Localization, Mapping and Exploration with Mobile Ground Robots in Disaster Environments

Zur Erlangung des akademischen Grades Doktor-Ingenieur (Dr.-Ing.)
Genehmigte Dissertation im Fachbereich Informatik von Kevin Daun
Tag der Einreichung: 26.02.2024, Tag der Prüfung: 08.04.2024

Erstreferent: Prof. Dr. Oskar von Stryk
Korreferent: Prof. Dr. Andreas Nüchter (Universität Würzburg)
Darmstadt, Technische Universität Darmstadt

TECHNISCHE UNIVERSITÄT DARMSTADT

Computer Science Department

Fachgebiet Simulation, Systemoptimierung und Robotik

Localization, Mapping and Exploration with Mobile Ground Robots in Disaster Environments

Accepted doctoral thesis in the department of Computer Science by Kevin Daun

Date of submission: 26.02.2024
Date of thesis defense: 08.04.2024

Darmstadt, Technische Universität Darmstadt

Abstract

Responding to disasters and threatening situations is a major challenge for first responders, authorities, and the public. The use of rescue and response robots can help to overcome the challenges by improving overall response capabilities, e.g., by providing valuable insights on dangerous areas through gathering data and creating 3D maps of the environment or performing remote physical actions while enabling first responders to maintain a safe distance from potential dangers.

However, the operating conditions for robots in disasters and threatening situations such as fire, flooding, collapse, or CBRNE are very difficult. Environmental conditions are usually very harsh, with challenging ground characteristics, versatile and irregular obstacles, and potentially disturbed visual conditions due to smoke, fog, and dust. Moreover, compared to industrial robot applications, missions, and environments have large variations and low repeatability and offer little prior knowledge and lead time, making applying common methods and approaches from mobile robot autonomy and artificial intelligence (AI) particularly challenging. At the same time, these conditions pose a challenge for remote teleoperation, as they increase the likelihood of fatal errors and mission failures for human operators.

This thesis focuses on developing mobile rescue robots with assistance functions motivated by advancing disaster response efficiency and safety, e.g., by contributing to autonomous robot exploration, that account for the specific requirements and challenges to support first responders and civil forces. Therefore, this work presents specific approaches for localization, mapping, and exploration with mobile ground robots in disaster response, addressing crucial challenges in three distinct areas.

Firstly, understanding the full range of specific requirements for (autonomous) assistance functions in rescue robots is crucial for research and development towards practical applicability. Previous analyses have primarily focused on general aspects, leaving a gap in the specific understanding of requirements for (autonomous) assistance abilities. We address this gap by deriving a novel model for an integrated function capability from established models for technology acceptance and derive a comprehensive, evidence-driven analysis of application requirements and research challenges for (autonomous) assistance abilities.

Secondly, sufficiently accurate and robust simultaneous localization and mapping (SLAM) in unknown environments without relying on GNSS support are essential for (semi-)autonomous operation. In particular, traversing uneven ground can lead to abrupt robot motions that existing SLAM methods cannot model accurately or efficiently enough. Furthermore, relevant environments are often unstructured and potentially visually degraded by smoke, dust, or fog. Therefore, we investigate new methods for robustly registering lidar scans, accurately estimating the trajectory in rough terrain, and efficiently mapping large-scale environments online on a mobile rescue robot system. The proposed approach gains accuracy and robustness by registering lidar data in a multi-resolution Truncated Signed Distance Function (TSDF) with a continuous-time trajectory representation. It enables the efficient mapping of large-scale environments by transferring a branch-and-bound-based loop closure detection approach for TSDF. Furthermore, we investigate extensions of the approach for the operation in visually degraded conditions with radar.

Thirdly, in response missions, robots might need to fulfill various tasks in a single mission. In such dynamic and versatile environments, first responders often have prior knowledge and better high-level decision-making skills than AI methods for the perception and reasoning of autonomous mobile robots. However, an operator's cognitive load is limited, and direct operator control is potentially error-prone, often inefficient, and not

always possible. Therefore, we investigate a new, efficient, and flexible method for multi-goal exploration that combines AI methods for perception with operator capabilities by extending a hierarchical planning approach for multi-goal scenarios and facilitating flexible operator assistance with an actionable environment representation based on affordances.

The innovations, methods, and implementations presented in this work have been successfully evaluated in various complex simulated and real-world robot experiments, demonstrating accuracy, robustness, and efficiency. Parts of the real-world evaluation are performed under the conditions of various international robotics competitions (RoboCup Rescue Robot League, EnRicH, World Robot Summit), demonstrating better accuracy and robustness than related approaches. In addition, the results from this thesis were used for their application in real missions and as input for two German consortium standards (DIN SPEC), which underline their impact in the field of disaster robotics.

Zusammenfassung

Die Bewältigung von Katastrophen und Gefahrenlagen stellt eine immense Herausforderung für Einsatzkräfte, Behörden und die Öffentlichkeit dar. Der Einsatz von Rettungsrobotern mit Assistenzfunktionen kann dazu beitragen, diese Herausforderungen zu mindern, indem sie die allgemeine Einsatzfähigkeit verbessern. Dies geschieht beispielsweise durch Aufklärung und die Erstellung von 3D-Karten der Umgebung sowie die Durchführung von Aktionen aus der Ferne. Damit tragen sie auch zur Sicherheit der Einsatzkräfte bei, indem sie diesen ermöglichen, einen sicheren Abstand zu potenziellen Gefahren zu wahren.

Die Einsatzbedingungen für Rettungsroboter in Szenarien wie Feuer, Überschwemmungen, Einstürzen oder CBRNE sind äußerst anspruchsvoll. Die Umweltbedingungen sind in der Regel sehr widrig, gekennzeichnet durch schwierige Bodenverhältnisse, vielfältige und unregelmäßige Hindernisse sowie potenziell eingeschränkte Sichtverhältnisse aufgrund von Rauch, Nebel und Staub. Im Vergleich zu industriellen Robotereinsätzen sind die Einsatzszenarien und Umgebungen äußerst variabel und wiederholen sich selten. Dies erschwert den Einsatz gängiger Methoden und Ansätze aus den Bereichen der autonomen mobilen Roboter und der künstlichen Intelligenz (KI) erheblich. Gleichzeitig stellen diese Bedingungen auch eine Herausforderung für die Fernsteuerung dar, da sie die Wahrscheinlichkeit erhöhen das Operatoren schwerwiegenden Fehler machen.

Diese Dissertation zielt darauf ab, mobile Rettungsroboter mit Assistenzfunktionen zu entwickeln, die die Effizienz und Sicherheit bei der Katastrophenhilfe verbessern sollen. Hierbei liegt ein besonderer Fokus auf der Entwicklung neuer Methoden zur autonomen Erkundung, welche den spezifischen Anforderungen und Herausforderungen bei Rettungseinsätzen gerecht werden. In dieser Arbeit werden daher Ansätze für die Lokalisierung, Kartierung und Exploration mit mobilen Bodenrobotern im Kontext der Bewältigung von Katastrophen und Gefahren vorgestellt und Forschungsthemen in drei verschiedenen Bereichen untersucht.

Erstens ist ein umfassendes Verständnis der spezifischen Anforderungen an (autonome) Assistenzfunktionen von Rettungsrobotern entscheidend für die Forschung und Entwicklung hin zur praktischen Anwendbarkeit. Bisherige Analysen haben sich in hauptsächlich auf allgemeine Aspekte konzentriert, sodass eine Lücke im spezifischen Verständnis der Anforderungen für (autonome) Assistenzfähigkeiten besteht. Diese Lücke adressieren wir, indem wir ein neus Modell für eine integrierte Funktionsfähigkeit aus etablierten Modellen zur Technologieakzeptanz ableiten und eine umfassende, evidenzbasierte Analyse der Anwendungsanforderungen und Forschungsherausforderungen für (autonome) Assistenzfähigkeiten erstellen.

Zweitens ist eine präzise und zuverlässige gleichzeitige Lokalisierung und Kartierung (SLAM) in unbekannten Umgebungen ohne GNSS-Unterstützung für den (teil-)autonomen Betrieb unerlässlich. Insbesondere das Überqueren von unebenem Gelände kann zu abrupten Roboterbewegungen führen, die mit typischen SLAM-Methoden nicht genau oder effizient modelliert werden können. Zudem sind relevante Umgebungen oft unstrukturiert und möglicherweise visuell beeinträchtigt. Daher untersuchen wir neue Methoden zur robusten Registrierung von Lidar-Scans, zur präzisen Schätzung der Trajektorie in unebenem Gelände und der effizienten Kartierung großer Umgebungen mit mobilen Rettungsrobotern. Der Ansatz erzielt hohe Genauigkeit und Robustheit durch die Registrierung von Lidar-Scans in einer multi-resolution Truncated Signed Distance Function (TSDF) mit zeitkontinuierlicher Trajektoriendarstellung und ermöglicht die effiziente Kartierung von großen Umgebungen durch die Übertragung eines branch-and-bound-basierten Ansatzes zur Erkennung von Loop Closures für TSDF. Darüber hinaus untersuchen wir Erweiterungen für die Anwendung mit Radar.

Drittens müssen Roboter bei Einsätzen möglicherweise verschiedene Aufgaben in einer einzigen Mission erfüllen. In solch dynamischen und vielseitigen Umgebungen verfügen Rettungskräfte oft über Vorwissen und potentiell bessere Entscheidungsfähigkeiten als KI-Methoden für die Perzeption und Logik von autonomen Robotern. Die kognitive Belastung des Bedieners ist jedoch begrenzt, und die direkte Kontrolle durch den Bediener ist fehleranfällig, potentiell ineffizient und nicht immer möglich. Daher präsentieren wir eine neue, effiziente und flexible Methode für die Multi-Ziel-Exploration vor, die KI-Methoden mit den Fähigkeiten des Operators kombiniert. Die geschieht durch die Erweiterung eines hierarchischen Planungsansatz für Multi-Ziel-Explorationund die Bereitstellung flexiblee Bedienerunterstützung durch eine Umgebungsrepräsentation auf der Grundlage von Affordanzen ermöglicht.

Die in dieser Arbeit vorgestellten Innovationen, Methoden und Implementierungen wurden in verschiedenen komplexen simulierten und realen Roboterexperimenten erfolgreich evaluiert, wobei Genauigkeit, Robustheit und Effizienz demonstriert wurden. Teile der realen Evaluierung wurden unter den Bedingungen verschiedener internationaler Robotikwettbewerbe (RoboCup Rescue Robot League, EnRicH, World Robot Summit) durchgeführt und zeigten eine bessere Genauigkeit und Robustheit als konkurrierende Ansätze. Darüber hinaus flossen Ergebnisse dieser Arbeit in zwei deutsche Konsortialnormen (DIN SPEC) ein und wurden bei einem Katastropheneinsatz eingesetzt, was die Bedeutung der Arbeit für den Bereich der Rettungsrobotik unterstreicht.

Contents

Acknowledgement

This thesis was written during my time as a research assistant at the Simulation, System Optimization, and Robotics Group (SIM) at the Department of Computer Science of the Technische Universität Darmstadt. Parts of my research have been supported by research grants from the German Federal Ministry of Education and Research (BMBF).

I am very grateful for the support I received from many people who accompanied me during the journey from my first robotic project in the SIM group as an undergraduate student until the completion of this thesis, even though not all are explicitly named here.

I appreciate the support from my supervisor, Prof. Dr. Oskar von Stryk, who gave me the opportunity and trust to work on exciting projects and research topics. I deeply value his support, patience, and expertise, which was crucial for creating this thesis. Furthermore, I am thankful for Prof. Dr. Andreas Nüchter agreeing to become my second referee.

I appreciate the support I received from Dr. Stefan Kohlbrecher, who first introduced me to rescue robotics, led Team Hector while I was a student, and supervised my Master's Thesis, thereby laying the foundation for my later role as a research assistant.

I appreciate the support I received from all members of Team Hector. I learned a lot from this team and have shared many exciting competitions and conferences. I especially want to thank Marius Schnaubelt for his support, insights, and fun during long (and late) car rides, especially during the many trips as part of the A-DRZ project. The joint disaster deployments were some of the most impressive experiences during my PhD phase.

I appreciate Felix Biemüller's support in proofreading parts of this thesis.

I would also like to thank the whole DRZ team for enabling parts of the evaluation of this thesis and fostering an environment for fruitful exchange and development. Furthermore, I would like to thank the former Google Cloud Robotics team for providing valuable insights into robotics development, research, and personal growth.

Most importantly, I want to thank my family for their constant support. I could not have finished the work of this thesis without them. I want to give a very special thanks to Nathalie for all her patience, advice, and encouragement.

This thesis was written independently and was grammatically revised with the help of Grammarly and ChatGPT. DeepL Translator was used to support the translation of the abstract from English into German.

The figures in this document use icons by Dewi Sari, juicy fish, freepik, smashicons, geotatah, good ware, eucalyp and berkahicon from flaticon.

List of Figures

List of Tables

List of Abbreviations

AI Artificial Intelligence
CBRNE Chemical, Biological, Radiological, Nuclear, and Explosives
CNN Convolutional Neural Network
EE Effort Expectancy
FC Facilitating Conditions
FOV field of view
GNSS global navigation satellite system
HRI human-robot interface
ICP Iterative Closest Points
IFAFRI International Forum to Advance First Responder Innovation
IMU inertial measurement unit
LiDAR light detection and ranging
NDT Normal Distribution Transform
NeRF Neural Radiance Field
NIST National Institute of Standards and Technology
PE Performance Expectancy
radar radio detection and ranging
RGB-D Red Green Blue-Depth
SDF Signed Distance Function
SI Social Influence
SLAM simultaneous localization and mapping
surfel Surface Element
TAM Technology Acceptance Model
TOED Theory of Optimal Experimental Design
TRL Technology Readiness Level
TSDF Truncated Signed Distance Function
TSP Traveling Salesperson Problem
UAV unmanned aerial vehicle
UI user interface
UTAUT Unified Theory of Acceptance and Use of Technology

1 Introduction

This chapter gives a general introduction to the topics and contents of this thesis. We first outline the motivation and goals of this work and give a compact overview of the proposed approaches for localization, mapping, and exploration in disaster environments.

1.1 Motivation

Responding to disasters and threatening situations is a major challenge for first responders, authorities, and the public. For example, in the United States in 2016, 69 firefighters died on duty, and 62085 suffered injuries requiring medical attention [17]. The authorities are faced with the challenge of making decisions while information is scarce, especially in the early phase of a disaster. Individuals can be affected directly by the disaster and be injured, trapped, or otherwise in need of urgent support. The use of rescue and response robots can help overcome the challenge by improving overall response capabilities. For example, they can provide valuable insights into a dangerous area by gathering data and creating 3D maps of the environment or performing remote physical actions while enabling first responders to maintain a safe distance from potential dangers.

In such situations, rescue and response robots can provide versatile abilities for a comprehensive situational overview, such as real-time data, including images and measurements. With assistance capabilities, these sensor modalities can be processed to generate insights from the data, such as semantic analyses (e.g., detecting a fire in a thermal image) or generating maps and 3D models. This can help in short- and long-term planning of further reaction measures. Furthermore, the robots can support the responders by initiating coping steps to limit the extent of a disaster (e.g., shutting down a machine or closing a valve) or directly supporting the responders (e.g., carrying a load). More complex assistance functions such as (semi-)autonomous navigation and exploration can further support the operator during the stressful and challenging task of remotely controlling a robot and even enable the operation in environments where no direct communication to the robot is possible.

While ground robots are already commonly applied in manufacturing, logistics, and inspection environments, this is not the case for rescue robotics. The tasks in the first three domains can be very complex, such as coordinating groups of logistic robots. However, the environment can typically be controlled, at least to some extent, and the tasks share a high repetitiveness. In contrast, disaster operations and environments are complex, vary widely, and have low repeatability with little prior knowledge and lead time. Furthermore, the environmental conditions are usually very harsh, with little structure in the environment, challenging ground characteristics, versatile and irregular obstacles, and potentially disturbed visual conditions due to smoke, fog, and dust. These conditions make the real-world application of disaster robots very challenging. Dealing with these challenges motivates various research fields such as system design, human factors, algorithms for sensing, planning and acting, and artificial intelligence.

While the application of ground robots in disaster response is still rare, several successful deployments include the reactions to the nuclear accident at the Fukushima Daiichi Nuclear Power Plant in 2011 and the Notre Dame fire in Paris in 2019. In Fukushima, one of the multiple ground robots was deployed under teleoperation to perform multiple tasks such as radiation and temperature measurements, collecting air

samples, investigating damaged piping by taking images and installing water measurement gauges [177]. In Paris, a teleoperated robot entered areas with high heat and risk of structural collapse of Notre Dame, extinguishing fires and removing debris[1]. However, both deployments also indicate the limitations of the applied robotic systems. The Fukushima deployment had a long preparation lead time, and the operators needed to be trained extensively. Finally, the robot was lost in the nuclear plant during the sixth mission due to teleoperation errors. In the Notre Dame deployment, the operators needed direct visual contact with the robot and the scene for the teleoperation of the robot, and the operation in an area without visual contact would not have been possible.

1.2 Goals and Content of this Work

This thesis focuses on the development of mobile rescue robots with assistance functions motivated by advancing disaster response efficiency and safety, e.g., by contributing to autonomous robot navigation, that account for the specific requirements and challenges to support first responders and civil forces in disaster environments. Therefore, this work presents specific approaches for localization, mapping, and exploration with mobile ground robots in disaster response relief, addressing crucial challenges in three distinct areas.

1.2.1 Requirements, Guidelines, and Challenges for Autonomous Capabilities and Assistance Functions for Ground Rescue Robots in Reconnaissance Missions

The remote operation of complex ground robots in disaster conditions is very demanding for human operators [116]. When deploying ground robots, state-of-the-art today is direct teleoperation of the robots [116]. Complex mission conditions can lead to operation errors and oversights, which can be catastrophic for the mission result. $50\,\%$ of mission failures are directly related to human errors, with workload overload as a main factor [116]. This emphasizes the need for assistance in robot operation.

The focus on teleoperation provides a large gap to the academic community, where competitions, trials, and experiments [92, 31, 87] demonstrate various autonomous assistance abilities, from assisted manipulation to large-scale autonomous exploration and mapping, that can help to improve performance, extend functionality and reduce workload overload of the operator. This gap indicates a potential lack of understanding of the full spectrum of application requirements and research challenges for assistance abilities and autonomy, which has so far not been fully addressed by the academic community. However, understanding the specific requirements for (autonomous) assistance abilities is crucial for research and development toward practical applicability.

To address this gap, we analyze application requirements, guidelines and scientific challenges for autonomy and assistance abilities in a holistic response mission context in Chapter 4. As part of the analysis, we derive a novel model for an integrated function capability from established models for technology acceptance, performed evaluations in disaster deployments as part of The Robotics Task Force established by the German Rescue Robotics Center DRZ and performed multiple requirements workshops as part of research projects with first responders from various emergency services.

1.2.2 Robust Simultaneous Localization and Mapping in Challenging Environments

A key challenge in supporting human operators with useful assistance abilities is the development of assistance abilities able to deal with the harsh environmental conditions at a disaster site. To enable (semi-)autonomous operation or support the operator with localization and mapping in an unknown environment, the robot must be able to sufficiently accurately localize itself in unknown environments without global navigation satellite

[1] https://spectrum.ieee.org/colossus-the-firefighting-robot-that-helped-save-notre-dame, accessed 16.01.2024

system (GNSS) reception, even on difficult terrain, and generate a map of the environment. In particular, traversing uneven ground can lead to abrupt robot motions that existing simultaneous localization and mapping (SLAM) methods cannot model accurately or efficiently enough. Furthermore, relevant environments are often unstructured and potentially visually degraded by smoke, dust, or fog.

In Chapter 5, we investigate novel methods for robustly registering scans, accurately estimating the trajectory in rough terrain, and efficiently mapping large space environments online on a mobile rescue robot system in disaster response conditions Furthermore, we investigate extensions for the operation in visually degraded conditions with radar.

1.2.3 Exploration and Data-Acquisition in Shared-Autonomy Multi-Goal Missions

In response missions, robots might need to fulfill various tasks in a single mission, such as creating a 3D map of an environment and also taking dose rate measurements for an accurate radiation map. Furthermore, during the mission, information dynamics can change. For example, hints of a missing person or the detection of a hazard might require immediate attention and modifications to the mission procedure. In such dynamic and versatile environments, first responders often have prior knowledge (e.g., rough location of missing persons), highly relevant context knowledge, and potentially better scene understanding and high-level decision-making skills than current Artificial Intelligence (AI) methods for perception and reasoning of autonomous mobile robots. However, the operator's cognitive load is limited, and direct operator control is error-prone, potentially inefficient, and often not possible (e.g., bandwidth, latency).

With current methods, fully autonomous robot operations are typically not possible due to the large complexity and variations in missions. Therefore, we investigate a new, efficient, and flexible method for multi-goal exploration that flexibly combines AI methods with operator abilities in Chapter 6. We investigate the aspects of efficient large-scale exploration for multi-goal missions, actionable environment representations, radiation mapping, and the embedding into a shared autonomy approach facilitating flexible operator assistance.

1.2.4 Approach Overview and Thesis Structure

In this thesis, we consider the topics of *localization, mapping, and exploration with mobile ground robots in disaster response*. Figure 1.1 provides an overview of the relation of the aspects. The approaches aim to enable a human operator with a mobile robot with assistance functions to efficiently perform reconnaissance or mitigation missions and tasks. Therefore, disaster conditions and operator needs and properties are analyzed in the *Requirements, Guidelines and Challenges for Autonomous Capabilities and Assistance Functions for Ground Rescue Robots in Reconnaissance Missions* in Chapter 4. Aspects of the resulting requirements and challenges are addressed in the approaches for *Robust Simultaneous Localization and Mapping in Challenging Environments* in Chapter 5, which covers novel approaches for registering lidar scans, accurately estimating the trajectory in rough terrain, and efficiently mapping large space environments, and *Exploration and Data-Acquisition in Shared-Autonomy Multi-Goal Missions* in Chapter 6, where we investigate a novel method for exploration and data acquisition in shared autonomy multi-goal missions. The robot with these assistance abilities provides the operator with abilities and assistance, while the operator can control the robot with these functions and perform supervision. The approaches are comprehensively evaluated in Chapter 7. Furthermore, we outline the background of this thesis and give a general overview of the state of the research in Chapter 2. Chapter 3 presents the robots with integrated function modules that were partly developed, integrated, and used in this thesis. Chapter 8 covers a final conclusion of the proposed approach and an outlook on future work.

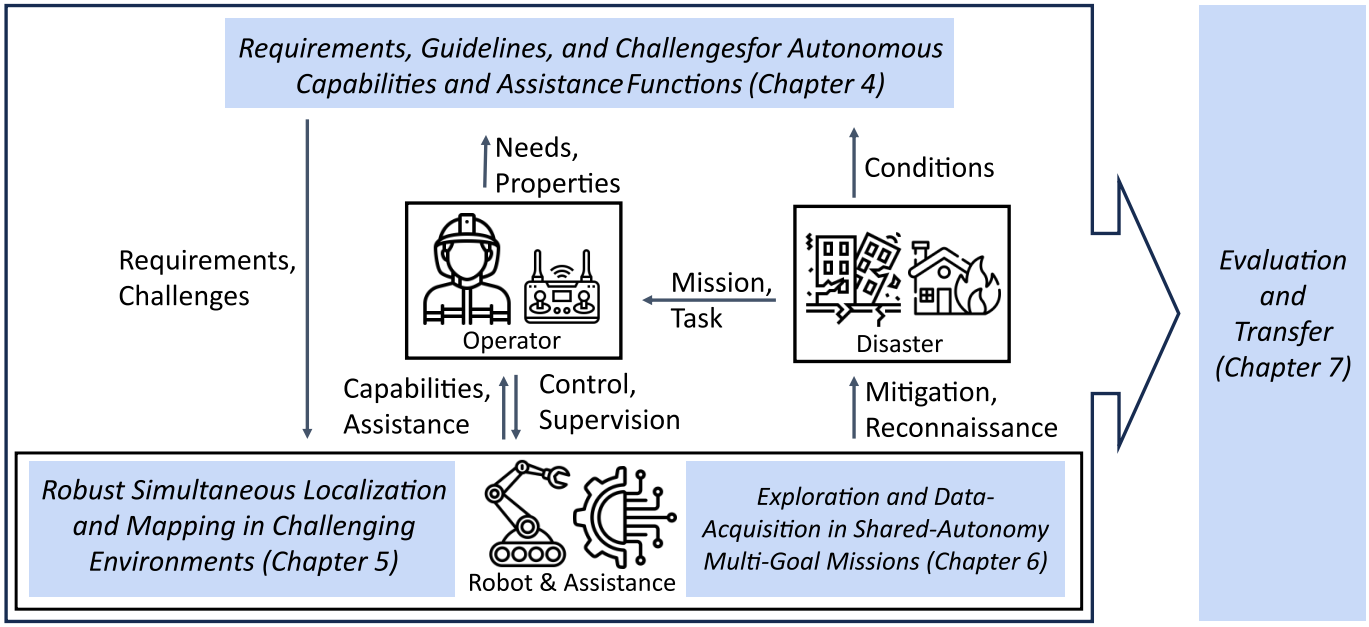

Figure 1.1: Overview of the relations of the proposed approaches for Localization, Mapping and Exploration with Mobile Ground Robots in Disaster.

1.2.5 Evaluation and Transfer

We perform comprehensive evaluations to investigate the performance and properties of the proposed methods in Chapter 7. We combine the use of simulation, benchmark data sets, and evaluations as part of a fully integrated robot system under conditions of multiple international robotic competitions. As typical benchmark data sets do not address the challenging motion characteristics when traversing obstacles or uneven terrain, we also introduce a novel data set.

Finally, we demonstrate the successful transfer of the proposed approach for practical applications such as evaluation as part of an actual disaster deployment and the contribution to the implementation of the German consortial standards (DIN SPEC) for robots in hazardous environments.

2 Background and State of Research

This chapter gives a general overview of the background and state of the art and research for topics discussed in this thesis. We first give a general introduction to the topic of robots and assistance capabilities for disaster response, simultaneous localization and mapping, and exploration. Each section also relates to design decisions made for the proposed approach and to the research questions addressed in Chapters 4 - 6. This chapter serves to give a general introduction to the topics. Additionally, Chapters 4 - 6 each cover a section providing a more focused and in-depth analysis and distinction to the related research publications of the proposed approach.

2.1 Robots and Assistance Capabilities for Disaster Response

This thesis considers rescue robots as robots that support first responders and other civil stakeholders in disasters or extreme situations. This covers the aspects of "sense and act at a distance from the site of a disaster or extreme incident" [116] but also the operation in direct, close collaboration with first responders, such as support in transportation. As such extreme events, we also consider deployments for civilian forces such as police operations with the potential presence of Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) threats.

Emergency and disaster management can be split into the phases of prevention, preparation, response, and recovery [13]. This work focuses on the application in the response phase, although the developed concepts and methods can also find applicability in the other phases.

Mobile robots can be classified based on the modality where they operate - maritime, aerial, and ground.

2.1.1 Aerial Robots

Aerial robots are an established and widespread tool in disaster response, e.g., in wilde fire prevention and response [109] or the search for missing persons[1]. A common use case is to provide improved situation awareness for fire brigades with live video of a color and/or thermal camera from an aerial perspective. Such systems are easy to deploy, cost-efficient (e.g., 2100 Euro for a DJI Mavic 3 Pro[2]), and provide a direct benefit by improving the situation awareness to the operation teams. Systems with complex and advanced sensing modalities such as light detection and ranging (LiDAR) [135] or manipulation capabilities [110] are currently more common in academic research. While the advantages of aerial robots are rapid deployment, data from an aerial perspective, low weight, easy transportation, and cost-effectiveness, there are also limitations. The limited payload restricts the type and amount of sensors that can be carried. Especially in indoor environments, the applicability is limited as navigation in narrow spaces is inherently difficult, the need for direct radio communication cannot be guaranteed, and the potential need for object interaction, such as opening doors, is not possible. Furthermore, the application is sensitive to weather conditions and the regulatory environment.

[1] `https://www.flyhound.com/`, accessed 25.11.2023
[2] `https://store.dji.com/`, accessed 25.11.2023

(a) AeroVironment telemax EVO Hybrid (Image: © AeroVironment, Inc.)

(b) Rosenbauer RTE Robot with a fire fighting monitor (Image: © Rosenbauer International AG)

Figure 2.1: Examples of ground robots applicable in disaster environments.

2.1.2 Ground Robots

The application of ground robots in disaster response is less widespread than aerial robots. Ground robots need to locomote over the ground structure, which can be very challenging with ground structures such as rubble and debris, potentially mud and water. In urban environments, the traversal of steps or stairs might be necessary. To account for this, ground robots, such as the Telerob Telemax or the Rosenbauer RTE Robot (see Figure 2.1), are typically larger and heavier than their aerial counterparts, requiring additional efforts for deployment logistics. Ground robots offer significant advantages, including the ability to transport larger and heavier payloads. Additionally, they allow direct interaction with the environment, such as manipulating objects with a robotic arm, e.g., opening doors, turning valves, taking samples, or removing debris. Ground robots are regularly applied when there are very high risks for first responders, such as the handling of potential explosive threats[3] or the handling of critical chemical containers[4]. In extreme environments they are also deployed for reconnaissance [117, 177], structural inspection [93] or manipulation tasks [177]. However, the operation mode in such environments is typically still in direct teleoperation, which leads to errors and oversights [116]. Further challenges are that radio connectivity in indoor environments cannot be guaranteed, requiring cables, sufficient repeaters, or autonomous operation capabilities. Overall, these conditions are also very hostile to human first responders and emphasize the need for robotic technologies to release human responders from such risk-intensive tasks.

2.1.3 Assistance Capabilities and Autonomy

The operation of rescue robots in disaster environments is very challenging. A large focus of the academic community is to develop autonomous and assistance capabilities for ground robots, which enable them to automate complex behaviors (e.g., autonomously exploring an environment) or single tasks (e.g., climbing stairs, opening a door). Such that an operator gains support in the control or does not need to control the robot at all.

In accordance with [47], we classify the control interaction as teleoperated, semi-autonomous, or autonomous:

[3] `https://www.avinc.com/ugv/telemax-evo/`, accessed 25.11.2023

[4] `https://www.basf.com/global/de/who-we-are/organization/locations/europe/german-sites/ludwigshafen/neighbor-basf/environment-and-safety/fire-department/about_us/fleet.html`, accessed 25.11.2023

- Teleoperated: "*Continuous, direct control of a robot by an operator using a user interface; assistance capabilities may be available and used to support operation*" [47]
- Semi-Autonomous: "*Control of the robot under continuous specification of intermediate goals with limited complexity by an operator using a user interface*" [47]
- Autonomous: "*Ability of a robot to independently fulfill a function, possibly defined by an operator, without the need for human intervention during execution*" [47]

Similar to the typical sense-plan-act-cycle of robotic systems [91], assistance capabilities also cover these three aspects. The sensing part covers using the sensor to measure internal (e.g., joint angle, orientation, battery voltage) or external (e.g., camera image, LiDAR scan, radiation dose rate, temperature) quantities. The plan part covers the fusion and reasoning from the sensor data. Depending on the function, this typically means maintaining a model of the robot's state and environment and reasoning about relevant actions. The act part then covers the execution of the planned actions, such as actuating a motor or providing the operator with a visualization.

2.1.4 Requirements and Challenges

To ensure that developed robotic systems apply to real-world needs, it is necessary to understand relevant requirements. The gathering of requirements and challenges started with the first deployments of rescue robots. The report on the deployment for the exploration of voids after the World Trade Center collapse in 2001 [117] discusses recommendations for rescue robot systems and autonomous control. The subsequent deployments [115] lead to a seminal and comprehensive analysis in [116] and [119] covering an overview of past robot deployments and provide insights on lessons learned, challenges, and practical guidelines. The authors emphasize the high maturity level required to operate successfully in the challenging conditions of actual disasters. They point out that with the given "operation envelopes" (operation conditions in disasters), robot control is very challenging and error-prone due to high cognitive load, perceptual impoverishment, and the choice of user interfaces and system designs. Further analyses focus on specific aspects such as human-robot interaction [118, 137] or general system design [49, 144].

Complementary are the insights from competitions that simulate aspects of a disaster deployment, such as the DARPA Robotics Challenge (DRC) or the DARPA Subterranean (SubT) Challenge. An analysis of the performances at the DRC [5] concludes that operator errors explain a large fraction of the overall errors and indicate the importance of human-robot interaction and early detection of errors in the robot or operator commands. Furthermore, they identify the need that the robots themselves need to be safe to operate in an environment close to humans. The authors also indicate the importance of using available sensing systems to support task execution. The evaluation of the team performances at the SubT Challenge [31] confirms the role of the human operator as a critical performance limiter and emphasizes the need for improved situational understanding, integrated interfaces, and reliable autonomous functions.

The works on requirements and challenges cover general deployment aspects and indicate the need for assistance capabilities. However, the understanding of requirements and challenges for the application of assistance capabilities is often anecdotal, and a comprehensive understanding is missing. We address this limitation and contribute to the understanding of challenges and requirements for the application of assistance capabilities for mobile rescue robots with a comprehensive analysis and share insights from evaluations in the scope training with first responders and deployments of a preliminary robotics task force in Chapter 4.

2.2 Simultaneous Localization and Mapping

A fundamental problem serving as an assistance capability itself and a prerequisite for many assistance capabilities is to estimate the location of the robot in an unknown environment and to provide a map of the environment.

Estimating the location requires a map to register the sensor data, and mapping requires the robot's location to integrate the sensor data consistently. Due to this cyclic dependency, localization and mapping are tightly coupled, and the problem is referred to as simultaneous localization and mapping (SLAM). Resolving the mutual dependency makes the SLAM problem a hard task and requires searching for a solution in a high-dimensional space.

As a fundamental capability in robotics, SLAM is an established research problem that has gained a lot of attention in the research community, resulting in various methods with varying assumptions and approaches. In the following subsections, we give a formal definition of the SLAM problem and subsequently discuss sensor types and map representations.

2.2.1 The SLAM Problem

Following [151], we define the SLAM problem as follows. We assume a mobile robot is operating in an unknown environment and perceiving its environment with noisy sensors. The goal is to estimate the robot's trajectory $\boldsymbol{x}_{0:t} = \{\boldsymbol{x}_0, \boldsymbol{x}_1, ..., \boldsymbol{x}_t\}$, where each $\boldsymbol{x}_i$ is a pose at time i, and the map of the environment $\boldsymbol{m}$. The map can have various representations, such as landmarks, objects, or surfaces.

Available data are relative motion estimates $\boldsymbol{u}_i$ for the motion between time i and $i-1$ as $\boldsymbol{u}_{1:t} = \{\boldsymbol{u}_1, \boldsymbol{u}_2, ...\boldsymbol{u}_t\}$, e.g., from integrating wheel encoders, and a sequence of measured pose observations relative to the environment $\boldsymbol{z}_{1:t} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, ...\boldsymbol{z}_t\}$, e.g., by detecting landmarks or matching laser scans. These estimates are suspect to uncertainty and measurement noise.

We follow the probabilistic formulation [151] to model the uncertainty and formulate the **full SLAM problem** as determining the posterior over the map $\boldsymbol{m}$ and the whole trajectory $\boldsymbol{x}_{0:t}$ given the sequences of observations $\boldsymbol{z}_{1:t}$ and controls $\boldsymbol{u}_{1:t}$

$$p(\boldsymbol{x}_{0:t}, \boldsymbol{m} | \boldsymbol{z}_{1:t}, \boldsymbol{u}_{1:t}). \tag{2.1}$$

However, for many applications, we are only interested in the most recent pose $\boldsymbol{x}_t$ of the robot, which is called **online SLAM**. Then the problem reduces to estimating the posterior over the map $\boldsymbol{m}$ and the latest pose $\boldsymbol{x}_t$

$$p(\boldsymbol{x}_t, \boldsymbol{m} | \boldsymbol{z}_{1:t}, \boldsymbol{u}_{1:t}). \tag{2.2}$$

Approaches to solving the full SLAM problem typically first need to gather data from the whole trajectory and then process it as a whole. Therefore, these approaches are referred to as **offline methods**. In contrast, approaches solving the online SLAM problem process the data as it is arriving. They can give an estimate of the current robot pose and map at runtime and are called **online methods**.

This general definition of the SLAM problem indicates various aspects that need to be considered. Depending on the application domain, choosing the right sensor type is crucial and has strong implications for the design of a suitable method. Furthermore, the SLAM problem covers two main aspects – estimating the robot's trajectory and mapping the environment. This requires suitable techniques to model the trajectory and the environment, which influence the choice of methods to optimize the trajectory and the map.

2.2.2 Sensor Types

An essential consideration is the choice of sensor to observe the environment to perform SLAM. Common choices are LiDAR [147], radio detection and ranging (radar) [74], monocular [114] or stereo vision [53], and

thermal stereo vision [113, 52]. Additionally, SLAM systems often incorporate internal sensors to support the state estimation such as an inertial measurement unit (IMU) or an odometer. Absolute reference sensors such as GNSS or external localization systems are typically not available in unknown indoor disaster environments and are, therefore, not considered in the scope of this work.

External Sensing

LiDAR ("light detection and ranging") sensors measure the distance to a surface by emitting laser beams and measuring the time for the reflected light to return. Advantages of LiDAR-based sensing, with state-of-the-art sensors for mobile applications, such as the Velodyne VLP-16 Puck[5] or the Ouster OS-0[6], are high-frequent and precise distance readings, large measurement range, wide horizontal field of view (FOV), average vertical FOV and average density, robustness to light conditions. Disadvantages are sensitivity to degraded visual conditions (smoke, fog, dust), sensitivity to reflective surfaces, and high price.

Radar ("radio detection and ranging ") sensors measure the distance and relative velocity to a surface by transmitting electromagnetic waves in the radio or microwave domain. The received reflections are then processed to estimate the surface's distance and relative velocity. Advantages are robustness to degraded visual conditions (smoke, fog, dust), obstacle penetration, observing surface velocity, and average cost. Disadvantages are low range and angular precision, interference of other radars, and multi-path return signals.

The experimental comparison of LiDAR and radar sensors by Gim et al. [64] also indicates the advantages of LiDAR with respect to measurement precision and coverage, while radar systems provide better environmental robustness and lower cost.

Monocular vision approaches use a sequence of images from a single camera to estimate the motion and 3D geometry of the environment. In contrast, **stereo vision approaches** typically use two cameras and leverage the disparity between the two images to reason about the 3D structure of the environment. A general challenge for vision-based sensing is that the sensor data does not provide direct range readings, and therefore, depth information needs to be inferred from a series of images in monocular approaches or by leveraging the image disparity for stereo approaches. The advantages are simple and cost-efficient hardware and color information. Disadvantages include sensitivity to lighting and visual conditions.

Stereo vision can also be performed with **thermal cameras** [113, 52], which improves the robustness to lighting conditions at the cost of an increased price, the lack of color information, and reduced resolution.

Internal Sensing

It is common to fuse an internal sensor, such as an **IMU**, which observes linear accelerations and angular velocities. Estimating the position and orientation from IMU observation requires a double integration of the linear acceleration and a single integration of the angular velocity [58]. Both the angular velocities and linear accelerations are suspect to noise and slowly varying biases. Therefore, the estimation of the position drifts strongly after short integration sequences. The linear acceleration allows for observation of the direction of gravity, which allows for filtering of the linear and angular observations and provides an accurate estimation of roll and pitch components even over longer periods. However, the yaw component cannot be observed by the direction of gravity and is still bound to drift, although as it incorporates only one instead of two integrations, the drift is much slower than for the position [160].

Odometers, which count the revolution of wheels or tracks, enable observation of components of the linear and angular velocity. As the odometer does not observe the slippage of the wheels/tracks, the accuracy is limited, and the estimates become very inaccurate on loose ground and rough terrain.

[5] `https://velodynelidar.com/products/puck/`, accessed 29.11.2023
[6] `https://ouster.com/products/hardware/os0-lidar-sensor`, accessed 29.11.2023

Sensor Choice for the Proposed Approach

To achieve a high accuracy and robustness in both the accuracy of the map and the pose estimate, we base the proposed approach in Chapter 5 on LiDAR fused with IMU and odometer readings. To account for degraded visual conditions, we investigate the incorporation of radar.

2.2.3 Map Representations for Lidar-based SLAM

As the accuracy of the map and the pose are tightly coupled in SLAM the choice of the map representation is crucial for the overall performance of the method. In the following, we first outline requirements and aspects that need to be considered for SLAM in disaster environments. Subsequently, we outline and discuss existing map representations.

Requirements

The **accuracy** of the environment representation is a crucial property, as it directly relates to the fidelity of the map and the accuracy of the localization. For LiDAR-based methods, this typically means the geometric accuracy of the surface representation or features and an estimate of the uncertainty, which directly relates to the accuracy of the registration of a LiDAR scan. For usage of the map for operator assistance or further assistance capabilities such as obstacle avoidance or path planning, it is important to achieve a **dense representation** of the environment with accurate surface representations and potentially additional information such as indications of free space, distance to surfaces or translucency. Especially for the application in disaster scenarios, **robustness** is crucial. Robustness, in this sense, means the property to achieve accurate representation even with noisy sensor data and to enable an accurate registration of LiDAR scans even with poor priors. We consider functions that operate online on a mobile robot with limited compute, therefore **memory and computational efficiency** are crucial for applicability. Accuracy, dense representation, and robustness are often properties that conflict with memory and computational efficiency. Therefore, finding representations that allow a sufficient trade-off between these aspects is essential for an overall well-performing approach.

Representations

Volumetric approaches sample the map at a given resolution to allow a reconstruction of the environment. Resulting in an usually high-dimensional map. Examples of volumetric maps are occupancy grid maps or Truncated Signed Distance Functions (TSDFs). **Feature-based** approaches extract sparse features from the sensor data and use these features as environment representation. Due to the reduced dimensionality, computation complexity is reduced, and feature-based methods tend to be more efficient. However, the extraction of sufficient, robust, and unique features is a hard problem. If the number of features in the environment is not sufficient, these methods are prone to mismatches, leading to inconsistent maps. Furthermore, they do not directly provide a dense 3D model of the environment that can be used for operator assistance. Therefore, we primarily focus on volumetric approaches.

Point Clouds are sets of typically 3D points from the range observations of LiDAR sensors. Naively accumulating all incoming scan observations into a single point cloud quickly leads to intractable models. Therefore, often only selected and/or filtered LiDAR scans are stored as so-called keyframes. The iterative closest points algorithm [10] enables accurate point cloud registration. With the availability of high-density and high-FOV LiDAR sensors such as the Ouster OS-0, Iterative Closest Points (ICP)-based LiDAR odometry methods gained again increased attention in the last years [163, 28, 171].

Occupancy Grids [69, 86, 71], also known as Probability Grids, represent the space as a 2D-pixel grid or 3D-voxel grid, where each cell stores the probability of being occupied. Maps can be updated by performing

a ray cast from the sensor origin to a range observation and using an inverse-sensor model to update the probability of all cells along the ray. The representation has the advantages of providing a sound probabilistic framework and an efficient map update scheme. The probabilities can be translated to free, occupied, and unknown space, are usable for planning with mobile robots, and are easily readable by humans. Disadvantages are that the accuracy is limited by the map resolution, and the computation of gradients is only possible directly at the surface. The flexible probabilistic framework enabled various successful applications such as GMapping [69], Hector SLAM [86], Cartographer [71] and Octomap [75].

Truncated Signed Distance Functions, also Truncated Signed Distance Fields, represent the space as a 2D-pixel grid or 3D-voxel grid, where each cell stores two values - the signed distance to the next surface and a weight indicating the confidence in the value. To gain efficiency, the distance function is only evaluated until a truncation threshold is reached. Thereby, values are only stored close to surfaces. The surface can be extracted by computing the zero isocontour, e.g., using the marching cubes algorithm [101]. Therefore, the surface accuracy is higher than the cell resolution and more accurate than for occupancy grids at the same resolution, enabling higher quality environment models and more accurate scan registration. A further advantage is the availability of gradients in the direction of the surface within the truncation distance. Disadvantages are that the map update induces more complexity than occupancy grids, projection errors for small angles are more notable, and thin objects cannot be represented without significant extensions [149]. While TSDF approaches typically apply some form of projective distance, Euclidean Distance Functions store the Euclidean Distance without truncation. This induces an increased computational cost but also enables efficient applications for motion planning [125]. ESDFs can be efficiently computed from TSDFs [125].

TSDFs gained popularity in the application with Red Green Blue-Depth (RGB-D) cameras with KinectFusion [120, 167]. However, applying these methods to laser scanners is challenging, as sensors also provide distance readings under highly slanted angles and are less dense, which leads to poor surface reconstructions. Initial works for applying TSDFs for 2D LiDAR SLAM are 2D SDF SLAM [59] and Ohm TSD SLAM [84].

Normal Distribution Transform (NDT) [12] and **Surface Elements (surfels)** [51, 168, 9] are two closely related methods which discretize space in a grid. Each cell stores a normal distribution of the 3D point measurements. NDTs represent this distribution as a mean and covariance matrix, whereas Surfels represent it by mean and the normal as the Eigenvector of the distribution. Such methods achieve high alignment accuracy and provide a direct estimate of the uncertainty. However, scan registration requires repeated alignment, and the representation does not provide direct gradients. The representation has been successfully applied to achieve highly accurate and large-scale 3D SLAM methods, such as ElasticFusion [168] or the approach by Behley and Stachniss [9].

Deep Learning-based environment representations have become more common in the last years, mainly for visual SLAM approaches [111]. Such systems either use Convolutional Neural Network (CNN) to extract a traditional environment representation, such as an Signed Distance Function (SDF) [173] or learn an implicit representation, such as implicit SDF encodings [130, 179] or **Neural Radiance Fields (NeRFs)** [107, 142] which produce high-quality visual results in room-scale environments.

However, due to the scale and diversity of scenes and the sparsity of LiDAR data, the applicability for LiDAR-based approaches is still an open research question. Recent results [46] indicate promising potentials with high-quality maps in urban environments. However, the proposed approach is not yet capable of real-time, and the performance in very unstructured disaster environments requires further investigation. Applications are loop closure detections [4] or object-based representations [166]. Approaches generating an implicit SDFs from point clouds, such as [54], produce high-quality results but are computationally intense and, until recently, limited to the extent of single objects. Recent results [183] demonstrate computationally expensive but memory-efficient mapping of large-scale environments.

Representation choice for the proposed approach

As discussed in this section, various environment representations can be utilized to enable accurate and efficient SLAM systems with dense maps. TSDFs directly provide accurate gradients within the truncation distance and are, therefore, well-suited for optimization approaches and robust registrations. Furthermore, the surface reconstruction is robust to noisy sensor integrations. Therefore, we focus on the application of TSDFs in the proposed approach. However, to enable the application of TSDFs for robust large-scale SLAM systems, various open research challenges remain. These cover the investigation of update schemes for 3D, robust and accurate scan registration in rough terrain, and efficient loop closure detection for large-scale SLAM. We propose solution strategies to these challenges in Chapter 5.

2.3 Robotic Exploration

In the exploration of initially unknown areas with mobile ground robots, the robots traverse the environment, aiming to fulfill a predefined exploration goal. Such goals can be versatile, for example, the creation of a 2D or 3D map of the environment, the coverage of the surface or space with a sensor to locate CBRNE threats such as radioactive materials or chemical hazards, or in the search for victims or objects of interest. Depending on the problem focus and scope, various terms are used to denote this problem. Common is the reference as active SLAM, where the SLAM problem is extended by the search for a control policy that aims to minimize the uncertainty in the map [26].

Exploration approaches typically need to solve three subproblems: "*1) Identification of potential actions, 2) Utility computation and 3) Action selection and execution*" [132]. The *identification* phase covers the representation of the progress towards the goal(s). For a spatial exploration, this is often a metric map split into the covered area, the uncovered area, and the frontiers. Potential actions can be derived from this representation, such as driving to a specific frontier. In the *utility computation step*, the identified actions are evaluated for the expected cost and gain, such as the size of the frontier and distance. Finally, in the *action selection and execution* step, a set of actions aiming at maximizing the utility is selected and executed.

For each of the steps, various approaches have been proposed. Comprehensive overviews are provided in the recent surveys by Placed et al. [132], Azpúrua et al. [6] and Ahmed et al. [3]. Based on the surveys, the following sections provide a brief and introductory overview of variations of approaches to solve the subproblems of the exploration problem. This serves to provide a general understanding related to the proposed approach in Chapter 6.

2.3.1 Identification of Potential Actions

The identification of potential actions typically covers two aspects: First, the creation of a map of the environment which is in a second step utilized to identify potential goal locations. Maps (see also Section 2.2.3) can be represented in various ways. Common is the representation as **metric maps**, such as occupancy grids, which allow a representation of the covered area, uncovered area, and the extraction of frontiers. Metric maps can also store other relevant quantities, such as the distance to obstacles, relevant measured quantities, sensor coverage, or semantic information. However, for complex and large-scale environments, the handling of naive metric maps can become computationally infeasible for update and planning. In contrast, **topological maps** represent the environment as a graph, which enables efficient planning even in large-scale environments. However, the generation and maintenance of a topological structure is a non-trivial task. To leverage the advantages of both representations **hybrid and hierarchical maps** combine both approaches. Such as small local metric maps connected by a global topological graph [25] or the combination of metric, semantic, and topological maps as 3D scene graphs [142].

The next step requires the identification of a set of potential actions. Typically, the potential navigation goal positions are identified, and the related actions are the navigation to the identified goal positions. Common are frontier-based approaches [172], which identify reachable positions at the border of known and unknown space. The naive evaluation of the whole map for frontiers becomes computationally infeasible for large-scale environments. Therefore, various approaches aim to find frontiers efficiently by applying incremental update schemes [82] or search-based strategies such as rapidly exploring random trees [159]. Other approaches aim to reduce the uncertainty in the map by identifying potential loop closure locations and adding these as goal positions [150]. Closely related to the frontier-based approaches are coverage-based approaches that identify positions that provide sensor coverage of areas of a model that are not sufficiently observed [25].

2.3.2 Utility Computation

To enable efficient exploration behaviors, it is necessary to not only identify potential actions but also predict the expected utility. Various utility estimates have been proposed, such as considering the Euclidean distance [172] to the goal, the size of the frontier [66], or the expected coverage [25]. A large class of approaches base the utility estimation on information theory and approximate the reduction of entropy of the map and the pose estimate by performing an action [15]. Furthermore, Theory of Optimal Experimental Design (TOED)-based approaches aim to minimize the expected covariance of the posterior state [57].

As the planning step can be computationally expensive and to achieve efficient global behaviors, it can be beneficial to extract not just a single action but a sequence of actions. However, actions are not independent as the execution of an action can influence the utility of other actions, e.g., by overlapping sensor coverages. This can be accounted for by predicting the gain of the selected action and updating the utility for all other action candidates [25].

2.3.3 Action Selection and Execution

Finally, an action or a sequence of actions needs to be selected from a set of identified and evaluated actions and be executed.

When the set of potential actions is discrete, the action with the highest utility can be selected. If a sequence of actions is considered, the interplay with the utility reduction, as outlined in the last section, needs to be accounted for.

Alternatively, methods that do not optimize for separate goal positions but directly optimize a whole trajectory can account more accurately for the submodularity in the problem. Examples are, belief-space planning approaches [7], optimal control formulations [98] or reinforcement learning-based appraoches [154]. However, enabling the applicability of these methods in versatile, complex, and previously unseen real-world environments online on systems with limited computing is still an ongoing research topic.

2.3.4 Relation to Proposed Approach

As we discussed in the last sections various approaches exist to enable efficient exploration of unknown spaces with mobile robots. In the proposed approach for shared-autonomy exploration in Chapter 6, we base the method on a coverage-based goal identification with a hierarchic environment representation [25] to enable the handling of multiple sensors in multi-goal missions. We formulate a modular utility computation and compute dense local and coarse global action sequences to achieve accurate and efficient exploration of large-scale environments.

3 Robots Used in this Thesis

This chapter presents the robots developed and used as evaluation and application platforms for the proposed approaches. We first provide a general overview of the considerations for the robots and then discuss the specific robots in detail.

3.1 Overview

The research and development of integrated assistance capabilities for mobile disaster robots requires suitable robots to integrate and evaluate the methods. However, commercially available platforms typically are either focused on direct teleoperation (e.g., Telerob Telemax) or do not have the required sensors and computing or sufficient open interfaces (e.g., Boston Dynamics Spot). Designing and constructing new robots is a challenging and time-consuming task. Therefore, most systems in this thesis are based on commercial platforms equipped with integrated modules developed at TU Darmstadt. These modules combine the required hardware (sensors, compute, electronics) and software.

In this thesis, mainly, three robots (see Figure 3.1) with varying locomotion, manipulation, and sensing capabilities are considered. The capabilities are summarized in Table 3.1.

The main navigation modules of all robots incorporate a LiDAR, an IMU, an Omni-Camera, and a computing unit. A key design consideration is the externally actuated spinning 3D-Lidar (see Figure 3.2) on Asterix and DRZ Telemax. The used sensor is a Velodyne VLP-16, which provides dense and comprehensive horizontal coverage of $360°$ with $0.1°$ resolution. In contrast, the vertical coverage is only $30°$ with a sparser resolution of $2°$. As the sensor is mounted tilted on a spinning axis, the overall vertical range increases by two times the tilting angle while also enabling dense coverage. On Asterix the mount is tilted by $45°$ yielding $120°$ effective FOV on Telemax the mounting angle is $30°$ yielding $90°$ effective FOV. However, the increased time of $1\,\mathrm{s}$ - $2\,\mathrm{s}$ for the coverage of a full accumulated scan leads to distortion in the scan due to the ego-motion of the robot, similar to the motion distortions in rolling shutter cameras. Therefore, the distortion needs to be considered and compensated in the state estimation and environment modeling. With the introduction of the Ouster

Table 3.1: Robot Capability Overview

	DRZ Telemax	KIARA Telemax	Asterix	eC Scout
Locomotion	4 adjustable tracked flippers		Whole-body tracks, pairwise adjustable flippers	4 wheels, skid steer
Manipulation	6-DoF Arm		6-DoF Arm	*none*
Perception				
Lidar	Actuated ($360°\times90°$)	Static ($360°\times90°$)	Actuated ($360°\times120°$)	Static ($360°\times90°$)
Cameras	RGBD, Tele, Omni, Wide, Thermal	RGBD, Tele, Omni, Thermal	RGBD, Omni, Thermal	RGBD, Tele, Omni, Thermal
Additional	IMU, GNSS, Radiation	IMU, GNSS, Radiation, Hazard Gas	IMU, GNSS, CO_2	IMU, GNSS, Radiation, Radar

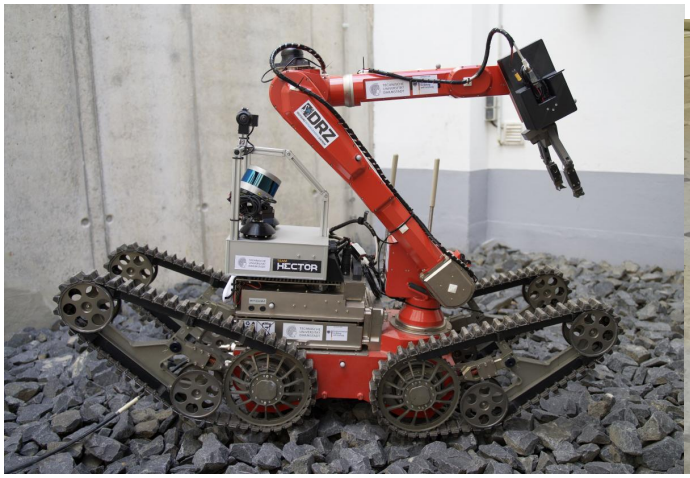

(a) DRZ Telemax

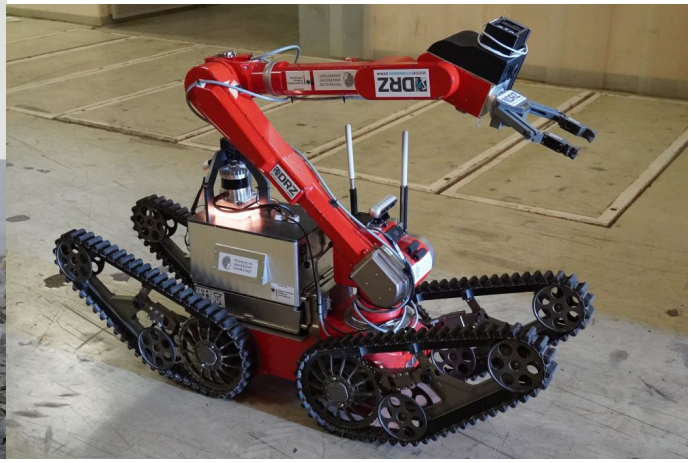

(b) KIARA Telemax

(c) Asterix

(d) emergenCITY Scout

Figure 3.1: Robotics platforms equipped with the developed integrated modules used in the development and evaluation of this thesis. (Images: Team Hector, TU Darmstadt)

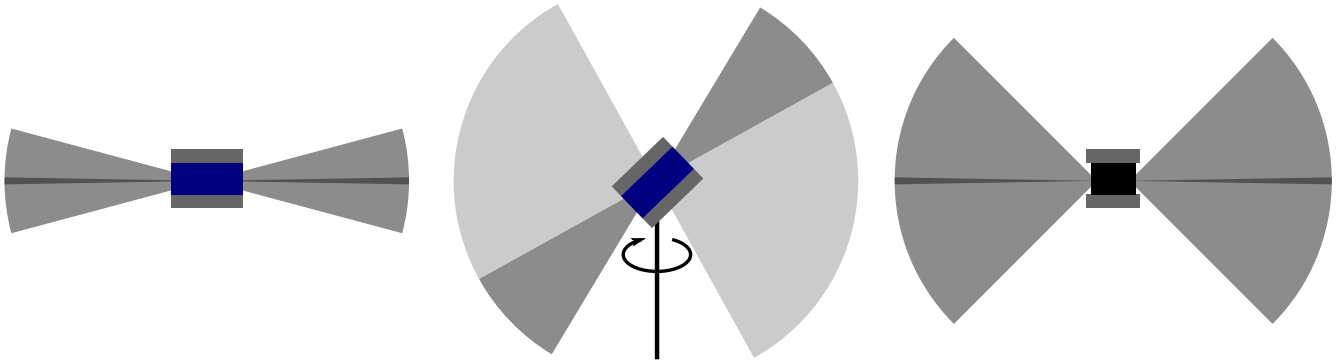

(a) Static Velodyne VLP-16 (b) Actuated Velodyne VLP-16 (c) Static Ouster OS0-128

Figure 3.2: Lidar vertical FOV comparison schematic. a) Static Velodyne VLP-16 with $30°$ opening angle and $2°$ sampling resolution. b) Actuated and tilted mounted Velodyne VLP-16. The coverage increases by two times the tilting angle, yielding an effective $120°$ coverage for a mounting angle of $45°$. c) Static Ouster OS0-128 with $90°$ opening angle and $1.4°$ sampling resolution.

OS0-128, which has a vertical coverage of $90°$ with an angular resolution of $1.4°$, no external actuation is used on KIARA Telmeax and eC Scout. However, even though the capture of a single scan takes only $0.05\,\mathrm{s}$ - $0.1\,\mathrm{s}$, scan distortions are notable for fast motions and need to be compensated.

3.2 DRZ and KIARA Telemax

The *Telerob Telemax* is a remote-controlled mobile search and rescue robot that combines high mobility and a highly accurate and strong manipulation arm manufactured by the "Telerob Gesellschaft für Fernhantierungstechnik mbH". The four independently controllable flippers enable the robot to traverse obstacles in urban environments, such as stairs, steps, or debris. The precise arm has a reach of $1.5\,\mathrm{m}$ and can carry $7.5\,\mathrm{kg}$ at maximum reach. This enables versatile and accurate manipulation capabilities such as opening doors or manipulating valves in real-world environments. The robot is owned by the "German Recue Robotics Center" (DRZ), which aims to encourage the development and transfer of robotic systems to assist in the rescue and protection of people and property. Within the scope of the research projects, "Development of the German Rescue Robotics Center" (A-DRZ) [94] and "AI assistance for robot-assisted reconnaissance and defense against acute radiological hazards" (KIARA) interacted function modules were developed to enable (autonomous) assistance functions and mapping capabilities in disaster environments.

The DRZ Version was continuously developed and improved in the period of 2018-2022, mainly by Marius Schnaubelt, with conceptual contributions of the author, and features two integrated modules. A main perception and navigation module is on the back side of the robot, with a tilted, spinning Velodyne VLP-16 PUCK 3D-Lidar, closely above a $360°$ Omnidirectional Camera. Inside the module is an XSENS MTi 100 inertial sensor and two compute units. LiDAR and IMU are the main sensors used for localization and mapping. The omnidirectional camera enables assistance functions with colored point clouds and supporting camera projections [124]. At the end-effector, the robot is equipped with a module to support inspection and manipulation containing multiple cameras (RGB-D, tele, wide-angle, thermal) and a low-power computer for processing the data. Additionally, the robot is equipped with a GNSS sensor at the elbow to enable a pose estimate with global reference in outdoor environments and a front-facing RGB-D camera for obstacle avoidance.

The KIARA Version was continuously developed and improved in the period of 2022-2023, mainly by the

author, Jonas Süß, and the TU Darmstadt Feinmechanik Werkstatt. It also features two integrated modules. The overall concept is similar to the DRZ Version with updated hardware components. The increased angular FOV and resolution of the Ouster OS0-128 resolves the need for the actuation of the sensor, allowing a more compact construction.

3.3 Asterix

Asterix is a mobile rescue robot designed to combine good sensor coverage, high mobility, and manipulation capabilities, enabling versatile autonomous assistance functions [143]. The movable flippers can be used to climb obstacles and stairs. The low center of mass enables the traversal of complex and uneven terrain. The compact arm enables basic manipulation tasks such as closing valves or opening doors in simple environments. The robot is used by *Team Hector*[1] for rescue robotics research and participation in the RoboCup Rescue Robot League [79].

In large parts, the robot follows a functional setup similar to *DRZ Telmax*, with a spinning 3D LiDAR, closely located IMU, omnidirectional camera, and RGB-D cameras. On the back side, the robot is equipped with a small pan-tilt sensor head with RGB-D and a thermal camera to provide sensor coverage even in narrow environments.

3.4 emergenCITY Scout

emergenCITY Scout is a wheeled mobile robot for research on assistance functions as part of the emgergenCITY research center [73]. The design is focused on the fast exploration of urban environments. The wheels enable faster locomotion than the tracked vehicles. However, the mobility is limited to simple steps, and the robot does not have manipulation capabilities.

The sensor setup follows a function setup similar to *KIARA Telemax*, with a static 3D lidar, IMU, omnidirectional camera, and RGB-D cameras. In the front, the robot is equipped with a large pan-tilt head with a thermal camera, zoom camera, and through-wall radar. The sensor-head enables various inspection and detection tasks, including the vital sign detection of people behind closed doors [145].

[1] `https://www.teamhector.de/`, accessed 16.01.2024

4 Requirements, Guidelines, and Challenges for Autonomous Capabilities and Assistance Functions for Ground Rescue Robots in Reconnaissance Missions

While rescue robots are becoming more established as part of disaster response, they are typically teleoperated in actual disasters. (Autonomous) assistance functions can improve performance, extend functionality, and reduce operator overload. It is necessary to understand relevant requirements to ensure that developed capabilities apply to real-world needs. Previous analyses focused on general aspects of rescue robots, leaving a gap in understanding requirements for (autonomous) assistance functions, which this chapter addresses.

This chapter provides a detailed, evidence-driven analysis of application requirements and guidelines, and research challenges for (autonomous) assistance functions for rescue robots in reconnaissance missions. The chapter is structured as follows: We first outline the related work and the key contributions of the approach, then we provide an overview of the proposed method, and we first investigate past deployments before investigating general and then capability-specific requirements and challenges.

As the case studies in Section 4.4 show, the deployment conditions, goals, and robotic capabilities are highly diverse. This also means that specific requirements for different robotic capabilities in diverse environments with different use cases strongly vary. Therefore, we provide a general understanding of the requirements on a functional level instead of quality requirements or attributes. Furthermore, we address multiple aspects by defining guidelines that cover concepts derived from successful deployments and studies but are less strict than the requirements.

This chapter is an extended and revised version of [41]. Insights from and concepts for training and mission deployments were published in [94] and [153].

4.1 Related Work

4.1.1 Rescue Robotics and Requirements

A comprehensive overview of the application, experiences, and requirements for rescue robotics is provided in [116] and [119]. The authors define classifications and characteristics of disasters and disaster robots, cover an overview of past robot deployments (until 2014), and provide insights on lessons learned, challenges, and practical guidelines. The authors emphasize the high maturity level required to operate successfully in the challenging conditions of actual disasters. They point out that with the given "operation envelopes" (operation conditions in actual disaster), robot control is very challenging and error-prone due to high cognitive load, perceptual impoverishment, and the choice of user interfaces and system designs. [118] focuses on interfaces for human-robot interaction, splitting the users into end-users, developers, and stakeholders / general public, and defines respective requirements.

The International Forum to Advance First Responder Innovation (IFAFRI) provides an analysis of ten

common global capability gaps for first responders[1]. Key gaps related to this work are "Remote Acquisition of Information," "Remote Operations," and "Actionable Intelligence." The authors mention the requirements of easy operation, operation in multiple environments, economic aspects, integration of data with response systems, and real-time visualization, analysis, extrapolation, and contextualization of data.

Delmerico et al. [45] survey the current state of the art in rescue robot systems and human-robot control interfaces, with statements of emergency response stakeholders and an assessment of technologies for future real-world disaster response and recovery. The authors consider robots as tools deployed by responders, which brings a benefit when they can be used to improve efficiency, perform a task that humans cannot perform (or only with intensive training, such as scuba diving), or allow them to operate remotely in high-risk environments. Regarding autonomy, it is emphasized that a human in the loop (a semi-autonomous mode in contrast to full autonomy) is strongly desired.

Doroftei, Matos, and Cubber [49] define requirements for rescue robots focusing on mechanical and structural aspects for teleoperated systems in realistic operation conditions. The requirements are gathered within the scope of the ICARUS project, which focuses on robots for data gathering. Schneider and Wildermuth [144] define various use cases for rescue robots and identify general requirements for robotic systems. The authors emphasize that robots should protect and expand the capability spectrum and should not introduce new potential hazards.

These works provide valuable insights into general requirements for rescue robots, which we build upon and extend for assistance functions and autonomy.

4.1.2 Standards and Norms

The IEEE Standard for Transparency of Autonomous Systems 7001 [169] outlines the need for transparency for autonomous systems for adjusting expectations, building confidence as a part of verification and validation, and ensuring accountability. The standard introduces a framework to measure and test transparency in six levels for five stakeholder groups: users, the general public and bystanders, safety certification agencies, incident/accident investigators, and lawyers/expert witnesses.

Resulting from the project for the "Establishment of the German Rescue Robotics Center (A-DRZ)" [94] a German consortium standard for the application of robots in hazardous applications [47] was introduced. The section on autonomy and assistance functions covers the aspects of transparency (by IEEE 7001), robustness to disturbances, safe operation, monitoring, and communication.

4.1.3 Technology Acceptance

Understanding why new technology is used and accepted is an established research field in cognitive science and information system theory.

An early established model is the Technology Acceptance Model (TAM) [42] by Davis. Based on the social-psychological "Theory of Reasoned Action," it postulates that the *Attitude Towards Using* a technology, is mainly influenced by its *Perceived Usefulness* and its *Perceived Ease of Use*. Based on the Attitude and Perceived Usefulness, an *Behavioral Intention to Use* is formed, leading to actual system use. In the subsequent years, more models were postulated, considering additional aspects.

A condensed model was created by Venkatesh et al. as the Unified Theory of Acceptance and Use of Technology (UTAUT) [161] (see Fig. 4.1) to explain and predict the acceptance of technology in organizational contexts.

UTAUT also assumes that system use is determined by a behavioral intention but explicitly models *facilitating conditions* as a dependency. It directly considers multiple predictors to form the behavioral intention:

[1] `https://www.internationalresponderforum.org/services/capability-gaps`

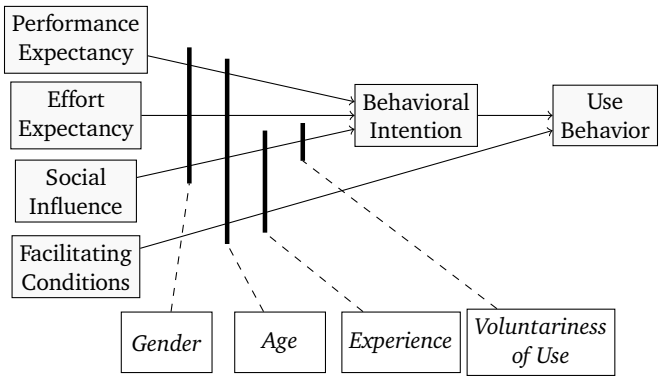

Figure 4.1: The ”Unified Theory of Acceptance and Use of Technology” (UTAUT) model predicts the acceptance and use of technology based on Performance Expectancy, Effort Expectancy, Social Influence and Facilitating Conditions moderated by Gender, Age, Experience and Voluntariness of Use. Image is based on [161].

- Performance expectancy: "the degree to which an individual believes that using the system will help him or her to attain gains in job performance" [161]
- Effort expectancy: "the degree of ease associated with the use of the system" [161]
- Social Influence: "the degree to which an individual perceives that important others believe he or she should use the new system" [161]

The effect of the predictors is moderated by the aspects of age, gender, experience, and voluntarines of use. As the model has been successfully applied to explain the acceptance and application of technology in various domains, we will apply it to the context of rescue robotics and use it as a basis for deriving requirements.

4.2 Contribution

While establishing rescue robots in disaster response is an ongoing process, multiple works identify primarily general requirements for rescue robots [116, 119, 45, 49, 144]. However, these analyses focus on general aspects of rescue robots, leaving a gap in understanding the full range of specific requirements for (autonomous) assistance functions. In this work, we address this gap and complement these analyses by providing a detailed, evidence-driven analysis of application requirements and guidelines, and research challenges for (autonomous) assistance functions for rescue robots in reconnaissance missions.

We base our reasoning on a general model for technology acceptance and propose to consider assistance functions together with their surrounding conditions holistically as *integrated function capabilities*. We define We define an integrated function capability as a capability that executes a specific function, implemented through software together with the necessary physical components (e.g., sensors, actuators, power sources, computing hardware) integrated into the robotic system to perform the function effectively.

We relate the requirements to the current state of the art and address challenges, limitations, needs, and recommendations to enable more functionality and efficiency in mobile rescue robotics. The key contributions are:

- Evidence-driven analysis of requirements and challenges for autonomy and assistance functions for rescue robots in reconnaissance missions

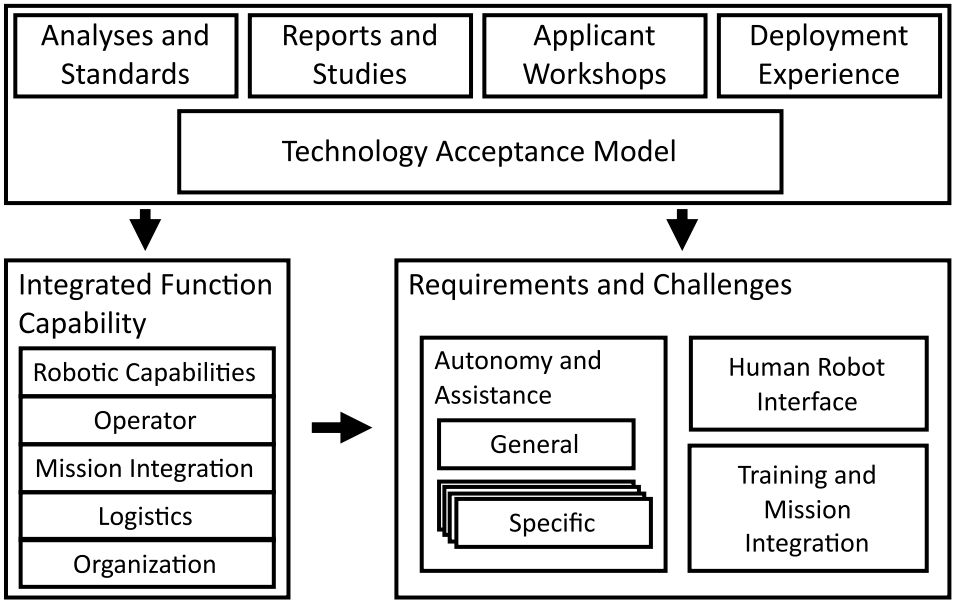

Figure 4.2: We base our analysis on a comprehensive model for technology acceptance and consider reports of past deployments, related analyses, our own deployment experience, and insights from workshops with first responders. We define relevant aspects of an integrated function capability and identify requirements and challenges.

- Analysis of rescue robot deployments on requirements and challenges for autonomy and assistance functions
- Recommendations for research, development, transfer, and application of (autonomous) assistance functions.

4.3 Overview

For our analysis (see Figure 4.2), we take into account:

- A condensed review of related analyses and standards
- A derivation of requirements from reports of past deployments
- Results from requirements workshops with first responders as part of the research projects "A-DRZ" [94] and "KIARA"[2]
- Insights and experience from the author as ”expert operators” and observers in the deployments as part of a trial German Robotic Task Force [153]

4.4 Case Studies

To understand the needs of first responders and the conditions in actual disasters, we investigate three different deployments: World Trade Center Collapse, Fukushima Daiichi Response, and Residential Complex Fire Essen.

[2] https://www.tu-darmstadt.de/kiara

The deployments were chosen based on the availability of sufficiently detailed reports and variances in the scenarios, conditions, and tasks.

4.4.1 Case Study: World Trade Center Collapse, 2001

On September 11, 2001, terrorists crashed two airplanes into the World Trade Center Twin Towers in New York, United States of America. The damage from the impact of the airplanes and the subsequent fires lead to the collapse of both skyscrapers within less than two hours.

The reaction to this incident is the first known application of robots for urban search and rescue [117]. Based on the report of Murphy [117], the Center for Robot-Assisted Search and Rescue utilized multiple rescue robots from September 11 until October 2, 2001, for searching for victims, searching excavation paths through the rubble, structural inspection and detection of hazardous material. In this period, the robots were on four deployments inspecting eight voids. A major challenge was the exploration of void spaces in the rubble pile. At this time, the default options were to use a camera on up to $2\,\mathrm{m}$ long pole or to send humans or dogs inside. In contrast to the pole, robots could go up to $20\,\mathrm{m}$ into the voids and, in contrast to humans or dogs, could also explore places with ongoing fires or risks of structural collapse.

The authors report on multiple challenges for the application of robots. A major difficulty was the lack of depth perception in all three dimensions in a heavily confined environment. This led to misjudgments of obstacles in front of the robot and the robot rolling over, as well as difficulties in judging the clearances to the sides and the top, and led to the robot getting stuck. Direct usage of odometry readings was not helpful as the readings were erroneous due to the rough terrain.

The post-analysis of camera footage revealed various robot operation errors and missed remains by the operators. The authors hypothesize that unaided human perception may not be sufficient for successful robot search and rescue.

4.4.2 Case Study: Fukushima Daiichi Response, 2011

On March 11, 2011, a devastating earthquake and subsequent tsunami hit eastern Japan, leading to a nuclear accident at the Fukushima Daiichi Nuclear Power Plant, resulting in very high radiation levels in and around the disaster site.

Following the report of Yoshida et al. [177], an iRobot Packbot was deployed on April 17, 2011, to perform a reconnaissance mission which confirmed very high levels of radiation in the building making human access impracticable. However, the system's operational readiness was strongly limited as the radio-based communication system only covered parts of the first floor, and the robot could not climb stairs. To perform further explorations, a joint research group from Tohoku University, Chiba Institute of Technology, and Kogakuin University built a new robot focused on reliability for hardware, communication, and radiation and the capability to perform multiple measurements. Another focus was the development of a system for easy teleoperation, as the operation had to be performed by a novice operator from the operator of the plant. From June 24, 2011, to October 20, 2011, the robot was deployed for six missions, performing multiple tasks such as radiation and temperature measurements, collecting air samples, investigating damaged piping by taking images, and installing water measurement gauges.

The authors report multiple challenges to the robot operation. The most severe limitation was a failure in the usage of a communication cable. The cable got jammed during the sixth mission, leading to the mission's failure and the loss of the robot. Furthermore, the terrain was rough, covering various obstacles such as steps, debris, and cables. Some environment conditions were not covered by tests, such as strong changes in illumination and high air temperature, which led to an overheating of motor drivers, almost leading to

the failure of the second mission. Furthermore, requirements such as additional sensors only came up after performing the first missions and had the be adopted at the site.

Regarding team integration, it was crucial to allow an easy operation for novice operators. The mission planning was challenging as there were differences between construction plans and the actual site (e.g., width of stairs).

4.4.3 Case Study: Residential Complex Fire Essen, 2022

On February 22, 2022, a residential apartment building in Essen, Germany, experienced a severe fire. As a result, 39 apartments on four floors were burned, and an entry ban was imposed for parts of the building. To assess the situation and gather information for the investigation of the fire's cause, the German Center for Rescue Robotics (DRZ)[3] deployed ground and air robots to create a 3D model and capture images of the interior.

Following the report by Surmann et al. [153] after an initial exploration with small unmanned aerial vehicles (UAVs), the terrain was evaluated to be suited for applying ground robots. A ground robot was deployed for three inspections the following day. For the first two missions on ground level, a person had visual contact with the robot and communicated with the operator. This was not feasible for the third mission at the upper level due to the entry ban. Therefore, a UAV was deployed to support the operator with visual support from the outside.

The authors report on challenging terrain conditions, with very narrow passages and loose debris on the ground. Visualizing registered point clouds with the 3D robot model helped the operator navigate through narrow environments. The usefulness of this assistance was limited by the availability of the WiFi network, which was only available in sufficient quality for an estimated $70\,\%$ of the mission time. The operator did not deploy more complex assistance functions (like waypoint control) due to concerns about the control system's performance in loose ground. The authors also emphasize the role of joint training of the academic staff with first responders for successful and efficient missions.

Overall, the authors conclude that rescue robots must provide robust functionalities that can be quickly deployed. These functionalities need to be exercised in joint training with first responders for value in disasters.

4.4.4 Case Studies Conclusion

The case studies show that perception and navigation are key capabilities described as difficult for the operators and motivate respective assistance functions. The challenging environments define implicit technical requirements for assistance functions to be able to handle the effects of robots operating in such conditions. Furthermore, the scenarios indicate the strong variance in potential scenarios, tasks, and environments in disasters, often unknown until the mission starts, inducing further challenges, especially for autonomous functions. Overall, assistance functions need to be integrated into an intuitive and easy-to-use/-learn user interface and trained together with first responders.

4.5 Analysis of Requirements, Guidelines and Challenges

In this section, we define application requirements and guidelines and identify research challenges for (autonomous) assistance functions as part of robotic capabilities for reconnaissance missions in disasters. We first discuss technology acceptance in the context of the UTAUT model and *integrated function capabilities*. Subsequently, we discuss human-robot interfaces, analyze assistance functions in general, and investigate specific assistance functions in detail.

[3]The first author was part of the ground robot response team.

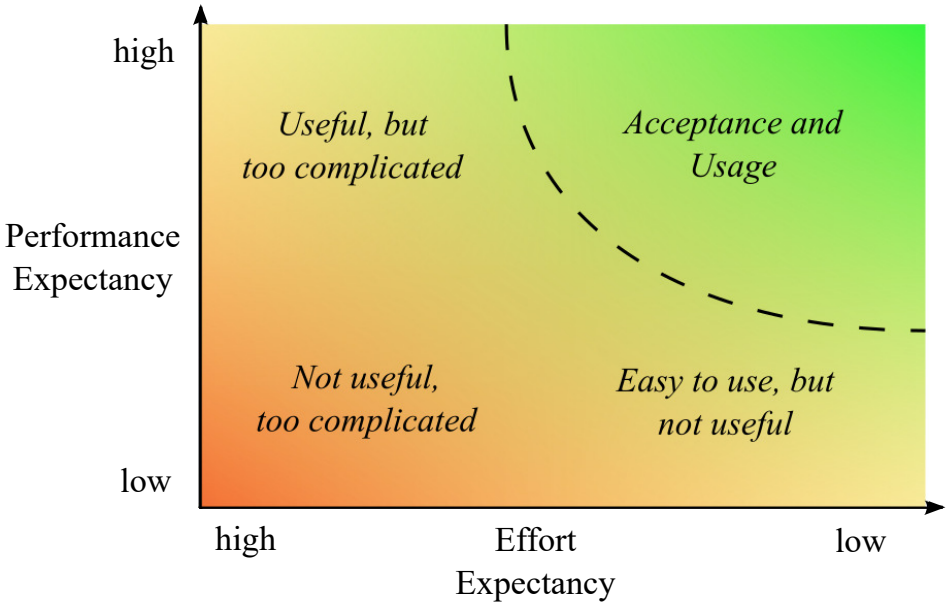

Figure 4.3: To fulfill the concept of effort expectancy of the UTAUT model, the application capability needs to be perceived as easy to learn and use. An assistance function will only be applied if it is expected to bring a benefit and is sufficiently easy to apply.

4.5.1 Technology Acceptance and Requirements for Integrated Robotic Capabilities in Disasters

Following the UTAUT model (see Section 4.1.3), central indicators for the acceptance and application of new technology are *Performance Expectancy (PE)*, *Effort Expectancy (EE)*, *Social Influence (SI)* and *Facilitating Conditions (FC)*. These aspects align well with the statements in the expert interviews in [45], where it is stated that robots need to provide capabilities better than established tools (PE) and should be simple with low training requirements (EE/FC).

PE and EE can be directly controlled by the design and communication of the robot and its capabilities (see Fig. 4.3), whereas SI and FC are environment conditions that need to be facilitated in communication and when introducing technology to first responders.

We consider rescue robots to be deployed as tools for first responders [45] that provide a set of capabilities relevant to the disaster response. To fulfill the aspects of PE and EE, this tool needs to provide a benefit compared to existing and established solutions. For example, in the deployment of the World Trade Center Collapse, robots provided a significant improvement in PE in comparison to a camera on a stick by being able to explore up to $20\,\mathrm{m}$ instead of $2\,\mathrm{m}$ into a void. As robot and assistance functions are deployed as part of a disaster response in a disaster response team and to include the aspects of the UTAUT model, we propose to consider the robot with its assistance functions as an *integrated function capability* (see Fig. 4.4), for which the following aspects need to be considered holistically:

- Robotic capabilities - the robot operating in the disaster site and providing capabilities ("operational envelope"[116]) (PE/EE)
- Operator - the operator controlling the robot with a human-robot interface at the control position ("operator's environment"[116]) (PE/EE)
- Mission integration - the integration in mission procedures with bidirectional communication of goals, states, and results to the operator and response team/organization (PE/EE/FC)

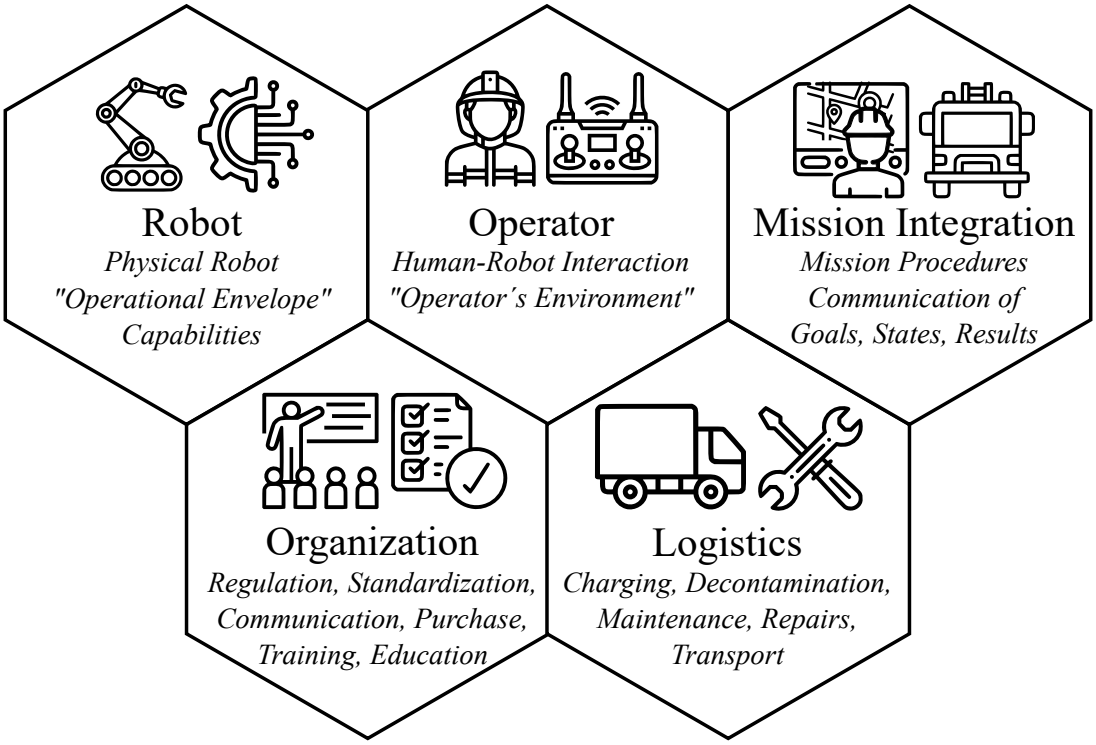

Figure 4.4: We propose to consider the robot with its assistance functions as an *Integrated Function Capability*, covering robot-, operator-, mission integration-, logistics- and organization-related aspects.

- Logistics - logistics of bringing robot and equipment to the disaster area and points of ingress and egress; maintenance; charging; decontamination and repairs (PE/EE/FC)
- Organization - purchase; regulation/standardization; training; communication and education (EE/SI/FC).

As we show in the following sections, robotic capabilities, operator, and mission integration provide various open research challenges. We see the role/potential (and ongoing success) in academic research to develop concepts, methods, and solutions as demonstrators with a medium Technology Readiness Level (TRL). Although the structures worldwide vary, we generally consider the hardening of these solutions to products (with relevant development, production, distribution, and support structures) at a higher TRL and the definition of logistics and organization structures rather suited for application-oriented organizations and companies, which emphasizes the need for cooperation of academic organizations, application-oriented organizations, companies, and first responders.

4.5.2 Human Robot Interaction, Training, and Mission Integration

An essential aspect for acceptance and performance is the interaction between robot and operator with a *human-robot interface (HRI)* [137]. The HRI should enable the operator to operate with the robot as a tool, providing the operator with good situation awareness and efficient robot control. A sufficient level of autonomous assistance can enable the operator and robot to perform like interdependent teammates [80]. There are in-depth analyses for requirements and limitations of HRI [118, 137]. [137] emphasizes a focus on user-driven design, which should be part of the whole development process, not just evaluation at the end. The HRI should reduce cognitive load for building and maintaining situation awareness. As robots often can be equipped with more sensors than an operator can simultaneously monitor, it is important to

provide the right information at the right time at the right level of abstraction instead of providing too much simultaneous information ("Add another camera syndrome" [116]) [137]. Rescue robots can be very complex to control, e.g., driving, arm movements, flippers, movable cameras..., and the HRI can help to simplify control complexity by reducing the control space the operator needs to think about [137]. Especially in the context of (semi-)autonomous assistance functions, the requirements of coactive design [80] observability, predictability, and directability are essential.

We consider the user-centered development of HRIs and the design of HRIs that enable synergies between the complementary capabilities of operator and robot as key challenges to enable the potential for efficient support of rescue robots with intelligent assistance functions in disasters.

Another aspect of human-robot interaction is learning to use the robot. Although an intuitive HRI can help to reduce the amount of training necessary, there is still the need for training for operators on learning how to control the robot and for training on general procedures to integrate the robot into disaster response missions. From the feedback of first responders, we learned that time and options for training are strongly limited due to a generally high work-load for first responders and a limited amount of available robots.

This motivates the need for efficient *training and education* concepts. To address the limit of available robots and reduce the effort for training, virtual training environments provide high potential [1]. Simulations have reached a high quality of realism (e.g., NVIDIA Isaac Sim[4]) and have become an established method in developing assistance functions. We consider concepts with empirical validation for efficient initial and regular training of first responders on robots with assistance functions as an open research question. Closely related is the investigation of metrics for the assessment of operator proficiency. An established framework for bench-marking the performance of operators and rescue robots are capability-focused test methods by National Institute of Standards and Technology (NIST) [79], which are also used in the RoboCup Rescue Robot League. For single capabilities, such as localization and mapping, complex data sets and evaluation methods are available [182]. However, we identify a need for performance measures for new autonomous functions, multi-modal perception and mapping tasks, and the consideration of mission integration and, therefore consider these as open research challenges.

Lastly, there is the aspect of *mission integration and procedures*. [116] covers guidelines on conducting fieldwork for disaster response and the organization of robots and team members. However, there is limited work on procedures for efficiently communicating goals and tasks from mission control to the operator/robot and status and results (maps, detections, measurements, ...) back to mission control in a meaningful and useful way. Furthermore, the robot might share its operation space with first responders. Therefore, it needs to indicate the robot's status (e.g., sound or light signals) to people operating in its vicinity.

4.5.3 General Requirements and Challenges for Assistance Functions and Autonomous Robot Capabilities

Assistance functions can enable additional functional capabilities, for instance, the automatic creation of maps and 3D models, thereby increasing the performance expectancy. Furthermore, assistance functions can help make the complex task of remote robot operation simpler and safer, for instance, by increasing situational awareness or by providing collision avoidance assistance as well as highly, thereby increasing the performance expectancy and reducing the effort expectancy.

We differ between *perception assistance*, which can support the operator and mission without exercising direct and active control of the robot, such as the creation of a map, and *(autonomous) control assistance*, which actively controls and moves the robot and thereby supports robot operation and provides (semi-)autonomous capabilities. In this section, we define requirements and challenges that generally apply to assistance functions,

[4] `https://developer.nvidia.com/isaac-sim`

which we extend for perception assistance in the next section and subsequently for (autonomous) control assistance.

Following [169], *transparency* is an integral requirement for adjusting expectations, building confidence for verification and validation, and ensuring accountability. In general terms, transparency means that it is always possible to understand why and how the robot behaves in a particular way presented in a meaningful level of abstraction. Furthermore, the operator should be able to understand the system's status and the principles by which decisions are made. This covers "*relevant goals; progress to those goals; models of its past, current and potential future environmental context (from sensors and other information)"; and relevant information about its current performance, such as reliability and error messages*" [169]. We also consider transparency a crucial aspect in expectation management to asses under which conditions assistance methods are expected to work and what limitations are.

Nevertheless, due to disasters' non-repetitive and chaotic nature, tools must be expected to be used in untested environments or conditions. This emphasizes the need for *robustness* and *modularity*. Both, on the level of a single function and the overall system. Assistance functions should be designed to work as well as possible with other components failing. For example, we consider the HeRO state estimation [112], which fuses multiple odometry estimates and achieves robustness through redundancy.

Compute on mobile platforms is limited, and the need for computing directly affects the size and power requirements of the physical systems. Therefore, we consider it imperative to aim for *computationally efficient* solutions.

We consider it a good practice to develop algorithms in an *environment agnostic* way, such that no to little manual scenario- or environment-specific parameter adaptation is necessary. KISS-ICP [163], which allows robust odometry with various sensors in various environments with the same parameters, is a positive example.

To provide value, the robot needs to perform the desired task at sufficient *speed*. From the statements by first responders, this was considered strongly task-specific. Typically, a reduction in speed compared to human responders was accepted when the task would impose a high risk to the responder, and no potential victim was assumed to be in imminent danger. Furthermore, the robots are expected to act *safely*, avoiding actions inducing unnecessary risks. While taking actions involving risks might be necessary, e.g., traversing an obstacle, this risk should be transparent to the operator.

As disaster scenarios are very complex, dynamic, and non-repetitive, with high levels of uncertainty, they are very challenging for high levels of autonomy. Human Operators can have extensive experience in disaster response and a good semantic understanding of the environment. However, they are not able to process all data at once, might have additional information (e.g., reports about the last known position of a missing worker) not directly available to the robot, and might have skills such as operating the robot to safely traverse unstructured obstacles or perform complex manipulation tasks highly accurate. The need for approaches for combining robot and operator skills was addressed in multiple workshops with different applicants. Therefore, we consider the efficient *combination of the complementary operator and assistance functions*, especially in the context of autonomous assistance functions, as an essential research challenge for developing efficient and useful rescue robots. This is especially relevant for complex behaviors such as (semi-)autonomous exploration or manipulation and closely related to the concept of coactive design [80]. Current research approaches often focus on high levels of autonomy, and there is a lack of focus on the efficient incorporation of operator capabilities.

4.5.4 Requirements and Challenges for Perception Assistance

We consider *Perception Assistance* as background capabilities, which can support the disaster response without exercising direct and active control of the robot. Such capabilities can support the operator in controlling the robot by providing supporting visualizations, guidance, or alerts. Furthermore, they can also support

the operator in reaching the mission goal safely and efficiently by providing mission-related information (robot pose, map, trajectory, points of interest). The provision of information (E.g., image, (live) video, (live) measurement, map) during the mission can be of direct value to first responders, enable decision-making for the next steps, and serve as part of the mission documentation. Additionally, high-fidelity maps and 3D models available during and after the mission can help with damage assessment or further planning of the general disaster response and recovery.

Real-Time Feedback and Visualization: The capability to provide real-time feedback and visualization of relevant, potentially processed, and fused sensor data to assist the operator in remote robot control. The requirements for human-robot interfaces (see Section 4.5.2) apply. Furthermore, communication aspects (bandwidth, latency) need to be considered to decide which data modalities (at which quality/resolution and frequency) are transmitted.

Status Monitoring: The capability to monitor the status of the robotic system and provide relevant information, guidance and alerts to the operator. This covers the state of the power supply/ battery, health of sensors, sensor data, actuators, communication, algorithms, and environment conditions [47], as well as monitoring (autonomous) assistance functions. An open research challenge is detecting if the required environment conditions of assistance functions are not met anymore. If needed, the methods should request a reduction in autonomy level or the operator to take over control [47].

Contextual Guidance and Alerts: The capability to provide contextual guidance and alerts based on detections in the environment, e.g., victims, hazards, thermal anomalies, relevant objects, or structures. As conditions in a disaster usually differ from conditions in training, detections should be linked to the data leading to detection (e.g., the assistance method detects a victim and adds an entry to a map, the operator should be provided with functions to see the data that led to the detection). The operator should be notified about the detections, if necessary, be provided with relevant guidance, and be able to discard or add their own detections.

Control/Navigation Guidance and Alerts: The capability to provide guidance and alerts for navigation and controlling the robot. E.g., automated alerts when a control input would lead (potentially/close) to an unstable state or a collision with the robot itself or the environment. As this might sometimes be intended, the operator should be able to suppress the alerts. Depending on the importance of the alert (e.g., imminent collision), this can also adhere to a stop or limit of control until cleared or suppressed by the operator. Guidance functions can help, e.g., by providing recommendations for trajectories or optimized robot configurations, e.g., for safe traversal of obstacles. We consider the stability assessment and pose guidance in disaster environments, considering both geometric and semantic aspects, an open research challenge.

State Estimation / SLAM: The capability to estimate the robot's pose. This is often tightly coupled to the simultaneous creation of a map of an unknown environment and the robot's location in the map (Simultaneous Localization and Mapping, SLAM). Such methods need to be robust to deal with the challenging environmental characteristics of disaster scenarios. Typically, GNSS is not or is only partially available. The operation in rough terrain and the traversal of obstacles can induce fast and abrupt rotational motions. Parts of the environment can be dynamic with other response actors operating in the vicinity. The scale of the environment can vary from narrow and small-scale environments to wide and open areas. Also, the overall scale can be large, inducing the need to account for loop closures. Illumination can change quickly and intensely, from well-lit rooms to dark rooms with no light and potentially dark walls. Dust, fog, smoke, dirt, and mud can obstruct sensors. A comprehensive study on open challenges is provided in [20]. For this work, we consider the robust fusion of multiple sensor modalities, the robustness against failures in data alignments and hardware failures, and the reduction of manual, environment-specific parameter tuning as highly relevant challenges.

Mapping: The capability to create a 2D/3D map or model of an environment. The map can cover single or multiple fused modalities. Maps should generally be easy to understand and interpret and consistent with domain-specific established guidelines (e.g., units, scale, representation, format). The coherence to established

formats is especially crucial for further usage and transfer. The shown quantities should be transparent, e.g., make model assumptions explicit, show measured values, and quantify uncertainties. Depending on the mapped modality and the intended utilization, requirements for accuracy vary. To ensure the success of the data captures for model generations that are not real-time capable, a lower quality preview provided in real-time can help to ensure that relevant data was captured.

4.5.5 Requirements and Challenges for (Autonomous) Control Assistance

We consider *control assistance* as capabilities that can help simplify the control of the robot, reduce the operator's cognitive load, reduce the amount of training needed, and improve safety, thereby improving the chance for a successful operation during a mission. Furthermore, they can enhance the operation space (e.g., areas with poor or no radio connection) or complex motions (e.g., traversal of an obstacle with simultaneous control of multiple joints), which would not be possible with direct teleoperation.

Locomotion Assistance The capability provides locomotion assistance by ensuring the safe locomotion of the robot given a continuous control input by the operator. While the operator provides a command goal (e.g., a direction by deflection of a joystick), the assistance function moves the robot in the general direction but modifies it locally to improve the trajectory (e.g., avoiding obstacles or instabilities or optimizing the robot configuration). Deviations from the command should be visualized to the operator and be able to be overwritten by the operator. We consider stability assessment and optimal trajectory design in disaster environments, taking into account both geometric and semantic aspects, as an open research challenge.

Navigation Assistance The capability to autonomously drive to a (or multiple) navigation goal(s) provided by the operator (e.g., a waypoint in a map), ensuring traversability along the path and avoiding obstacles. The planned trajectory should be visualized to the operator and be modifiable. Similar challenges as for locomotion assistance apply.

(Supervised-)Autonomous Exploration The capability to autonomously explore an area with a goal defined by the operator. The goal can vary depending on the use case and cover various aspects, such as creating a complete geometric map or searching for potential victims, dangerous goods, or heat sources. Termination conditions (e.g., quality, distance, area) and technical limits (e.g., min. battery, distance, time) need to be defined. As the robot might move out of areas with radio coverage, handling for connection losses needs to be considered, such as a return-to-home functionality. The current exploration goal (and potential sub-goals), the queue of considered goals, and the planned trajectory should be visualized for the operator. The operator should be able to add, modify, and prioritize the goals. [20] identifies the fast and accurate prediction of future state and performance guarantees as open research questions. As detailed in Sec. 4.5.3, we consider synergies between assistance functions and operator a key challenge.

Manipulation Assistance The capability to assist with manipulation tasks such as opening a door or taking a sample. Manipulation can be very versatile, with various levels of assistance, from moving the arm close to an object where the operator performs the object interaction to full autonomous task execution. As manipulation potentially involves direct interaction with the environment (for touching objects), a close feedback loop with the operator is necessary, such as a visualization of the planned motion, which needs to be confirmed by the operator or suited HRIs providing force feedback. Technical details, such as inverse kinematics or singularities, should be hidden from the operator. As detailed in Sec. 4.5.3, we consider synergies between assistance functions and operators a key challenge.

5 Robust Simultaneous Localization and Mapping in Challenging Environments

To perform missions within unknown, degraded, and GNSS-denied environments, autonomous mobile rescue robots need to localize themselves in such environments and create a map of it using a simultaneous localization and mapping (SLAM) approach. The capability to create accurate maps and precisely locate the robot's pose in the map are key prerequisites for many higher-level autonomous functions, such as navigation or exploration, while also directly providing operator assistance. During search and rescue missions, the motion characteristics can be highly challenging. For example, traversing uneven ground or obstacles induces aggressive roll-pitch motions and poor odometry estimates as tracks/wheels slip. Environments are typically unstructured and might contain narrow indoor transits in addition to wide, open outdoor spaces with translucent vegetation. Smoke, dust, and fog can further impair the data quality of visual sensors such as cameras or LiDAR. These characteristics make both state estimation and mapping highly challenging.

In this chapter, we introduce a novel continuous-time-based approach for robust and efficient SLAM in challenging terrain to address the key challenges of robustly registering scans, accurately estimating the trajectory in rough terrain, and consistently mapping large space environments online on a mobile rescue robot system.

The chapter is structured as follows. We first outline the related work and the key contributions of the approach. Then, we provide an overview of the proposed method before investigating the wheel-inertial odometry, lidar-inertial odometry, and the large-scale SLAM with pose graph back-end and loop closure detection. Finally, we transfer and extend the proposed approach for the application of radar in visually degraded environments. An overview of experiments, applications, and results is provided in Chapter 7.

This chapter extends, revises, and combines multiple publications. The wheel-inertial and lidar-inertial odometry (Section 5.3 to Section 5.5) have been published in [39], the large-scale SLAM approach (Section 5.6) has been published in [37, 39] and the extension for visually degraded environments (Section 5.7) summarizes [157]. An earlier version for lidar-inertial odometry and the large-scale SLAM have been published in the author's Master Thesis [36]. Therefore, the formulation of the large-scale SLAM (Section 5.6) is in parts also based on [36].

5.1 Related Work

This section discusses the related work with respect to the aspects relevant to the proposed approach. A general introduction and background on SLAM is provided in Section 2.2. The SLAM problem covers two aspects – estimating the robot's trajectory and mapping the environment. This requires suitable techniques to model the trajectory and the environment, which influence the choice of methods to optimize the trajectory and the map.

Discrete- and Continuous-Time SLAM Discrete-time SLAM approaches represent the estimate of the trajectory by sampling estimates at discrete times, whereas continuous-time SLAM (CT-SLAM) approaches represent the estimate of the trajectory as a continuous function defined by a discrete set of control points. In contrast to discrete approaches, continuous approaches easily enable fusing high-frequency data as the

dimensionality of the state only depends on the number of control points and is independent of the frequency of the sensor data. Therefore, CT-SLAM approaches are well suited for modeling rough motions, which we observe with mobile robots in disaster environments.

Early work towards CT-SLAM approaches was introduced by Bosse and Zlot [14], who propose a linear interpolation-based registration scheme to estimate the continuous trajectory of spinning 2D LiDARs by matching the geometric structure of local point clusters. Thereby, they are able to compensate for distortions in the scan cloud induced during the scan acquisition. Following up, LOAM [181] proposes a two-fold approach: one part continuously performs a low accuracy registration to achieve a high-frequency velocity update while another part performs a less frequent higher accuracy registration to correct for drift and update the map. Point clouds are matched by extracting edge and plane features to perform efficient scan registration. LeGO-LOAM [146] extends the LOAM approach by separating the ground for scan matching and gaining further efficiency by splitting the optimization in solving different components of the six-degree-of-freedom transformation separately. LIO-SAM [147] applies the LOAM registration scheme in a smoothing and mapping context, which fuses preintegrated IMU measurements [58] and LiDAR registration jointly in a pose graph framework.

Elastic LIDAR Fusion [129] uses a linear interpolation scheme to correct for scan distortions Instead of linear interpolation, various works leverage more complex trajectory representations such as B-Splines [102, 50, 134]. Nüchter et al. [122] apply a global continuous time formulation to improve the registration results of Cartographer [71].

Recent time-discrete approaches demonstrate efficient, accurate, and robust results for modern 3D-LiDAR systems by performing ICP-based registrations. KISS-ICP [163] leverages a combination of adaptive thresholding for correspondence matching, a robust kernel, motion compensation approach, and point cloud subsampling, demonstrating accurate results over various configurations in real-world scenarios without the need for specific parameter adaptations. FAST-LIO2 [171] maintains the map as an iterative kd-tree as a map and an iterative Kalman Filter for sensor fusion. Chen et al. [28] use an efficiency-optimized variant of the generalized ICP-approach with adaptive keyframe sampling to achieve efficient and accurate results in complex real-world environments. In [29], the authors extend the approach by a continuous-time motion correction, yielding more accurate results than the initial approach, which is especially significant for fast motions.

Aiming to provide robust localization in uncertain settings, LOCUS [126] proposes to leverage a multi-stage scan matching scheme fusing multi-modal odometry sources to achieve robust SLAM in cave exploration scenarios.

We build upon the registration concept of registering range data against linearly interpolated poses by [14], transfer it for multi-resolution TSDF, and stabilize it by incorporating fused inertial-odometry observations.

TSDF as Map Representation in SLAM Representing the environment using TSDFs, a volumetric environment representation storing the truncated signed distance to the next surface in each cell was introduced in the seminal work of Curless and Levoy [34]. TSDF gained further attention with the introduction of KinectFusion [120], a method for live 3D tracking and mapping of room-scale environments. By performing a point-to-plane ICP optimized for efficient usage of GPU parallelization, they were able to generate high-resolution 3D maps in real-time. Bylow et al. [19] propose an alternative optimization scheme by directly minimizing the depth error of the RGB-D image on the TSDF. Thereby improving the accuracy of the tracked pose. Complementing the previous works, Slavcheva et al. [148] propose representing the RGB-D Image as TSDF and performing direct TSDF to TSDF registration. This yields further improvements in the size of the convergence basin, rotational motion estimation, and reconstruction quality. However, explicit representation of the scan as TSDF requires significant computation operations. Therefore, [178] extends the approach and overcomes the need for an explicit representation by introducing a pseudo point set.

May et al. [105] generalize the KinectFusion approach to jointly register 2D and 3D data. In [123], the approach is integrated into a mobile mapping solution that can be carried as a back-pack.

LiDAR systems differ from RGB-D cameras, as the range and field of view are significantly larger, leading to more measurements with steep incident angles, which induce projection errors in TSDFs. To correct these errors, Fossel, Tuyls, and Sturm [59] compute regression lines in the scan. In contrast, we estimate scan normals to approximate the Euclidean distance. Koch et al. propose Ohm TSD SLAM [84, 85] a 2D TSDF-based SLAM pipeline for combined mapping with multiple robots.

Conventional TSDF approaches suffer from overwriting artifacts when objects are thinner than the truncation distance. Splietker and Behnke [149] overcome this issue by storing the SDF value for multiple surface orientations separately.

An established technique to improve the scalability and robustness of grid-based methods is multi-resolution approaches. Hector SLAM [89] performs robust 2D scan-to-map matching against a pyramid of occupancy grids, starting at the coarsest resolution and forwarding each result as initialization for the next finer resolution. Quenzel and Behnke [134] apply an adaptive resolution selection scheme to perform efficient surfel-based scan registration. Chen, Bautembach, and Izadi [27] leverage a hierarchical data structure reconstruction of large-scale scenes with fine geometric details from depth cameras on GPUs. Vespa et al. [162] adaptively choose the octree-resolution based on depth image resolution and the distance to the object for SLAM in room-scale environments and demonstrate up to six-fold execution speed-ups to single resolution grids.

To achieve robust and accurate scan registration, we build upon the direct registration scheme of [19] and extend it for a direct multi-resolution scheme in a continuous time registration formulation.

Large-scale SLAM and Loop Closure Detection Noise in sensor measurements and inaccuracies in the environment representation induce uncertainties in the registration of range data, leading to the accumulation of errors in the pose estimate and the map, which become notable over large distances. The error can be bound by fusing absolute references, such as GNSS [127] or local reference systems. However, such systems are typically not available in unknown indoor disaster environments.

The error can also be bound by detecting the revisit of areas, creating a geometric constraint between the respective poses, and leveraging this information to correct the past trajectory estimate. This process is commonly referred to as *Loop Closure*. The representation and optimization of the trajectory is a well-investigated research problem with established, efficient, and robust solution strategies, such as sparse factor graphs [68, 44, 43]. However, detecting loop closures is a challenging research problem. With increasing uncertainty in the pose, the search space can become large and, due to self-similarities in the environment, highly non-convex with many local optima. While for small search areas, ICP-based approaches can be feasible [147], they become computationally intractable for larger areas. To reduce the computational cost, approaches to extract geometrical descriptors [97, 104] or to perform histogram-based matching [72] have been proposed. Wang, Wang, and Xie [165] improve the efficiency and precision by incorporating intensity readings in global descriptors. Hess et al. [71] propose a branch-and-bound scheme to efficiently match scans in Occupancy Grids.

Recently, learning-based approaches gained increased focus [4], such as using a CNNs for scan association [30, 100, 176] or the application of deep descriptors yielding robust results [184]. However, the application and evaluation of such systems require dedicated hardware, and they are often limited to room-scale environments [184] or focused on autonomous driving-focused datasets [30], which are less complex and less degraded than the disaster environments this work is focusing on.

To achieve robust and accurate loop closure detections on mobile systems with limited computing, we build upon the branch-and-bound scheme proposed in [71] and transfer it to the application with TSDF.

Relation to Cartographer, Hector SLAM and the author's Master Thesis The proposed approach is related to Hector SLAM [89] and Cartographer [71]. Hector SLAM [89] is a robust 2D LiDAR SLAM system based on scan-to-map matching with a pyramid of occupancy grids. The system gained popularity and enabled various applications [88], such as 2D localization and mapping in the RoboCup Rescue Robot League [87]. However, the system does not incorporate loop-closure detection, so the applicable environment size is limited. Cartographer [71] builds upon the occupancy grid-based registration scheme, but instead of matching against a grid pyramid, it simultaneously matches against a coarse and a fine grid by summing the respective residual terms in the optimization process. Furthermore, Cartographer extends the optimization problem for 3D and large-scale mapping by introducing an efficient branch-and-bound-based loop closure detection with a pose graph back-end. Additionally, it extends the approach to support multi-robot trajectories.

Within the author's master thesis [36], a discrete-time optimization-based lidar-odometry and TSDF as map representation were investigated building on the Cartographer implementation. Furthermore, a generalized formulation of the loop closure detection was proposed.

The proposed approach in this thesis builds upon the Cartographer implementation. It extends it by a continuous-time registration scheme with a multi-resolution TSDF to achieve robust SLAM in disaster environments. In contrast to the formulation in Cartographer, the proposed approach does not sum the residuals but uses the highest resolution information available, thereby combining the benefits of high accuracy of the high resolution and robustness of the coarse resolution. Furthermore, the loop closure detection formulation is transferred for the application with TSDF as a specific implementation of the general formulation in [36].

5.2 Contribution

Dealing with these harsh conditions in disasters requires SLAM methods that are robust to the fast motions and potentially poor initialization induced by them, robust to degraded visual conditions, accurate to enable autonomous robot navigation, and efficient to run on mobile robot systems with limited compute. We propose multiple methodological novelties to improve the performance of SLAM methods for these requirements.

As localization and mapping are tightly coupled, we investigate suited environment representation and inference methods. We base our approach on TSDFs, where every cell models the distance to the nearest object surface, enabling sub-pixel accuracy. As TSDF provide gradients around the surface, they are naturally suited for optimization-based approaches. The representation provides gradients towards the surface in a larger area than, e.g., Occupancy Grids, leading to larger convergence basins and enabling scan matching with robustness to poor initialization. TSDF-based approaches became popular with RGB-D cameras [120], but so far, their utility for laser-based SLAM has not been fully exploited. A main challenge in applying these methods to laser scanners is that those sensors also provide distance readings under highly slanted angles, which leads to poor surface reconstruction. To overcome this issue, we investigate novel update and inference schemes. Another established environment model is occupancy grids. We perform detailed comparisons of TSDF and Occupancy Grids for LiDAR scan registration and demonstrate that TSDFs yield improvements in the accuracy and robustness of the registration.

Discrete trajectory representation limits the accuracy of fast and abrupt motions, which occur when traversing rough terrain or obstacles. Continuous-time SLAM approaches represent the pose as a time-continuous estimate that provides high accuracy and allows correcting for distortions induced by motion during the scan capture. We propose to combine continuous-time pose estimation and robust scan registration based on multi-resolution signed distance functions to achieve robust, accurate, and efficient SLAM.

Efficient loop closure detection approaches are missing for TSDF, limiting the applicability of TSDF-based methods for large-scale mapping. Therefore, we propose to transfer an efficient branch-and-bound-based loop closure detection for Occupancy Grids [71] for TSDFs.

The proposed approach is implemented based on Cartographer [71], and available open source as *HectorGrapher* [39]. In Chapter 7, we evaluate the proposed methods in multiple publicly available real-world data sets, own data sets, and under conditions of robotic competitions, demonstrating improvements in accuracy and robustness compared to related approaches. Furthermore, the approach was a core element for achieving the Best-in-Class "*Exploration and Mapping*" *Award* at the *RoboCup 2021 Rescue Robot League* with Team Hector.

Lidar-based methods achieve robust and accurate results in complex environments but require sufficiently good visual conditions. As visually degraded conditions need to be considered for various applications in disasters, such as fire or building collapse, it is crucial to enable mobile robots to localize and navigate in degraded visual conditions. In contrast to LiDAR sensors, radar sensors provide reduced accuracy and density and have a lower signal-to-noise ratio under good visual conditions but are less affected by degraded visual conditions. Therefore, radar sensors are well suited for navigation in environments with degraded visual conditions. Key challenges for the successful adaptation of the proposed LiDAR approach for radar are achieving accurate scan registration despite the reduced signal-to-noise ratio and enabling an accurate map update, although radar can return observations from behind walls. We address these challenges and extend the proposed lidar-based approach for visually degraded conditions to radar. Covering a novel forward sensor model for TSDF map update and robust scan registration scheme, enabling the navigation in room-size environments even in degraded visual conditions such as dense smoke.

5.3 System Overview

The proposed pipeline for solving the SLAM problem consists of three main components: wheel-inertial odometry, lidar-inertial odometry, and the pose graph back-end with loop closure detection. The wheel-inertial odometry fuses wheel odometry and IMU observations to gain a low latency and high frequency pose estimate. The result is forwarded to lidar-inertial odometry, where LiDAR point clouds are registered in multi-resolution TSDF submaps to achieve accurate and robust lidar-inertial odometry. Lidar-inertial odometry still induces small errors, leading to drift in the pose estimate over time. To maintain large-scale consistency of the map, the poses from the lidar-inertial odometry are stored in a pose graph, where we perform loop-closure detection and large-scale optimization. Optimizing the large-scale pose graph yields a large-scale consistent map and pose estimate.

5.4 Wheel-Inertial Odometry

Track and wheel encoders enable a high-frequency, low-latency estimate of the current robot motion state. Due to slippage and model errors, this estimate is only a rough approximation. Such errors are particularly strong for rotational motions in slippage-based drive kinematics such as skid-steer-kinematics. Inertial sensors provide high-frequency estimates of the linear acceleration and the rotational velocity. Estimating the position requires a double integration of the acceleration, which amplifies even small errors and induces unbound drift after short periods. In contrast, integration of the angular components requires only a single integration and only induces a small drift.

To reach an accurate motion estimate at a high frequency, we fuse inertial and wheel odometry measurements in an optimization problem, which can be stated as a pose graph, as shown in Fig. 5.2. Given the last known state $\boldsymbol{x}_i = \{\boldsymbol{p}_i, \boldsymbol{v}_i, \boldsymbol{R}_i, \boldsymbol{b}_i\}$ in the IMU frame, with the position vector $\boldsymbol{p}_i$, linear velocity vector $\boldsymbol{v}_i$, orientation matrix $\boldsymbol{R}_i$ and the IMU biases $\boldsymbol{b}_i$, and the unknown state $\boldsymbol{x}_j$ at τ_j we apply the IMU preintegration method introduced in [58] to estimate the changes in state $\Delta\boldsymbol{v}_{ij}^{pre}$, $\Delta\boldsymbol{p}_{ij}^{pre}$ and $\Delta\boldsymbol{R}_{ij}^{pre}$ in the time interval $\Delta\tau_{ij}$ to define

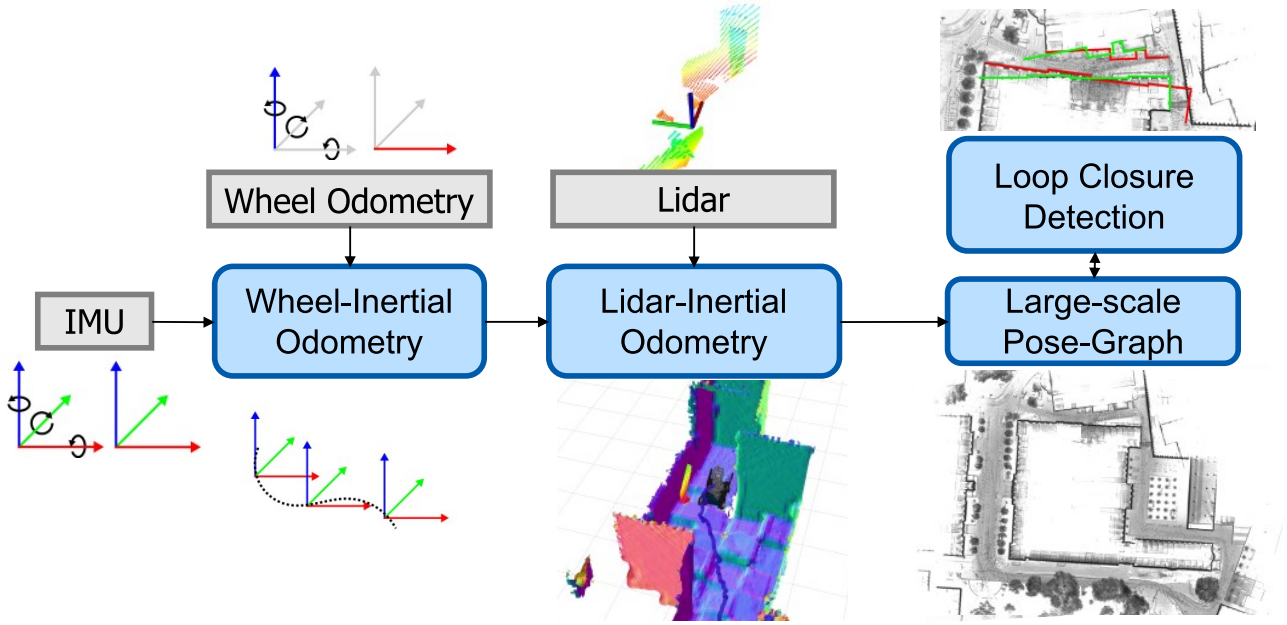

Figure 5.1: SLAM System Overview. The arrows indicate the direction of data flow.

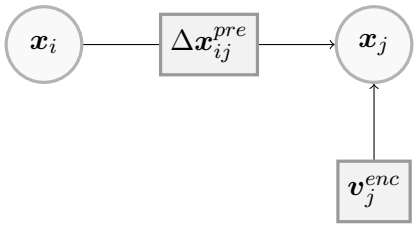

Figure 5.2: We fuse preintegrated IMU measurements Δx_{ij}^{pre} and linear velocity estimates v_j^{enc} from track or wheel encoders in a small factor graph to compute a high-frequency, low-latency estimate of the robot motion.

the residuals:

$$\boldsymbol{r}_{ij}^{v} = \boldsymbol{R}_i^T(\boldsymbol{v}_j - \boldsymbol{v}_i - \boldsymbol{g}\Delta\tau_{ij}) - \Delta\boldsymbol{v}_{ij}^{pre} \tag{5.1}$$

$$\boldsymbol{r}_{ij}^{p} = \boldsymbol{R}_i^T(\boldsymbol{p}_j - \boldsymbol{p}_i - \boldsymbol{v}_i\Delta\tau_{ij} - \frac{1}{2}\boldsymbol{g}\Delta\tau_{ij}^2) - \Delta\boldsymbol{p}_{ij}^{pre} \tag{5.2}$$

$$\boldsymbol{r}_{ij}^{R} = \Delta\boldsymbol{R}_{ji}^{pre}\boldsymbol{R}_i^T\boldsymbol{R}_j \tag{5.3}$$

$$\boldsymbol{r}_{ij}^{b} = \boldsymbol{b}_j - \boldsymbol{b}_i \tag{5.4}$$

with the gravity vector $\boldsymbol{g}$. Additionally, we add the estimated linear velocity from the track/wheel odometry converted into the IMU frame $\boldsymbol{v}_j^{enc}$ as unary constraints

$$\boldsymbol{r}_{j}^{v} = \boldsymbol{v}_j^{enc} - \boldsymbol{v}_j. \tag{5.5}$$

The residuals of Equation (5.1) - Equation (5.5) are added to a non-linear least squares problem which we solve using the Levenberg-Marquardt implementation in GTSAM[1]. Furthermore, we assume to be in a steady state when both IMU and odometry indicate no motion, which improves the estimate of the biases.

5.5 Lidar-Inertial Odometry

In the next step, the lidar-inertial odometry registers LiDAR point clouds in multi-resolution TSDF submaps to achieve an accurate and robust estimate of the local trajectory and map. In the following, we first outline the notation, then discuss the multi-resolution TSDF approach, and finally describe the optimization scheme for time-continuous trajectories.

5.5.1 Notation

We model the robot state as a time-continuous trajectory, similar to the formulation in [129]. The trajectory is represented by a linear Lie-group valued spline, which is defined by a set of timestamped control points C. Each control point $c_i = [\boldsymbol{T}_i, \tau_i]$ is given by a rigid transformation $\boldsymbol{T}_i \in SE(3) = [\boldsymbol{R}|\boldsymbol{p}]$ at time τ_i, which transforms from the robot body frame into the map frame. For simplicity, we assume that the IMU frame, robot body frame, and LiDAR frame coincide. Poses at time τ between two consecutive control points c_i and c_j with $\tau_i < \tau < \tau_j$ are evaluated by linear interpolation

$$\begin{aligned}
\alpha &= (\tau - \tau_i)/(\tau_j - \tau_i) \\
\overline{\boldsymbol{p}} &= (1-\alpha)\boldsymbol{p}_i + \alpha\boldsymbol{p}_j \\
\overline{\boldsymbol{R}} &= slerp(\alpha, \boldsymbol{R}_i, \boldsymbol{R}_j)
\end{aligned}$$

with the spherical linear interpolation operator $slerp$. We denote the interpolated transform as $\overline{\boldsymbol{T}}(\tau) = [\overline{\boldsymbol{R}}|\overline{\boldsymbol{p}}]$.

We denote the observations of single structured LiDAR scan $\boldsymbol{H} = \{\boldsymbol{h}_i\}_{i=1,\dots,N}$ consisting of N timestamped range observations $\boldsymbol{h} = (h_x, h_y, h_z, h_t)$ in the scan coordinate frame with the sensor pose as the origin of the scan frame.

5.5.2 Environment Representation as Truncated Signed Distance Function

We model the map as a multi-resolution TSDF. In the following, we first outline the principle of a single-resolution TSDF and then extend it to the multi-resolution formulation in the next section.

[1] https://gtsam.org/

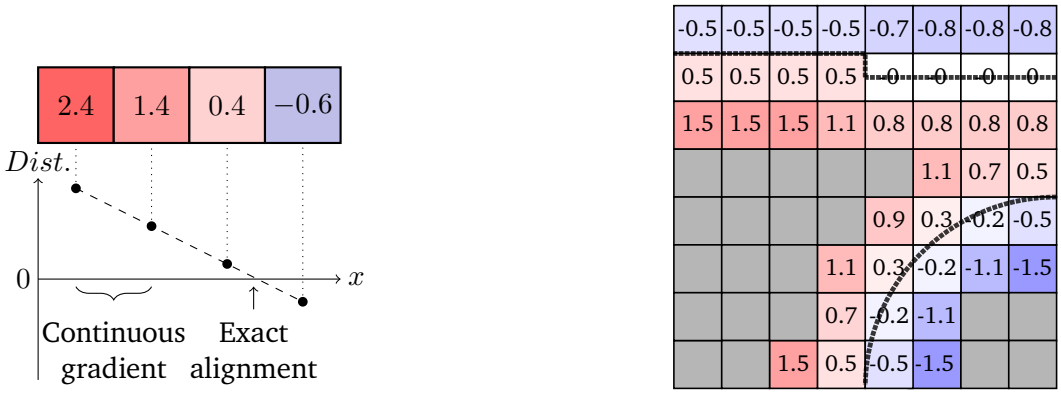

(a) TSDF 1D

(b) TSDF 2D

Figure 5.3: 1D- and 2D-TSDF examples. The underlying surface of the scene is indicated by the dashed lines. Undefined cells are marked in gray.

(a) Projective Distance Update

(b) Approximate Euclidean Distance Update

Figure 5.4: To update the TSDF with a scan observation $\boldsymbol{h}_i$ we evaluate two distance methods. (a) updates the cells in the direction of measurement $\boldsymbol{h}_i$ to origin $\boldsymbol{x}$. (b) updates the cells in the direction of the scan normal $\boldsymbol{n}_i$.

The Signed Distance Function [34] $\Phi : \mathbb{R}^3 \to \mathbb{R}$ maps from for each position in space to the scalar, signed distance to the nearest surface as shown in Figure 5.3. Φ is positive outside of objects and negative inside of objects. Therefore, object surfaces are encoded as the zero isocontour ($\Phi = 0$). As an exhaustive evaluation of Φ quickly becomes computationally intractable, TSDFs evaluate Φ only close to the surface. Each evaluation of a position further away from the closest surface than the truncation distance τ is truncated

$$\Phi_\tau(\boldsymbol{x}) = \begin{cases} -\tau, & \text{if } \Phi(\boldsymbol{x}) < -\tau \\ \Phi(\boldsymbol{x}), & \text{if } |\Phi(\boldsymbol{x})| \leq \tau \\ \tau, & \text{if } \Phi(\boldsymbol{x}) > \tau. \end{cases} \tag{5.6}$$

To represent a 3D scene as a TSDF, space is discretized into a regular grid. Each grid cell c contains the current estimate of the TSDF $M_\Phi(c)$, and a scalar weight $M_w(c)$ indicating the confidence in the TSDF value.

Projective Distance Update

Estimating the proper Euclidean distance for every cell is computationally intense. Therefore, most approaches approximate it. A common approach for updating the TSDF from depth images is the projective distance cell updates [120] (see Fig. 5.4a). To update a TSDF from a new range observation, we model each measurement as a ray from the sensor origin located at the origin of the sensor frame to the position of the measurement $\boldsymbol{h}_i$. Thus, the direction is $\hat{\boldsymbol{h}} = \frac{\boldsymbol{h}_i}{||\boldsymbol{h}_i||}$. We parameterize the scan based on an interpolation parameter u as a ray

$$\boldsymbol{v}_{projective}(u) = \boldsymbol{h}_i + u\hat{\boldsymbol{h}}. \tag{5.7}$$

All cells c along the ray $\boldsymbol{v}(u)$ with $u \in [-||\boldsymbol{h}_i||, \tau]$ are updated by taking a weighted moving average of the distance measurements

$$M_\Phi(c) := \frac{M_W(c)M_\Phi(c) + \omega(u)\Phi_\tau(u)}{W(c) + \omega(u)} \tag{5.8}$$

$$M_W(c) := M_W(c) + \omega(u) \tag{5.9}$$

with the update weighting function $\omega(u)$. In 2D, we apply an exponential weighting function as proposed in [19]. Additionally, we reduce the weight outside the truncation distance in front of the observation to reduce overwriting effects when the ray passes by close to other obstacles

$$\omega(u) = \begin{cases} e^{-\sigma(u+\epsilon)^2}, & \text{if } u \leq -\epsilon \\ 1, & \text{if } u \leq \tau \text{ and } u > -\epsilon \\ w_{free}, & u > \tau. \end{cases} \tag{5.10}$$

In our 2D experiments, values in the range of $0.1 \leq w_{free} \leq 0.5$ yielded accurate tracking results. In 3D, we did not observe significant improvements in the exponential weighting function in comparison to the constant weighting function $\omega(u) = const$. Therefore, we apply the constant weighting function in our 3D applications, as it reduces computational complexity and the number of parameters that need to be adapted.

Approximate Euclidean Distance Update

The projective distance function is only accurate for observations with viewing angles orthogonal to the surface or exactly at the surface.

The less orthogonal the viewing angle becomes, the more inaccurate the update distances. To compensate for the distance biases induced by using the projective distance, we leverage scan normals to approximate the Euclidean distance to the surface, see Figure 5.4b. Therefore, instead of directly using the projective distance, we propose to approximate the distance along a ray from the observation along the scan normal $\boldsymbol{n}$:

$$\boldsymbol{v}_{euclidean}(u) = \boldsymbol{h}_i - u\boldsymbol{n}. \tag{5.11}$$

Note that $\boldsymbol{v}_{euclidean} \approx \boldsymbol{v}_{projective}$ if the angle of incidence is large, which is typically the case for RGB-D cameras as they cannot observe surfaces at oblique angles. However, laser scanners have a high signal-to-noise ratio, which leads to observations even under small angles.

Update Rule Comparison

To qualitatively compare the TSDF update rules (for an in-depth quantitative analysis, see Chapter 7), we simulate a long, narrow hallway with a square and a circle obstacle inside. The robot is a simulated Pioneer

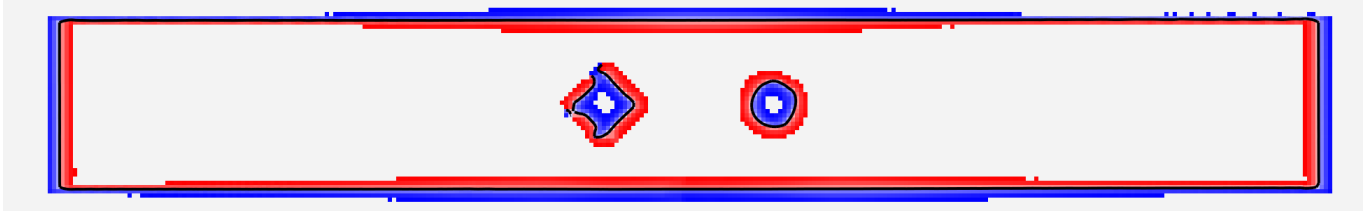

(a) Projective Distance Update

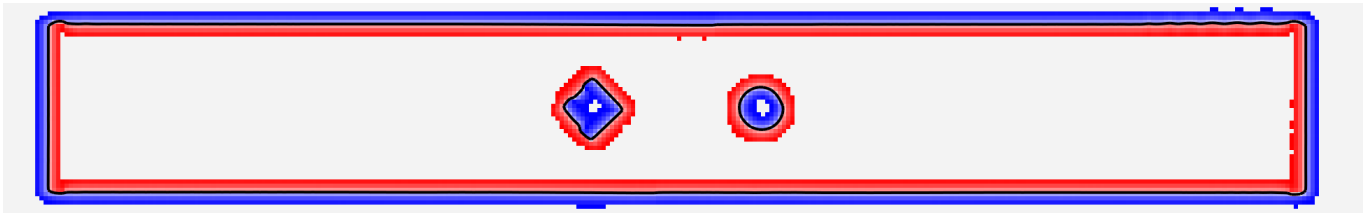

(b) Approximated Euclidean Distance Update

Figure 5.5: Comparison of TSDF maps with different update rules in a simulated environment. Red indicates a positive TSDF value, and blue encodes negative values. The reconstructed surfaces are indicated by black lines.

3DX with a LiDAR with $180°$ FOV and a measurement noise with a standard deviation $\sigma^2 = 0.01\,\mathrm{m}^2$. The robot drives a trajectory similar to an eight around the two obstacles. Therefore, the outer ends of the hallways are mainly observed from skewed angles. Fig. 5.5 shows TSDF maps with projective distance and the approximated Euclidean distance updates.

The projective approach shows artifacts modeling the square, which is reconstructed more consistently with the Euclidean approximation. In the top right corner, the surface reconstruction with the Euclidean approach shows small jittering.

The projective approach shows a strong overestimation of the absolute distance values close to the corners on the upper and lower edges of the outer rectangle and provides only a narrow truncation band. These errors are to be expected as these cells have only been observed under oblique angles. In contrast, the approximate Euclidean update yields a more consistent estimation of the Euclidean distance within the truncation band around the surface.

5.5.3 Optimization

We register the LiDAR scans and wheel-inertial odometry in a joint optimization. For registering scans in TSDFs, two paradigms are prevalent: ICP-based registration [120] and direct optimization of the pose on the TSDF [19]. As [19] indicates benefits in accuracy and efficiency for the direct optimization, we use the direct formulation as a basis and extend it by the multi-resolution registration and wheel-inertial odometry terms.

For the original direct optimization approach, the convergence basin is limited by the truncation distance, as outside the truncation distance, the gradient is constant. Furthermore, it requires an interpolation of the discrete TSDF grid to compute gradients, which might not be possible if the scan data is sparse. Finally, expressiveness and, thereby, accuracy depend on the grid resolution.

To leverage the precision of the high grid resolution and the robustness of a lower resolution, we apply a multi-resolution scan matching approach (see Fig. 5.6). We maintain two TSDF grids with different resolutions, a high-resolution grid M_Φ^{high} and a low-resolution grid M_Φ^{low}, both with the same relative truncation distance with respect to the resolution. For scan matching, we evaluate SDF at the highest resolution, providing a valid gradient. This leads to coarse and robust gradients in large areas around the surface and precise gradients

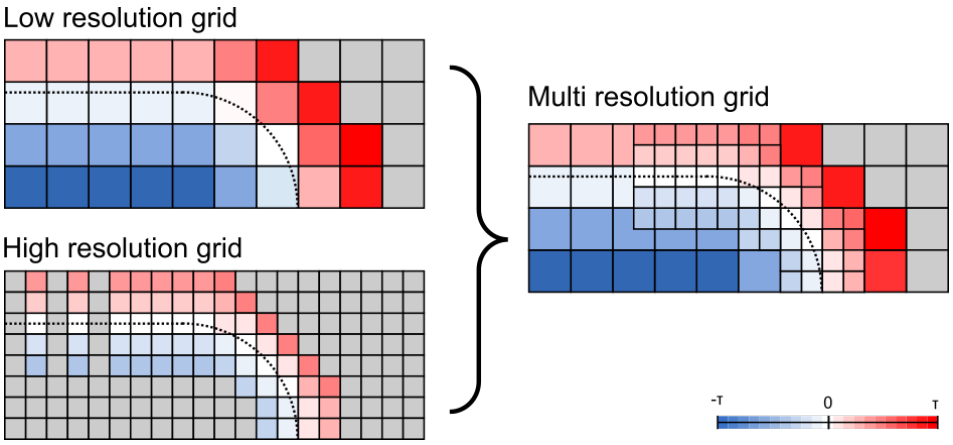

Figure 5.6: We maintain two grids, a high-resolution grid and a low-resolution grid with the same relative truncation distance. During scan matching, for each point, we evaluate the SDF at the highest resolution, providing a valid interpolation.

close to the surface.

Following the derivations in [19], we phrase the registration as a Nonlinear Least Squares Problem

$$\min_{\boldsymbol{C}} \sum_{i=1}^{N} \left(\Phi_I^{MR}(\overline{\boldsymbol{T}}(\tau_i)\boldsymbol{h}_i)\right)^2 + \sum_{i=1}^{|C|-1} (r(T_{i+1}^{-1} T_i \Delta T_{i;i+1}^{odom}))^2, \tag{5.12}$$

$$r(\boldsymbol{T}_i) := [\boldsymbol{p} | RPY(\boldsymbol{R})] \tag{5.13}$$

Φ_I^{MR} is the tri-linear interpolation of the highest resolution grid available, providing a valid interpolation. We consider an interpolation to be valid when none of the eight neighboring cells is uninitialized. $RPY(\boldsymbol{R})$ is the extraction of the Cardanian angles from the rotation matrix $\boldsymbol{R}$. We solve the optimization with the Levenberg-Marquardt method using the Ceres Solver [2] and compute gradients with automatic differentiation.

5.6 Large-scale SLAM

Scan matching induces small registration errors, resulting in an accumulation of errors over large distances and leading to global inconsistencies in the pose estimate and the map. To correct these errors, we follow the approach of Cartographer [71] and generate many small, locally consistent *submaps*. The spatial relation of the submaps is represented as a pose graph, where each node represents a location and edges relate to constraints between poses, which are derived from sensor observations. If sufficient valid constraints have been found, solving an optimization problem to minimize the constraint violations yields consistent pose estimates and allows the combination of the submaps into a consistent map.

The pose graph is constructed by two types of constraints - *odometry constraints* and *loop closure constraints*. *Odometry constraints* are relative constraints between two scan matches and are added for each new optimized pose of the lidar-inertial-odometry. *Loop closure constraints* are relative constraints between poses that are geometrically close but temporally distant and arise when a previous location is revisited, e.g., after performing a loop. The detection of such constraints, in the following referred to as "*constraint search*", is typically computationally intense, as it involves the search in a large and highly non-convex space with many local optima. Therefore, it is a key challenge to efficiently and accurately identify loop closure constraints online on mobile systems.

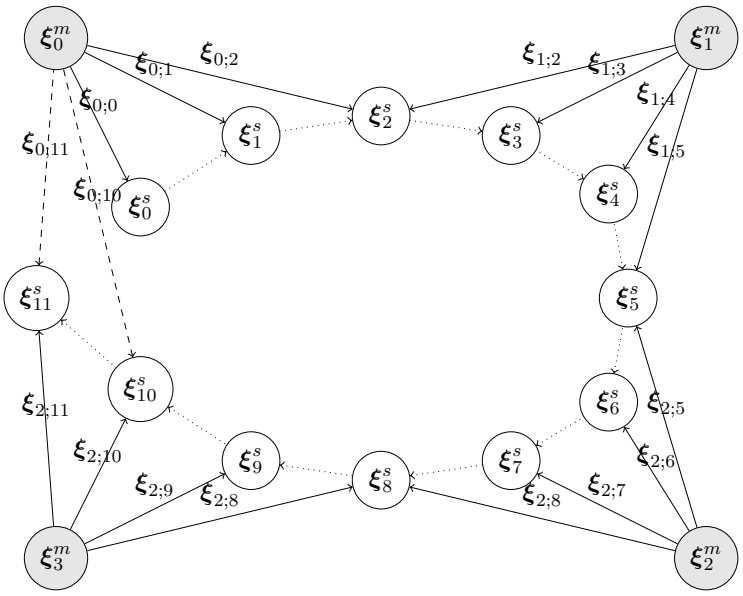

Figure 5.7: Topology and notation of a pose graph, with the scan nodes $\boldsymbol{\xi}_i^s$ (white circles), the submap nodes $\boldsymbol{\xi}_i^m$ (gray circles) and the constraints $\boldsymbol{\xi}_{ij}$. The solid lines indicate lidar-inertial odometry constraints, and the dashed lines indicate loop closure constraints. The trajectory is marked by the dotted line.

Cartographer [71] proposes an efficient branch-and-bound-based approach to compute loop closure constraints for occupancy grids, enabling the search for loop closures on mobile systems in large areas in real-time. Comparable efficient methods are lacking for TSDFs. We address this gap and extend the Cartographer method to efficiently detect loop closures in TSDFs by transferring the scan matching problem and bounding function for TSDFs. For the overall formulation and implementation of the pose graph and its optimization, we use Cartographer, which builds upon [68] and [90].

5.6.1 Pose Graph Optimization

We use the pose graph formulation and implementation proposed in Cartographer [71]. The pose graph (see Figure 5.7) is constructed iteratively by adding each scan match pose as a node and the relative pose to the matched submap as odometry constraints. Furthermore, all scans are considered candidates for loop closures. Whenever a submap is finished, a subset of the scans inserted in the finished submap is matched with all other submaps within a search window around the finished submap. If a good match is found (see Section 5.6.2), the corresponding relative pose is added to the pose graph as a loop closure constraint.

The pose graph optimization follows [90] and is formulated as a Nonlinear Least Squares Problem. Given the the submap poses $\boldsymbol{\Xi}^m = \{\boldsymbol{\xi}_i^m\}_{i=1,\dots,m}$, the scan poses $\boldsymbol{\Xi}^s = \{\boldsymbol{\xi}_i^s\}_{i=1,\dots,n}$, the respective constraints for a pair of a submap i and a scan j as the relative poses $\boldsymbol{\xi}_{ij}$ and the corresponding covariance matrices $\boldsymbol{\Sigma}_{ij}$, the pose graph optimization problem can be formulated as

$$\underset{\boldsymbol{\Xi}^m,\boldsymbol{\Xi}^s}{\operatorname{argmin}} \frac{1}{2} \sum_{ij} \rho(E^2(\boldsymbol{\xi}_i^m, \boldsymbol{\xi}_j^s; \boldsymbol{\Sigma}_{ij}, \boldsymbol{\xi}_{ij})). \tag{5.14}$$

The residual E for a constraint computes as

$$E^2(\boldsymbol{\xi}_i^m, \boldsymbol{\xi}_j^s; \boldsymbol{\Sigma}_{ij}, \boldsymbol{\xi}_{ij}) = \boldsymbol{e}(\boldsymbol{\xi}_i^m, \boldsymbol{\xi}_j^s, \boldsymbol{\xi}_{ij})^T \boldsymbol{\Sigma}_{ij}^{-1} \boldsymbol{e}(\boldsymbol{\xi}_i^m, \boldsymbol{\xi}_j^s, \boldsymbol{\xi}_{ij}), \tag{5.15}$$

$$\boldsymbol{e}(\boldsymbol{\xi}_i^m, \boldsymbol{\xi}_j^s, \boldsymbol{\xi}_{ij}) = \boldsymbol{\xi}_{ij} - \begin{pmatrix} \boldsymbol{R}_{\boldsymbol{\xi}_i^m}^{-1}(\boldsymbol{t}_{\boldsymbol{\xi}_i^m} - \boldsymbol{t}_{\boldsymbol{\xi}_j^s}) \\ \boldsymbol{\xi}_{i;r}^m - \boldsymbol{\xi}_{j;r}^s \end{pmatrix} \tag{5.16}$$

with $\boldsymbol{\xi}_{i;r}^m$ and $\boldsymbol{\xi}_{j;r}^s$ denoting the rotational components of $\boldsymbol{\xi}_i^m$ and $\boldsymbol{\xi}_j^s$. Furthermore, Huber loss ρ is applied to improve the robustness against incorrect constraints. The problem is solved with the Levenberg-Marquardt method using Ceres [2] with automatic differentiation.

5.6.2 Efficient Constraint Search with Branch and Bound Scan Matching

The detection of loop closure constraints between scans and submaps requires the search for an optimal scan match $\boldsymbol{\xi}^\star$ within a large search window $\mathcal{W}$. If the search window is smaller than the accumulated error in the pose estimate, the constraint will be missed. A naive, exhaustive search of the search window is theoretically possible. However, for large search windows, it is computationally intractable.

The efficient branch-and-bound-based approach proposed in Cartographer [71] enables to solve this problem on mobile systems in large areas in real-time for occupancy grids. Comparable efficient methods are lacking for TSDFs. We address this gap and extend the Cartographer method to efficiently detect loop closures in TSDFs by transferring the scan matching problem and bounding function for TSDFs. In the following, we first present the constraint detection problem for TSDF, then summarize the Branch and Bound approach [71], and finally formulate the TSDF specific bounding function proposed in [37].

Constraint Detection Scan Matching for TSDF

To detect pose graph constraints, we search a match $\boldsymbol{\xi}^\star$ for a scan $\boldsymbol{H}$ within the search window $\mathcal{W}$. This can be formulated by transferring the Occupancy Grid-based formulation of [71] for TSDF as:

$$\boldsymbol{\xi}^\star = \underset{\boldsymbol{\xi} \in \mathcal{W}}{\operatorname{argmin}} \sum_{k=1}^{K} |\Phi_N(\boldsymbol{T}_{\boldsymbol{\xi}} \boldsymbol{h}_k)| \tag{5.17}$$

$$s.t. \sum_{k=1}^{K} |\Phi_N(\boldsymbol{T}_{\boldsymbol{\xi}} \boldsymbol{h}_k)| < e_{max}. \tag{5.18}$$

Φ_N is the nearest neighbor interpolation of the TSDF grid map M_Φ. The nearest neighbor interpolation allows an efficient evaluation of the bounding function.

As we are only interested in good matches, we require the solution to have an alignment error smaller than e_{max}. If no solution satisfies this requirement, no constraint is added to the pose graph.

Branch and Bound

The constraint detection problem can be efficiently solved by utilizing branch and bound approaches, which are optimization methods that represent the discretized search space as a tree, with each node corresponding to a subspace of the search space. Efficiency in the evaluation of the optimal solution is gained by estimating bounds for the optimal solution within the subtrees. Thereby, the evaluation of large parts of the tree can be omitted, and only a small part of the search space has to be evaluated.

The formulation of a branch and bound approach requires a branching rule, a node selection strategy, and a bounding function. The *branching rule* defines how to represent the discretized search space as a tree. The *node selection strategy* defines how to traverse the tree efficiently, and the *bounding function* provides a bound on the best obtainable solution in a subtree. The branch and bound algorithm then performs a search over the tree for the best solution, omitting subtrees if their bounding function is worse than the currently best

candidate or a required score threshold. Therefore, a good bounding function is crucial for the overall method performance.

Cartographer [71] proposes strategies for each of the steps. The approach discretizes the search space in fixed rotations and separate translations for the location of each map cell in the search window $\mathcal{W}$. For the*branching rule*, the approach combines in each node eight neighboring translations and a single rotation. Thereby, a node at height h combines up to $2^h \times 2^h \times 2^h$ translations while representing a single rotation. As *node selection strategy*, the approach uses a depth-first strategy to reach a good solution candidate early in the search for omitting further subtrees.

Bounding function for TSDF

As a bounding function, we transfer the Occupancy Grid-based formulation of [71] for TSDF. An accurate and efficient evaluation of a lower bound of the optimal solution at an inner node c is crucial to prune large parts of the tree. Approximating the optimal solution within the search window by the optimal match for each scan point within the potential maximum size of the search window yields the bounding function bound(c) as

$$\min_{\boldsymbol{\xi}_c \in \mathcal{W}_c} \sum_{k=1}^{K} |\Phi_N(\boldsymbol{T}_{\boldsymbol{\xi}_c} \boldsymbol{h}_k)| \tag{5.19}$$

$$\geq \sum_{k=1}^{K} \min_{\boldsymbol{\xi}_c \in \mathcal{W}_c} |\Phi_N(\boldsymbol{T}_{\boldsymbol{\xi}_c} \boldsymbol{h}_k)| = \text{bound}(c). \tag{5.20}$$

The bounding function can be efficiently precomputed as a grid $\Phi^h_{precomp}$ for each height h

$$\Phi^h_{precomp}([x, y, z]^T) = \min_{\substack{x' \in [x, x+r(2^h-1)] \\ y' \in [y, y+r(2^h-1)] \\ z' \in [z, z+r(2^h-1)]}} |\Phi_N([x', y', z']^T)|. \tag{5.21}$$

Each cell in the precomputed grid contains the minimum values of the $2^h \times 2^h \times 2^h$ cube around it.

With the precomputed grid $M^h_{precomp}$ we can evaluate the bounding function efficiently as

$$\text{bound}(c) = \sum_{k=1}^{K} \Phi^{c_h}_{precomp}(\boldsymbol{T}_{\boldsymbol{\xi}_c} \boldsymbol{h}_k). \tag{5.22}$$

5.7 SLAM in Degraded Visual Conditions

Lidar-based methods achieve robust and accurate results in complex environments but require sufficiently good visual conditions. Dense smoke, dust, and fog impair the quality of LiDAR sensor measurements and can lead from a reduction in performance up to complete failure (see Figure 5.8). As such conditions need to be considered for various applications in disasters, such as fire or building collapse, it is crucial to enable mobile robots to localize and navigate in degraded visual conditions.

In contrast to LiDAR sensors, radar sensors provide reduced accuracy and density and have a lower signal-to-noise ratio under good visual conditions but are less affected by degraded visual conditions. Therefore, radar sensors are well suited to enable navigation in environments with degraded visual conditions.

Key challenges for the successful adaptation of the proposed LiDAR approach for radar are achieving accurate scan registration despite the reduced signal-to-noise ratio and enabling an accurate map update,

(a) Camera image at fog insertion start.

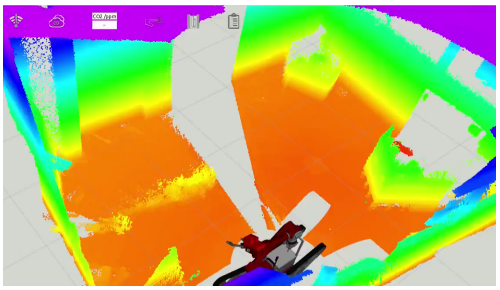

(b) Lidar cloud at fog insertion start.

(c) Camera image in a fog-filled room.

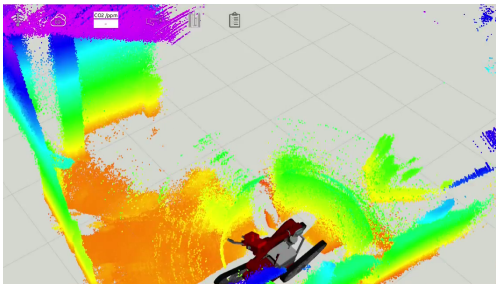

(d) Lidar cloud in a fog-filled room.

Figure 5.8: Lidar and camera data degrade in poor visual conditions such as fog or smoke and are not suitable for navigation. a) and b) show camera and LiDAR data in a smoke chamber at the start of the insertion of disco fog. c) and d) show the same environment after it is filled with fog.

although radar can return observations from behind walls. In [157], we address these challenges and extend the proposed lidar-based approach for visually degraded conditions by extending it to radar. The author's main contributions are work on the general conceptualization of the robust scan registration, map update model, and evaluation. Most of the detailed work, especially for implementation and evaluation, was performed by Moritz Torchalla and Marius Schnaubelt and builds on the author's implementation in [37].

The proposed approach covers three main aspects *radar processing*, a *forward sensor model for TSDF map update* and *robust scan registration*.

Radar Processing

To optimize the quality of the final radar signal, the raw radar data is processed by a flexible and modular processing pipeline covering multiple Fast Fourier Transformations to gain accurate distance and velocity estimates and Constant False Alarm Rates detection to enable the extraction of discrete range observations, similar to a point cloud.

TSDF Map Update

In contrast to lidar, radar can also penetrate objects and yield range observations for multiple stacked surfaces. While TSDF map update approaches often assume free space between each observation and the sensor origin, this does not hold true for radar. Therefore, we propose an adapted method that only updates free space up to the first detection. The method gains efficiency by leveraging angular binning for the range observations.

Robust Scan Registration

The increased number of outliers in the scan processing makes the registration infeasible for a native non-linear least squares formulation. The outliers induce high residuals, reducing the overall accuracy of the optimization. Therefore, the optimization is adapted for the robust optimization method Graduated Non-Convexity [175], which reduces the impact of outliers during the optimization and thereby increases the robustness and accuracy of the solver for high outlier ratios.

Application Example

As an application example, we evaluated the performance of the radar SLAM in comparison to LiDAR SLAM in the DRZ challenge in the DRZ Living Lab. The chamber is S-shaped with drywalls and can be filled with artificial fog. Figure 5.9 shows comparisons for LiDAR and radar SLAM without fog. Both methods accurately estimate the trajectory of the robot. While the LiDAR map is very sharp and consistent, the radar map is more noisy but also overall consistent, and walls can be distinctly indicated. When fog is added, the radar SLAM generates results of similar quality, while the LiDAR SLAM fails. A visual comparison of the radar and LiDAR data in the fog environment in Figure 5.9c also indicates that radar provides sensor reading relating to the geometry of the environment, while this is not the case for the LiDAR data. Overall, the quality of the map and trajectory estimation indicate the applicability for navigation in degraded visual conditions.

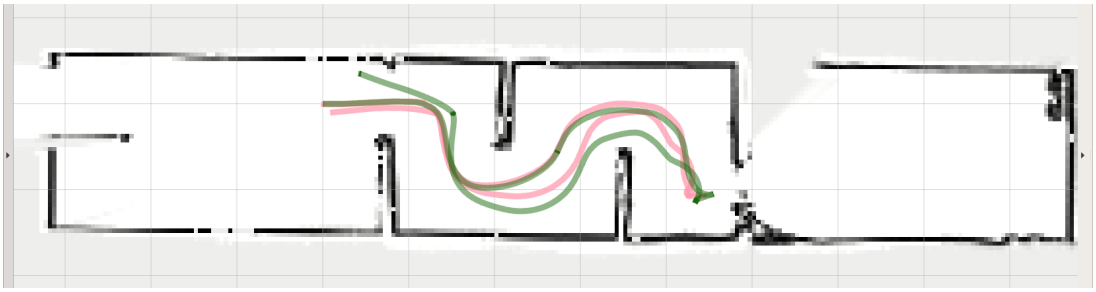

(a) LiDAR SLAM without fog.

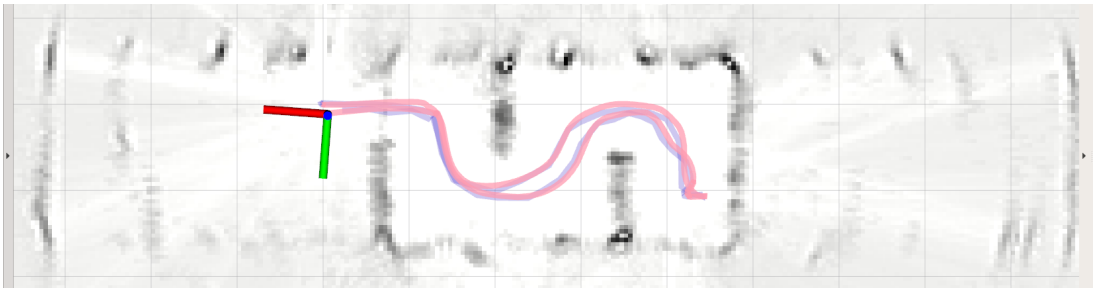

(b) Radar SLAM without fog.

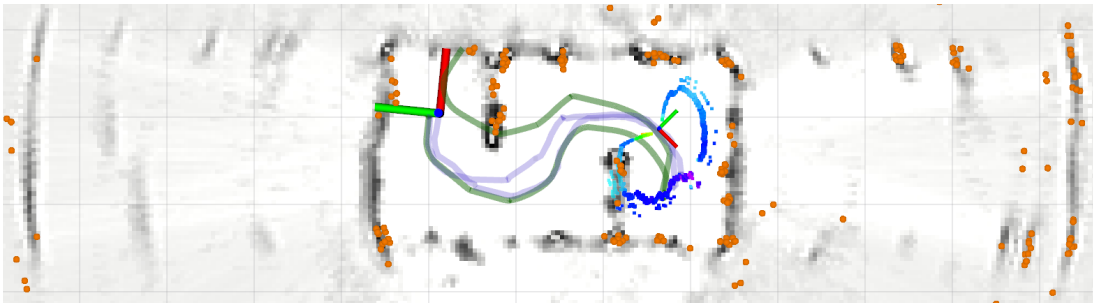

(c) Radar SLAM with smoke. The sample point clouds indicate the sensor quality. The radar cloud (orange) accurately captures the environment characteristics, whereas the LiDAR cloud (rainbow) is strongly impaired by the fog.

Figure 5.9: Comparison of maps and trajectories created by LiDAR and radar SLAM with the Asterix robot in the DRZ challenge. Green: Odometry, Blue: radar trajectory, Pink: LiDAR trajectory. One square corresponds to 1 m^2. (Images: [157] ©2021 IEEE)

6 Exploration and Data-Acquisition in Shared-Autonomy Multi-Goal Missions

Exploration and data-acquisition missions with mobile ground robots in disaster response missions often involve multiple goals with versatile and dynamically changing situations. As a stable radio connection to the robot can often not be guaranteed and the operator's cognitive capacity is limited, enabling (semi-)autonomous assistance for such missions is important. However, the low repetitiveness, high complexity, and situation dynamics are challenging for current autonomous methods.

While related state-of-the-art methods often focus on fully autonomous approaches for single-goal missions [25], this chapter proposes a shared-autonomy approach for multi-goal missions, which allows fully autonomous operation but strongly benefits and incorporates by design the capabilities of the operator for scene understanding and decision making. Key elements of the framework are the coverage-based multi-goal exploration, which enables the efficient exploration of large-scale environments, and the affordance-based actionable environment representation, which enables complex object interaction and planning sequences.

The chapter is structured as follows: we first outline the related work and the key contributions of the approach, then we provide an overview of the proposed method before investigating the aspects of environment representations, exploration planning, operator interaction, and model generation in detail. An overview of experiments and results is provided in Chapter 7.

Parts of this chapter have been previously published. The affordance-based actionability concept has been published in [8], and the radiation mapping approach in [152].

6.1 Related Works

Exploration and Active SLAM methods typically need to solve three subproblems: 1) the identification of the potential actions, 2) the estimation of the action utility, and 3) the selection and execution of actions [132]. A general introduction to the Exploration and Active SLAM problem is provided in Chapter 2. This section focuses on the related works directly related to the proposed method. We first outline related approaches for exploring large-scale environments, then discuss the approaches for multi-goal exploration, and finally, discuss actionability and user interaction.

6.1.1 Exploration in Large-Scale Environments

With the introduction of the first concepts for efficient exploration of unknown environments, such as frontier-based [172] approaches for metric maps, the efficient exploration of large-scale environments became a major research interest. Challenges include planning efficient trajectories and action sequences and efficiency in the planning methods themselves, as naive planning and mapping schemes can quickly become computationally infeasible for large-scale environments.

Therefore, various approaches aim to find frontiers efficiently by applying incremental update schemes [82] or search-based strategies such as rapidly exploring random trees [159]. Other approaches aim to reduce the uncertainty in the map by identifying potential loop closure locations and adding these as goal positions [150].

Furthermore, topological and hybrid map representations improve planning efficiency for large-scale environments. Various recent approaches for the exploration of complex, large-scale 3D environments have been proposed and successfully evaluated in the context of the DARPA SubT Challenge.

Cao et al. propose the *TARE planner* [24] focusing on the efficient surface coverage-based exploration of large-scale environments. The approach gains computational efficiency through a hierarchical framework with dense local planning and coarse global planning. Efficient routes are computed by frequently solving traveling salesperson problems over high-utility viewpoints. In [25], the authors extend the approach in an integrated planning and exploration system together with a planning approach for unknown environments [174]. Furthermore, in [23], the authors extend the approach for multi-robot exploration. Overall, the approach is evaluated in various simulated environments and in complex underground environments by application by the CMU Team in the DARPA SubT Challenge, achieving the "Most Sectors Explored Award."

Dang et al. propose *GBPlanner* [35], a graph-based exploration approach for aerial and legged robots. Similar to TARE, they also use a combined local and global planning architecture. The local planner is based on a rapidly exploring random graph, which optimizes a local exploration gain based on volumetric coverage while accounting for dynamic feasibility, traversability, and collision avoidance. Globally, the approach utilizes a sparse graph to maintain large-scale relations and global goals. In contrast to TARE, the global graph is only utilized when a local exploration goal is completed. In [95], the approach is extended for heterogeneous multi-robot exploration. The approach is successfully evaluated in various complex environments and applied by Team CERBERUS in the DARPA SubT Challenge, enabling the team to win the overall competition [158].

Best et al. [11] propose an approach for exploration with a team of aerial robots leveraging both range and vision sensing modalities in confined environments focusing on resilience. The approach is based on a behavior tree architecture, which switches between different behaviors, including exploration or the response to adverse events, such as degraded perception due to dust. The approach combines a local and global planner. The global planner plans coarse global graph-based trajectories related to the current behavior. Whereas, the local planner performs motion-primitive-based local planning to the next subgoal of the global trajectory. The author proposes combining lidar and camera coverage to estimate the utility of potential viewpoints. The approach is successfully evaluated in complex real-world environments by Team Explorer in the DARPA SubT challenge.

To enable efficient exploration assistance for large-scale environments, the proposed approach follows the dense local and coarse global paradigm proposed by the previous methods [24, 35, 11]. To enable globally and locally efficient trajectories, the proposed approach is based on the overall structure of the TARE planner [24] but differs in the specific approaches for local planning, graph generation, and environment coverage mapping and extends for multi-goal exploration by incorporating spatial coverage.

6.1.2 Multi-Goal Exploration

While common exploration approaches often focus on maximizing a single goal, such as covering a map with a single sensor, multi-goal approaches focus on combining multiple goals, such as covering multiple modalities, or considering constraints, such as staying close to relevant agents in the environment.

Mandischer et al. [103] propose a novel approach for finding moving operators in firefighting operations under the constraint of staying close to human firefighters in 2D. The authors propose a next-best view-based exploration method that combines multiple modalities, covering prior information, detected targets, the direction of the operator, and progress. The overall utility is determined by multiplying factors for each utility. Calisi et al. [22] combine the creation of a map and victim detection. The authors propose to use Petri nets to model behaviors that switch between mapping and investigating victim hypotheses.

Butzke and Likhachev [18] consider multi-robot exploration with an extended frontier-based exploration approach that combines utility terms for information gain, per-robot regions, and distance to other robots. The

results demonstrate efficient distributed exploration with multiple robots. Bramblett, Peddi, and Bezzo [16] also addresses the coordination of multiple robots for exploration and incorporates the aspect of robot-to-robot communication by explicitly modeling behaviors for exploration, rendevous, and task allocation. Further approaches incorporate energy consumption [136] or the availability of communication links [32].

The proposed approach combines multiple sensor modalities in the coverage-based exploration. Furthermore, the proposed approach allows the switch between different behaviors, similar to [16]. However, the switches are initiated by a human operator instead of an autonomous agent.

6.1.3 Actionability and Operator Interaction

It is crucial to consider actionability in the environment representation to enable meaningful operator interactions and complex planning sequences. Common approaches model the semantics of the environment by maintaining a 3D model of the environment and assigning instances and semantic classes to geometric segments [140, 67]. Such structures can be embedded in 3D scene graphs, which are hierarchical mixed topological graph structures, allowing for reasoning on spatial and semantic relations in scenes [141]. An example of the real-time creation of 3D scene graphs is the Hydra framework [77], which combines a semantic 3D mesh with a hierarchical topological graph structure of objects and places, rooms, and buildings. Semantic classes can also be assigned to exploration-related instances. Gomez, Hernandez, and Barber [65] propose to assign semantic classes to exploration frontiers, incorporating the information in the utility estimation.

Semantic 3D scene graphs enable reasoning about the environment. To allow the planning of complex actions on such data structures requires additional information such as a knowledge base [61]. Alternatively, actionability can be achieved by directly modeling potential actions in the environment representation.

A concept to describe actions that can be performed on an object is affordances [180]. The concept is common in the object manipulation context, e.g., to model grasps [139], to model supportability and leanability [81], or for object recognition by combining visual appearance and grasp affordances [63]. The representation is less common in the navigation context. Qi et al. [133] describe the navigability of surfaces with spatial affordances.

In the proposed approach, we consider a combination of topological graph-based environment representation [141], with object affordances [180] for object interaction tasks to allow flexible interactions with the environment.

6.2 Contribution

While related state-of-the-art methods often focus on fully autonomous approaches for single-goal missions [24, 35], we propose an operator-related approach for multi-goal missions, which allows fully autonomous operation but strongly benefits and incorporates by design the capabilities of the operator for scene understanding and decision making. We base the proposed method on the concept of the hierarchical surface-coverage-based approach by TARE [24] and extend it by a novel formulation to allow for combined spatial (e.g., radiation or hazard sensor) and surface coverage goals missions. We embed the approach in a proposed shared-autonomy framework that follows the shared autonomy principle and allows flexible changes in the autonomy level from assisted teleoperation to full autonomy. The interaction concept follows the ideas of coactive design [80] and ensures observability, predictability, and directability. The operator can request complex environment interaction with a novel, integrated, actionable, affordance-based environment representation, which enables complex object interaction and planning sequences. Furthermore, we propose a novel method to accurately map dose rates in 2D and 3D based on Gaussian Processes.

The evaluation in Chapter 7 demonstrates the efficiency of the proposed planner and demonstrates improved efficiency and coverage compared to a frontier-based planner in complex simulation environments.

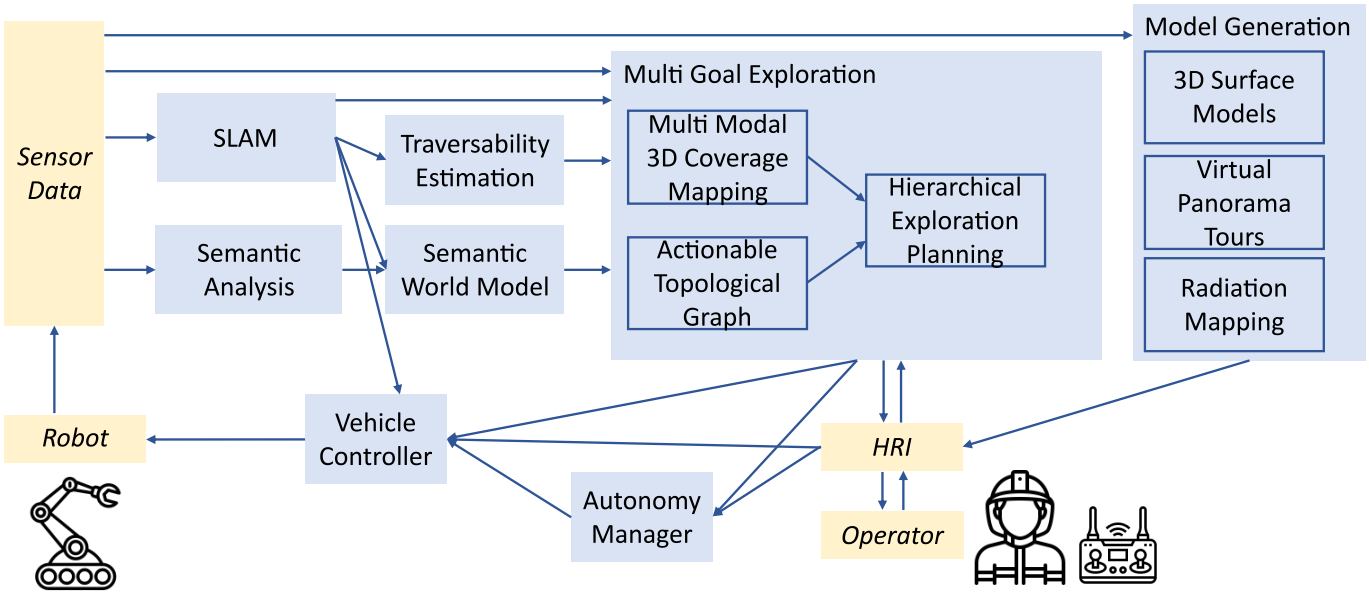

Figure 6.1: Overview of the shared-autonomy exploration approach. The approach combines hierarchical coverage-based mapping and planning with actionable object representations. The arrows indicate the directions of the main information flow. Although most components directly provide information to the HRI, for clarity, this is only indicated for a subset.

Furthermore, we demonstrate the capability to perform complex missions with autonomous and interactive interactions.

6.3 System Overview

Dealing with the challenging application conditions of low repetitiveness and high complexity makes the application of fully autonomous systems currently infeasible for complex real-world disaster conditions. On the other hand, direct teleoperation leads to errors in control, resulting in mission failures and oversights [116]. Therefore, we propose a shared-autonomy exploration concept that creates synergies between the capabilities of autonomous assistance functions and the scene understanding and decision-making skills of a human operator. The multi-goal exploration framework follows the principle of shared autonomy and allows flexible changes in the autonomy level from assisted teleoperation to full autonomy. The interaction concept embeds the ideas of coactive design [80] by ensuring observability, predictability, and directability.

Key elements of the framework (see Figure 6.1) are the coverage-based multi-goal exploration, which enables the efficient exploration of large-scale environments, and the affordance-based actionable environment representation, which enables complex object interaction and planning sequences.

For the exploration assistance function, we apply a hierarchical approach with dense local and coarse global planning. The overall concept is inspired by the TARE planner [24] but differs in the several representations of the components, the extension for multi-goal exploration, the focus on shared autonomy, and the operator interaction concept.

For the coverage-based exploration, we follow the hierarchical environment representation paradigm of [25] and segment the environment into equally sized segments. Potential exploration goals are evaluated frequently and densely for the segments in the vicinity of the robot. For all other segments, we maintain the status of observed, partially observed, or unobserved. The observation goals for partially observed segments are stored

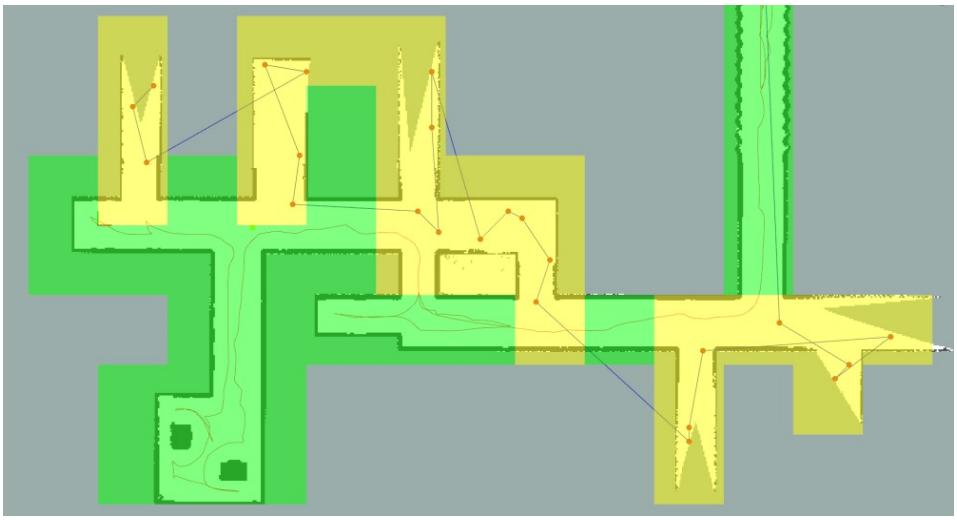

Figure 6.2: For efficiently exploring large areas, we use a hierarchical approach. The environment is divided into equal-sized segments, with those in close proximity to the robot being densely and frequently sampled. All relevant viewpoints of these segments are added to the planning problem. Distant segments are only reevaluated when alterations occur. For partially explored segments (yellow), only the viewpoint with the highest score is added to the planning problem. Fully explored segments (green) are no longer considered for the planning problem.

in a global topological graph. The trajectory planning accounts for the local and global goals by repeatedly solving a Traveling Salesperson Problem (TSP) to compute the optimal exploration path. The operator can directly interact with the exploration approach and prioritize exploration regions or modify the coverage accuracy for regions. The visualization of the explored and unexplored areas and the proposed exploration route can also serve as assistance for the mission procedure planning in a teleoperated control mode.

To allow the incorporation of object interactions in the planning, objects are stored with affordances representing potential actions. Thereby, they are incorporated for direct operator interactions or complex planning queries with multi-stage requirements. As shown in [8], the operator can provide a complex and abstract goal, and the assistance function can automatically deduct the respective steps to reach the goal.

6.4 Environment Representations for Multi-Goal Exploration

Enabling efficient planning and exploration assistance functions for large-scale environments requires suitable environment representations. A crucial component of the exploration algorithm is accurately and efficiently predicting the utility of an action.

In the environment representation, we apply a modular formulation and account for both *surface coverage* (see Section 6.4.1), the observation of surface coverage of the 3D environment, such as a geometric representation of a 3D map based on lidar observations and *spatial coverage* (see Section 6.4.2), the observation of locally measured scalar quantities in the 2D-constrained observation space, such as radiation dose rate, radio connectivity, or temperature. Furthermore, we maintain a topological graph for efficient planning in large-scale environments (see Section 6.4.3) and affordance-based instance representations (see Section 6.4.4) to allow flexible interactions with the environment.

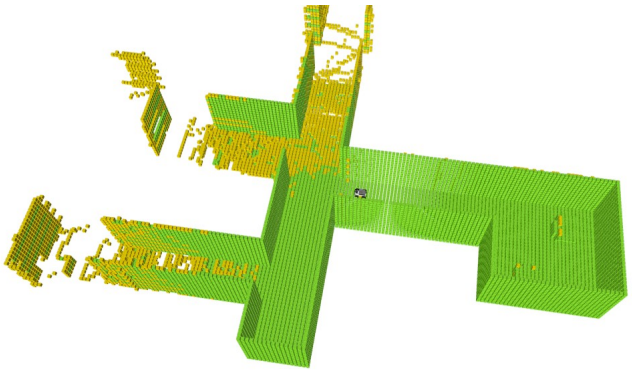

(a) 3D surface coverage map. The fully observed cells are indicated in green, and the partially observed cells are indicated in yellow.

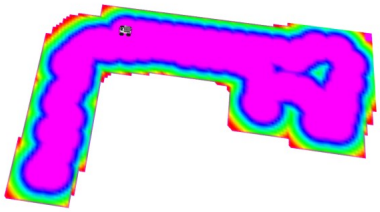

(b) 2D spatial coverage map. The rainbow scheme indicates the coverage score, with magenta indicating full coverage and red/white indicating no coverage.

Figure 6.3: We use 3D and 2D map structures to monitor the progress of the coverage mapping and the basis for predicting the viewpoint utility.

6.4.1 3D Surface Coverage Mapping

A common exploration goal is to create a dense, geometric 3D map of the environment by accumulating range measurements, e.g., from a LiDAR or an RGB-D camera. Range sensors allow the observation of surfaces from a distance, which makes predictions about the coverage from an unvisited viewpoint non-trivial as the sensor FOV and occlusions constrain the observation space of a given surface segment. Therefore, efficient surface coverage approaches need to estimate the expected surface coverage for potential viewpoints. Furthermore, the position can impact the quality of the surface observation. Lidar measurements close to the robot can be assumed to be more accurate than far-away readings. Uncertainties in the calibration or the orientation of the pose estimate propagate over long distances and induce larger modeling errors further away from the robot.

Coverage Mapping We maintain a 3D coverage grid G (see Figure 6.3a) to store the observation status of the 3D environment. As naively storing dense 3D data for large environments induces infeasible memory requirements, we store the 3D coverage grid as a spatially hashed 3D map [121, 83]. Each cell stores a floating value representing the observation score for the cell. A higher score indicates a better coverage and more accurate observation of the part of the environment. Each cell is initialized with zero and increased for each LiDAR point observation falling into the cell. Therefore, cells with a value of zero indicate unknown segments, cells with a value larger than zero and smaller than a threshold value $\tau_{observed}$ indicate partially observed segments, and cells surpassing $\tau_{observed}$ indicate a fully observed segment.

To update the coverage grid with a SLAM registered point cloud $H = \{\boldsymbol{h}_1, \boldsymbol{h}_2, \ldots, \boldsymbol{h}_n\}$ with the scan points $\boldsymbol{h}_i \in \mathbb{R}^3$ in the map coordinate frame, we increase the coverage score for each point falling into a cell. To account for the depth uncertainty, we apply an update scheme that reduces the coverage gain with increasing distance between observation and sensor. With a minimum and maximum distance thresholds d_{min} and d_{max}, the updated value for the cell $G(\boldsymbol{h})$ enclosing the scan point $\boldsymbol{h}$ can be computed as

$$G(\boldsymbol{h}) := \begin{cases} G(\boldsymbol{h}), & range(\boldsymbol{h}) > d_{max} \\ G(\boldsymbol{h}) + \frac{d_{min}}{range(\boldsymbol{h})}, & d_{min} < range(\boldsymbol{h}) < d_{max} \\ G(\boldsymbol{h}) + 1, & \text{else} \end{cases} \tag{6.1}$$

with $range(\boldsymbol{h})$ indicating the distance between the sensor and the observation. The update scheme is applied

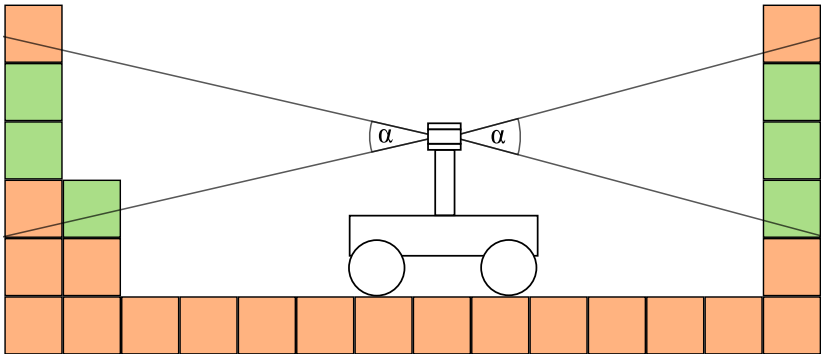

Figure 6.4: The approximated lidar visibility estimation considers the field of view α and the maximum range r_{max} of the sensor to predict which parts of the coverage grid can be observed from a viewpoint. The observable cells are indicated in green, and the unobservable cells are indicated in orange.

for all points in the point cloud.

Utility Estimation We apply an approximated lidar-visibility estimation to predict the utility u_v of visiting a viewpoint $\boldsymbol{v}$. The visibility estimation accounts for the vertical sensor field of view α and the maximum visibility distance threshold r_{max}. The locations of all partially observed cells P are stored in a kd-tree, which allows efficient range-based queries, such as a radius search, which gives all cells within a radius around a query position. Furthermore, we define a value n_{max}, the number of cells in the range needed to achieve the maximum utility score. We first use the radius search to compute the subset of cells $\boldsymbol{P}_r$ within visibility distance r_{max} to the sensor origin p_{sensor} as

$$\boldsymbol{P}_r = \{\boldsymbol{p} \in \boldsymbol{P} \mid \quad || \, \boldsymbol{p} - \boldsymbol{p}_{sensor} \, || < r_{max}\}. \tag{6.2}$$

In the second step, we remove all points outside the vertical field of view and perform a simplified occlusion test for neighboring cells to remove occluded cells. We denote the remaining points as $\hat{\boldsymbol{P}}$. The raw utility score $u_{\boldsymbol{v}}^{raw}$ for the viewpoint computes as the ratio between the expected number of perceived cells and the maximum number of perceived cells threshold n_{max} as

$$u_{\boldsymbol{v}}^{raw} = | \, \hat{\boldsymbol{P}} \, | \, / n_{max}. \tag{6.3}$$

To avoid repeatedly visiting the same position, we apply a penalty factor to viewpoints close to the position already visited. For each visited position in the distance, we apply a penalty proportional to the distance to the visited viewpoint. We denote the penalty distance $d_{penalty}$ and the robot trajectory $\boldsymbol{T} = \{\boldsymbol{t}_1, \boldsymbol{t}_2, \ldots, \boldsymbol{t}_n\}$ as a set of discrete positions $\boldsymbol{t}_i \in \mathbb{R}^3$ sampled with constant time intervals. The penalty factor $c_{penalty}$ for a viewpoint location p_v computes as:

$$\boldsymbol{T}_v = \{\boldsymbol{t} \in \boldsymbol{T} \mid \quad || \, \boldsymbol{t} - \boldsymbol{p}_v \, || < d_p\}, \tag{6.4}$$

$$c_{penalty} = \prod_{\boldsymbol{t}_i \in \boldsymbol{T}_p} \frac{||\boldsymbol{t}_i - \boldsymbol{p}_v||}{d_p}. \tag{6.5}$$

The term is zero if the viewpoint was exactly visited, which leads to zero utility, and therefore, the viewpoint will not be considered for planning anymore. If there is no viewpoint within the penalty distance, no discount is applied. The expected surface coverage utility $u_{\boldsymbol{v}}$ of a viewpoint $\boldsymbol{v}$ computes as

$$u_{\boldsymbol{v}} = c_{penalty} u_{\boldsymbol{v}}^{raw}. \tag{6.6}$$

6.4.2 2D Spatial Coverage Mapping

Another common goal is the creation of maps with sensors observing quantities directly at the sensor location, such as a dosimeter measuring the dose rate at the sensor location. We assume a sensor is rigidly mounted at a robot. Therefore, the observable space is locally 2D-constrained by the ground surface.

For applications such as dose rate monitoring or the search for hazardous substances, it is crucial to sample space at a sufficient spatial density.

Coverage Mapping We maintain a 2D-grid map to store the observation status of the environment. Each grid cell contains a floating value indicating the truncated distance to the closest observation. We define the truncation distance d_τ. We choose the target distance between two sensor observations as the truncation distance.

For each measurement, we update all cells in the truncation distance around the measurement location $\boldsymbol{p}_m$ to compute the minimum truncated distance to a sensor reading as

$$M(\boldsymbol{p}_c) := \begin{cases} M(\boldsymbol{p}_c), & range(\boldsymbol{p}_m) >= M(\boldsymbol{p}_c) \\ range(\boldsymbol{p}_m), & else. \end{cases} \tag{6.7}$$

Utility Estimation The utility of a viewpoint can be directly inferred from the coverage map distance value at the viewpoint location. Given the target sampling distance d_s. We assign no score to all viewpoints with a map value smaller than $\frac{d_s}{2}$. Between $\frac{d_s}{2}$ and d_s, we apply a linear relation to the utility. All viewpoints with a distance large d_s yield the maximum utility score, leading to the following formula for the computation of the spatial utility $u_{\boldsymbol{v}}$ of a viewpoint $\boldsymbol{v}$ at position $\boldsymbol{p}_v$

$$u_{\boldsymbol{v}} := \begin{cases} 0, & M(\boldsymbol{p}_v) <= \frac{d_s}{2} \\ \frac{2(M(\boldsymbol{p}_v) - \frac{d_s}{2})}{d_s}, & \frac{d_s}{2} <= M(\boldsymbol{p}_v) <= d_s \\ 1, & else. \end{cases} \tag{6.8}$$

6.4.3 Topological Graph

Efficiently planning over large distances requires a suitable abstraction of traversable space. Topological graphs represent the traversability of the environment as a sparse graph structure (see Figure 6.5), with nodes indicating positions and the edges indicating a connection to another node typically annotated with a distance metric.

An established method to compute such graphs is the usage of Voroni diagrams [156, 170]. We use a 2D traversability map [56] to derive the graph by applying a series of image operations.

Given the 2D traversability map represented as a binary 2D image, perform the following operations:

- Inflation: Inflate the occupied space by the radius of the robot to only account for traversable space.
- Distance Transform: Compute the distance transform, indicating the distance to the next surface, for the free space
- Skeletonization: Skeletonize the binarized distance transform image to gain a sparse representation of the environment.
- Graph extraction: Iterate over the skeletonized image to extract nodes and edges.

The result is a sparse graph of the environment, which allows efficient planning of trajectories even over large distances with standard graph search algorithms.

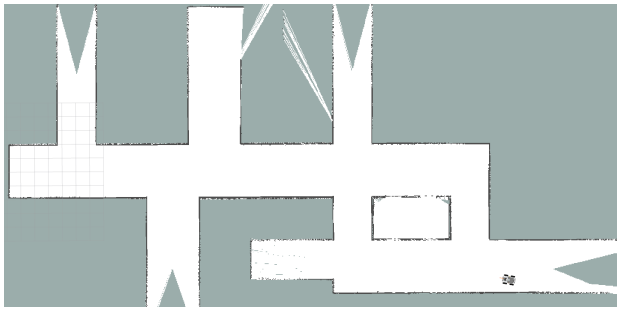

(a) 2D traversability grid map

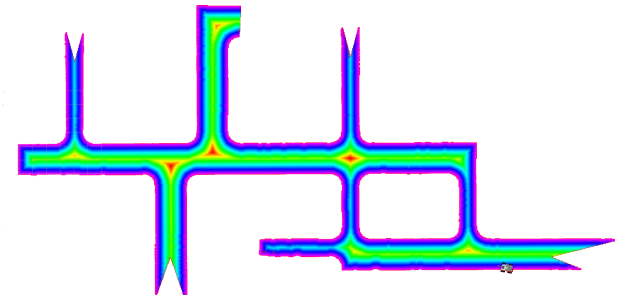

(b) Distance Transform

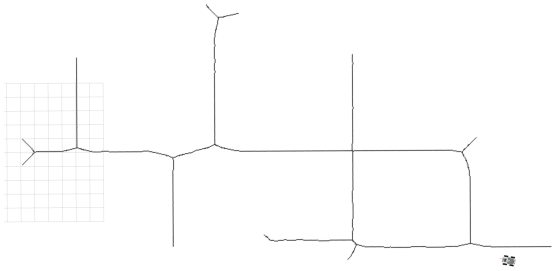

(c) Skeletonized Binarized Distance Transform

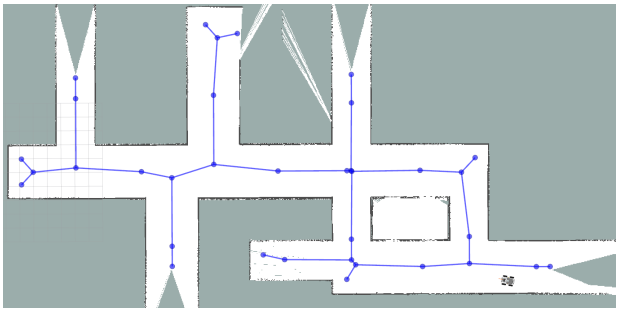

(d) Topological graph (blue) overlaying the occupancy grid map

Figure 6.5: Topological graph generation process. We compute a distance transform (b) from the 2D traversability map (a). Then, we skeletonize the binarized traversibility map (c) and extract the topological structure (d).

6.4.4 Actionable Instance Representation

During an exploration mission, the robot might encounter objects that allow for interaction, such as a door that can be opened or stairs that can be traversed. Such interactions can also require dependencies, such as a detected fire can be extinguished, but in order to perform the fire extinguishing action, the robot needs to acquire this capability first by picking up a fire extinguishing device. This requires a suitable, actionable representation of the environment semantics.

To allow complex interactions, we maintain a 3D map of semantic object instances and inscribe the instances with affordances. An affordance represents an action that can be performed on an object instance (e.g., the action *pick-up* on an object instance of a *tool* in the environment). We define an affordance by its requirements and effects. As we show in [8], additionally, taking priorities and costs into account allows the autonomous planning and execution of complex missions.

To generate the semantic 3D map, we follow the approach implemented in *hector_worldmodel* [106]. The point clouds registered by the SLAM system are accumulated in an Occupancy Grid represented as Octree. Camera images are semantically segmented (e.g., by using YOLO [164]). The masks of relevant objects are then projected onto the occupancy grid and either merged with existing object instance observations or added as a new object instance. The affordances are pre-defined per class and added to each instance.

The overall method was initiated and high-level conceptualized by the author. Most of the detailed work, especially for detailed conceptualization, implementation, and evaluation, was performed by Frederik Bark [8]. The author's main contributions are work on the general conceptualization of affordance-based mapping and planning approach and the integration into the exploration concept proposed in this thesis.

6.5 Online Multi-Goal Exploration

The proposed approach for efficient online multi-goal exploration is based on the hierarchical coverage exploration approach TARE [25]. Similar to TARE, we perform a hierarchical approach that samples viewpoints in the vicinity of the robot, selects the best viewpoints according to the utility function, and then solves a TSP problem over the local and global viewpoints.

6.5.1 Local Viewpoint Sampling

We sample the segment containing the robot and the adjacent segments with a regular grid pattern in equal distance d_{vp} viewpoints. Each viewpoint is checked for reachability with the distance transform map used for the topological graph generation. All unreachable points are sorted out. For each reachable view point $\boldsymbol{v}$, we compute the utility function scores $\mathbf{u}(\mathbf{v}) = (u_1(\mathbf{v}), u_2(\mathbf{v})...)$ as the element-wise minimum of the spatial and surface coverage utilities and maintain the maximum utility as viewpoint score $u^*(\mathbf{v}) = \max \mathbf{u}(\mathbf{v})$. We assume a maximum number of points n_{max} to consider in the planning and a minimum utility threshold u_{min} required to add a point in the planning. We add the viewpoint with the highest score to the planning horizon. Then, reduce the score of the remaining viewpoints to account for the expected coverage as

$$u^*(\boldsymbol{v}) := \begin{cases} u^*(\boldsymbol{v}), & ||\boldsymbol{v} - \boldsymbol{v}_{inserted}|| > d_{dropoff} \\ \frac{||\boldsymbol{v}-\boldsymbol{v}_{inserted}||}{d_{dropoff}} u^*(\boldsymbol{v}), & \text{else} \end{cases} \tag{6.9}$$

The insertion and utility reduction are repeated until the maximum number of points n_{max} is reached or no point surpasses the minimum utility threshold u_{min}. In the case that none of the initial utility scores surpass u_{min}, the local area is considered fully explored, and only global waypoints from other segments are considered for the current planning step.

6.5.2 Global Viewpoint Sampling

For each segment that is not part of the local search horizon, the viewpoints are sampled similarly to the local sampling. However, only the best viewpoint for each segment is added to the global planning horizon, presuming it surpasses the minimum utility threshold u_{min}. If no viewpoint surpasses the utility threshold, no viewpoint is added, and the segment is considered to be fully explored. As the results for each segment only change when new observations are added, we cache the result and only reevaluate it after new observations have been added to the segment.

6.5.3 Route Planning

Given both the local viewpoints from the local sampling and the global viewpoints from the global sampling, we are looking for the fastest route to cover all viewpoints. This is an instance of the NP-hard TSP. Following [128], we formulate the problem as an integer linear programming problem such as in the Miller–Tucker–Zemlin formulation [108]. Given the set of viewpoint indices $V = \{v_1, v_2, \dots, v_n\} \in \mathbb{R}^n$ referencing to n viewpoints v and the distance matrix $\boldsymbol{D} \in \mathbb{R}^{n \times n}$ with each scalar entry d_{ij} indicating the distance between the ith and jth viewpoint the problem can be formulated as

$$\min \sum_{i=1}^{n} \sum_{j \neq i, j=1}^{n} d_{ij} x_{ij} \qquad i, j = 1, \dots, n; \tag{6.10}$$

$$s.t. \sum_{i=1, i \neq j}^{n} x_{ij} = 1 \qquad j = 1, \dots, n; \tag{6.11}$$

$$\sum_{j=1, j \neq i}^{n} x_{ij} = 1 \qquad i = 1, \dots, n; \tag{6.12}$$

$$u_i - u_j + 1 \leq (n-1)(1 - x_{ij}) \qquad 2 \leq i \neq j \leq n; \tag{6.13}$$

$$2 \leq u_i \leq n \qquad 2 \leq i \leq n. \tag{6.14}$$

With $x_{ij} \in \{0, 1\}$ indicating whether the edge between i and j is part of the solution and the helper variables u indicating the order of the viewpoints and guaranteeing a single route as a solution. We solve the problem with the guided local search method implemented in the *Google OR-Tools Routing Library* [60]. We approximate the cost of traveling from one viewpoint to another by the shortest path between the points in the topological graph. The distance matrix $\boldsymbol{D}$ containing the shortest path distance between all viewpoints can be efficiently precomputed with the Floyd-Warshall algorithm [33].

6.6 User-Interaction Concepts for Efficient Exploration and Data-Acquisition

The proposed approach is based on the concept of shared autonomy (see Figure 6.6), where an operator can flexibly choose and switch the degree of autonomy of the robot. As the default setting, we consider an operator controlling the system in supervised autonomy mode. The operator provides a high-level goal, such as exploring an environment, with further specifications that should be considered, such as the coverage of the environment's surface for creating a 3D model or the spatial coverage with a sensor, e.g., for creating a map of the dose rate or thermal distribution in the environment. The system then transparently represents the state towards these goals and the currently planned action sequence. At all times, the operator can switch continuously between autonomy levels. In a semi-autonomous control mode, for example, to provide a waypoint navigation goal, the operator considers relevant or switches to assisted teleoperation and commands

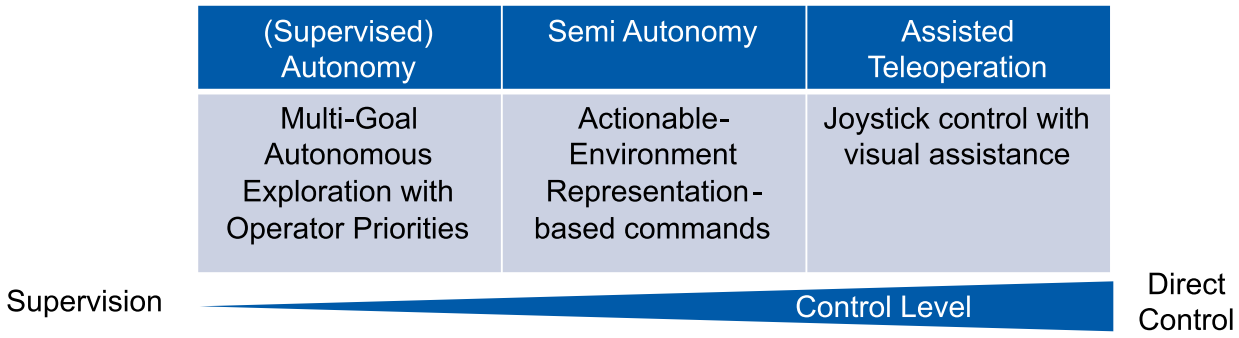

Figure 6.6: The user-related exploration and data-acquisition approach follows the shared autonomy principle. Allowing the operator to choose the level of autonomy and switch as needed flexibly.

the system via joystick commands, e.g., for difficult obstacle traversals or manipulation tasks. Afterward, the operator can always switch back to an increased autonomy level and continue the autonomous exploration. This allows a trade-off between the operator's cognitive load and the system's safety and action space.

The implementation into a User Interface is shown in Figure 6.7. The interface is based on RViz[1] and the QML-based user interface (UI) tools [55] by Fabian and Stryk. Transparent cuboids visualize the segments and provide interactions via context menus to prioritize or ignore a segment in the exploration. The color indicates the exploration status. Furthermore, the operator can set the quality requirements for spatial and surface mapping per segment. This allows scenarios with varying coverage requirements, such as a general search of the building with a more detailed search of a room of interest. The current path and the selected local and global viewpoints are visualized as 3D markers. The coverage exploration can be paused and resumed by a large play/pause button. When the operator manually starts a behavior, such as driving to a waypoint, performing an object affordance action, or providing a joystick input, the coverage exploration is paused automatically. In pause mode, the coverage assistance still provides updated visual assistance. The operator can then continue in assisted teleoperation or semi-autonomous operation. Pressing the play/pause button can activate the autonomous operation of the coverage exploration at any time, which also accounts for the progress during the pause. Detected objects are indicated by markers, and potential affordance-related actions are indicated by context menus.

6.7 Online Model Generation

To accumulate sensor data and transfer it in a readable and understandable depiction, we consider three different types of models, which can be directly generated online on the robot.

6.7.1 2D and 3D Radiation Mapping

Monitoring potential radiation sources with mobile robots requires the creation of accurate maps of the dose rate distribution. Shadowing, multi-source environments, and sensor errors can lead to implausible results for mapping methods, which make strong model assumptions. In [152], we propose a Gaussian Processed-based method that requires only weak model assumptions and gains efficiency by leveraging pre-sampling and local map update schemes. The author's main contributions are work on the general conceptualization of the mapping scheme, sensor modeling, and evaluation. Most of the detailed work, especially for implementation and detailed evaluation, was performed by Jonas Süß and Martin Volz. The resulting method can predict the

[1] `http://wiki.ros.org/rviz`

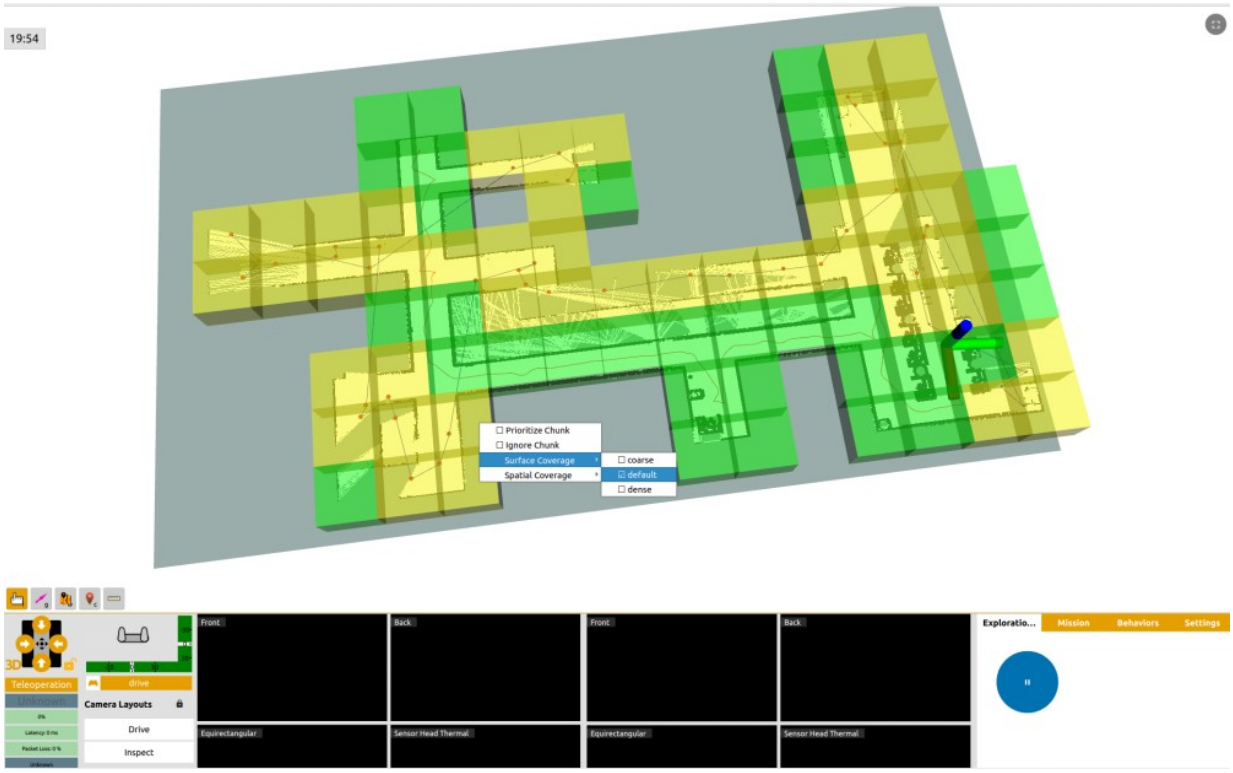

Figure 6.7: Visualization of the UI integration of the user interaction concept.

dose rate distribution in complex indoor environments with multiple sources and quantify the uncertainty in their estimates.

6.7.2 Point Cloud Accumulator

3D point cloud models of the environment provide accurate assessments of the geometric relation of the environment. However, with modern lidar sensors providing 5.2 million points per second[2], naively accumulating the registered point cloud observations quickly becomes computationally infeasible. Therefore, we use *pointcloud_accumulator*[3] developed by Jasper Süß, which uses an incremental kd-tree [21] to compute a globally voxel-grid filtered point cloud efficiently.

6.7.3 Panorama Tour Generator

While 3D point clouds can provide accurate geometric representations of the environments, camera images can often capture higher-resolution visual information about the environment. Furthermore, camera images are easy to interpret for human operators. Therefore, we automatically create panorama tour models of the environment. The overall method was initiated and high-level conceptualized by the author. Most of the detailed work, especially for implementation and evaluation, was performed by Jonas Süß. The method uses a created 2D map of the environment and combines it with panorama images taken by the robot. The images are filtered for quality measures such as brightness and sharpness. The web-based user interface allows flexible and natural interactions between the viewpoints.

[2]Ouster OS0-128, `https://ouster.com/products/hardware/os0-lidar-sensor`, accessed 18.02.24
[3]`https://github.com/tu-darmstadt-ros-pkg/pointcloud_accumulator`, accessed 18.02.24

7 Evaluation and Transfer

This chapter provides an overview of the proposed methods' evaluation results and their transfers to practical applications.

The reproducible evaluation of capabilities for mobile robots is challenging, as mobile robots are complex systems containing many hardware and software components that interact with each other. This introduces many potential sources for measurement uncertainties and errors and requires large integration efforts. Therefore, we consider versatile evaluation techniques with different levels of realism. We combine evaluation with external and own benchmark data sets, simulations, and trials under the conditions of robotics competitions to reach a comprehensive result on the performance of the proposed approach with respect to the targeted application and related approaches. In the following, this chapter provides a comprehensive overview of evaluations of the proposed SLAM approach, followed by evaluations of the proposed exploration approach. It concludes with examples of practical transfer of the research results.

Parts of this chapter have been previously published. The evaluations of the proposed SLAM approach have been published in [37, 39] and have been revised for the presentation in this chapter.

7.1 Evaluation - Robust Simultaneous Localization and Mapping in Challenging Environments

To comprehensively evaluate the developed SLAM system, we consider a wide variety of environments, data sets, and applications. We first investigate the 2D SLAM accuracy in large-scale environments on two data sets of the Radish dataset [76], comparing the proposed approach to related methods. Subsequently, we analyze the accuracy and robustness of the 3D mapping in rough terrain with aggressive motions in the ground truth annotated set data of the created DRZ Living Lab Motion Capture Dataset. Then, we investigate the behavior of large-scale 3D mapping in a challenging mixed indoor-outdoor environment. Afterward, we evaluate the applicability of the SLAM system under the conditions of operation in robotic competitions. Finally, in the last section of this chapter, we provide an outlook on evaluation as a transfer for an actual disaster deployment of the preliminary robotic task force of the German Center of Rescue Robotics.

7.1.1 Radish Dataset - 2D Large Scale SLAM

To evaluate the large-scale mapping capability in 2D environments and compare the proposed approach to other methods, we evaluate on benchmarks from the radish data set [76] using the relative displacement error metric suggested in [96]. The measure compares the error in relative poses with respect to manually annotated ground truth relations. We choose two data sets with different characteristics. The MIT Killian Court provides more than two hours of data with loop closures over large distances and long hallways, whereas the Freiburg Bldg. 79 is a more narrow and cluttered environment with loop closures on smaller scales.

As the data sets differ in sensor configurations and characteristics, we adjust the parameters for each data set individually. For MIT Killian Court we use a grid resolution $r = 0.075\,\mathrm{m}$ and a truncation distance $\tau = 0.15\,\mathrm{m}$. For Freiburg Bldg. 79 we use $r = 0.1\,\mathrm{m}$ and $\tau = 0.15\,\mathrm{m}$. On both data sets, we use odometry information additional scan matcher with an angular search window of $0.1\,\mathrm{rad}$ to improve the scan matching initialization.

Table 7.1: Radish Dataset - Quantitative Error Comparison

	TSDF Projective	TSDF Euclidean	Cartographer	Graph Mapping [96]
MIT Killian Court				
Absolute translational [m]	**0.0276 ± 0.0232**	**0.0276 ± 0.0235**	0.0324 ± 0.0270	0.050 ± 0.056
Squared translational [m^2]	**0.0013 ± 0.0089**	**0.0013 ± 0.0095**	0.0018 ± 0.0099	0.006 ± 0.029
Absolute rotational [°]	0.2807 ± 0.2462	**0.2802 ± 0.2435**	0.3183 ± 0.2883	0.5 ± 0.5
Squared rotational [deg^2]	0.1394 ± 0.26865	**0.1378 ± 0.2591**	0.1844 ± 0.4912	0.9 ± 0.9
Freiburg Bldg. 79				
Absolute translational [m]	**0.0382 ± 0.0292**	0.0391 ± 0.0298	0.0395 ± 0.0306	0.056 ± 0.042
Squared translational [m^2]	**0.0023 ± 0.0044**	0.0024 ± 0.0045	0.0025 ± 0.0048	0.005 ± 0.011
Absolute rotational [°]	0.4245 ± 0.4610	**0.4204 ± 0.4606**	0.4333 ± 0.4735	0.6 ± 0.6
Squared rotational [deg^2]	0.3926 ± 1.2308	**0.3887 ± 1.1806**	0.4118 ± 1.2475	0.7 ± 1.7

Table 7.2: Radish Dataset - Runtime Comparison

	TSDF Projective	TSDF Euclidean	Cartographer
MIT Killian Court			
Wall Clock Time [s]	81.0	74.8	103.5
CPU Time [s]	177.6	162.4	220.8
Memory [MB]	1201.4	1171.2	886.0
Freiburg Bldg. 79			
Wall Clock Time [s]	20.7	18.9	19.1
CPU Time [s]	68.1	107.3	71.7
Memory [MB]	228.6	245.8	215.9

As we were able to achieve better results with the current occupancy grid implementation in Cartographer than the results in the original publication, we use the current cartographer implementation performance, as measured by us. We choose the same settings for resolution and the additional scan matcher for occupancy grids as for TSDF.

Table 7.1 shows the results of our TSDF approach in comparison to the occupancy grids in Cartographer and Graph Mapping [96]. The two TSDF approaches achieve similar accuracy on both benchmarks. In comparison to occupancy grids, the TSDF approaches reduce the absolute translational error by $13.6\,\%$, the squared error by $27.8\,\%$ and the rotational errors by similar margins in the MIT Killian Court data set. In the Freiburg Bldg. 79 data set, the TSDF approaches perform slightly better than the occupancy grids with margins between $1\,\%$ and $5\,\%$.

A potential reason for the small improvement in the Freiburg data set is that it contains many small objects. TSDFs suffer from overwriting effects, which degrade the map fidelity when objects smaller than the truncation distance are observed from multiple viewpoints. This is a known issue also described in [84], a solution strategy to store multiple distances for multiple normal orientations per cell is outlined in [149].

A comparison of the runtimes and memory loads is shown in Table 7.2. For the MIT Killian Court data set, the TSDF requires less run time than the occupancy grid. In contrast, Cartographer is slightly faster on the Freiburg Bldg. 79 data set. Overall, the TSDF is faster in scan matching, but the map update requires more computations and the bounds for the loop closure branch-and-bound are slightly worse, yielding an overall similar runtime. For both data sets the TSDF approach needs $10\,\%$ to $20\,\%$ more memory. An increase in memory is to be expected as TSDF maps store two values, signed distance, and weight, per grid cell instead of one for occupancy grids.

7.1.2 DRZ Living Lab Motion Capture Dataset - Aggressive Motions on Rough Terrain

To evaluate the performance of the proposed approach in challenging terrains, we captured four sequences on challenging terrain. Each is tracked with a high-performance Qualisys optical motion capture system. We use the data of the motion capture as ground truth for the evaluation. The four sequences (see Figure 7.1) contain 1) double pitch ramps, which induce a fast pitch motion when traversing 2) a loose woodpile that slips when traversed 3) the RoboCup Rescue Robot League "Maneuvering 3 - Traverse" lane which contains a 2.4m long $30°$ incline and 4) the RoboCup Rescue Robot League "Mobility 4 - Elevated Ramps" lane containing a diagonal hill terrain consisting of $60\,\mathrm{cm}$ ramps with sloped tops. The sequences are between $59\,\mathrm{s}$ and $149\,\mathrm{s}$ in duration. The dataset is published in [38].

We compare the accuracy of the proposed HectorGrapher and Cartographer[71] with reference to the measurements from a visual motion capture system. Other state-of-the-art systems, such as LIO-SAM[147], are not included as, to our best knowledge, and after unsuccessful integration attempts, they do not support the spinning lidar configuration. We use the relative displacement benchmark measure suggested in [96], splitting the motion capture trajectory in $0.5\,\mathrm{s}$ sequences and comparing the errors for the translation and rotation component of each sequence.

The error metrics are shown in Figure 7.2. For the rotation component, both methods achieve comparable results, with outliers in the same value region. HectorGrapher performs slightly worse for small errors. In contrast, for the translation, significant differences are notable. The median error of Cartographer is more than two times the median error of Hectorgrapher in the woodpile and elevated ramps scenarios, and even in the continuous ramps and traverse scenario $40\,\%$-$60\,\%$ larger. Notable are also the outliers with high errors with $12\,\mathrm{cm}$-$14\,\mathrm{cm}$ in three scenarios, whereas the errors for HectorGrapher are no larger than $6\,\mathrm{cm}$. Both HectorGrapher and Cartographer are able to generate qualitatively comparable maps of the environment for moderate motions. For fast motions, such as pitching when traversing ramps, the improved localization accuracy can also be observed in the scan registration quality. Figure 7.3 shows the registered point clouds for

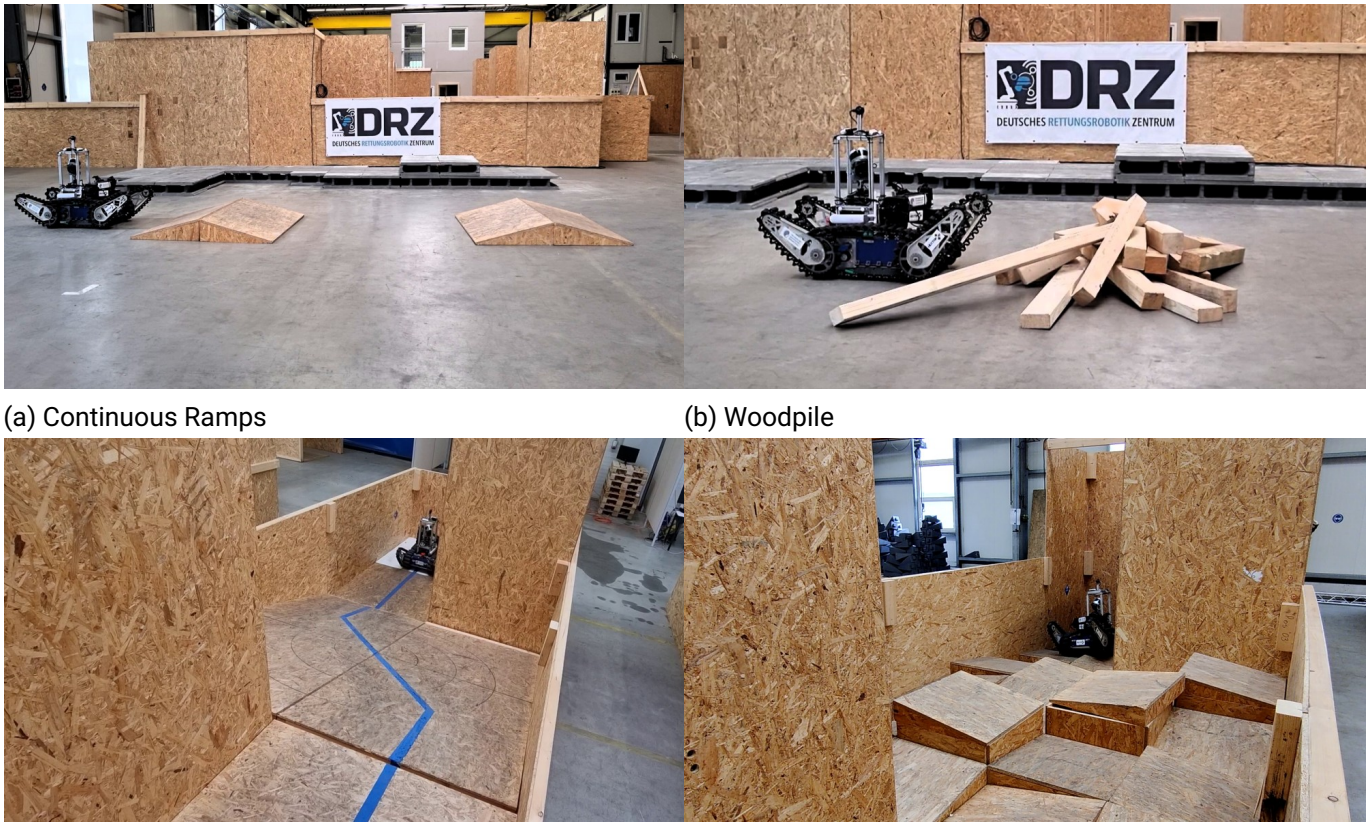

(a) Continuous Ramps

(b) Woodpile

(c) RoboCup Rescue Robot League: Maneuvering 3 - Traverse

(d) RoboCup Rescue Robot League: Mobility 4 - Elevated Ramps

Figure 7.1: Overview of the evaluation scenarios with ground truth data.

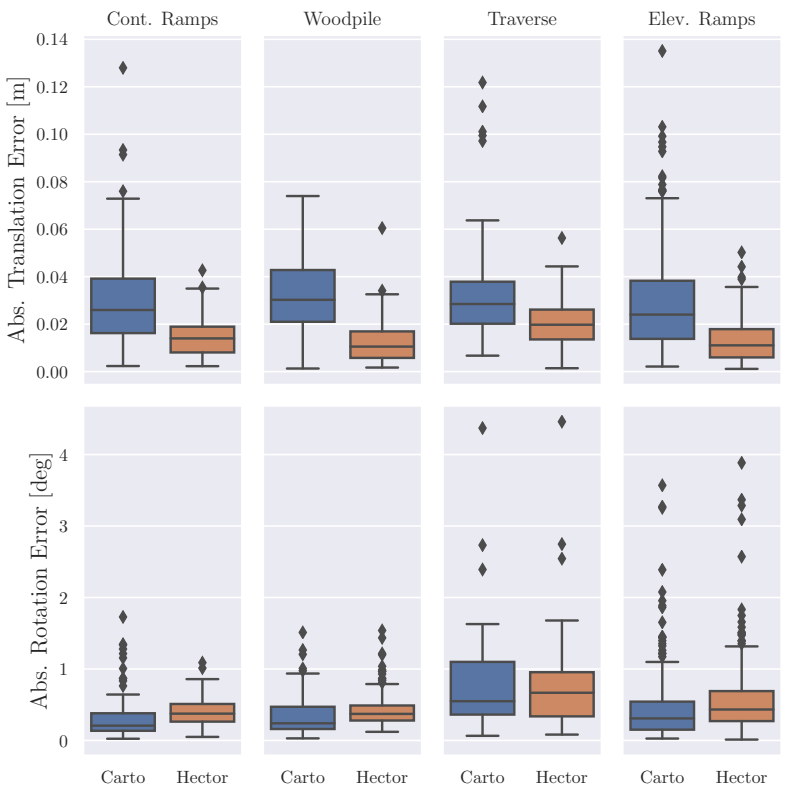

Figure 7.2: Quantitative Error Comparison of Cartographer (Carto) and HectorGrapher (Hector) in the four evaluation scenarios.

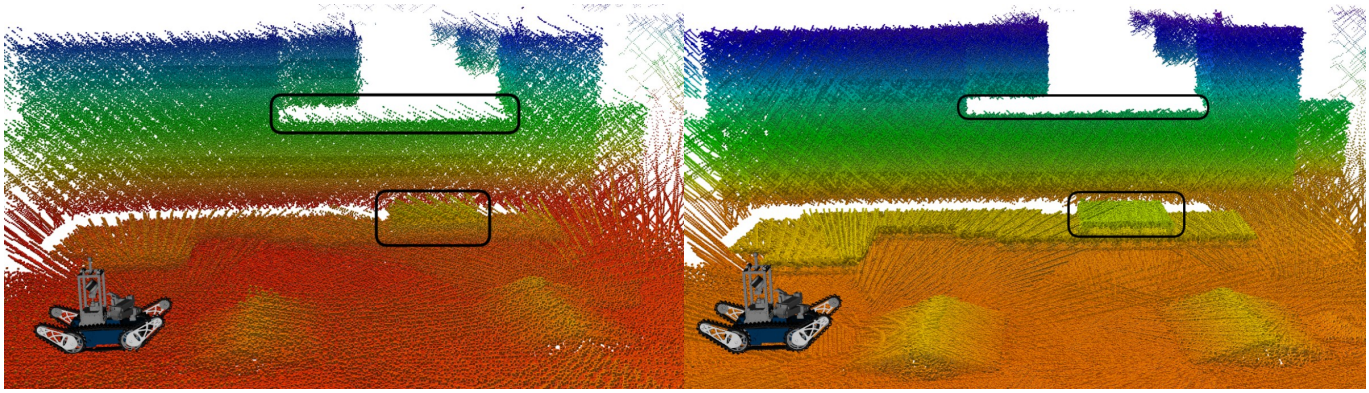

(a) Cartographer (b) HectorGrapher

Figure 7.3: Comparison of the registered point clouds (colored by height) in the "Continuous ramps" scenario. The Cartographer result shows shift artifacts at the pallet stack and the wall (marked by the black boxes), which are less notable in the HectorGrapher result.

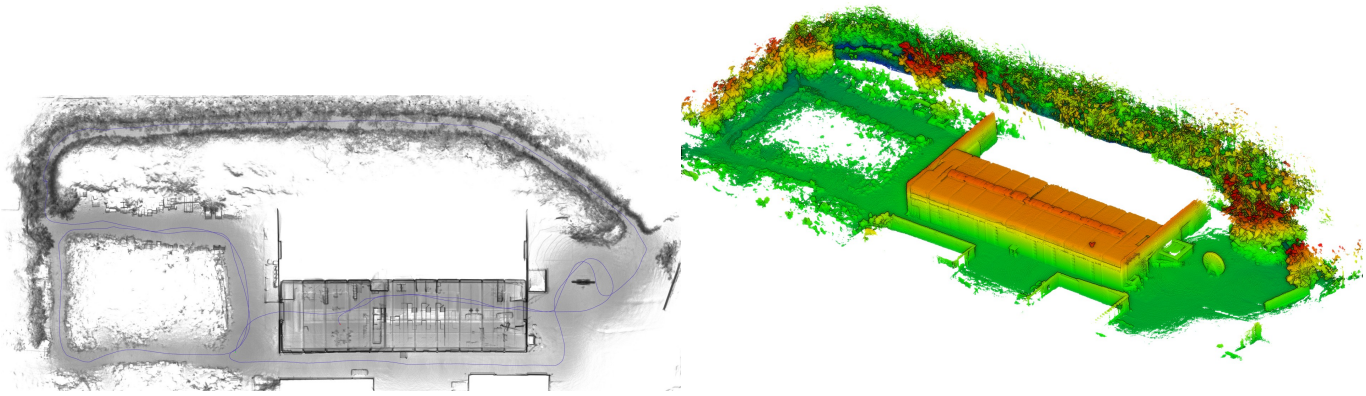

(a) X-Ray of 3D Point Cloud (b) 3D Point Cloud

Figure 7.4: Map, trajectory and point cloud for the DRZ Living Lab Loops scenario.

the "Continuous ramps" scenario. While the Cartographer result shows shift artifacts at the pallet stack and the wall, such artifacts are less notable in the HectorGrapher result.

Runtimes for HectorGrapher were 4-5.5 times higher than for Cartographer, which seems plausible as the optimization problem becomes significantly larger and the update of the TSDF requires more computation than the update of the occupancy grid in Cartographer.

7.1.3 Scout DRZ Loops - 3D Large Scale SLAM

To evaluate the proposed approach for large-scale mapping capabilities, we captured a mixed indoor-outdoor dataset with three loops at the DRZ Living Lab in Dortmund with the emergenCITY Scout robot. In contrast to the previous evaluations, this dataset also covers three loops and outdoor terrain, including an unpaved path through a scrapyard and an unpaved trail through a small forest. The unstructuredness and translucency of these environments make scan registration and mapping challenging. The dataset covers $704\,\mathrm{s}$ of data and a distance of approx. $920\,\mathrm{m}$.

As the data set is large in scale, we run the full SLAM pipeline with lidar odometry, loop closure detection, and pose graph optimization. emergenCITY Scout is equipped with an Ouster OS-0 128 scanner capturing data at more than 10x the rate of the Velodyne VLP-16. To keep computations manageable, we downsample the point cloud and only insert $10\,\%$ of the points in the TSDF.

The resulting map, trajectory, and point cloud are shown in Figure 7.4. The map overall shows a high consistency, with all three loops closed accurately. Geometry inside the building and outdoors on the scrapyard are sharp, and the forest track is mapped consistently. Minor registration artifacts are notable in the lower left side of the DRZ Living Lab building, indicating a small angular offset in the submap registration.

While setting up the configuration for the data set, we noted the tendency of the lidar odometry to erroneous rotational motions in the forest part, leading to a warping of the map. This issue could be resolved by increasing the weight of the wheel-inertial odometry rotation component in lidar-odometry optimization.

We performed the computations with an AMD Ryzen 7 3800X processor, taking $648.7\,\mathrm{s}$ wall time and $1832.6\,\mathrm{s}$ CPU time. In comparison to the RoboCup dataset, the increased point cloud data leads to increased time of the map update, while the loop closure detection is executed in separate threads and thereby mainly increases the CPU time but not the wall time. The peak memory usage was $2.03\,\mathrm{GB}$, and the real-time factor was 1.09, demonstrating the real-time capability of the proposed approach with current hardware.

Competition	Results
RoboCup 2018	Best in Class Exploration, Outdoor CarryBot and Small Robot Awards
World Robot Summit 2018	1st place in the Plant Disaster Challenge
RoboCup 2019 German Open	2nd place, Best in Class Exploration
EnRicH 2019	Radiation Mapping Award
RoboCup 2019	3rd place, Best in Class Exploration and Outdoor CarryBot
RoboCup 2021	Best in Class Dexterity and Best in Class Exploration and Mapping
RoboCup 2021 German Open	1st place and Best in Class Exploration
EnRicH 2021	3D Mapping and Radiation Mapping Award
RoboCup 2022 German Open	1st place and Best in Class Autonomy Award
RoboCup 2022	3rd place, Best in Class Exploration
AIRA Challenge 2022	4th place
RoboCup 2023 German Open	1st place, Best in Class Autonomy and Mobility
EnRicH 2023	Radiation Mapping Award
RoboCup 2023	2nd place, Best in Class Autonomy, Technology Challenge Award, and Best Team Description Paper

Table 7.3: Evaluations of the proposed SLAM system in competitions with Team Hector in chronological order. No competitions took place in 2020 due to the COVID-19 pandemic.

7.1.4 Evaluation in Robotic Competitions

Evaluations with data sets, such as in the previous sections, provide comparability as different methods can be evaluated with the same input. However, such data sets are limited in validity for real applications for multiple reasons. Biases can be induced by the selection of considered data sets, as the authors can choose which data sets to use for their evaluation and try their methods multiple times until they reach a satisfying result. Furthermore, published data sets are often cleaner than actual missions, where calibrations might not be optimal, operators might make errors, tasks might vary, the environment location and conditions cannot be chosen beforehand, and other processes might be running on the computer during the same time. In many aspects, evaluations under competition conditions are closer to real-world applications than evaluations with data sets, as the competition organizers create an environment with only limited knowledge by the operating team. Furthermore, there is typically only one or a few runs where the method needs to work with no or only limited chances for repetitions, emphasizing the focus of robustness on the approach as part of an overall integrated system.

The proposed SLAM system has been successfully evaluated by Team Hector in various development stages in a wide variety of competitions, enabling autonomous functions and mapping capabilities. An overview of the competitions and the team results is provided in Table 7.3.

RoboCup Rescue Robot League

At the RoboCup Rescue Robot League competitions, the robot needs to perform multiple tasks from the categories of maneuvering, mobility, dexterity, and exploration, which are derived from capabilities needed for actual disasters. While some tasks, such as exploration, require an autonomous operation (depending on the rule version), other tasks gain a score bonus for autonomous operation. In the exploration tasks, the submitted 2D and 3D maps are the main evaluation criteria. In the nine RoboCup (4x German Open, 5x World Championship) participations, the proposed SLAM system was applied in all competitions as the main SLAM system, enabling autonomous and semi-autonomous functions. The proposed system provided reliable

localization and accurate mapping even in narrow exploration arenas and on challenging terrain. Thereby, it played a crucial role in achieving high scores for good overall results and Best in Class prices for exploration and autonomy. A detailed evaluation of the data of the RoboCup Rescue Robot League 2021 follows in the next section.

European Robotics Hackathon (EnRicH)

The European Robotics Hackathon (EnRicH) takes place every two years at the inactive Nuclear Power Plant (NPP) Zwentendorf in Austria. The competition mimics nuclear accidents with active radiation sources. The robots need to search a part of the NPP and perform multiple tasks: 1) create a 3D map of the environment, 2) locate the radiation sources and create a map of the dose rate distribution, 3) perform manipulation tasks, and 4) find a missing worker and take him to a safe position. In 2023, the rule to autonomously explore at least parts of the environment was introduced. In the three participations, the proposed SLAM system enabled the team to create accurate 3D maps of complex, real-world environments and provided exploration assistance for the operator. The resulting artifacts were rewarded with the Award for Best 3D Mapping in 2021 and Best Radiation Mapping in all years, where the SLAM system provided the localization reference and geometric 3D map. Therefore, the proposed SLAM system was successfully evaluated and played a crucial role in the successful competition participation.

World Robot Summit and AIRA Challenge

The Plant Disaster Challenge at World Robot Summit 2018 and the AIRA Challenge 2022 evaluated the performance of mobile robots in industrial settings. In contrast to RoboCup or EnRicH, the exact structure of the environment with the location of relevant objects was known beforehand. Therefore, a map of the environment could be captured before the competition, allowing the SLAM system to operate in a localization mode against a static map. This allowed us to perform autonomous behaviors at the World Robot Summit, although during the competition, the operator teleoperated the robot using the localization estimate for assistance, as this was faster than the autonomous operation. At the AIRA challenge, it was required that the robot perform the mission fully autonomously. The localization information enabled the successful deployment of assistance functions to drive to waypoints and traverse stairs at predefined positions.

7.1.5 Evaluation at RoboCup Rescue Robot League 2021

We evaluated the proposed SLAM system for the RoboCup Rescue Robot League 2021, winning the best-in-class "Exploration and Mapping" award. Each team had to set up a scenario following the same rules, following the NIST guidelines for evaluation of rescue robots[1], distributing 10 barrels at two different heights as visual fiducials. The fiducials appear as circles in the 2D projection of the map and are utilized to measure accuracy and completeness. To make the terrain challenging, every $4.8\,\mathrm{m}$ had to contain a small obstacle (see Figure 7.5a), such as a wooden bar or a ramp. As part of the scenario we traversed the RoboCup German Open - "EXP 1 Map on Continuous Ramps" Arena (see Figure 7.5b) which is a narrow, $1.2\,\mathrm{m}$ wide corridor in waveform in a $7.2\,\mathrm{m}\times2.4\,\mathrm{m}$ area continuously paved with ramps. The ramps induce fast roll-pitch motions. In combination with the narrow environment, both tasks - navigation and mapping - are challenging.

The data set covers $622\,\mathrm{s}$ of data. As the data set is rather small in scale and does not cover large loops, we only perform the lidar odometry part and do not need to check for loop closures.

The resulting map, trajectory, and point cloud are shown in Figure 7.6. The map overall shows a high consistency, even in the narrow parts such as the EXP1 arena. The fiducials are clearly visible in the 2D

[1] `https://rrl.robocup.org/forms-guides-labels/`

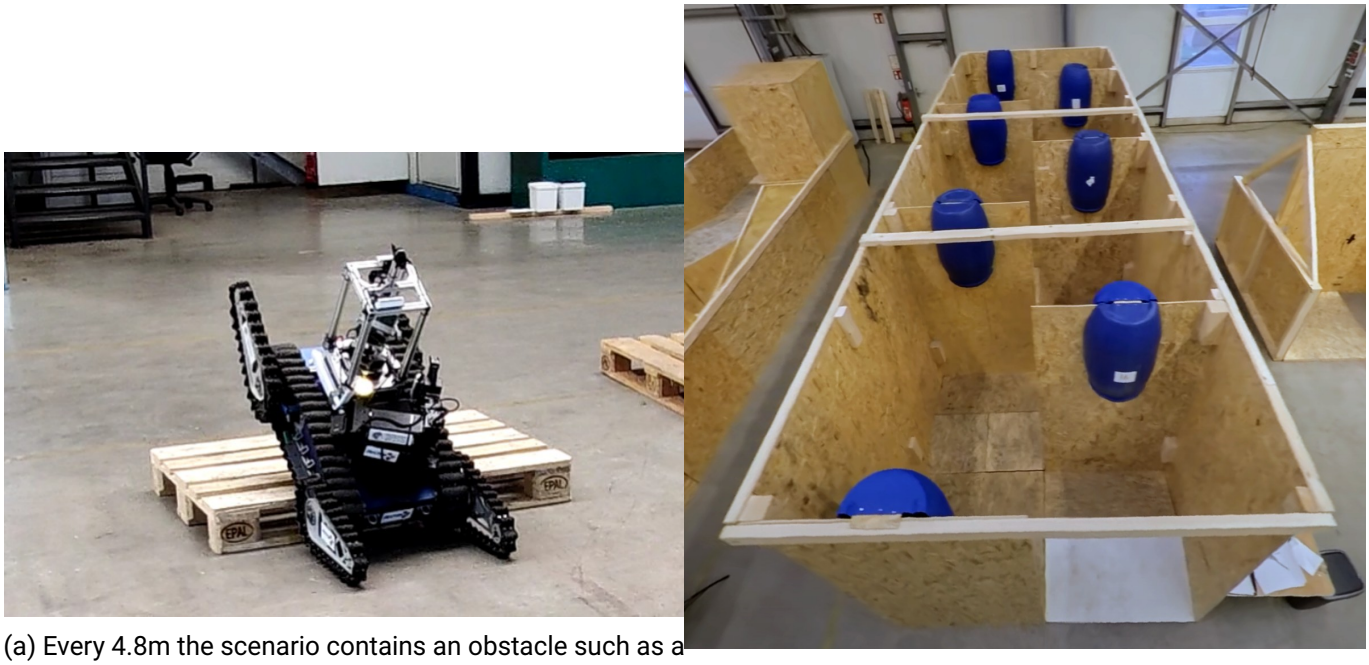

(a) Every 4.8m the scenario contains an obstacle such as a pallet, ramp or wooden bar.

(b) Aerial Photo of the EXP1 Arena with barrel fiducials.

Figure 7.5: The RoboCup Rescue Robot League 2021 scenario contains multiple small obstacles and the EXP1 exploration arena.

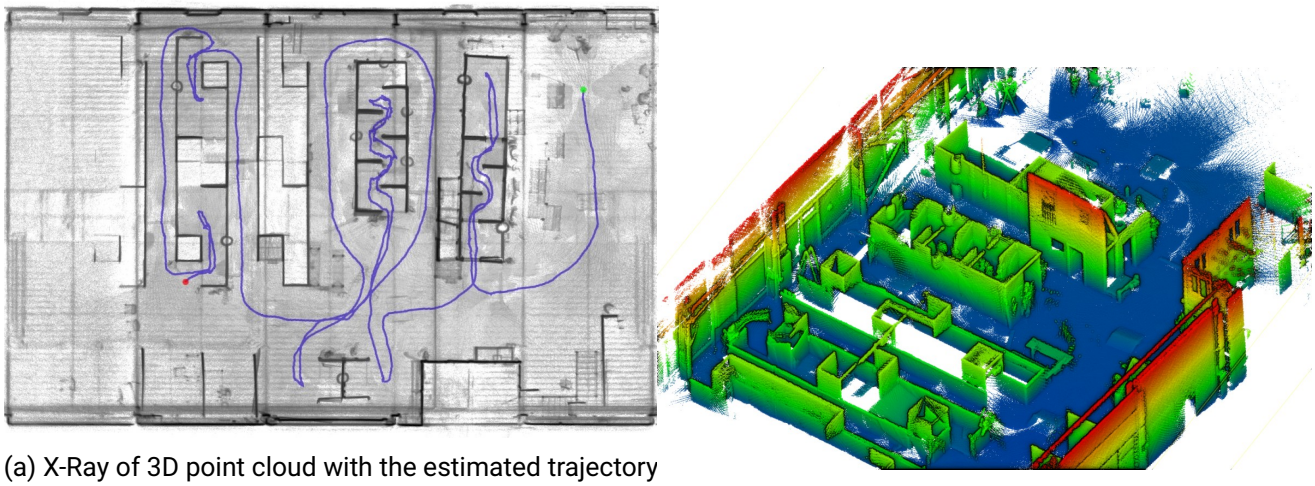

(a) X-Ray of 3D point cloud with the estimated trajectory indicated in blue

(b) 3D point cloud - Roof removed for better visibility

Figure 7.6: Map, trajectory, and point cloud for the RoboCup 2021 scenario.

Figure 7.7: Workshop exploration scenario. The scenario simulates a workshop exploration with the search for a potential radiation source. In the scenario, the simulated KIARA Telemax starts at the top left corner, the radiation source is placed on the rear table at the right end, and the door to the small room in the rear can be opened for further investigation. The scenario mimics an actual workshop and training area at the Kerntechnische Hilfsdienst GmbH in Karlsruhe, Germany.

projection and demonstrate accuracy and coverage. On a desktop computer with an AMD Ryzen 9 3900X processor, computations took $193.4\,\mathrm{s}$ wall time and $191.9\,\mathrm{s}$ CPU time on a single core with a peak memory usage of $321.39\,\mathrm{MB}$ yielding a real-time factor of 3.24. The low computational cost makes the approach well-suited for localization and mapping on mobile robots with limited hardware.

7.2 Evaluation - Operator-related Exploration and Data-Acquisition in Multi-Goal Missions

To investigate the proposed exploration concept, we consider two complex simulated scenarios. The first scenario mimics a workshop and investigates the approach in the context of the shared autonomy concept. The second scenario covers a large-scale environment and investigates the large-scale exploration performance.

7.2.1 Workshop Scenario - Shared autonomy exploration for radiation source localization

To investigate the suitability of the proposed approach for explorations in complex missions with operator interactions, we consider the scenario of an investigation of a workshop with a presumed radiation source. Figure 7.7 shows an overview of the scenario environment.

In the presumed scenario, the goal is to capture a 3D model of the environment and a 2D map of the dose rate with the simulated KIARA Telemax robot. As part of that, the radiation source should be localized as accurately as possible. Furthermore, potential environmental interactions such as a detailed investigation of relevant objects or a door opening to reach the rear room are required to explore the environment.

The overall workshop, which mimics an actual workshop and training area at the Kerntechnische Hilfsdienst GmbH in Karlsruhe, Germany, has a size of $23\,\mathrm{m} \times 10\,\mathrm{m}$ and covers various objects and obstacles. In the scenario, the robot starts at the top left corner, the radiation source is placed on the rear table at the right end, and the door to the small room in the rear can be opened for further investigation. The environment is simulated with the Webots simulator, which provides accurate modeling of physics and sensor data. To focus

on the exploration aspect, the ground truth data provided by the simulator is utilized for localization and semantic analysis of images.

Scenario Procedure

An example scenario procedure is shown in Figure 7.8a and the respective operator interactions and autonomy levels in Figure 7.8b. The robot starts in the rear-left corner. The remote operator initiates an autonomous exploration of the environment. During the exploration, the progress is visualized by remaining viewpoints and the coverage indication of the submaps. Furthermore, the environment model is continuously updated, and relevant semantic objects such as hazard signs or doors are added to the environment model. As the radiation dose rate increases in the right part of the scenario, the operator chooses an increased spatial coverage resolution by clicking on the marker for the respective submaps. Close to the first table, the measurements reach a high level and decrease again after the robot diverges from the table. Therefore, the operator decides to take manual control of the robot, driving close to the table and inspecting it in detail with the arm. The operator identifies the radiation source and continues the autonomous exploration assistance of the right part to check for additional sources. The exploration assistance continues to cover the right part of the workshop. The exploration assistance terminates, and there is no indication of further radiation sources. The operator decides to explore a door that was passed. Therefore, the operator clicks on the door marker and chooses the *go to* action. At the door, the operator takes manual control and uses the manipulator to open the door. As the room behind is very narrow, the operator manually navigates the robot through the room. The room does not contain suspicious objects or notable dose rate readings. Therefore, the operator manually drives the robot out of the room and sets a waypoint at the starting position, using the navigation assistance to drive back to the initial position.

Resulting maps of the dose rate and the room are shown in Figure 7.9. The radiation map clearly indicates the position of the radiation source. The localization error of the estimated radiation source 3D position with respect to ground truth is $0.052\,\mathrm{m}$. This clearly indicates where the source is located on the table and allows detailed planning for further reaction measures. The visual of the 3D map is complete, indicating a good surface coverage of the environment.

The scenario demonstrates the proposed approach's applicability to a scenario with flexible switches between operator and autonomous capabilities.

Exploration Strategy Comparison

To investigate the effect of different exploration strategies on the source localization accuracy, duration, and covered distance, we consider an autonomous exploration of the main hall of the workshop environment. As the arm is not used in the autonomous exploration and the exploration algorithm uses the robot body as a reference link, we simulate the radiation sensor at the center of the robot body. We consider four exploration strategies: 1) only surface coverage - "surface," 2) surface and coarse spatial coverage - "default," 3) surface and dense coverage - "dense" and 4) surface and coarse coverage in the left half and dense in the right half, similar to the previous scenario - "adaptive."

For each strategy, the scenario is repeated five times. The resulting radiation source localization errors, exploration durations, and covered distance are shown in Figure 7.10. The mean errors in position for the surface strategy are $2.6\,\mathrm{m}$, $1.8\,\mathrm{m}$ for default, $1.18\,\mathrm{m}$ for dense, and $1.26\,\mathrm{m}$ for adaptive. Thereby, dense and adaptive show comparable results, while default and surface each significantly worsen the source localization estimate. A denser spatial coverage should enable a more accurate estimate of the radiation source position, which is supported by the results. As the surface strategy does not directly define spatial sampling constraints, the spatial coverage is the coarsest, leading to the highest error. Exploration duration and covered distance

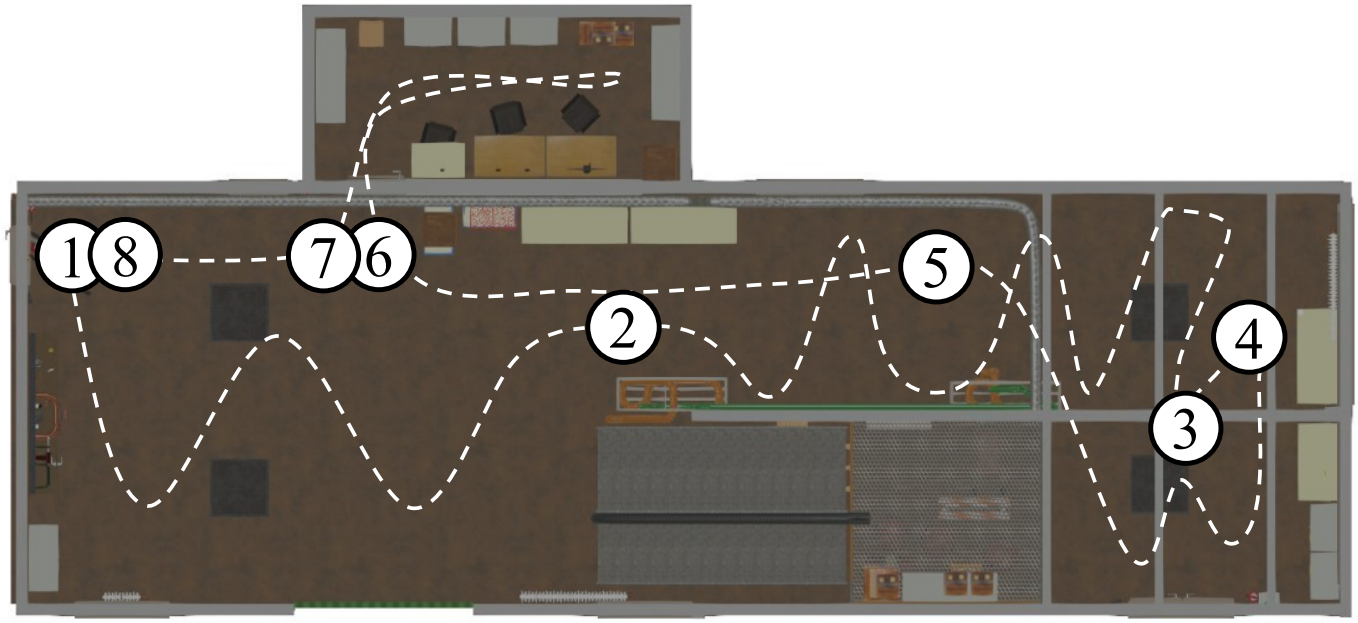

(a) Scenario procedure schematic. Simplified robot trajectory indicated by white dashes.

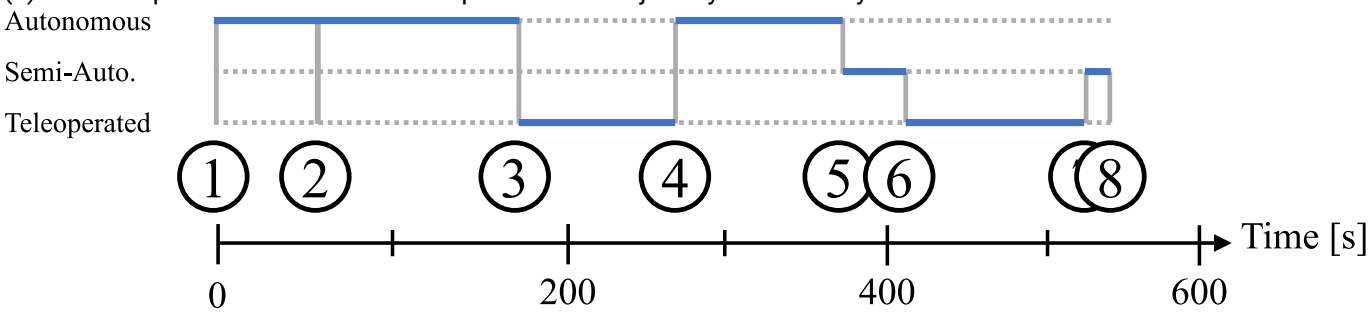

(b) Autonomy levels during the scenario procedure.

Figure 7.8: Workshop scenario procedure. Robot locations for relevant environment interactions are indicated by the numbers. (1) Initiation of the exploration. (2) Increase the spatial coverage resolution for the right half. (3) Pause of the autonomous exploration and manual inspection of the table. (4) Continuation of autonomous exploration (5) Exploration terminates, operator command to navigate to door. (6) Manual opening of the door and inspection of the small room. (7) Waypoint navigation to start. (8) End of mission.

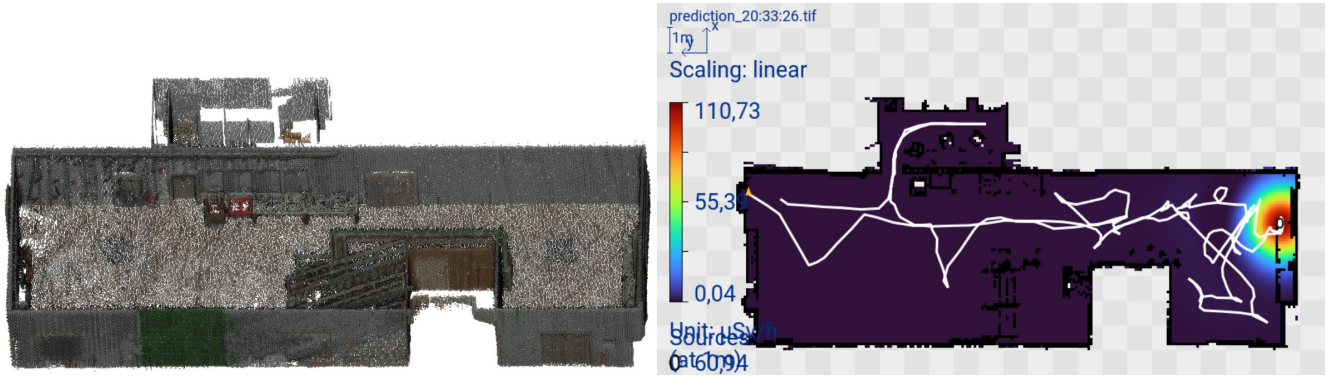

(a) Captured 3D point cloud of the scenario. The roof was removed for better visibility.

(b) 2D Radiation map with estimated source location. The white line indicates the sensor trajectory.

Figure 7.9: Resulting 3D point cloud and 2D radiation map of the workshop scenario procedure.

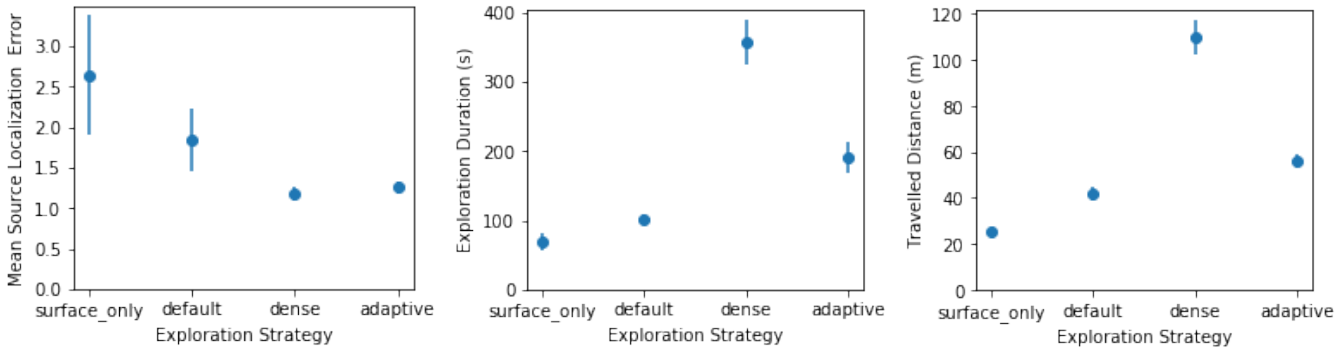

(a) Source localization error

(b) Exploration duration

(c) Covered distance

Figure 7.10: Evaluation of the source localization error, duration, and covered distance for different exploration strategies. Mean and standard deviation of 5 runs per strategy.

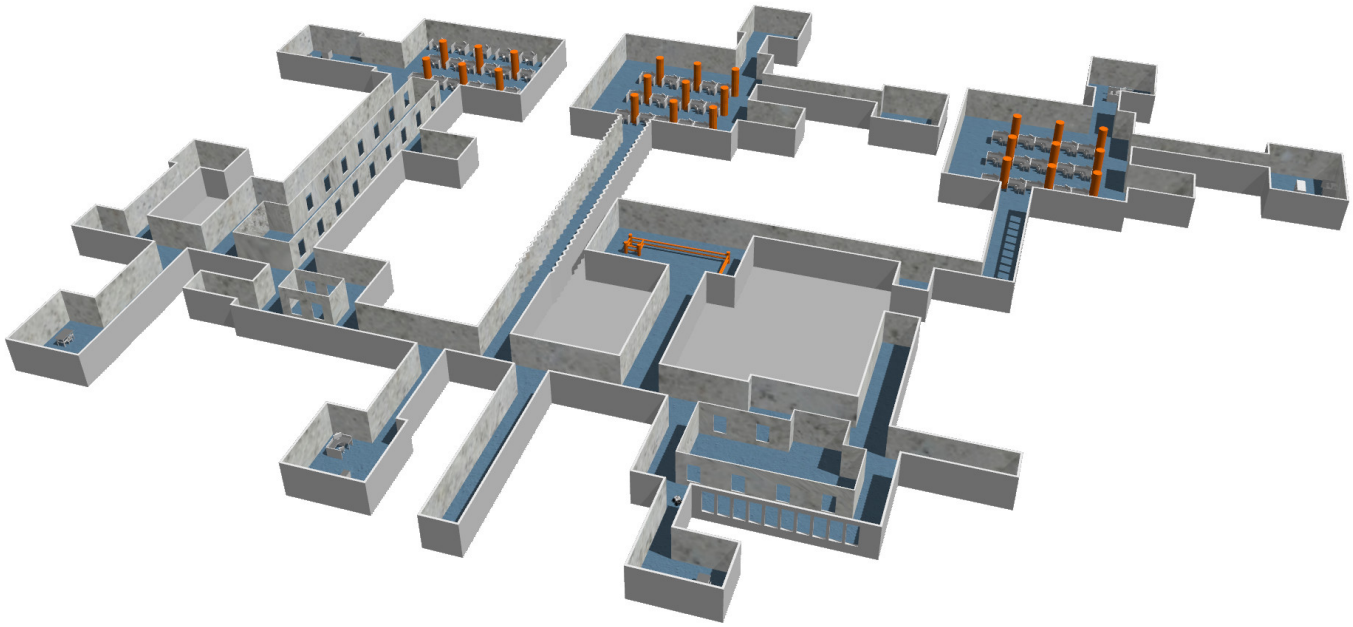

Figure 7.11: CMU indoor scenario. The scenario covers outer dimensions of $130\,\mathrm{m} \times 100\,\mathrm{m}$.

show very comparable results. The surface approach covers a distance of $25.2\,\mathrm{m}$ in $69.3\,\mathrm{s}$, default $41.5\,\mathrm{m}$ in $101.3\,\mathrm{s}$, dense $109.7\,\mathrm{m}$ in $356.8\,\mathrm{s}$ and adaptive $56.0\,\mathrm{m}$ in $191.9\,\mathrm{s}$. For surface, default, and dense, the results are reversed to the localization error. The dense approach is slower, by a factor of 5.1, compared to the surface approach and 3.5 to the default approach. The adaptive approach is still slower than the surface and default approach but 1.9 times faster than the dense approach, achieving a similar localization accuracy as only the part relevant to the source localization is densely mapped. The adaptive approach compromises between high localization accuracy and exploration duration, leading to increased exploration efficiency.

7.2.2 CMU Scenario - Autonomous exploration in large-scale environments

To investigate the behavior of the exploration assistance in large-scale environments, we consider the large-scale "Indoor Corridors" environment of the CMU Autonomous Exploration Development Environment [23] as shown in Figure 7.11. The scenario is large-scale with outer dimensions of $130\,\mathrm{m} \times 100\,\mathrm{m}$ and covers a convoluted environment with obstacles and thin structures.

We are considering a fully autonomous exploration with the emergenCITY Scout robot simulated in Gazebo. For the exploration, we consider three strategies: the proposed method with only surface coverage "*surface_only,*" the proposed method with combined surface and spatial coverage "*combined*" and a frontier-based exploration approach [88] "*frontier.*" To measure spatial coverage, we integrate the visited area with a distance of maximum $2\,\mathrm{m}$ to the robot trajectory by maintaining a coverage grid map with $0.1\,\mathrm{m}$ resolution. For the surface coverage, we maintain a 3D coverage grid with $0.1\,\mathrm{m}$ resolution and count the number of fully observed cells, which approximate the covered 3D surface of the environment.

For each strategy, the spatial and surface coverage during the exploration for three samples per strategy are shown in Figure 7.12. The combined approach achieves the highest spatial coverage of $3954\,\mathrm{m}^2$, the surface_only strategy follows with $3648\,\mathrm{m}^2$, $7.7\,\%$ less and lastly the frontier-based approach with $3140\,\mathrm{m}^2$, $20.6\,\%$ less. For surface coverage, the result is similar, with smaller relative margins of $4.2\,\%$ and $10.5\,\%$, respectively. For the required time, the order is reversed: explorations with surface_only are the quickest,

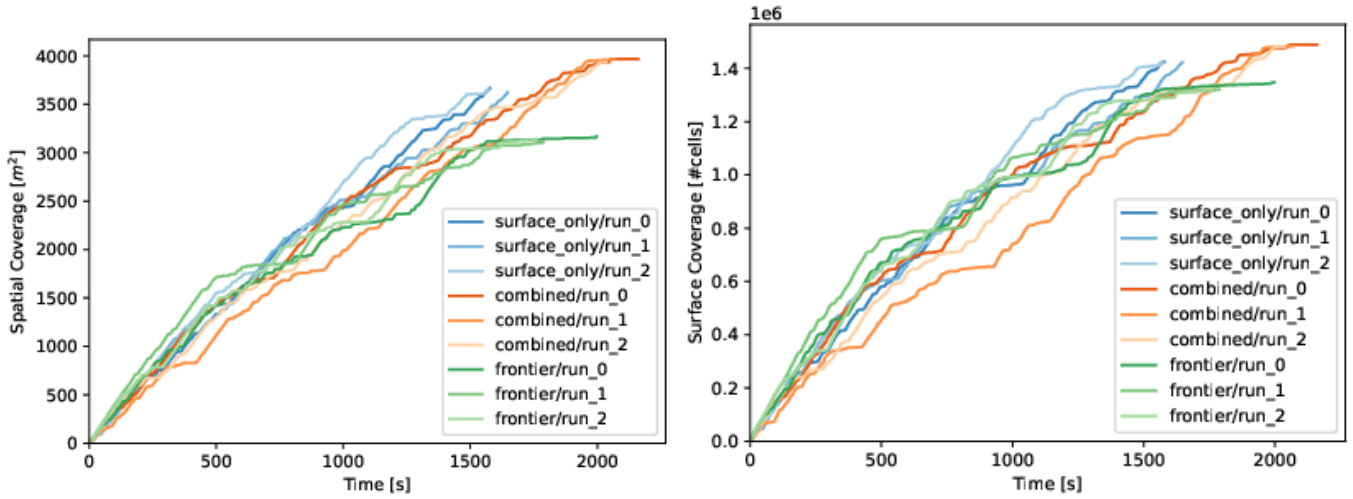

(a) Spatial Coverage

(b) Surface Coverage

Figure 7.12: Coverage over time for the CMU Indoor scenario with three exploration strategies: the proposed method with only surface coverage "surface_only," the proposed method with combined surface and spatial coverage "combined" and a frontier-based exploration approach "frontier." For each strategy, three samples are depicted.

taking an average $1599\,\mathrm{s}$, frontier takes on average $1848\,\mathrm{s}$, additional $15.6\,\%$ and combined $2088\,\mathrm{s}$, additional $30.6\,\%$.

For both coverages, the frontier strategy explores slightly faster in the first approx. $700\,\mathrm{s}$, when the surface strategy takes a slight lead. Notable are plateaus in the coverage for the frontier-based approach starting at $1400\,\mathrm{s}$ - $1500\,\mathrm{s}$, where the method catches up on small remaining frontiers. The combined approach explores slightly slower than the other approaches in the early phase, with small plateaus in the surface coverage, and then it consistently continues until small plateaus arise before termination. The combined approach also accounts for the spatial coverage, so the slightly slower coverage is reasonable, as the other approaches only optimize for direct coverage of the surface. Conversely, this also explains the better coverage of the combined approach. The faster and better coverage of the surface_only strategy in comparison to the frontier approach indicates improved efficiency of the viewpoint-based approach with global planning. Cao et al. [25] observe even larger margins for their approach in comparison to a frontier-based approach, indicating the need for further investigations.

7.3 Transfer

Concepts, insights, and systems developed in this thesis were transferred to contribute to practical applications. In the following, we outline the contribution of requirements to German consortial standards (DIN SPEC) for robotic systems for use in hazardous applications and the application of the robotic demonstrator DRZ Telemax with the developed SLAM system for inspection and mapping in response to a residential complex fire.

(a) DRZ Telemax robot in front of the building ruins.

(b) The operator teleoperates the robot with SLAM localization and registered point cloud as assistance.

Figure 7.13: Residential Complex Fire Deployment in Essen in February 2022. (Images: Nils Heidemann)

7.3.1 German consortial standards - DIN SPEC 91477 Robotic systems for use in hazardous applications

The creation of two German consortial standards (DIN SPEC) for robotic systems for use in hazardous applications was initiated by a consortium originating from the creation of the German Center for Rescue Robotics. The final consortium also included multiple other companies and research organizations. The goal of the first standard *DIN SPEC 91477-1:2023-05 Robotic systems for use in hazardous applications - Part 1: General requirements* [47] was to provide developers and applicants with a categorization of robotic systems and a definition of general requirements. The second standard *DIN SPEC 91477-2:2023-05 Robotic systems for use in hazardous applications - Part 2: Requirements for firefighting robots* [48] provided further specifications specifically for firefighting robots. Parts of the identified requirements served as input for the implementation of the first standard, with specific input to standard chapters on environmental conditions and requirements for autonomy and assistance functions. Further specific requirements for autonomy and assistance functions for firefighting robots were contributed to the second standard.

7.3.2 Disaster Deployment - Residential Complex Fire Essen

As described in Section 4.4.3, on February 22, 2022, a residential apartment building in Essen, Germany, experienced a severe fire. As a result, 39 apartments on four floors were burned, and an entry ban was imposed for parts of the building. To assess the situation and gather information for the investigation of the fire's cause, the German Center for Rescue Robotics (DRZ)[2] deployed ground and air robots to create a 3D model and capture images of the interior.

The DRZ Telemax ground robot was deployed for three inspections the following day. For the first two missions on ground level, a person had visual contact with the robot and communicated with the operator.

[2]The author was part of the ground robot response team and teleoperated the robot.

This was not feasible for the third mission at the upper level due to the entry ban. Therefore, a UAV was deployed to support the operator with visual support from the outside.

During all missions, the robot was teleoperated by the operator. The navigation was challenging as the environment covered very narrow passages, where the robot could collide with the walls and loose debris on the ground, which led to track slippage and ground contact with the lower part of the chassis. Visualizing the robot's position and state with registered point clouds with the 3D robot model helped assess distances to the environment and navigate through narrow environments. First, 3D models could already be computed live during the deployment missions with the proposed SLAM approach. To further improve the quality, the data was post-processed after the deployment, and the final models were transmitted one week after the deployment.

After inspecting the data, the investigators expressed the need not only for geometric data but also for high-quality visual models. This request led to the development of the automatic panorama tour tool outlined in Chapter 6. Experiences, observations, and derived requirements are accounted for in Chapter 4.

Overall, the deployment demonstrated the suitability of the proposed SLAM approach for the targeted applications in disaster response. However, the deployment also indicated that for practical applicability, not only research questions need to be addressed, but also robust, practical integrated systems and streamlined processing workflows are necessary.

8 Conclusion

This thesis investigates novel approaches for localization, mapping, and exploration as assistance abilities for mobile ground robots in disaster environments, aiming to advance disaster response efficiency and safety by contributing, e.g., to autonomous robot navigation, that account for the specific requirements and challenges to support first responders and civil forces.

8.1 Summary of Contributions

Requirements and Challenges for Autonomy and Assistance Functions for Ground Rescue Robots in Reconnaissance Missions Understanding the specific requirements for (semi-)autonomous assistance functions in rescue robots is crucial for research and development toward practical applicability. Previous analyses have primarily focused on general aspects, leaving a gap in the specific understanding of the full spectrum of requirements for (autonomous) assistance functions. In Chapter 4, we address this gap by deriving a novel model for an integrated function capability from established models for technology acceptance and derive a comprehensive, evidence-driven analysis of application requirements and research challenges for (autonomous) assistance abilities.

The analysis of requirements and challenges has been published in [41], which was nominated as *Best Paper Award Finalist*. Insights from training and deployments have been published in [153] and [94]. Parts of the identified requirements served as input for the implementation of the German consortial standards *DIN SPEC 91477-1:2023-05 Robotic systems for use in hazardous applications - Part 1: General requirements* [47] and *DIN SPEC 91477-2:2023-05 Robotic systems for use in hazardous applications - Part 2: Requirements for firefighting robots* [48].

Robust Simultaneous Localization and Mapping in Challenging Environments Sufficiently accurate and robust SLAM in unknown environments without relying on GNSS support are essential for (semi-)autonomous operation. In particular, traversing uneven ground can lead to abrupt robot motions that existing SLAM methods cannot model accurately or efficiently enough. Furthermore, relevant environments are often unstructured and potentially visually degraded by smoke, dust, or fog. To account for these properties, we present novel approaches to improve the accuracy, robustness, and efficiency of position estimation and map generation. Additionally, we propose extensions to transfer the approach for the operation in visually degraded conditions with radar.

The evaluation results in Chapter 7 demonstrate improvements in accuracy and robustness with respect to a state-of-the-art method. We successfully evaluated the approach in various international competitions as part of integrated robot systems (RoboCup Rescue Robot League, World Robot Summit, EnRicH, AIRA), demonstrating more robust and accurate results than competing approaches. The approach enabled the creation of accurate maps and the execution of complex autonomous behaviors, such as autonomous obstacle traversal. Moreover, we demonstrated the applicability in the evaluation in an actual disaster environment as part of a deployment of the Robotic Task Force established by German Rescue Robotic Center DRZ after a large building complex fire.

The developed approach for robust 3D mapping is published as HectorGrapher [39], the large-scale loop closure and pose graph approach in [37], the extension for degraded visual conditions in [157] and insights on the evaluation in the Essen disaster in [153]. The implementation is available as open source[1] and the created DRZ Living Lab data set as open data[2].

Operator-related Exploration and Data-Acquisition in Multi-Goal Missions In response missions, robots might need to fulfill various tasks in a single mission. In such dynamic and versatile environments, first responders often have prior knowledge and potentially better high-level decision-making skills than existing AI methods for the perception and reasoning of autonomous robots. However, an operator's cognitive load is limited, and direct operator control is often error-prone, inefficient, and not always possible.

While related state-of-the-art methods often focus on fully autonomous approaches for single-goal missions, we propose a shared-autonomy approach for multi-goal missions, which allows fully autonomous operation but strongly benefits and incorporates by design the capabilities of the operator for scene understanding and decision-making. We propose a novel method to allow for spatial (e.g., radiation or hazard sensor) and surface coverage goals missions. The method is embedded in a proposed framework that follows the shared autonomy principle and allows flexible changes in the autonomy level from assisted teleoperation to full autonomy. The operator can request complex environment interaction with a novel, integrated, actionable, affordance-based environment representation, which enables complex object interaction and planning sequences. Furthermore, we propose a novel method to efficiently and accurately map dose rates in 2D and 3D based on Gaussian Processes.

The evaluation in Chapter 7 demonstrates the proposed planner's efficiency and improved efficiency and coverage compared to a frontier-based planner in complex simulation environments. Furthermore, the evaluation demonstrates the capability to perform complex missions with autonomous and interactive interactions.

The actionability concept has been published in [8], and the radiation mapping approach in [152] together with a data set with multiple radiation sources [40].

Evaluation and Transfer We performed comprehensive evaluations to investigate the performance and properties of the proposed methods in Chapter 7. We evaluated the proposed SLAM method on established benchmark data, demonstrating improvements in accuracy for large-scale 2D mapping compared to related methods. As typical benchmarks do not address the challenging motion characteristics when traversing obstacles or uneven terrain, we introduce the DRZ Living Lab dataset [38]. Our evaluation of that data set demonstrates improvements in accuracy and robustness for 3D mapping to a related method. Furthermore, we evaluated the proposed SLAM method as part of a fully integrated robot system under conditions of multiple international robotic competitions, contributing to multiple prices and demonstrating better accuracy and robustness than competing methods. We evaluated the exploration method in two complex simulated environments, demonstrating improved efficiency and coverage compared to a frontier-based planner and the capability to perform complex missions with autonomous and interactive interactions.

Finally, we demonstrate the successful transfer of the proposed approach for practical applications such as evaluation as part of an actual disaster deployment and the contribution to implementing the German consortial standards (DIN SPEC) for robots in hazardous environments.

[1] `https://github.com/tu-darmstadt-ros-pkg/hectorgrapher`
[2] `https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3973`

8.2 Outlook

The advancements and contributions presented in this thesis lay a solid foundation for further research and development in the field of disaster robotics and provide versatile aspects for future exploration and improvements toward capabilities for real-world disasters.

A key challenge is robustness, resilience, and reliability under adverse disaster conditions, encompassing all components from hardware to control, perception, reasoning, operator interaction, and novel robotic assistance abilities. The proposed SLAM approach addresses different sensing modalities, lidar, and radar, demonstrating robustness to fast motions and degraded visual conditions. While these sensors are considered separately, combined approaches incorporating further sensors, such as color or thermal cameras, have the potential to yield accurate and robust results across various conditions. Initial works for such a system are proposed by the NeBula framework [112], considering separate sources and fusing them subsequently. However, achieving a flexible and computationally efficient approach with a unified environment representation is an open research question.

Scaling from a single robot to a fleet of heterogeneous cooperating robots can enhance the capabilities significantly. Results from the DARPA SubT Challenge demonstrate large potentials for the cooperative exploration of large-scale environments [31] but also indicate limitations in the function capabilities and robustness. Concepts such as sparse robot swarms indicate the potential to enable versatile distributed capabilities [155].

Another key challenge is the robot-operator interaction. The shared autonomy framework proposed in this thesis represents a step towards effective collaboration between humans and robots in disaster response tasks. Further investigation into intuitive and efficient user interfaces and interaction paradigms combining the operator and assistance functions is essential to facilitate seamless collaboration and enhance overall efficiency.

Learning-based approaches demonstrate impressive results in many relevant tasks, such as environment modeling [107], and semantic understanding [164], control [78], enabling complex behaviors [99] and skills [131]. Recent progress [70] towards the robustness and safety of reinforcement learning-based approaches indicates promising directions to enable easier object interactions with complex robots. These approaches provide a large potential for further powerful robotic assistance functions. However, the combination of typically low training data, generalization, transparency, potential out-of-distribution data during deployment, limited computing power, and real-time requirements on mobile robotics systems pose versatile research and application challenges for successful transfer and application.

Bibliography

[1] Pooya Adami, Patrick B Rodrigues, Peter J Woods, Burcin Becerik-Gerber, Lucio Soibelman, Yasemin Copur-Gencturk, and Gale Lucas. “Impact of VR-based training on human–robot interaction for remote operating construction robots”. In: *Journal of Computing in Civil Engineering* 36.3 (2022), p. 04022006.

[2] Sameer Agarwal, Keir Mierle, et al. *Ceres Solver*. http://ceres-solver.org.

[3] Muhammad Farhan Ahmed, Khayyam Masood, Vincent Fremont, and Isabelle Fantoni. “Active slam: A review on last decade”. In: *Sensors* 23.19 (2023), p. 8097.

[4] Saba Arshad and Gon-Woo Kim. “Role of deep learning in loop closure detection for visual and lidar slam: A survey”. In: *Sensors* 21.4 (2021), p. 1243.

[5] Christopher G Atkeson, PW Babu Benzun, Nandan Banerjee, Dmitry Berenson, Christoper P Bove, Xiongyi Cui, Mathew DeDonato, Ruixiang Du, Siyuan Feng, Perry Franklin, et al. “What happened at the DARPA robotics challenge finals”. In: *The DARPA robotics challenge finals: Humanoid robots to the rescue* (2018), pp. 667–684.

[6] Héctor Azpúrua, Maíra Saboia, Gustavo M Freitas, Lillian Clark, Ali-akbar Agha-mohammadi, Gustavo Pessin, Mario FM Campos, and Douglas G Macharet. “A Survey on the autonomous exploration of confined subterranean spaces: Perspectives from real-word and industrial robotic deployments”. In: *Robotics and Autonomous Systems* 160 (2023), p. 104304.

[7] Haoyu Bai, David Hsu, and Wee Sun Lee. “Integrated perception and planning in the continuous space: A POMDP approach”. In: *The International Journal of Robotics Research* 33.9 (2014), pp. 1288–1302.

[8] Frederik Bark, Kevin Daun, and Oskar von Stryk. “Affordance-based Actionable Semantic Mapping and Planning for Mobile Rescue Robots”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 53–60.

[9] Jens Behley and Cyrill Stachniss. “Efficient Surfel-Based SLAM using 3D Laser Range Data in Urban Environments.” In: *Robotics: Science and Systems*. Vol. 2018. 2018, p. 59.

[10] Paul J Besl, Neil D McKay, et al. “A method for registration of 3-D shapes”. In: *IEEE Transactions on pattern analysis and machine intelligence* 14.2 (1992), pp. 239–256.

[11] Graeme Best, Rohit Garg, John Keller, Geoffrey A Hollinger, and Sebastian Scherer. “Resilient multi-sensor exploration of multifarious environments with a team of aerial robots”. In: *Robotics: Science and Systems (RSS)*. 2022.

[12] Peter Biber and Wolfgang Straßer. “The normal distributions transform: A new approach to laser scan matching”. In: *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*. Vol. 3. IEEE. 2003, pp. 2743–2748.

[13] B Wayne Blanchard. “Guide to emergency management and related terms, definitions, concepts, acronyms, organizations, programs, guidance, executive orders & legislation: A tutorial on emergency management, broadly defined, past and present”. In: *United States. Federal Emergency Management Agency*. United States. Federal Emergency Management Agency. 2008.

[14] Michael Bosse and Robert Zlot. “Continuous 3D scan-matching with a spinning 2D laser”. In: *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE. 2009, pp. 4312–4319.

[15] Frederic Bourgault, Alexei A Makarenko, Stefan B Williams, Ben Grocholsky, and Hugh F Durrant-Whyte. “Information based adaptive robotic exploration”. In: *IEEE/RSJ international conference on intelligent robots and systems*. Vol. 1. IEEE. 2002, pp. 540–545.

[16] Lauren Bramblett, Rahul Peddi, and Nicola Bezzo. “Coordinated multi-agent exploration, rendezvous, & task allocation in unknown environments with limited connectivity”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2022, pp. 12706–12712.

[17] David T Butry, David Webb, Stanley Gilbert, and Jennifer Taylor. *The economics of firefighter injuries in the United States*. US Department of Commerce, National Institute of Standards and Technology ..., 2019.

[18] Jonathan Butzke and Maxim Likhachev. “Planning for multi-robot exploration with multiple objective utility functions”. In: *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2011, pp. 3254–3259.

[19] Erik Bylow, Jürgen Sturm, Christian Kerl, Fredrik Kahl, and Daniel Cremers. “Real-time camera tracking and 3D reconstruction using signed distance functions.” In: *Robotics: Science and Systems*. Vol. 2. 2013.

[20] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age”. In: *IEEE Transactions on robotics* 32.6 (2016), pp. 1309–1332.

[21] Yixi Cai, Wei Xu, and Fu Zhang. “ikd-Tree: An Incremental KD Tree for Robotic Applications”. In: *arXiv preprint arXiv:2102.10808* (2021).

[22] Daniele Calisi, Alessandro Farinelli, Luca Iocchi, and Daniele Nardi. “Multi-objective exploration and search for autonomous rescue robots”. In: *Journal of Field Robotics* 24.8-9 (2007), pp. 763–777.

[23] C Cao, H Zhu, Z Ren, H Choset, and J Zhang. “Representation granularity enables time-efficient autonomous exploration in large, complex worlds”. In: *Science Robotics* 8.80 (2023).

[24] Chao Cao, Hongbiao Zhu, Howie Choset, and Ji Zhang. “TARE: A Hierarchical Framework for Efficiently Exploring Complex 3D Environments.” In: *Robotics: Science and Systems*. Vol. 5. 2021.

[25] Chao Cao, Hongbiao Zhu, Fan Yang, Yukun Xia, Howie Choset, Jean Oh, and Ji Zhang. “Autonomous exploration development environment and the planning algorithms”. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE. 2022, pp. 8921–8928.

[26] Henry Carrillo, Ian Reid, and José A Castellanos. “On the comparison of uncertainty criteria for active SLAM”. In: *2012 IEEE International Conference on Robotics and Automation*. IEEE. 2012, pp. 2080–2087.

[27] Jiawen Chen, Dennis Bautembach, and Shahram Izadi. “Scalable real-time volumetric surface reconstruction”. In: *ACM Transactions on Graphics (ToG)* 32.4 (2013), pp. 1–16.

[28] Kenny Chen, Brett T Lopez, Ali-akbar Agha-mohammadi, and Ankur Mehta. “Direct lidar odometry: Fast localization with dense point clouds”. In: *IEEE Robotics and Automation Letters* 7.2 (2022), pp. 2000–2007.

[29] Kenny Chen, Ryan Nemiroff, and Brett T Lopez. “Direct lidar-inertial odometry: Lightweight lio with continuous-time motion correction”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2023, pp. 3983–3989.

[30] X. Chen, T. Läbe, A. Milioto, T. Röhling, J. Behley, and C. Stachniss. “OverlapNet: A Siamese Network for Computing LiDAR Scan Similarity with Applications to Loop Closing and Localization”. In: *Autonomous Robots* 46 (2021), pp. 61–81. ISSN: 1573-7527. DOI: 10.1007/s10514-021-09999-0.

[31] Timothy H Chung, Viktor Orekhov, and Angela Maio. “Into the Robotic Depths: Analysis and Insights from the DARPA Subterranean Challenge”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 6 (2023), pp. 477–502.

[32] Micah Corah, Cormac O’Meadhra, Kshitij Goel, and Nathan Michael. “Communication-efficient planning and mapping for multi-robot exploration in large environments”. In: *IEEE Robotics and Automation Letters* 4.2 (2019), pp. 1715–1721.

[33] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. 3rd. The MIT Press, 2009. ISBN: 0262033844.

[34] Brian Curless and Marc Levoy. “A volumetric method for building complex models from range images”. In: *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM. 1996, pp. 303–312.

[35] Tung Dang, Marco Tranzatto, Shehryar Khattak, Frank Mascarich, Kostas Alexis, and Marco Hutter. “Graph-based subterranean exploration path planning using aerial and legged robots”. In: *Journal of Field Robotics* 37.8 (2020), pp. 1363–1388.

[36] Kevin Daun. “Robust 3D SLAM for Mobile Search and Rescue Robots in Challenging Environments”. MA thesis. Technische Universität Darmstadt, Fachbereich Informatik, 2017.

[37] Kevin Daun, Stefan Kohlbrecher, Jürgen Sturm, and Oskar von Stryk. “Large scale 2d laser slam using truncated signed distance functions”. In: *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2019, pp. 222–228.

[38] Kevin Daun, Schnaubelt Marius, and Oskar von Stryk. *DRZ Living Lab Tracked Robot SLAM Dataset*. 2021. URL: https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3973.

[39] Kevin Daun, Marius Schnaubelt, Stefan Kohlbrecher, and Oskar von Stryk. “HectorGrapher: Continuous-time lidar SLAM with multi-resolution signed distance function registration for challenging terrain”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 152–159.

[40] Kevin Daun, Jonas Süß, Martin Volz, and Oskar von Stryk. *Hector Enrich 2023 Radiation Mapping Dataset*. 2023. URL: https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3974.

[41] Kevin Daun and Oskar Von Stryk. “Requirements and challenges for autonomy and assistance functions for ground rescue robots in reconnaissance missions”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 83–90.

[42] Fred D Davis. “Perceived usefulness, perceived ease of use, and user acceptance of information technology”. In: *MIS quarterly* (1989), pp. 319–340.

[43] Frank Dellaert. “Factor Graphs: Exploiting Structure in Robotics”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 4 ().

[44] Frank Dellaert, Michael Kaess, et al. “Factor graphs for robot perception”. In: *Foundations and Trends® in Robotics* 6.1-2 (2017), pp. 1–139.

[45] Jeffrey Delmerico, Stefano Mintchev, Alessandro Giusti, Boris Gromov, Kamilo Melo, Tomislav Horvat, Cesar Cadena, Marco Hutter, Auke Ijspeert, Dario Floreano, et al. “The current state and future outlook of rescue robotics”. In: *Journal of Field Robotics* 36.7 (2019), pp. 1171–1191.

[46] Junyuan Deng, Qi Wu, Xieyuanli Chen, Songpengcheng Xia, Zhen Sun, Guoqing Liu, Wenxian Yu, and Ling Pei. “Nerf-loam: Neural implicit representation for large-scale incremental lidar odometry and mapping”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, pp. 8218–8227.

[47] *Robotic systems for use in hazardous applications - Part 1: General requirements*. Standard. 10772 Berlin, May 2023.

[48] *Robotic systems for use in hazardous applications - Part 2: Requirements for firefighting robots*. Standard. 10772 Berlin, May 2023.

[49] Daniela Doroftei, Anibal Matos, and Geert de Cubber. “Designing search and rescue robots towards realistic user requirements”. In: *Applied mechanics and materials* 658 (2014), pp. 612–617.

[50] David Droeschel and Sven Behnke. “Efficient continuous-time SLAM for 3D lidar-based online mapping”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 5000–5007.

[51] David Droeschel, Jörg Stückler, and Sven Behnke. “Local multi-resolution representation for 6D motion estimation and mapping with a continuously rotating 3D laser scanner”. In: *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE. 2014, pp. 5221–5226.

[52] Raimund Edlinger, Gabriel Himmelbauer, Gerald Zauner, and Andreas Nüchter. “Visual odometry and mapping under poor visibility conditions using a stereo infrared thermal imaging system”. In: *Electronic Imaging* 35 (2023), pp. 1–7.

[53] Jakob Engel, Jörg Stückler, and Daniel Cremers. “Large-scale direct SLAM with stereo cameras”. In: *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE. 2015, pp. 1935–1942.

[54] Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J Mitra, and Michael Wimmer. “Points2surf learning implicit surfaces from point clouds”. In: *European Conference on Computer Vision*. Springer. 2020, pp. 108–124.

[55] Stefan Fabian and Oskar von Stryk. “Open-source tools for efficient ROS and ROS2-based 2D human-robot interface development”. In: *2021 European Conference on Mobile Robots (ECMR)*. IEEE. 2021, pp. 1–6.

[56] Péter Fankhauser, Michael Bloesch, and Marco Hutter. “Probabilistic Terrain Mapping for Mobile Robots with Uncertain Localization”. In: *IEEE Robotics and Automation Letters (RA-L)* 3.4 (2018), pp. 3019–3026. DOI: 10.1109/LRA.2018.2849506.

[57] Hans Jacob S Feder, John J Leonard, and Christopher M Smith. “Adaptive mobile robot navigation and mapping”. In: *The International Journal of Robotics Research* 18.7 (1999), pp. 650–668.

[58] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. “On-Manifold Preintegration for Real-Time Visual–Inertial Odometry”. In: *IEEE Transactions on Robotics* 33.1 (2016), pp. 1–21.

[59] Joscha-David Fossel, Karl Tuyls, and Jürgen Sturm. “2D-SDF-SLAM: A signed distance function based SLAM frontend for laser scanners”. In: *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE. 2015, pp. 1949–1955.

[60] Vincent Furnon and Laurent Perron. *OR-Tools Routing Library*. Version v9.8. Google, Nov. 15, 2023. URL: https://developers.google.com/optimization/routing/.

[61] Cipriano Galindo, Juan-Antonio Fernández-Madrigal, Javier González, and Alessandro Saffiotti. “Robot task planning using semantic maps”. In: *Robotics and autonomous systems* 56.11 (2008), pp. 955–966.

[62] Gregor HW Gebhardt, Kevin Daun, Marius Schnaubelt, and Gerhard Neumann. “Learning robust policies for object manipulation with robot swarms”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 7688–7695.

[63] Arjan Gijsberts, Tatiana Tommasi, Giorgio Metta, and Barbara Caputo. “Object recognition using visuo-affordance maps”. In: *IEEE/RSJ IROS*. 2010, pp. 1572–1578. DOI: 10.1109/IROS.2010.5649238.

[64] Haeyeon Gim, Seungmin Baek, JeongKi Park, Hoyong Lee, Chiwon Sung, KyungTae Kim, and Soohee Han. “Suitability of Various Lidar and Radar Sensors for Application in Robotics: A Measurable Capability Comparison”. In: *IEEE Robotics & Automation Magazine* (2022).

[65] Clara Gomez, Alejandra C Hernandez, and Ramon Barber. “Topological frontier-based exploration and map-building using semantic information”. In: *Sensors* 19.20 (2019), p. 4595.

[66] Héctor H González-Banos and Jean-Claude Latombe. “Navigation strategies for exploring indoor environments”. In: *The International Journal of Robotics Research* 21.10-11 (2002), pp. 829–848.

[67] Margarita Grinvald, Federico Tombari, Roland Siegwart, and Juan Nieto. “TSDF++: A Multi-Object Formulation for Dynamic Object Tracking and Reconstruction”. In: *IEEE ICRA*. 2021, pp. 14192–14198. DOI: 10.1109/ICRA48506.2021.9560923.

[68] Giorgio Grisetti, Rainer Kummerle, Cyrill Stachniss, and Wolfram Burgard. “A tutorial on graph-based SLAM”. In: *IEEE Intelligent Transportation Systems Magazine* 2.4 (2010), pp. 31–43.

[69] Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. “Improved techniques for grid mapping with rao-blackwellized particle filters”. In: *IEEE transactions on Robotics* 23.1 (2007), pp. 34–46.

[70] Shangding Gu, Alap Kshirsagar, Yali Du, Guang Chen, Jan Peters, and Alois Knoll. “A human-centered safe robot reinforcement learning framework with interactive behaviors”. In: *Frontiers in Neurorobotics* 17 (2023).

[71] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. “Real-Time Loop Closure in 2D LIDAR SLAM”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. 2016, pp. 1271–1278.

[72] Marian Himstedt, Jan Frost, Sven Hellbach, Hans-Joachim Böhme, and Erik Maehle. “Large scale place recognition in 2D LIDAR scans using geometrical landmark relations”. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2014, pp. 5030–5035.

[73] Matthias Hollick, Anne Hofmeister, Jens Ivo Engels, Bernd Freisleben, Gerrit Hornung, Anja Klein, Michèle Knodt, Imke Lorenz, Patrick Lieser, Max Mühlhäuser, et al. “Emergencity: A paradigm shift towards resilient digital cities”. In: *World Congress on Resilience, Reliability and Asset Management (WCRRAM)*. 2019, pp. 383–406.

[74] Ziyang Hong, Yvan Petillot, and Sen Wang. “Radarslam: Radar based large-scale slam in all weathers”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2020, pp. 5164–5170.

[75] Armin Hornung, Kai M Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. “OctoMap: An efficient probabilistic 3D mapping framework based on octrees”. In: *Autonomous robots* 34 (2013), pp. 189–206.

[76] Andrew Howard and Nicholas Roy. *The Robotics Data Set Repository (Radish)*. 2003. URL: http://radish.sourceforge.net/.

[77] N. Hughes, Y. Chang, and L. Carlone. “Hydra: A Real-time Spatial Perception System for 3D Scene Graph Construction and Optimization”. In: (2022).

[78] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. “Learning agile and dynamic motor skills for legged robots”. In: *Science Robotics* 4.26 (2019), eaau5872.

[79] Adam Jacoff, Raymond Sheh, Ann-Marie Virts, Tetsuya Kimura, Johannes Pellenz, Sören Schwertfeger, and Jackrit Suthakorn. “Using competitions to advance the development of standard test methods for response robots”. In: *Proceedings of the Workshop on Performance Metrics for Intelligent Systems*. ACM. 2012, pp. 182–189.

[80] Matthew Johnson, Jeffrey M Bradshaw, Paul J Feltovich, Catholijn M Jonker, M Birna Van Riemsdijk, and Maarten Sierhuis. “Coactive design: Designing support for interdependence in joint activity”. In: *Journal of Human-Robot Interaction* 3.1 (2014), pp. 43–69.

[81] Peter Kaiser, Markus Grotz, Eren E. Aksoy, Martin Do, Nikolaus Vahrenkamp, and Tamim Asfour. “Validation of whole-body loco-manipulation affordances for pushability and liftability”. In: *IEEE-RAS Humanoids*. 2015, pp. 920–927. DOI: 10.1109/HUMANOIDS.2015.7363471.

[82] Matan Keidar and Gal A Kaminka. “Robot exploration with fast frontier detection: Theory and experiments”. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. 2012, pp. 113–120.

[83] Matthew Klingensmith, Ivan Dryanovski, Siddhartha Srinivasa, and Jizhong Xiao. “Chisel: Real Time Large Scale 3D Reconstruction Onboard a Mobile Device using Spatially Hashed Signed Distance Fields.” In: *Robotics: Science and Systems*. Vol. 4. 2015.

[84] Philipp Koch, Stefan May, Michael Schmidpeter, Markus Kühn, Christian Pfitzner, Christian Merkl, Rainer Koch, Martin Fees, Jon Martin, Daniel Ammon, et al. “Multi-robot localization and mapping based on signed distance functions”. In: *Autonomous Robot Systems and Competitions (ICARSC)*. IEEE. 2015, pp. 77–82.

[85] Philipp Koch, Stefan May, Michael Schmidpeter, Markus Kühn, Christian Pfitzner, Christian Merkl, Rainer Koch, Martin Fees, Jon Martin, Daniel Ammon, et al. “Multi-robot localization and mapping based on signed distance functions”. In: *Journal of Intelligent & Robotic Systems* 83.3-4 (2016), pp. 409–428.

[86] Stefan Kohlbrecher. “A Scalable, Platform-Independent SLAM System for Urban Search and Rescue”. MA thesis. Technische Universität Darmstadt, Department of Computer Science, 2009.

[87] Stefan Kohlbrecher, Florian Kunz, Dorothea Koert, Christian Rose, Paul Manns, Kevin Daun, Johannes Schubert, Alexander Stumpf, and Oskar von Stryk. “Towards highly reliable autonomy for urban search and rescue robots”. In: *RoboCup 2014: Robot World Cup XVIII 18*. Springer. 2015, pp. 118–129.

[88] Stefan Kohlbrecher, Johannes Meyer, Thorsten Graber, Karen Petersen, Uwe Klingauf, and Oskar Von Stryk. “Hector open source modules for autonomous mapping and navigation with rescue robots”. In: *RoboCup 2013: Robot World Cup XVII 17*. Springer. 2014, pp. 624–631.

[89] Stefan Kohlbrecher, Oskar Von Stryk, Johannes Meyer, and Uwe Klingauf. “A flexible and scalable slam system with full 3d motion estimation”. In: *Safety, Security, and Rescue Robotics (SSRR), 2011 IEEE International Symposium on*. IEEE. 2011, pp. 155–160.

[90] Kurt Konolige, Giorgio Grisetti, Rainer Kümmerle, Wolfram Burgard, Benson Limketkai, and Regis Vincent. “Efficient sparse pose adjustment for 2D mapping”. In: *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE. 2010, pp. 22–29.

[91] David Kortenkamp, Reid Simmons, and Davide Brugali. “Robotic Systems Architectures and Programming”. In: *Springer Handbook of Robotics*. Ed. by Bruno Siciliano and Oussama Khatib. Cham: Springer International Publishing, 2016, pp. 283–306. ISBN: 978-3-319-32552-1. DOI: 10.1007/978-3-319-32552-1_12. URL: https://doi.org/10.1007/978-3-319-32552-1_12.

[92] Eric Krotkov, Douglas Hackett, Larry Jackel, Michael Perschbacher, James Pippine, Jesse Strauss, Gill Pratt, and Christopher Orlowski. “The DARPA robotics challenge finals: Results and perspectives”. In: *The DARPA robotics challenge finals: Humanoid robots to the rescue* (2018), pp. 1–26.

[93] Ivana Kruijff-Korbayová, Luigi Freda, Mario Gianni, Valsamis Ntouskos, Václav Hlaváč, Vladimír Kubelka, Erik Zimmermann, Hartmut Surmann, Kresimir Dulic, Wolfgang Rottner, et al. “Deployment of ground and aerial robots in earthquake-struck amatrice in italy (brief report)”. In: *2016 IEEE international symposium on safety, security, and rescue robotics (SSRR)*. IEEE. 2016, pp. 278–279.

[94] Ivana Kruijff-Korbayová, Robert Grafe, Nils Heidemann, Alexander Berrang, Cai Hussung, Christian Willms, Peter Fettke, Marius Beul, Jan Quenzel, Daniel Schleich, Sven Behnke, Janis Tiemann, Johannes Güldenring, Manuel Patchou, Christian Arendt, Christian Wietfeld, Kevin Daun, Marius Schnaubelt, Oskar von Stryk, Alexander Lel, Alexander Miller, Christof Röhrig, Thomas Straßmann, Thomas Barz, Stefan Soltau, Felix Kremer, Stefan Rilling, Rohan Haseloff, Stefan Grobelny, Artur Leinweber, Gerhard Senkowski, Marc Thurow, Dominik Slomma, and Hartmut Surmann. “German rescue robotics center (drz): A holistic approach for robotic systems assisting in emergency response”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 138–145.

[95] Mihir Kulkarni, Mihir Dharmadhikari, Marco Tranzatto, Samuel Zimmermann, Victor Reijgwart, Paolo De Petris, Huan Nguyen, Nikhil Khedekar, Christos Papachristos, Lionel Ott, et al. “Autonomous teamed exploration of subterranean environments using legged and aerial robots”. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE. 2022, pp. 3306–3313.

[96] Rainer Kümmerle, Bastian Steder, Christian Dornhege, Michael Ruhnke, Giorgio Grisetti, Cyrill Stachniss, and Alexander Kleiner. “On measuring the accuracy of SLAM algorithms”. In: *Autonomous Robots* 27.4 (2009), p. 387.

[97] Mathieu Labbé and François Michaud. “Appearance-Based Loop Closure Detection for Online Large-Scale and Long-Term Operation”. In: *IEEE Transactions on Robotics* 29.3 (2013), pp. 734–745. DOI: 10.1109/TRO.2013.2242375.

[98] Jerome Le Ny and George J Pappas. “On trajectory optimization for active sensing in Gaussian process models”. In: *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE. 2009, pp. 6286–6292.

[99] Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection”. In: *The International journal of robotics research* 37.4-5 (2018), pp. 421–436.

[100] Jiaxin Li, Huangying Zhan, Ben M Chen, Ian Reid, and Gim Hee Lee. “Deep learning for 2D scan matching and loop closure”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 763–768.

[101] William E Lorensen and Harvey E Cline. “Marching cubes: A high resolution 3D surface construction algorithm”. In: *Seminal graphics: pioneering efforts that shaped the field*. 1998, pp. 347–353.

[102] Steven Lovegrove, Alonso Patron-Perez, and Gabe Sibley. “Spline Fusion: A continuous-time representation for visual-inertial fusion with application to rolling shutter cameras.” In: *BMVC*. Vol. 2. 5. 2013, p. 8.

[103] Nils Mandischer, Marius Gürtler, Sebastian Döbler, Mathias Hüsing, and Burkhard Corves. “Finding Moving Operators in Firefighting Operations Based on Multi-Goal Next-Best-View Exploration”. In: *2022 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2022, pp. 47–52.

[104] Fernando Martín, Rudolph Triebel, Luis Moreno, and Roland Siegwart. “Two different tools for three-dimensional mapping: DE-based scan matching and feature-based loop detection”. In: *Robotica* 32.1 (2014), pp. 19–41.

[105] Stefan May, Philipp Koch, Rainer Koch, Christian Merkl, Christian Pfitzner, and Andreas Nüchter. “A Generalized 2D and 3D Multi-Sensor Data Integration Approach based on Signed Distance Functions for Multi-Modal Robotic Mapping.” In: *VMV*. 2014, pp. 95–102.

[106] J. Meyer, P. Schnitzspan, S. Kohlbrecher, K. Petersen, O. Schwahn, M. Andriluka, U. Klingauf, S. Roth, B. Schiele, and O. von Stryk. “A Semantic World Model for Urban Search and Rescue Based on Heterogeneous Sensors”. In: *RoboCup 2010: Robot Soccer World Cup XIV*. Ed. by Javier Ruiz-del-Solar, Eric Chown, and Paul G. Ploeger. Lecture Notes in Computer Science, Lecture Notes in Artificial Intelligence. 2011, pp. 180–193.

[107] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. “Nerf: Representing scenes as neural radiance fields for view synthesis”. In: *Communications of the ACM* 65.1 (2021), pp. 99–106.

[108] Clair E Miller, Albert W Tucker, and Richard A Zemlin. “Integer programming formulation of traveling salesman problems”. In: *Journal of the ACM (JACM)* 7.4 (1960), pp. 326–329.

[109] Ankita Mohapatra and Timothy Trinh. “Early wildfire detection technologies in practice—a review”. In: *Sustainability* 14.19 (2022), p. 12270.

[110] Abdullah Mohiuddin, Taha Tarek, Yahya Zweiri, and Dongming Gan. “A survey of single and multi-UAV aerial manipulation”. In: *Unmanned Systems* 8.02 (2020), pp. 119–147.

[111] Saad Mokssit, Daniel Bonilla Licea, Bassma Guermah, and Mounir Ghogho. “Deep Learning Techniques for Visual SLAM: A Survey”. In: *IEEE Access* 11 (2023), pp. 20026–20050.

[112] Benjamin Morrell, Rohan Thakker, Àngel Santamaria Navarro, Amanda Bouman, Xiaoming Lei, Jeffrey Edlund, Torkom Pailevanian, Tiago Stegun Vaquero, Yun Lin Chang, Thomas Touma, et al. “NeBula: TEAM CoSTAR’s robotic autonomy solution that won phase II of DARPA subterranean challenge”. In: *Field robotics* 2 (2022), pp. 1432–1506.

[113] Tarek Mouats, Nabil Aouf, Lounis Chermak, and Mark A Richardson. “Thermal stereo odometry for UAVs”. In: *IEEE Sensors Journal* 15.11 (2015), pp. 6335–6347.

[114] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. “ORB-SLAM: a versatile and accurate monocular SLAM system”. In: *IEEE transactions on robotics* 31.5 (2015), pp. 1147–1163.

[115] Robin R Murphy. “A decade of rescue robots”. In: *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE. 2012, pp. 5448–5449.

[116] Robin R Murphy. *Disaster robotics*. MIT press, 2014.

[117] Robin R Murphy. “Trial by fire”. In: *IEEE Robotics & Automation Magazine* 11.3 (2004), pp. 50–61.

[118] Robin R Murphy and Satoshi Tadokoro. “User interfaces for human-robot interaction in field robotics”. In: *Disaster Robotics: Results from the ImPACT Tough Robotics Challenge* (2019), pp. 507–528.

[119] Robin R. Murphy, Satoshi Tadokoro, and Alexander Kleiner. “Disaster Robotics”. In: *Springer Handbook of Robotics*. Ed. by Bruno Siciliano and Oussama Khatib. Cham: Springer International Publishing, 2016, pp. 1577–1604. ISBN: 978-3-319-32552-1. DOI: 10.1007/978-3-319-32552-1_60. URL: https://doi.org/10.1007/978-3-319-32552-1_60.

[120] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. “KinectFusion: Real-time dense surface mapping and tracking”. In: *Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on*. IEEE. 2011, pp. 127–136.

[121] Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Marc Stamminger. “Real-time 3D reconstruction at scale using voxel hashing”. In: *ACM Transactions on Graphics (ToG)* 32.6 (2013), pp. 1–11.

[122] Andreas Nüchter, M Bleier, Johannes Schauer, and Peter Janotta. “Improving Google’s Cartographer 3D mapping by continuous-time slam”. In: *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* 42 (2017), p. 543.

[123] Andreas Nüchter, Dorit Borrmann, Philipp Koch, Markus Kühn, and Stefan May. “A man-portable, IMU-free mobile mapping system”. In: *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 2 (2015), pp. 17–23.

[124] Martin Oehler and Oskar von Stryk. “A flexible framework for virtual omnidirectional vision to improve operator situation awareness”. In: *2021 European Conference on Mobile Robots (ECMR)*. IEEE. 2021, pp. 1–6.

[125] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. “Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 1366–1373.

[126] Matteo Palieri, Benjamin Morrell, Abhishek Thakur, Kamak Ebadi, Jeremy Nash, Arghya Chatterjee, Christoforos Kanellakis, Luca Carlone, Cataldo Guaragnella, and Ali-akbar Agha-mohammadi. “LOCUS: A Multi-Sensor Lidar-Centric Solution for High-Precision Odometry and 3D Mapping in Real-Time”. In: *IEEE Robotics and Automation Letters* 6.2 (2021), pp. 421–428. DOI: 10.1109/LRA.2020.3044864.

[127] Stefano Panzieri, Federica Pascucci, and Giovanni Ulivi. “An outdoor navigation system using GPS and inertial platform”. In: *IEEE/ASME transactions on Mechatronics* 7.2 (2002), pp. 134–142.

[128] Christos H Papadimitriou and Kenneth Steiglitz. *Combinatorial optimization: algorithms and complexity*. Courier Corporation, 1998.

[129] Chanoh Park, Peyman Moghadam, Soohwan Kim, Alberto Elfes, Clinton Fookes, and Sridha Sridharan. “Elastic lidar fusion: Dense map-centric continuous-time slam”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1206–1213.

[130] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. “Deepsdf: Learning continuous signed distance functions for shape representation”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 165–174.

[131] Jan Peters and Stefan Schaal. “Reinforcement learning of motor skills with policy gradients”. In: *Neural networks* 21.4 (2008), pp. 682–697.

[132] Julio A Placed, Jared Strader, Henry Carrillo, Nikolay Atanasov, Vadim Indelman, Luca Carlone, and José A Castellanos. “A survey on active simultaneous localization and mapping: State of the art and new frontiers”. In: *IEEE Transactions on Robotics* (2023).

[133] William Qi, Ravi Teja Mullapudi, Saurabh Gupta, and Deva Ramanan. “Learning to Move with Affordance Maps”. In: (2020).

[134] Jan Quenzel and Sven Behnke. “Real-time multi-adaptive-resolution-surfel 6D LiDAR odometry using continuous-time trajectory optimization”. In: (2021), pp. 5499–5506.

[135] Jan Quenzel, Malte Splietker, Dmytro Pavlichenko, Daniel Schleich, Christian Lenz, Max Schwarz, Michael Schreiber, Marius Beul, and Sven Behnke. “Autonomous fire fighting with a UAV-UGV team at MBZIRC 2020”. In: *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE. 2021, pp. 934–941.

[136] Micha Rappaport. “Energy-aware mobile robot exploration with adaptive decision thresholds”. In: *Proceedings of ISR 2016: 47st International Symposium on Robotics*. VDE. 2016, pp. 1–8.

[137] Daniel J Rea and Stela H Seo. “Still not solved: A call for renewed focus on user-centered teleoperation interfaces”. In: *Frontiers in Robotics and AI* 9 (2022), p. 704225.

[138] Adrian Rebmann, Jana-Rebecca Rehse, Mira Pinter, Marius Schnaubelt, Kevin Daun, and Peter Fettke. “Iot-based activity recognition for process assistance in human-robot disaster response”. In: *Business Process Management Forum: BPM Forum 2020, Seville, Spain, September 13–18, 2020, Proceedings 18*. Springer. 2020, pp. 71–87.

[139] A. Romay, S. Kohlbrecher, D. C. Conner, A. Stumpf, and O. von Stryk. “Template-based manipulation in unstructured environments for supervised semi-autonomous humanoid robots”. In: *IEEE-RAS Humanoids*. 2014, pp. 979–986.

[140] Antoni Rosinol, Marcus Abate, Yun Chang, and Luca Carlone. “Kimera: an Open-Source Library for Real-Time Metric-Semantic Localization and Mapping”. In: *IEEE ICRA*. 2020.

[141] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone. “3D dynamic scene graphs: Actionable spatial perception with places, objects, and humans”. In: *Robotics: Science and Systems (RSS)* (2020).

[142] Antoni Rosinol, John J Leonard, and Luca Carlone. “Nerf-slam: Real-time dense monocular slam with neural radiance fields”. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2023, pp. 3437–3444.

[143] Marius Schnaubelt, Tobias Ullrich, Moritz Torchalla, Jonas Diegelmann, Matthias Hoffmann, and Oskar von Stryk. “Entwicklung eines autonomiefokussierten hochmobilen Bodenrobotersystems für den Katastrophenschutz”. In: *Digital-Fachtagung VDI-MECHATRONIK 2021*. Ed. by Torsten Bertram, Burkhard Corves, Klaus Janschek, and Stephan Rinderknecht. Universitäts- und Landesbibliothek Darmstadt, Mar. 2021, pp. 20–25.

[144] Frank E Schneider and Dennis Wildermuth. “Using robots for firefighters and first responders: Scenario specification and exemplary system description”. In: *2017 18th International Carpathian Control Conference (ICCC)*. IEEE. 2017, pp. 216–221.

[145] Christian A Schroth, Christian Eckrich, Ibrahim Kakouche, Stefan Fabian, Oskar von Stryk, Abdelhak M Zoubir, and Michael Muma. “Emergency response person localization and vital sign estimation using a semi-autonomous robot mounted SFCW radar”. In: *IEEE Transactions on Biomedical Engineering* (2024).

[146] Tixiao Shan and Brendan Englot. “Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain”. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018, pp. 4758–4765.

[147] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Rus Daniela. “LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2020, pp. 5135–5142.

[148] Miroslava Slavcheva, Wadim Kehl, Nassir Navab, and Slobodan Ilic. “Sdf-2-sdf registration for real-time 3d reconstruction from RGB-D data”. In: *International Journal of Computer Vision* (2018), pp. 1–22.

[149] Malte Splietker and Sven Behnke. “Directional TSDF: Modeling Surface Orientation for Coherent Meshes”. In: Nov. 2019. DOI: 10.1109/IROS40897.2019.8968264.

[150] Cyrill Stachniss, Dirk Hahnel, and Wolfram Burgard. “Exploration with active loop-closing for FastSLAM”. In: *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*. Vol. 2. IEEE. 2004, pp. 1505–1510.

[151] Cyrill Stachniss, John J. Leonard, and Sebastian Thrun. “Simultaneous Localization and Mapping”. In: *Springer Handbook of Robotics*. Ed. by Bruno Siciliano and Oussama Khatib. Cham: Springer International Publishing, 2016, pp. 1153–1176. ISBN: 978-3-319-32552-1. DOI: 10.1007/978-3-319-32552-1_46. URL: https://doi.org/10.1007/978-3-319-32552-1_46.

[152] Jonas SüB, Martin Volz, Kevin Daun, and Oskar von Stryk. “Online 2D-3D Radiation Mapping and Source Localization using Gaussian Processes with Mobile Ground Robots”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 39–46.

[153] Hartmut Surmann, Kevin Daun, Marius Schnaubelt, Oskar von Stryk, Manuel Patchou, Stefan Böcker, Christian Wietfeld, Jan Quenzel, Daniel Schleich, Sven Behnke, Robert Grafe, Nils Heidemann, Dominik Slomma, and Ivana Kruijff-Korbayová. “Lessons from robot-assisted disaster response deployments by the German Rescue Robotics Center task force”. In: *Journal of Field Robotics* (2023).

[154] Lei Tai and Ming Liu. “A robot exploration strategy based on q-learning network”. In: *2016 IEEE International Conference on Real-Time Computing and Robotics (RCAR)*. IEEE. 2016, pp. 57–62.

[155] Danesh Tarapore, Roderich Groß, and Klaus-Peter Zauner. “Sparse robot swarms: Moving swarms to real-world applications”. In: *Frontiers in Robotics and AI* 7 (2020), p. 83.

[156] Sebastian Thrun. “Learning metric-topological maps for indoor mobile robot navigation”. In: *Artificial Intelligence* 99.1 (1998), pp. 21–71.

[157] Moritz Torchalla, Marius Schnaubelt, Kevin Daun, and Oskar von Stryk. “Robust Multisensor Fusion for Reliable Mapping and Navigation in Degraded Visual Conditions”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 110–117.

[158] Marco Tranzatto, Takahiro Miki, Mihir Dharmadhikari, Lukas Bernreiter, Mihir Kulkarni, Frank Mascarich, Olov Andersson, Shehryar Khattak, Marco Hutter, Roland Siegwart, et al. “CERBERUS in the DARPA Subterranean Challenge”. In: *Science Robotics* 7.66 (2022), eabp9742.

[159] Hassan Umari and Shayok Mukhopadhyay. “Autonomous robotic exploration based on multiple rapidly-exploring randomized trees”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 1396–1402.

[160] Roberto G Valenti, Ivan Dryanovski, and Jizhong Xiao. “Keeping a good attitude: A quaternion-based orientation filter for IMUs and MARGs”. In: *Sensors* 15.8 (2015), pp. 19302–19330.

[161] Viswanath Venkatesh, Michael G Morris, Gordon B Davis, and Fred D Davis. “User acceptance of information technology: Toward a unified view”. In: *MIS quarterly* (2003), pp. 425–478.

[162] Emanuele Vespa, Nils Funk, Paul HJ Kelly, and Stefan Leutenegger. “Adaptive-resolution octree-based volumetric SLAM”. In: *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, pp. 654–662.

[163] Ignacio Vizzo, Tiziano Guadagnino, Benedikt Mersch, Louis Wiesmann, Jens Behley, and Cyrill Stachniss. “Kiss-icp: In defense of point-to-point icp–simple, accurate, and robust registration if done the right way”. In: *IEEE Robotics and Automation Letters* 8.2 (2023), pp. 1029–1036.

[164] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. “YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, pp. 7464–7475.

[165] Han Wang, Chen Wang, and Lihua Xie. “Intensity scan context: Coding intensity and geometry relations for loop closure detection”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 2095–2101.

[166] Jingwen Wang, Martin Rünz, and Lourdes Agapito. “DSP-SLAM: Object oriented SLAM with deep shape priors”. In: *2021 International Conference on 3D Vision (3DV)*. IEEE. 2021, pp. 1362–1371.

[167] Thomas Whelan, Michael Kaess, Maurice Fallon, Hordur Johannsson, John Leonard, and John McDonald. “Kintinuous: Spatially extended kinectfusion”. In: (2012).

[168] Thomas Whelan, Stefan Leutenegger, Renato Salas-Moreno, Ben Glocker, and Andrew Davison. “ElasticFusion: Dense SLAM without a pose graph”. In: Robotics: Science and Systems. 2015.

[169] Alan Winfield, Eleanor Watson, Takashi Egawa, Emily Barwell, Iain Barclay, Serena Booth, Louise A Dennis, Helen Hastie, Ali Hossaini, Naomi Jacobs, et al. “IEEE Standard for Transparency of Autonomous Systems”. In: (2022).

[170] Kai M Wurm, Cyrill Stachniss, and Wolfram Burgard. “Coordinated multi-robot exploration using a segmentation of the environment”. In: *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2008, pp. 1160–1165.

[171] Wei Xu, Yixi Cai, Dongjiao He, Jiarong Lin, and Fu Zhang. “Fast-lio2: Fast direct lidar-inertial odometry”. In: *IEEE Transactions on Robotics* 38.4 (2022), pp. 2053–2073.

[172] Brian Yamauchi. “A frontier-based approach for autonomous exploration”. In: *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA’97.’Towards New Computational Principles for Robotics and Automation’*. IEEE. 1997, pp. 146–151.

[173] Chen Yang, Qi Chen, Yaoyao Yang, Jingyu Zhang, Minshun Wu, and Kuizhi Mei. “SDF-SLAM: A deep learning based highly accurate SLAM using monocular camera aiming at indoor map reconstruction with semantic and depth fusion”. In: *IEEE Access* 10 (2022), pp. 10259–10272.

[174] Fan Yang, Chao Cao, Hongbiao Zhu, Jean Oh, and Ji Zhang. “FAR planner: Fast, attemptable route planner using dynamic visibility update”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2022, pp. 9–16.

[175] Heng Yang, Pasquale Antonante, Vasileios Tzoumas, and Luca Carlone. “Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 1127–1134.

[176] Huan Yin, Li Tang, Xiaqing Ding, Yue Wang, and Rong Xiong. “Locnet: Global localization in 3d point clouds for mobile vehicles”. In: *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE. 2018, pp. 728–733.

[177] Tomoaki Yoshida, Keiji Nagatani, Satoshi Tadokoro, Takeshi Nishimura, and Eiji Koyanagi. “Improvements to the rescue robot quince toward future indoor surveillance missions in the Fukushima Daiichi nuclear power plant”. In: *Field and Service Robotics: Results of the 8th International Conference*. Springer. 2014, pp. 19–32.

[178] Yijun Yuan and Andreas Nüchter. “Indirect point cloud registration: Aligning distance fields using a pseudo third point set”. In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 7075–7082.

[179] Yijun Yuan and Andreas Nüchter. “Uni-fusion: Universal continuous mapping”. In: *IEEE Transactions on Robotics* (2024).

[180] Philipp Zech, Simon Haller, Safoura R.L. Lakani, Barry Ridge, Emre Ugur, and Justus Piater. “Computational models of affordance in robotics: a taxonomy and systematic classification”. In: *Adaptive Behavior* 25.5 (2017), pp. 235–271.

[181] Ji Zhang and Sanjiv Singh. “LOAM: Lidar Odometry and Mapping in Real-time.” In: *Robotics: Science and Systems*. Vol. 2. 2014.

[182] Shibo Zhao, Damanpreet Singh, Haoxiang Sun, Rushan Jiang, YuanJun Gao, Tianhao Wu, Jay Karhade, Chuck Whittaker, Ian Higgins, Jiahe Xu, et al. “SubT-MRS: A Subterranean, Multi-Robot, Multi-Spectral and Multi-Degraded Dataset for Robust SLAM”. In: *arXiv preprint arXiv:2307.07607* (2023).

[183] Xingguang Zhong, Yue Pan, Jens Behley, and Cyrill Stachniss. “Shine-mapping: Large-scale 3d mapping using sparse hierarchical implicit neural representations”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2023, pp. 8371–8377.

[184] Youjie Zhou, Yiming Wang, Fabio Poiesi, Qi Qin, and Yi Wan. “Loop closure detection using local 3D deep descriptors”. In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 6335–6342.

Own Publications

Journal Publications

Hartmut Surmann, Kevin Daun, Marius Schnaubelt, Oskar von Stryk, Manuel Patchou, Stefan Böcker, Christian Wietfeld, Jan Quenzel, Daniel Schleich, Sven Behnke, Robert Grafe, Nils Heidemann, Dominik Slomma, and Ivana Kruijff-Korbayová. “Lessons from robot-assisted disaster response deployments by the German Rescue Robotics Center task force”. In: *Journal of Field Robotics* (2023).

Conference Publications

Frederik Bark, Kevin Daun, and Oskar von Stryk. “Affordance-based Actionable Semantic Mapping and Planning for Mobile Rescue Robots”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 53–60.

Kevin Daun, Stefan Kohlbrecher, Jürgen Sturm, and Oskar von Stryk. “Large scale 2d laser slam using truncated signed distance functions”. In: *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2019, pp. 222–228.

Kevin Daun, Marius Schnaubelt, Stefan Kohlbrecher, and Oskar von Stryk. “HectorGrapher: Continuous-time lidar SLAM with multi-resolution signed distance function registration for challenging terrain”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 152–159.

Kevin Daun and Oskar Von Stryk. “Requirements and challenges for autonomy and assistance functions for ground rescue robots in reconnaissance missions”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 83–90.

Gregor HW Gebhardt, Kevin Daun, Marius Schnaubelt, and Gerhard Neumann. “Learning robust policies for object manipulation with robot swarms”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 7688–7695.

Stefan Kohlbrecher, Florian Kunz, Dorothea Koert, Christian Rose, Paul Manns, Kevin Daun, Johannes Schubert, Alexander Stumpf, and Oskar von Stryk. “Towards highly reliable autonomy for urban search and rescue robots”. In: *RoboCup 2014: Robot World Cup XVIII 18*. Springer. 2015, pp. 118–129.

Ivana Kruijff-Korbayová, Robert Grafe, Nils Heidemann, Alexander Berrang, Cai Hussung, Christian Willms, Peter Fettke, Marius Beul, Jan Quenzel, Daniel Schleich, Sven Behnke, Janis Tiemann, Johannes Güldenring, Manuel Patchou, Christian Arendt, Christian Wietfeld, Kevin Daun, Marius Schnaubelt, Oskar von Stryk, Alexander Lel, Alexander Miller, Christof Röhrig, Thomas Straßmann, Thomas Barz, Stefan Soltau, Felix Kremer, Stefan Rilling, Rohan Haseloff, Stefan Grobelny, Artur Leinweber, Gerhard Senkowski, Marc Thurow, Dominik Slomma, and Hartmut Surmann. “German rescue robotics center (drz): A holistic approach for robotic systems assisting in emergency response”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 138–145.

Adrian Rebmann, Jana-Rebecca Rehse, Mira Pinter, Marius Schnaubelt, Kevin Daun, and Peter Fettke. “Iot-based activity recognition for process assistance in human-robot disaster response”. In: *Business Process Management Forum: BPM Forum 2020, Seville, Spain, September 13–18, 2020, Proceedings 18*. Springer. 2020, pp. 71–87.

Jonas SüB, Martin Volz, Kevin Daun, and Oskar von Stryk. “Online 2D-3D Radiation Mapping and Source Localization using Gaussian Processes with Mobile Ground Robots”. In: *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2023, pp. 39–46.

Moritz Torchalla, Marius Schnaubelt, Kevin Daun, and Oskar von Stryk. “Robust Multisensor Fusion for Reliable Mapping and Navigation in Degraded Visual Conditions”. In: *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2021, pp. 110–117.

Data Sets

Kevin Daun, Schnaubelt Marius, and Oskar von Stryk. *DRZ Living Lab Tracked Robot SLAM Dataset*. 2021. URL: `https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3973`.

Kevin Daun, Jonas Süß, Martin Volz, and Oskar von Stryk. *Hector Enrich 2023 Radiation Mapping Dataset*. 2023. URL: `https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3974`.

Erklärungen laut Promotionsordnung

§ 8 Abs. 1 lit. d PromO

Ich versichere hiermit, dass zu einem vorherigen Zeitpunkt noch keine Promotion versucht wurde. In diesem Fall sind nähere Angaben über Zeitpunkt, Hochschule, Dissertationsthema und Ergebnis dieses Versuchs mitzuteilen.

§ 9 Abs. 1 PromO

Ich versichere hiermit, dass die vorliegende Dissertation – abgesehen von den in ihr ausdrücklich genannten Hilfen – selbstständig verfasst wurde und dass die „Grundsätze zur Sicherung guter wissenschaftlicher Praxis an der Technischen Universität Darmstadt“ und die „Leitlinien zum Umgang mit digitalen Forschungsdaten an der TU Darmstadt“ in den jeweils aktuellen Versionen bei der Verfassung der Dissertation beachtet wurden.

§ 9 Abs. 2 PromO

Die Arbeit hat bisher noch nicht zu Prüfungszwecken gedient.

Darmstadt, 26.02.2024

K. Daun