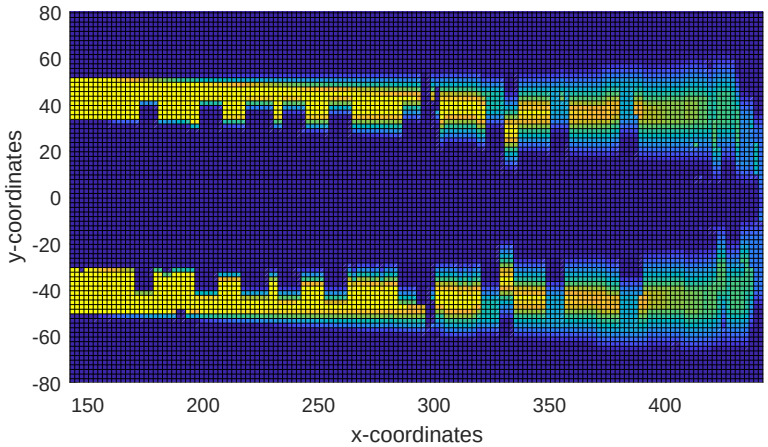

Figure A-12: Probability that the POD $= 0.8$ based on the uncertain trajectory.

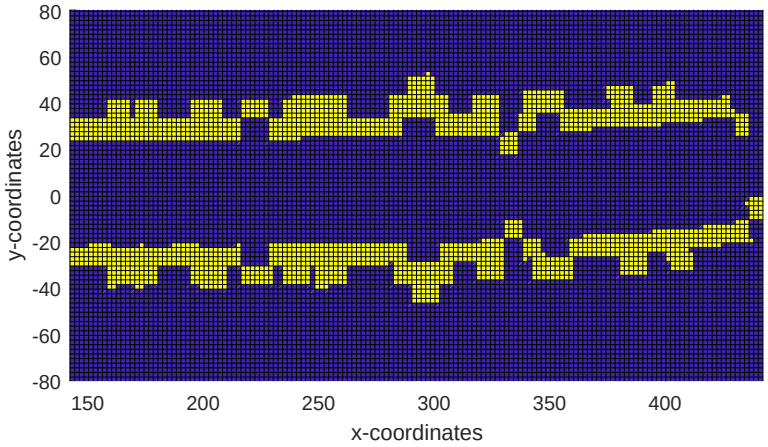

Figure A-13: Probability that the POD $= 0.9$ based on the real trajectory.

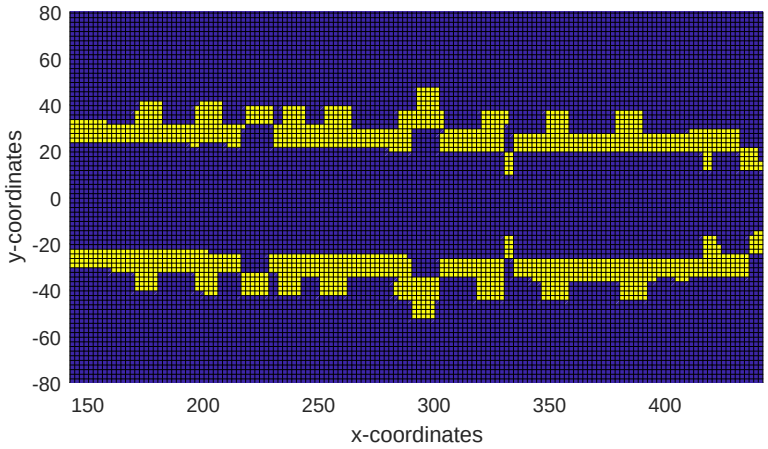

Figure A-14: Probability that the POD $= 0.9$ based the measured trajectory.

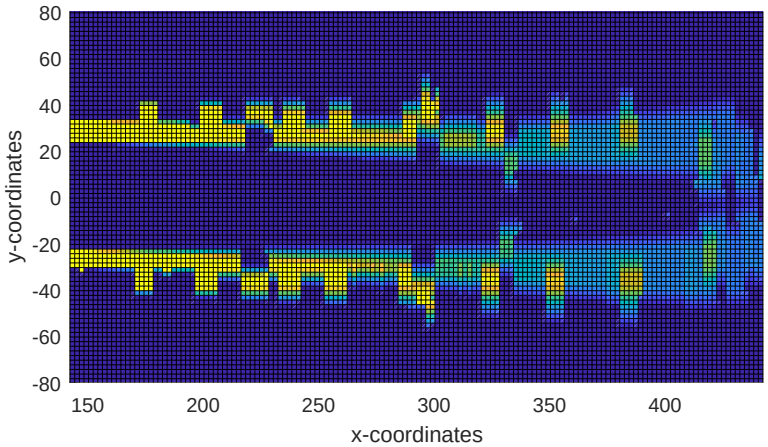

Figure A-15: Probability that the POD $= 0.9$ based on the uncertain trajectory.

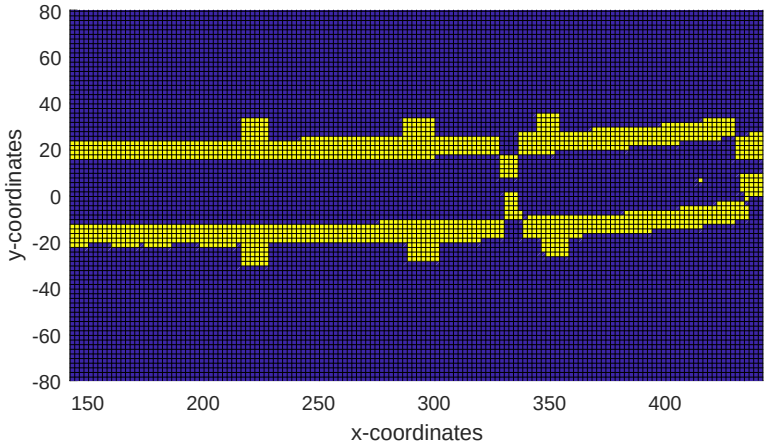

Figure A-16: Probability that the POD $= 0.95$ based on the real trajectory.

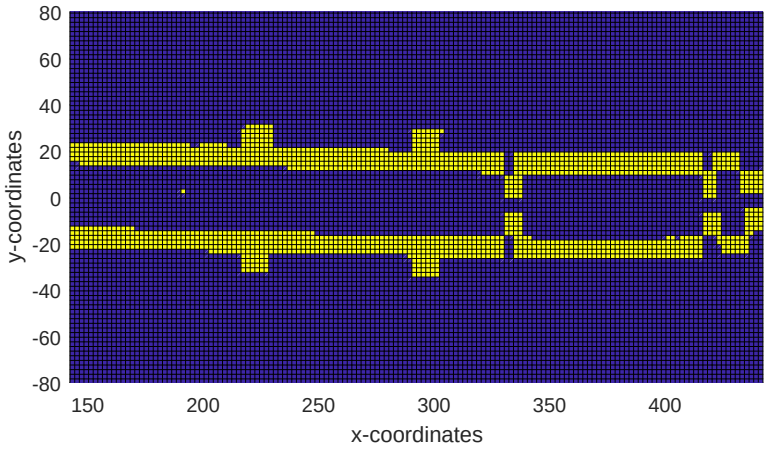

Figure A-17: Probability that the POD $= 0.95$ based the measured trajectory.

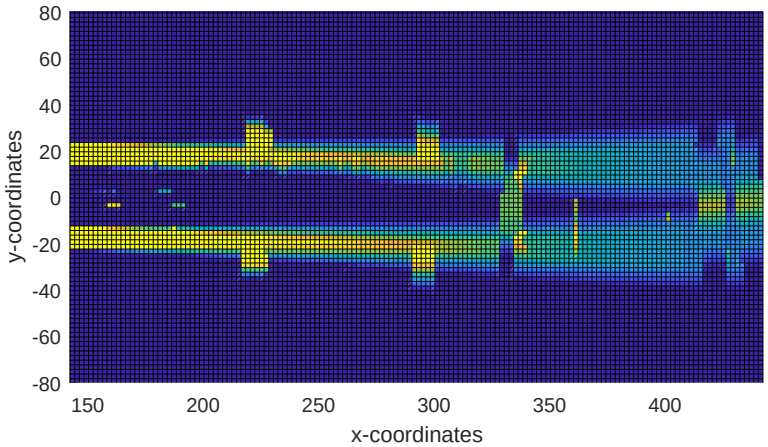

Figure A-18: Probability that the POD $= 0.95$ based on the uncertain trajectory.

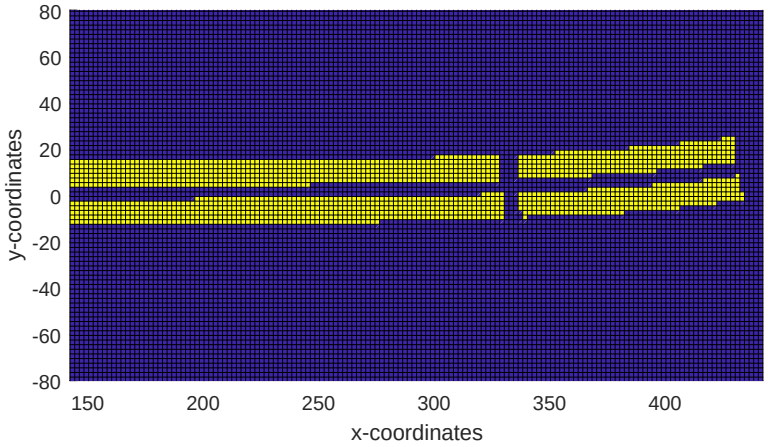

Figure A-19: Probability that the POD $= 1$ based on the real trajectory.

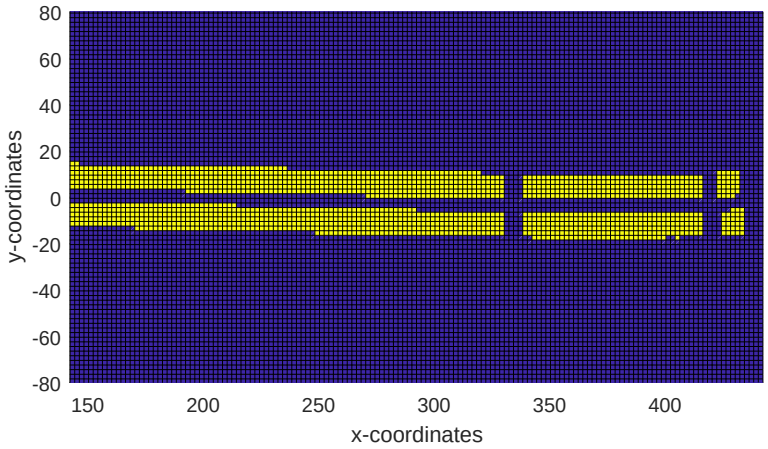

Figure A-20: Probability that the POD $= 1$ based the measured trajectory.

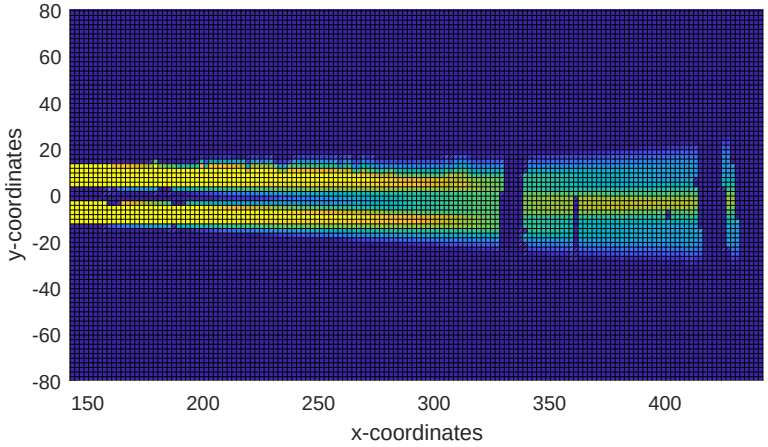

Figure A-21: Probability that the POD $= 1$ based on the uncertain trajectory.

Appendix B

Simulation Results Tables

This appendix contains four tables, which contain the exact values that have been used to create the figures used in the evaluation in Chapter 5. On the subsequent pages, the figures used in the evaluation of the coverage performance models are shown again, with a table underneath showing the exact values used to create these figures.

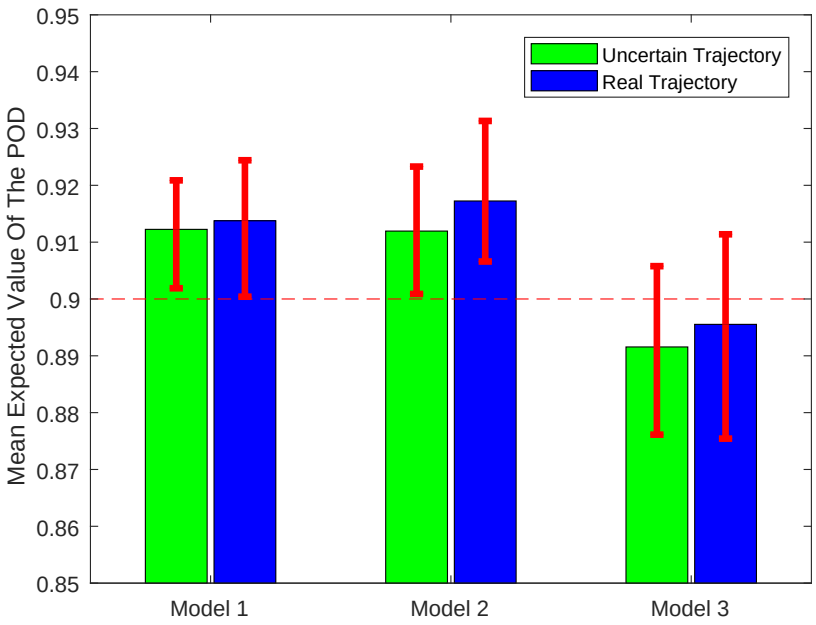

Figure B-1: Mean POD expectations for the different models and trajectories ($c_2 = 0.286$).

Figure B-1 is the same as Figure 5-1, shown in Chapter 5. The minimum, mean and maximum values used to create these figures are shown in Table B-1.

Table B-1: Mean POD expectations for the different models and trajectories ($c_2 = 0.286$).

	Trajectory	Min	Mean	Max
Model 1	Uncertain	0.902	0.912	0.921
	Real	0.900	0.914	0.924
Model 2	Uncertain	0.901	0.912	0.923
	Real	0.907	0.917	0.931
Model 3	Uncertain	0.876	0.892	0.906
	Real	0.875	0.896	0.911

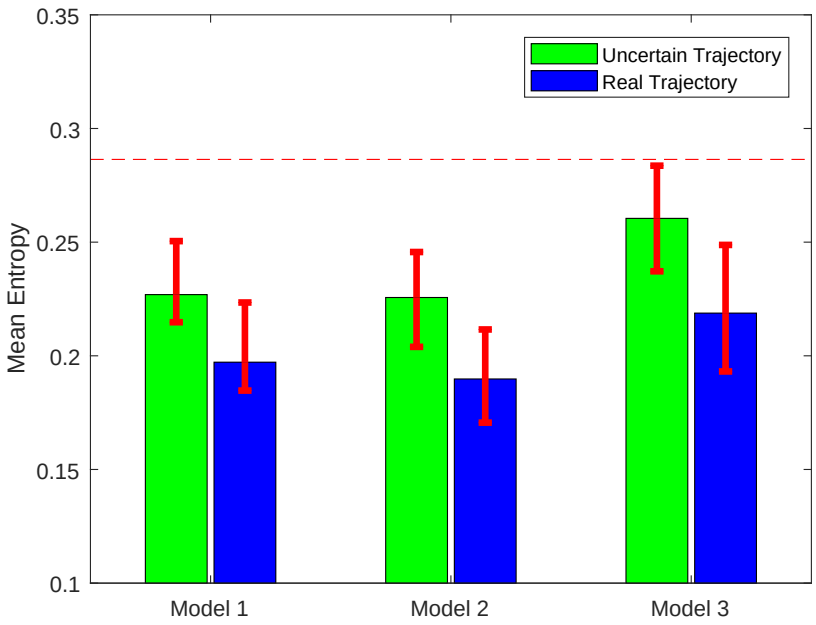

Figure B-2: Mean entropies for the different models and trajectories ($c_2 = 0.286$).

Figure B-2 is the same as Figure 5-2, shown in Chapter 5. The minimum, mean and maximum values used to create these figures are shown in Table B-2.

Table B-2: Mean entropies for the different models and trajectories ($c_2 = 0.286$).

	Trajectory	Min	Mean	Max
Model 1	Uncertain	0.215	0.227	0.251
	Real	0.185	0.197	0.223
Model 2	Uncertain	0.204	0.226	0.246
	Real	0.171	0.190	0.212
Model 3	Uncertain	0.237	0.261	0.284
	Real	0.193	0.219	0.249

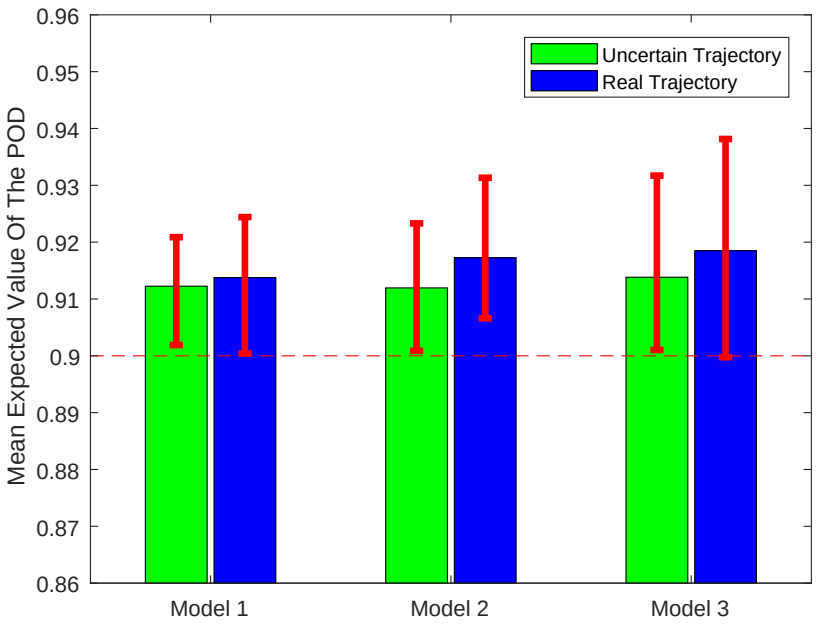

Figure B-3: Mean POD expectations for the different models and trajectories ($c_2 = 0.242$).

Figure B-3 is the same as Figure 5-3, shown in Chapter 5. The minimum, mean and maximum values used to create these figures are shown in Table B-3.

Table B-3: Mean POD expectations for the different models and trajectories ($c_2 = 0.242$).

	Trajectory	Min	Mean	Max
Model 1	Uncertain	0.902	0.912	0.921
	Real	0.900	0.914	0.924
Model 2	Uncertain	0.901	0.912	0.923
	Real	0.907	0.917	0.931
Model 3	Uncertain	0.901	0.914	0.932
	Real	0.900	0.919	0.938

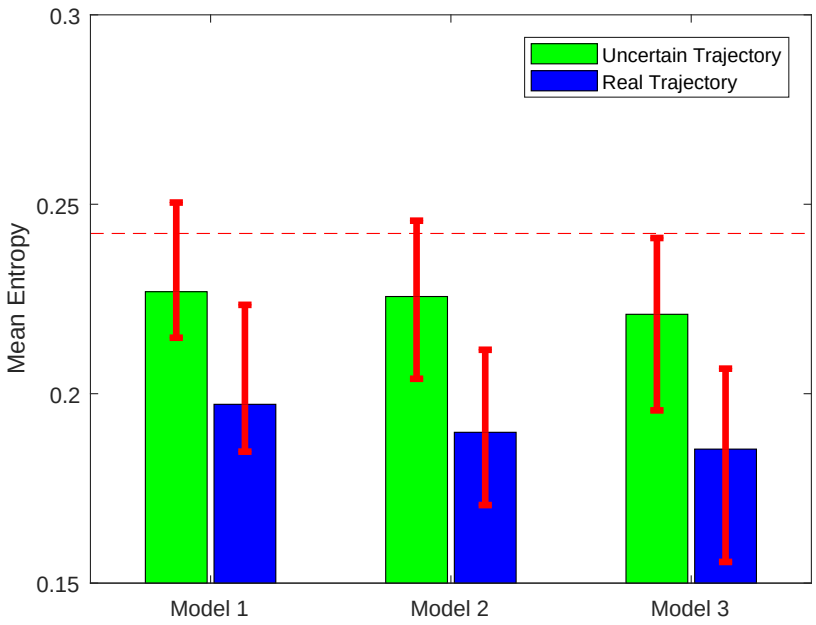

Figure B-4: Mean entropies for the different models and trajectories ($c_2 = 0.242$).

Figure B-4 is the same as Figure 5-4, shown in Chapter 5. The minimum, mean and maximum values used to create these figures are shown in Table B-4.

Table B-4: Mean entropies for the different models and trajectories ($c_2 = 0.242$).

	Trajectory	Min	Mean	Max
Model 1	Uncertain	0.215	0.227	0.251
	Real	0.185	0.197	0.223
Model 2	Uncertain	0.204	0.226	0.246
	Real	0.171	0.190	0.212
Model 3	Uncertain	0.196	0.221	0.241
	Real	0.156	0.185	0.207

Bibliography

[1] F. Gutierrez, M. Gutowski, S. Ganther, P. Hogarth, and C. Wallace, "Deep rippled bedforms in loch ness: Evidence from an auv bathymetry survey," in *Autonomous Underwater Vehicles (AUV), 2014 IEEE/OES*, pp. 1–3, IEEE, 2014.

[2] Y. Tipsuwan and P. Hoonsuwan, "Design and implementation of an auv for petroleum pipeline inspection," in *2015 7th International Conference on Information Technology and Electrical Engineering (ICITEE)*, pp. 382–387, IEEE, 2015.

[3] S. Giodini, A. Hunter, H. Naus, B. Bakker, D. Bekers, M. Ditzel, R. v. Vossen, S. Dugelay, F. Baralli, and A. Beckers, "Decision making on auvs for adaptive minehunting surveys," in *Proceedings of the European Conference on Undersea Defence Technology, UDT Europe 2014, 10-12 June, 2014, Liverpool, UK*, Nexus Media, Ltd, 2014.

[4] R. B. Wynn, V. A. Huvenne, T. P. Le Bas, B. J. Murton, D. P. Connelly, B. J. Bett, H. A. Ruhl, K. J. Morris, J. Peakall, D. R. Parsons, *et al.*, "Autonomous underwater vehicles (auvs): Their past, present and future contributions to the advancement of marine geoscience," *Marine Geology*, vol. 352, pp. 451–468, 2014.

[5] A. Sousa, L. Madureira, J. Coelho, J. Pinto, J. Pereira, J. B. Sousa, and P. Dias, "Lauv: The man-portable autonomous underwater vehicle," *IFAC Proceedings Volumes*, vol. 45, no. 5, pp. 268–274, 2012.

[6] D. P. Williams, "On optimal auv track-spacing for underwater mine detection," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 4755–4762, IEEE, 2010.

[7] P. E. Hagen, "Gap filler sensors for sas," in *OCEANS 2011*, pp. 1–6, IEEE, 2011.

[8] A. Pizzo, "Overview of navigation, obstacle avoidance and localisation of auvs," 2015.

[9] D. P. Williams, F. Baralli, M. Micheli, and S. Vasoli, "Adaptive underwater sonar surveys in the presence of strong currents," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pp. 2604–2611, IEEE, 2016.

[10] T. R. Krogstad and M. S. Wiig, "Autonomous survey and identification planning for auv mcm operations," *Proc. Undersea Defence Technology Europe, Liverpool, UK*, 2014.

[11] T. Huntsberger, H. Aghazarian, T. Estlin, and D. Gaines, "Control architecture for robotic agent command and sensing," 2008.

[12] J. Albus, H.-M. Huang, A. Lacaze, M. Schneier, M. Juberts, H. Scott, S. Balakirsky, P. W. Shackleford, T. Hong, J. Michaloski, *et al.*, "4d/rcs: A reference model architecture for unmanned vehicle systems version 2.0," 2002.

[13] E. Miguelanez, P. Patron, K. E. Brown, Y. R. Petillot, and D. M. Lane, "Semantic knowledge-based framework to improve the situation awareness of autonomous underwater vehicles," *IEEE Transactions on Knowledge and Data Engineering*, vol. 23, no. 5, pp. 759–773, 2011.

[14] F. Maurelli, M. Carreras, J. Salvi, D. Lane, K. Kyriakopoulos, G. Karras, M. Fox, D. Long, P. Kormushev, and D. Caldwell, "The pandora project: A success story in auv autonomy," in *OCEANS 2016-Shanghai*, pp. 1–8, IEEE, 2016.

[15] J. R. Boyd, "Organic design for command and control," *A discourse on winning and losing*, 1987.

[16] G. Davies and E. Signell, "Espresso scientific user guide," *NATO underwater research centre, Nurc-sp-2006-003*, 2006.

[17] L. Paull, M. Seto, and H. Li, "Area coverage planning that accounts for pose uncertainty with an auv seabed surveying application," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 6592–6599, IEEE, 2014.

[18] Z. Zeng, L. Lian, K. Sammut, F. He, Y. Tang, and A. Lammas, "A survey on path planning for persistent autonomy of autonomous underwater vehicles," *Ocean Engineering*, vol. 110, pp. 303–313, 2015.

[19] N. Abreu, N. Cruz, and A. Matos, "Accounting for uncertainty in search operations using auvs," in *Underwater Technology (UT), 2017 IEEE*, pp. 1–8, IEEE, 2017.

[20] L. Paull, S. Saeedi, M. Seto, and H. Li, "Sensor-driven online coverage planning for autonomous underwater vehicles," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 6, pp. 1827–1838, 2013.

[21] M. Cashmore, M. Fox, T. Larkworthy, D. Long, and D. Magazzeni, "Auv mission control via temporal planning," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 6535–6541, IEEE, 2014.

[22] D. M. Lane, F. Maurelli, P. Kormushev, M. Carreras, M. Fox, and K. Kyriakopoulos, "Pandora-persistent autonomy through learning, adaptation, observation and replanningâŸĚ," *IFAC-PapersOnLine*, vol. 48, no. 2, pp. 238–243, 2015.

[23] F. M. Dekking, *A Modern Introduction to Probability and Statistics: Understanding why and how*. Springer Science & Business Media, 2005.

[24] B. Heck and E. Kamen, *Fundamentals of signals and systems using the web and matlab*. Prentice-Hall, third edition edition, 2007.

[25] D. C. Lay, *Linear algebra and its applications*. Addison Wesley, Boston, 2006.

[26] A. Papoulis and S. U. Pillai, *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education, 2002.

[27] C. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.

[28] J. Stewart, *Calculus*. Cengage Learning, 2008.

[29] J. L. W. V. Jensen, "Sur les fonctions convexes et les inégalités entre les valeurs moyennes," *Acta mathematica*, vol. 30, no. 1, pp. 175–193, 1906.

Glossary

List of Acronyms

4D/RCS	Four-dimensional Real-time Control System
ANN	Artificial Neural Network
ATR	Automatic Target Recognition
AUV	Autonomous Underwater Vehicle
CARACaS	Control Architecture For Robotic Agent Command And Sensing
CPP	Coverage Path Planning
DVL	Doppler Velocity Log
FLS	Forward-looking Sonar
GA	Genetic Algorithm
GPS	Global Positioning System
INS	Inertial Navigation System
LAUV	Light Autonomous Underwater Vehicle
MCM	Mine Countermeasures
NATO	North Atlantic Treaty Organization
OODA	Observe, Orient, Decide & Act
POD	Probability Of Detection
PRNG	Pseudorandom Number Generator
SAS	Synthetic Aperture Sonar
SSS	Side-scan Sonar

List of Symbols

μ	Mean of a normal distribution
σ	Standard deviation of a normal distribution
θ_m	Measured rotation of the AUV about the z-axis with respect to the positive x-axis in the global coordinate frame
$\bar{C}$	Average coverage of the search area (General)
$\bar{E}_e$	Average coverage of the search area (Coverage Performance Model 2)
$\bar{E}_p$	Average coverage of the search area (Coverage Performance Model 1)
$\bar{H}_h$	Average coverage of the search area (Coverage Performance Model 3)
$\boldsymbol{C}$	Discrete random variable describing the POD in a grid cell
$\boldsymbol{C}_{\text{uncovered}}$	Discrete random variable describing the POD in an uncovered grid cell
$\boldsymbol{D}$	Binary random variable describing the coverage of a grid cell
$\mathcal{P}(r, o)$	Function mapping r and o into a POD
$\max(\boldsymbol{X}, \boldsymbol{Y})$	Random variable describing the maximum of the random variables $\boldsymbol{X}$ and $\boldsymbol{Y}$
$\mathrm{E}[\boldsymbol{X}]$	Expected value of the random variable $\boldsymbol{X}$
$\mathrm{F}_{\boldsymbol{X}}$	Cumulative distribution function of the random variable $\boldsymbol{X}$
$\mathrm{H}(\boldsymbol{X})$	Function calculating the entropy of a discrete random variable $\boldsymbol{X}$
$\mathrm{H}(\mathrm{E}[\boldsymbol{C}])$	Function calculating the shifted entropy of the expected value of the POD
$\mathrm{P}(O)$	Probability of the event O
${}^{b}_{a}T$	Homogeneous transformation matrix transforming from coordinate frame a to b
A	Set of optimization parameters describing the coverage path
A_d	Distance between adjacent straight tracks in the coverage path
A_l	Location of the first straight track of the coverage path in the search area
A_n	Number of straight tracks in the coverage path
c	Required average coverage of the search area
c_1	Required average expected value of the POD in the search area
c_2	Required average entropy in the search area
D	Event that there is an object detected in a grid cell
D^c	Event that there is no object detected in a grid cell
$f_{\boldsymbol{X},\boldsymbol{Y}}$	Joint probability density function of the random variables $\boldsymbol{X}$ and $\boldsymbol{Y}$
g_x	The x-axis in the global coordinate frame
g_y	The y-axis in the global coordinate frame
$H_{R,e}$	Function calculating the entropy reduction (Coverage Performance Model 2)
$H_{R,h}$	Function calculating the entropy reduction (Coverage Performance Model 3)
$H_{R,p}$	Function calculating the entropy reduction (Coverage Performance Model 1)
H_R	Function calculating the entropy reduction (General)
N_c	Number of grid cells in the search area
N_l	Number of possible AUV locations
N_o	Number of outcomes of the discrete random variable describing the POD

N_p	Number of sonar image resolution polygons in the sensor coordinate frame
O	Event that there is an object located in a grid cell
o	Ocean floor type
O^c	Event that there is no object located in a grid cell
r	Sonar image resolution
r_x	The x-axis in the robot coordinate frame
r_y	The y-axis in the robot coordinate frame
s_x	The x-axis in the sensor coordinate frame
s_y	The y-axis in the sensor coordinate frame
w	Weight factor associated with the cost of a straight track
x_m	Measured x-coordinate of the AUV's location in the global coordinate frame
y_m	Measured y-coordinate of the AUV's location in the global coordinate frame