Robust large-scale mapping and localization

—

Combining robust sensing and introspection

Örebro Studies in Technology 100

Daniel Adolfsson

Robust large-scale mapping and localization

—

Combining robust sensing and introspection

Cover image: Estimated trajectory and map using the method TBV Radar SLAM, proposed in this thesis. The trajectory of TBV Radar SLAM is indicated in blue and can be compared to the ground truth indicated in red. The sequence is part of the Oxford Radar RobotCar Dataset.

Title: Robust large-scale mapping and localization—Combining robust sensing and introspection

Publisher: Örebro University, 2023
www.oru.se/forskning/publikationer-och-avhandlingar

Printer: Örebro University, Repro 10/2023

ISSN 1650-8580
ISBN 978-91-7529-524-4

Abstract

Daniel Adolfsson (2023): Robust large-scale mapping and localization—Combining robust sensing and introspection. Örebro Studies in Technology 100.

The presence of autonomous systems is rapidly increasing in society and industry. To achieve successful, efficient, and safe deployment of autonomous systems, they must be navigated by means of highly robust localization systems. Additionally, these systems need to localize accurately and efficiently in real-time under adverse environmental conditions, and within considerably diverse and new previously unseen environments.

This thesis focuses on investigating methods to achieve robust large-scale localization and mapping, incorporating robustness at multiple stages. Specifically, the research explores methods with *sensory robustness*, utilizing radar, which exhibits tolerance to harsh weather, dust, and variations in lighting conditions. Furthermore, the thesis presents methods with *algorithmic robustness*, which prevent failures by incorporating introspective awareness of localization quality. This thesis aims to answer the following research questions:

How can radar data be efficiently filtered and represented for robust radar odometry? How can accurate and robust odometry be achieved with radar? How can localization quality be assessed and leveraged for robust detection of localization failures? How can self-awareness of localization quality be utilized to enhance the robustness of a localization system?

While addressing these research questions, this thesis makes the following contributions to large-scale localization and mapping: A method for robust and efficient radar processing and state-of-the-art odometry estimation, and a method for self-assessment of localization quality and failure detection in lidar and radar localization. Self-assessment of localization quality is integrated into robust systems for large-scale Simultaneous Localization And Mapping, and rapid global localization in prior maps. These systems leverage self-assessment of localization quality to improve performance and prevent failures in loop closure and global localization, and consequently achieve safe robot localization.

The methods presented in this thesis were evaluated through comparative assessments of public benchmarks and real-world data collected from various industrial scenarios. These evaluations serve to validate the effectiveness and reliability of the proposed approaches. As a result, this research represents a significant advancement toward achieving highly robust localization capabilities with broad applicability.

Keywords: SLAM, Localization, Robustness, Radar

Daniel Adolfsson, School of Science and Technology
Örebro University, SE-701 82 Örebro, Sweden, dla.adolfsson@gmail.com

Acknowledgements

I would like to express my deepest gratitude to my supervisors, Henrik, Martin, and Stephanie, for their invaluable guidance and support throughout my doctoral journey. Henrik, thank you for all your support and patience, especially during the earlier and more challenging periods. I am grateful for your guidance and understanding, allowing me the time and freedom to pursue my research curiosities. Martin, I am immensely grateful for all our fruitful unplanned meetings. Your thorough reviews and honest critique of my work have been vital to my improvement. Stephanie, thank you for being readily available when I needed assistance, particularly during the earlier and later stages of my Ph.D. Furthermore, I would like to acknowledge the encouragement provided by Achim, who, although not my former supervisor, played a crucial role in my research journey.

I would also like to express my gratitude to my coauthors. Anas, Qianfang and Mattias, thank you for your collaboration and valuable contributions to my research projects. Manuel, Vladimír and Maximilian – a special thanks for your willingness to join my research projects shortly after arriving at the university. Your dedication and involvement have been remarkable.

I would like to extend my thanks to all my friends, colleagues, and visiting students at AASS. Your friendship and encouragement have made this time truly memorable.

Thanks to my family and friends. Most importantly, I want to express my heartfelt gratitude to Sofia. Her belief in me and constant encouragement were crucial in my decision to pursue a doctoral degree. Sofia's valuable feedback and careful review of my papers played a vital role in completing this thesis. I feel incredibly lucky to have her with me, especially for our new adventures together with Ellie.

Sofia, you provide meaning and inspiration – this achievement is ours.

“The reliable cartographer”

A robot, sleek and strong and bold,
Perceives the world, with senses untold,
Estimating motion, smooth as silk,
Robustly mapping, with perfect skill.

Revisiting places it has seen,
Its sensors keen, and evergreen,
Correcting estimates, with such ease,
Robustly mapping, with such expertise.

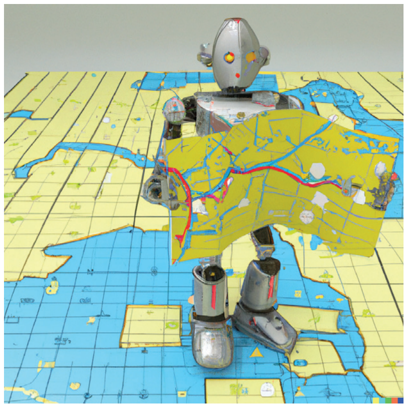

Each correction made, with careful care,
Maps become clearer, beyond compare,
A masterpiece, a sight to behold,
Robustly mapping, with stories untold.

Verifying every correction,
Sensors scanning, with perfect detection,
The maps it builds, a work of art,
Robustly mapping, with all its heart.

From tempest’s roar to blizzard’s blow,
The robot braves, with perfect glow,
Through storms unleashed, a world to know,
Robustly mapping, where harsh winds throw.

A map in hand, with landmarks bold,
The robot navigates with sight so cold,
Pose regression, its guiding star,
Helps it find its place, no matter how far.

ChatGPT, DALL·E2 and Daniel Adolfsson – 2023

Contents

List of Figures

List of Footnotes

List of publications

The work within this thesis has been published in a series of articles. For completeness, all versions of evolved articles–workshop, conference and journal–are included in this list.

Papers included in this thesis

Paper I Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim Lilienthal, and Henrik Andreasson. Oriented surface points for efficient and accurate radar odometry. In *Radar Perception for All-Weather Autonomy, a Half-Day Workshop, ICRA*, 2021. URL `https://arxiv.org/abs/2109.09994`. arXiv

Paper II Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. CFEAR radarodometry - conservative filtering for efficient and accurate radar odometry. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5462–5469, 2021

Paper III Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Lidar-level localization with radar? the CFEAR approach to accurate, fast, and robust large-scale radar odometry in diverse environments. *IEEE Transactions on Robotics*, 39(2):1476–1495, 2023

Paper IV Anas Alhashimi, Daniel Adolfsson, Martin Magnusson, Henrik Andreasson, and Achim J. Lilienthal. BFAR – bounded false alarm rate detector for improved radar odometry estimation, 2021. URL `https://arxiv.org/abs/2109.09669`

Paper V Daniel Adolfsson, Martin Magnusson, Qianfang Liao, Achim J. Lilienthal, and Henrik Andreasson. CorAl – are the point clouds correctly aligned? In *2021 European Conference on Mobile Robots (ECMR)*, pages 1–7, 2021

Paper VI Daniel Adolfsson, Manuel Castellano-Quero, Martin Magnusson, Achim J. Lilienthal, and Henrik Andreasson. CorAl: Introspection for robust radar and lidar perception in diverse environments using differential entropy. *Robotics and Autonomous Systems*, 155:104136, 2022. ISSN 0921-8890. URL `https://www.sciencedirect.com/science/article/pii/S0921889022000768`

Paper VII Daniel Adolfsson, Mattias Karlsson, Vladimír Kubelka, Martin Magnusson, and Henrik Andreasson. TBV radar SLAM – trust but verify loop candidates. *IEEE Robotics and Automation Letters*, 8(6):3613–3620, 2023

Paper VIII Li Sun, Daniel Adolfsson, Martin Magnusson, Henrik Andreasson, Ingmar Posner, and Tom Duckett. Localising faster: Efficient and precise lidar-based robot localisation in large-scale environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4386–4392, 2020

Paper IX Zhicheng Zhou, Cheng Zhao, Daniel Adolfsson, Songzhi Su, Yang Gao, Tom Duckett, and Li Sun. NDT-transformer: Large-scale 3d point cloud localisation using the normal distribution transform representation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5654–5660, 2021

Paper X Daniel Adolfsson, Stephanie Lowry, and Henrik Andreasson. Improving localisation accuracy using submaps in warehouses. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Workshop on Robotics for Logistics in Warehouses and Environments Shared with Humans*, 2018

Paper XI Daniel Adolfsson, Stephanie Lowry, Martin Magnusson, Achim Lilienthal, and Henrik Andreasson. A submap per perspective - selecting subsets for SuPer mapping that afford superior localization quality. In *2019 European Conference on Mobile Robots (ECMR)*, pages 1–7, 2019

Paper XII Henrik Andreasson, Daniel Adolfsson, Todor Stoyanov, Martin Magnusson, and Achim J. Lilienthal. Incorporating ego-motion uncertainty estimates in range data registration. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1389–1395, 2017

Author contributions

For most articles, D. Adolfsson contributed to the majority of the design, implementation, evaluation, analysis, writing, and presentation. For the remaining articles, the contributions of D. Adolfsson are stated for each paper. Additionally, a select set of coauthor contributions, which deserve additional recognition, are stated here.

Paper II A. Alhashimi contributed to the evaluation.

Paper IV D. Adolfsson contributed to the literature survey and provided the odometry estimation module which the article builds upon.

Paper V Q. Liao contributed to the evaluation.

Paper VI M. Castellano-Quero contributed to the implementation and evaluation of the system.

Paper VII M. Karlsson contributed to the evaluation, and implementation of the place recognition module. V. Kubelka contributed to the implementation and provided a thorough analysis of odometry uncertainty which was briefly discussed in the paper.

Paper VIII D. Adolfsson contributed with the framework for topometric mapping and localization. Additionally. D. Adolfsson did a majority of the design, implementation, evaluation, and writing related to the integration of deep pose predictions with MCL localization.

Paper IX D. Adolfsson contributed with a method that condense a 3d point cloud into a fixed size *NDT* representation, such that geometrical features are preserved. A fixed and accurate NDT representation is required to learn descriptors for robust place recognition.

Paper XII D. Adolfsson contributed to writing, revising, and presenting of the paper.

Chapter 1
Introduction

Autonomous systems are becoming increasingly more present in society and industry, operating in the presence of or in collaboration with humans. Applications range from delivery services, self-driving vehicles, and transportation within large-scale urban environments, to industrial forklifts and machines operating within mining and construction scenarios. These systems can enhance societal efficiency and productivity, and play a crucial role in freeing humans from hazardous – but necessary – work. However, a profound level of robustness and safety is required to successfully integrate autonomous systems into our society with a high degree of public acceptance. A fundamental aspect of robust operation lies in the system's ability to navigate by localizing safely and effectively, even under challenging conditions.

Safe and uninterrupted autonomous localization commonly requires all-weather robust sensing in harsh conditions. For example, self-driving vehicles must perceive their surroundings effectively during the day and night, and during adverse weather conditions such as heavy snow, rain, and fog. Similarly, mining and construction machines are expected to operate in the presence of dust. These conditions may reduce the visibility of exteroceptive sensors such as lidar, and, in particular, camera. As a consequence, subsequent methods for localization can be compromised.

Fortunately, radar – which has recently become compact and accurate with long range – remains largely unaffected under these conditions. Therefore, the sensor modality presents promising advantages for robust and uninterrupted autonomy even under adverse conditions.

However, system-wide robustness or resilience cannot be achieved solely from environment-tolerant sensing, but requires careful algorithmic consideration at all stages. In particular, localization algorithms need to be enhanced with introspective capabilities, to evaluate their success in executing their tasks, and act accordingly to improve their performance and prevent hazardous failures.

The primary goal of this thesis is to develop autonomous localization methods that emphasize robustness through the use of reliable sensors – such as radar – and robust algorithms for filtering, pose tracking, and place recognition. Furthermore, the thesis explores methods for introspective self-assessment of localization quality, which is leveraged to enhance the overall system-level robustness for both Simultaneous Localization And Mapping (SLAM) and global localization in prior maps.

1.1 Thesis outline

This chapter introduces the research topics and presents the research questions and contributions of this thesis. The articles that constitute the core of this thesis are appended to this thesis and briefly summarized in Sec. 1.5. Chapters 2, 3 and 4 summarize the work carried out in this thesis per research topic. For each topic, the related work and published articles are presented, followed by a discussion on select topics and findings. Finally, Chapter 5 concludes with a summary of the thesis' contributions to the research questions and suggests future research topics.

1.2 Problem statement

This thesis focuses on robust systems for large-scale Simultaneous Localization And Mapping, and global localization and tracking in prior maps. In the scope of this thesis, a system is considered robust if performing consistently well, with a low failure rate, in a wide range of environments and over a long period of time [35, 112]. For that reason, the thesis investigates localization and mapping methods with a focus on understanding how performance generalizes across diverse environments.

A localization system requires prior maps that must be created at some point, for example, using SLAM. We distinguish between SLAM, which focuses on building new maps, and localization in prior maps, that is, the task of localizing within these previously built maps. However, both tasks can have the following subtasks in common: *pose tracking* and *place recognition*. Pose tracking for SLAM is more commonly referred to as *odometry estimation*, and generally does not rely on prior maps. Within a SLAM session, place recognition is more commonly referred to as *loop closure*, or rather takes place as one of several subtasks within a loop closure subsystem. This thesis investigates methods for odometry and pose tracking, loop closure in SLAM, and global localization.

1.2.1 Radar interpretation and odometry estimation

Long-term failure-free localization requires a sensor modality that can operate well during day and night, is unaffected by dust and precipitation, and provides long-range, accurate, and information-rich data. For that reason, the thesis is

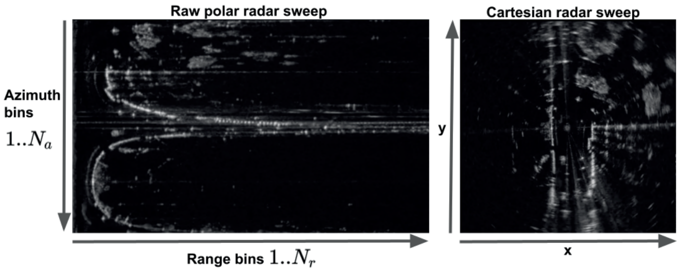

Figure 1.1: A full 360° sweep from a spinning radar, presented in its original polar form (left) and computed Cartesian form. Scan the QR code or follow the link found in the footnote for a video visualization of spinning radar data.

mainly concerned with spinning radar and, to a lesser extent, lidar. The benefits of robust radar sensing do, however, come at the price of more data to process and challenging data interpretation. Radar has challenging noise characteristics, including speckle noise, receiver saturation, and multipath reflections. The radar considered in this thesis is a mechanically spinning single-beam Frequency-Modulated Continuous Wave (FMCW) radar, specifically the model series "CIR" developed by Navtech. The radar is continuously spinning while separating measurements into 360° sweeps. An example of spinning radar data is presented in Fig. 1.1 [1]. A visualization of radar and lidar data is presented in Fig. 1.2, and a comparison of their resilience to dust is presented in Fig. 1.3 [2].

Lidar generally provides sparse point clouds with a single reading per beam. On the contrary, the radar considered in this thesis provides dense and detailed range-intensity measurements according to its operating resolution. In other words, radar intensity measurements are provided for every selection of range and bearing around the sensor. Given the considerable amount of noisy measurements and the high amount of information, filtering and creating lightweight representations for real-time localization is investigated in this thesis. Specifically, the thesis explores interpreting measurement sweeps as a set of sparse oriented surface points, which serves as an accurate and compact condensation of radar sweeps, and is suitable for efficient localization. These efficiency and noise considerations are essential for enabling the promising radar modality to be used as a primary source of localization.

In addition to radar interpretation, estimating radar odometry involves *scan matching* (or *registration*), local mapping, and considerations of the motion kinematics and sensing principles of the mechanically spinning radar. These

[1] https://tinyurl.com/SpinningRadarData
[2] https://tinyurl.com/RadarVsLidar

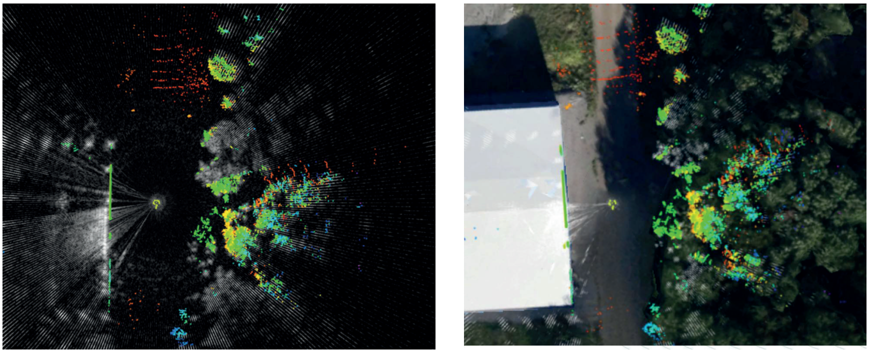

Figure 1.2: A visual comparison between lidar and radar. Range-intensity measurements are visualized in color (lidar) and grayscale (radar), on top of a black background (left), and on an overhead image (right).

aspects are studied with the aim of robustness and low drift, for accurate mapping of new previously unseen environments.

1.2.2 Localization quality self-assessment

Additionally, this thesis aims to explore self-assessment of localization quality for improving localization robustness. [**Paper V-VI**] study localization quality self-assessment in an isolated context. In contrast, [**Paper VII**] and [**Paper VIII**] integrate localization quality assessment into full-fledged methods for SLAM and localization in prior maps, and consequently, improve robustness.

Localization quality is explored from multiple perspectives. In [**Paper VI**], the quality of localization is assessed and represented as a *likelihood*, i.e. a single scalar. The quality is assessed between a point cloud pair that is expected to have been aligned by the localization system. In this context, the quality represents the likelihood of points being correctly aligned, or as assumed to be analogous, a pose estimate being successful. [**Paper VII**] fuses the previous alignment quality with other sources that are suitable for assessing loop closure quality. Similarly, the loop quality assessment is represented as the likelihood of correct alignment, although it is applied specifically for verifying that loop closures are correct.

Another formulation of localization quality is adopted in [**Paper VIII**]. In the article, the prediction of localization quality is learned and represented as a pose with uncertainty, that is, a *pose distribution*. The prediction, or quality, is then leveraged during the integration of corresponding global pose estimates to improve robustness. Both these representations, i.e. the scalar likelihood, and

Figure 1.3: A visual comparison of sensors' resiliency to dust. Range-intensity measurements are visualized for lidar (top image, in color) and radar (top, grayscale). The lidar cannot penetrate the dust, caused by a nearby machine that performs a brushing task. The lidar outputs a cloud of point cloud of false detections in close proximity to the sensor. Follow the QR code or click the link in the footnote for a video visualization.

the distribution, have advantages and are used for different tasks within this thesis.

In summary, the thesis aims to improve the algorithmic robustness of these previously mentioned subtasks and integrate these techniques into reliable and safe methods for SLAM in diverse environments and for localization in previous maps.

1.2.3 Limitations

Today, there exist multiple types of radar, including spinning automotive, SoC, and 4D radar, of which some provide Doppler information to measure the speed of objects. Each radar type poses its own challenges and opportunities for various applications. For an overview of the various types, the reader is referred to literature survey articles [55, 113]. The radar considered in this thesis is a spinning radar without Doppler information.

1.3 Research questions

Based on the objectives of this thesis, the following research questions were formulated.

RQ1: ***How can radar data be efficiently filtered and represented for robust radar odometry?***

Spinning radar has great range, accuracy, and detail level. However, it produces a large amount of data with challenging noise characteristics. Hence, for efficient and accurate localization, data needs to be filtered and processed into a lightweight representation that is suitable for the task. Processing range data is currently an important focus in radar localization research and requires that quality is maintained over diverse environments without model or parameter changes.

This question is addressed in [**Paper I-IV**].

RQ2: ***How can accurate and robust odometry be achieved with radar?***
Radar-based odometry requires algorithmic robustness within diverse, possibly feature-poor environments (with varied scales) while maintaining consistently high accuracy. This necessitates consideration of numerous aspects, including initial filtering and processing, but also scan- and map representation, registration, and sensor fusion.

This question is addressed in [**Paper I-IV,XII**].

RQ3: ***How can localization quality be assessed and leveraged for robust detection of localization failures?***
Introspective self-assessment of localization quality is crucial for safe navigation. Failures need to be prevented or detected and recovered from. Since pose estimation commonly relies on point cloud registration, evaluating alignment quality – after the alignment of point clouds – could provide a means to assess localization quality. This research aims to investigate how and to which extent alignment quality can be assessed and utilized to detect point cloud misalignment, and in turn, detect localization failures. This question is addressed in [**Paper V-VI**].

RQ4: ***How can self-awareness of localization quality be utilized to enhance the robustness of a localization system?*** Localization quality needs to be leveraged by the localization system to improve its robustness. This thesis investigates the use of quality assessment for global localization within prior maps, and for radar SLAM by enhancing loop closure verification. This question is addressed in [**Paper VII, VIII**].

1.4 Thesis contributions

The most important contributions of the appended articles are listed below.

1. A robust, efficient, and accurate real-time pipeline for radar processing and state-of-the-art radar odometry estimation that generalizes across environments. A detailed ablation study gives insight into the most important aspects of radar odometry. A sensor fusion approach is presented which integrates an external source of odometry with uncertainty into registration, hence achieving robustness within highly feature-poor environments. [**Paper I-IV,XII**] and contributes to **RQ1** and **RQ2**.

2. A method for fault detection in lidar and radar localization that generalizes across environments. Localization quality is assessed from alignment quality that is sensitive to small misalignments. This is addressed in [**Paper V-VII**] and contributes to **RQ3**.

3. A framework for introspective and accurate real-time radar SLAM that generalizes well across environments. The framework integrates a novel unified loop-retrieval and verification module that combines multiple sources of information to improve loop quality assessment for highly robust loop closure. This is addressed in [**Paper VII**] and contributes to **RQ4**.

4. A hybrid metric-topological localization framework that improves localization utilizing specialized submaps. The proposed technique groups nearby measurements into a set of submaps. During localization, the robot selects the most suitable submap from the current robot's perspective. The evaluation shows that a set of local maps is preferred to a single monolithic map for accurate pose tracking. This is addressed in [**Paper X,XI**] and contributes to pose tracking more generally.

5. A method for rapid global localization in large-scale environments by integrating localization estimates including measures of uncertainty – acquired from deep Gaussian pose regression – together with Monte Carlo localization, via importance resampling. This is addressed in [**Paper VIII**] and contributes to **RQ4**.

1.5 Paper summary

The following section summarizes the articles included in this thesis.

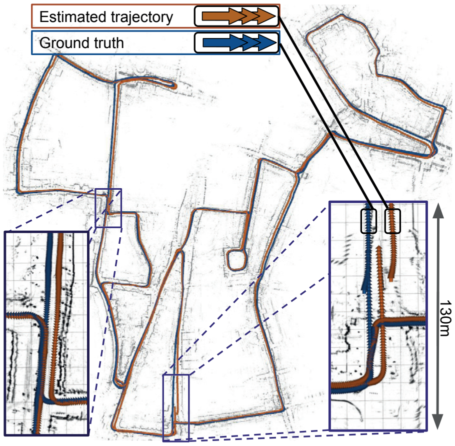

Figure 1.4: Overview and demo of CFEAR radar odometry

1.5.1 Abstract of Paper I-III

Papers I-III propose "CFEAR radar odometry" (see Fig. 1.4 [3]), a method for spinning radar odometry estimation in diverse large-scale environments. CFEAR processes radar data into a lightweight representation for efficient and accurate registration. Motion is estimated by register the latest scan to a history of previous scans jointly. **[Paper III]** provides an ablation study that quantifies the importance of each component in the odometry pipeline, including radar filtering, surface point resolution, registration cost, loss function, local keyframe history, weighted surface point estimation, weighted correspondences, and motion compensation. Contributions include:

- A pipeline for radar processing (detailed on Fig. 2.3 [4]), for efficient state-of-the-art odometry estimation.
- A thorough ablation study that quantifies the importance of components included in the processing and odometry pipeline.
- An evaluation that quantitatively demonstrates how well the method generalizes across environments. Multi-keyframe registration and residual quality weighting play important roles in robustness within feature-poor environments, while generalizing better to new environments.
- Dataset, C++ code, and evaluation have been released to the community.

[3] https://tinyurl.com/CFEARDemoT-RO
[4] https://tinyurl.com/CfearOverview

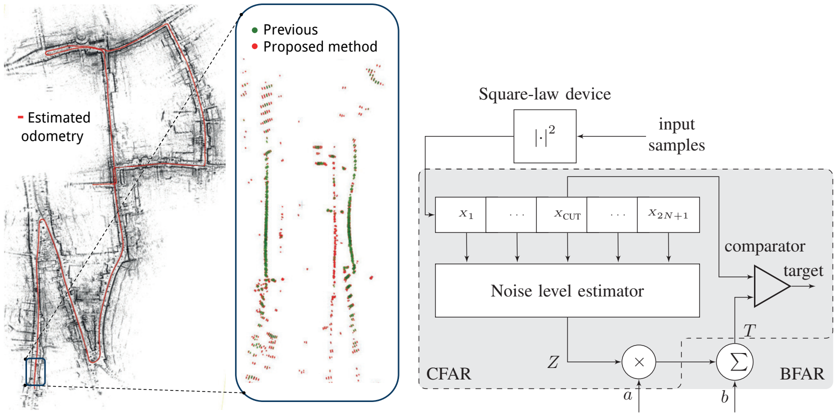

Figure 1.5: Left: Estimated odometry, and filtered radar data using the proposed method “BFAR”. Right: Operating principle of BFAR.

1.5.2 Abstract of Paper IV

[Paper IV] presents BFAR – a detector for filtering FMCW radar data presented in Fig. 1.5. Filtering of radar data is considered challenging for applications with autonomous ground vehicles, due to its heterogeneous and context-dependent noise characteristics. The classical radar filter Cell Averaging Constant False Alarm Rate (CA-CFAR) operates by adaptively estimating the noise level, from which a landmark can be distinguished. BFAR extends CA-CFAR by applying an affine transformation on the adaptive noise level, thus combining fixed and adaptive filtering. The fixed noise level can be tuned or learned from data. BFAR has been evaluated on the task of radar odometry, replacing the default k-*strongest* filter from the CFEAR pipeline in Papers I-III" (the version of CFEAR presented in **[Paper II]**). Evaluation demonstrated an improvement with 12% reduced drift. The insight that a combination of adaptive and fixed estimation of background noise level is preferred over fully adaptive, is in line with the finding in **[Paper III]** (Sec IV.A.1).

To summarize, **[Paper IV]** present the following contributions:

- A novel radar target detection method named “BFAR” that improves the detection of features useful for radar odometry.
- An evaluation demonstrating the effectiveness of BFAR in the radar odometry pipeline “CFEAR” with an improvement in odometry quality.

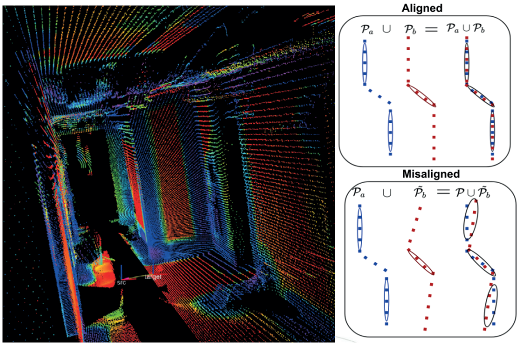

Figure 1.6: Left: per-point differential entropy. Middle: principle of CorAl alignment quality assessment.

1.5.3 Abstract of Paper V-VI

[Paper V-VI] presents CorAl [5]. The name CorAl comes from *are the point clouds correctly aligned?*. CorAl is a method for fault detection within lidar and radar localization systems. Fault detection in localization is formulated as a misalignment problem, following point cloud registration. If misalignment can be detected between point cloud pairs, it is an indication that a failure has occurred. This principle can be applied at various stages in a localization pipeline, e.g. during loop closing – after registering a scan with a potential loop candidate, or during pose tracking – to consecutively aligned scans, or between scan and map when localizing in prior maps. CorAl makes use of dual per-point differential entropy measurements, before and after joining the pair of scans as depicted in Fig. 1.6. These dual measurements can be used to classify misalignment and highlight segments where misalignment is evident. The decision boundary between aligned and misaligned point clouds is automatically learned from a set of aligned point cloud pairs. The evaluation shows that CorAl excels at detecting small misalignments. Additionally, CorAl can be tuned for classification in multiple environments jointly and even generalize well to new unseen environments. The contributions are:

- A simple and intuitive measure of alignment *correctness*.
- A method for classification of alignment correctness which generalizes across environments without retraining.
- Automatic learning of classification from aligned point clouds.
- An ablation study that investigates how practical aspects such as distance and error magnitude affect performance.

[5] `https://tinyurl.com/CorAl-demo`

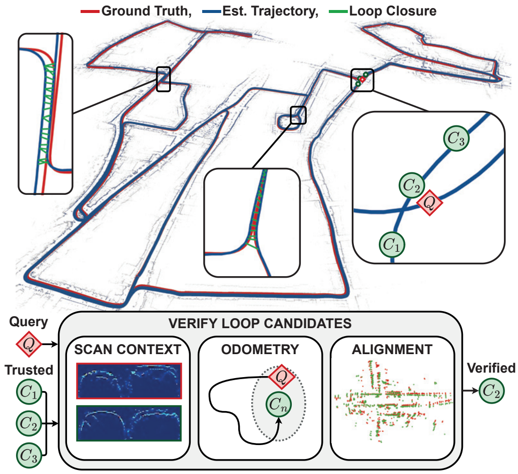

demo presentation

Figure 1.7: Verification of loop constraints within TBV Radar SLAM.

1.5.4 Abstract of Paper VII

Robust and safe SLAM requires fault detection and awareness at multiple stages, from sensing and odometry estimation to loop closure. [**Paper VII**] presents TBV (Trust But Verify) Radar SLAM, a method that explores the integration of robustness for large-scale radar-only SLAM. TBV is briefly presented here [6], with the full system demonstrated in Fig. 1.7 [7].

TBV combines the dust and precipitation-resistant radar modality, together with robust odometry estimation ([**Paper I-III**]), and a loop closure module that introspectively checks loop candidates. A high rate of correct-loop-retrievals with rare false positives is achieved, making SLAM both accurate and safe. Key to the results is the introspective evaluation of loop closure candidates. The evaluation combines an analysis of place matching quality, an alignment verification (based on [**Paper V-VI**]), and accordance with odometry uncertainty. TBV Radar SLAM achieves state-of-the-art performance with 64% reduced trajectory error while generalizing to substantially different environments. In summary, the following contributions are presented.

- A combination of techniques for a high level of correct loop retrievals.
- A loop verification step that retrieves, registers, verifies, and selects between multiple candidates for robust and safe loop closure.

[6] `https://tinyurl.com/TbvIros23Short`
[7] `https://tinyurl.com/TBVRadarSLAM`

- A framework for radar-only, real-time SLAM in large-scale environments that extends the state-of-the-art, and generalizes across environments without the need to change parameters.

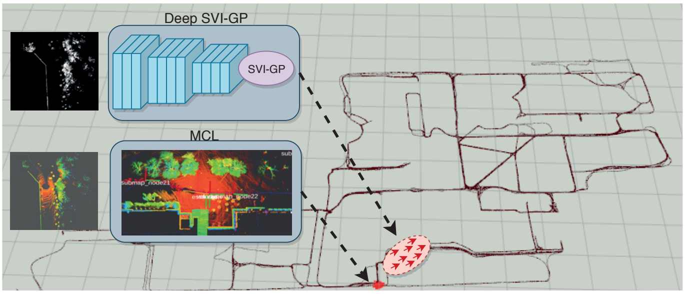

Figure 1.8: Overview of global localization using “Localizing faster”

1.5.5 Abstract of Paper VIII

[Paper VIII] proposes a novel approach for global localization in large-scale environments. Classical map-based localization methods such as Monte Carlo Localization (MCL) have demonstrated the feasibility of accurate and robust localization in prior maps, but become intractable for large-scale environments. However, recent learning methods can rapidly estimate the 6-DOF pose albeit at a lower level of accuracy. This article proposes a system that combines fast learning-based localization with accurate MCL – thus achieving both fast and accurate localization. Specifically, the learning method estimates a pose distribution (using only a bird’s-eye view image) using Gaussian process regression with a deep kernel. The pose distribution is used to seed MCL during each importance sampling step. The particle weights are calculated based on a scan-to-map alignment within a topological localization system (see **[Paper XI]**). The pose particles in MCL are then refined iteratively. Experiments demonstrate that a localization accuracy of 0.75 m can be reached within 1.94 s (median of 0.8 s) in a large-scale environment of approximately 0.5 km^2.

The article, which is demonstrated in Fig. 1.8 [8], presents the following contributions:

- A hybrid deep one-shot / Monte Carlo localization framework for fast and smooth global localization in large-scale environments.
- Integration of rough pose distribution estimates into a localization system. This is done by sampling particles from the pose distribution that are integrated into MCL. The notion of localization uncertainty (pose distribution) allows for faster and more reliable localization.
- A deep non-parametric model for estimating the distribution of global 6-DOF pose from range data.

[8] `https://tinyurl.com/LocalisingFaster`

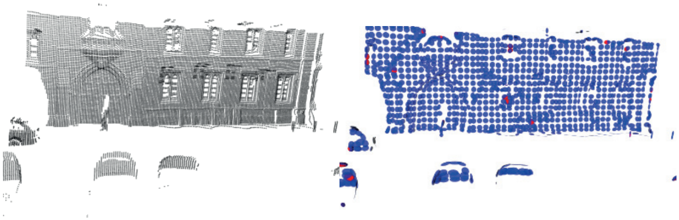

Figure 1.9: The proposed method NDT Transformers condenses a 3d point cloud into a fixed-size NDT representation from which place recognition is learned. Left (in gray): Raw point cloud within an urban street scene, Center (in blue): Fix size NDT set that effectively retains the geometric features within the scene. Center (in red): NDT cells that contain redundant geometrical information are removed using the proposed algorithm.

1.5.6 Abstract of Paper IX

3D point cloud-based place recognition is a crucial component of lidar-based SLAM systems and is highly demanded in autonomous driving within large-scale GPS-challenged environments. **[Paper IX]** presents "NDT-Transformer" – a method for learning place recognition from 3d point clouds. The method is demonstrated in presented in Fig. 1.9 [9]. The proposed approach employs a 3D Normal Distribution Transform (NDT) representation to condense the raw, dense 3D point clouds into probabilistic distributions (NDT cells), providing the geometrical shape description. The NDT-Transformer network then learns a global descriptor, from a set of 3D NDT cell representations, enriched with both geometrical and contextual information. Descriptor retrieval is achieved through a query database for place recognition. The evaluation shows that NDT-Transformer scores a 7.52% improvement in average top 1 recall and 2.73% in average top 1% recall on the Oxford Robotcar benchmark. The combination of NDT representation and NDT-Transformer network enables the learned global descriptors to have a better representation of both geometrical and contextual information. The main contributions are:

- A computationally efficient method for large-scale point cloud recognition. The proposed method condenses a 3D point cloud into a fixed-size NDT representation, such that geometrical features are preserved.
- A neural network architecture named: "NDT Transformers" learns global descriptors with contextual clues from the NDT representation.

[9] `https://tinyurl.com/NdtTransformersPresent`

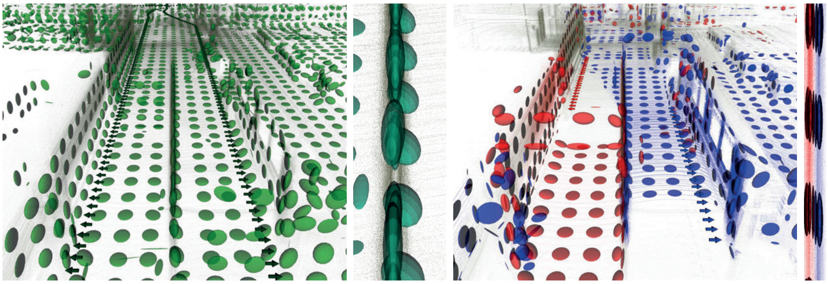

Figure 1.10: Green: Monolithic map of the environment. The center wall between the corridors is imprecisely mapped, with high uncertainty (inflated ellipsoids). Red/Blue: The same environment is mapped using two submaps, centered on each side of the center wall. When the wall is mapped once from each side, the mapping uncertainty is greatly reduced (flat ellipsoids), and localization is consequently improved.

1.5.7 Abstract of Paper X-XI

Precise localization requires highly accurate geometric maps of the environment. A challenge with monolithic voxel-based map representations is that the expressiveness of the map is limited by the resolution and that integration of measurements from different perspectives introduces imprecisions, thereby reducing localization accuracy. **[Paper X-XI]** present SuPer (Submap per Perspective) – an extra map pre-processing step which runs after SLAM for improved localization accuracy. SuPer uses a pose graph to compute a hybrid metric-topological (HMT) map, i.e. a set of local submaps. Each submap represents the environment as observed around its origin, thus alleviating the imprecisions and the need for a lower map resolution. After preprocessing, the robot navigates and selects the submap that best explains the environment from its current perspective. The proposed method was evaluated on real-world and simulated data from an industrial scenario. The results demonstrate significantly improved localization accuracy, with up to 46% improvement compared to localization in global maps and up to 25% better compared to alternative submapping approaches. The following contributions are presented in the articles.

- An insight into the challenge with monolithic maps, and an evaluation that demonstrates that: for localization, a set of smaller local maps is preferred over a single monolithic map.
- A strategy for clustering scans into submaps that outperform the commonly adopted incremental submap strategy.
- A fast submap selection strategy that retries the submap that best fits the current perspective of the robot.

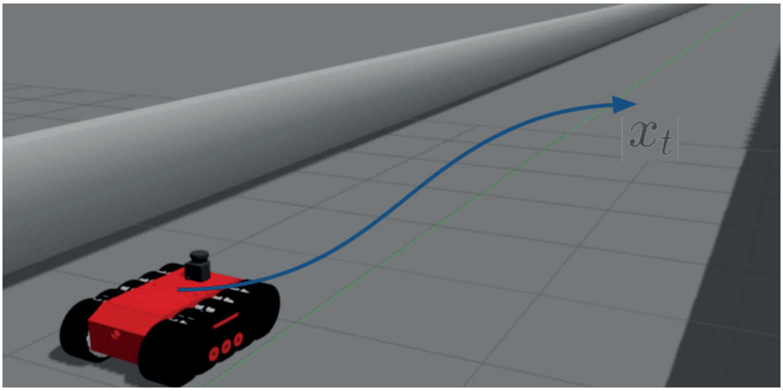

Figure 1.11: Localizing in feature-poor environments with range measurements is challenging and requires careful fusion of external motion estimates.

1.5.8 Abstract of Paper XII

In robotics navigation, point cloud registration is essential for the robot to track its pose. Pose tracking can be improved by the use of additionally available sources of ego-motion estimates, such as wheel odometry. This is often done by using the ego-motion estimate to predict an initial alignment for point cloud registration, or by fusing the two separate sources based on maximum likelihood estimation. [**Paper XII**] describes an improved strategy for incorporating ego-motion estimates (with uncertainty) into point-cloud registration. This is done by penalizing solutions during alignment optimization, which violates ego-motion uncertainty. As a consequence, registration robustness is greatly increased, especially within feature-poor and self-similar environments where previous fusion strategies are likely to fail. The experimental evaluation demonstrates a significant improvement in accuracy within ambiguous robotics working scenarios, including an industrial production site and warehouse environments. In summary, the presented contribution is:

- A method for incorporating ego-motion uncertainty measurements into point cloud registration for improved robustness in pose tracking within feature-poor environments.

Chapter 2
Radar interpretation and odometry estimation

A fundamental challenge to achieving robust localization is the interpretation of sensor data and the creation of suitable representations. This chapter discusses the interpretation and creation of radar data representations with a primary focus on *odometry estimation* – a key challenge in localization. The representations presented in this chapter are, however, not only suitable for odometry. For example, the articles presented in chapter 3 make use of the proposed representations for usage in radar fault detection, loop closure and SLAM.

2.1 Related work – radar filtering and representation

Radars' resilience to adverse conditions brings favorable properties for robust localization, which has motivated localization and mapping research in various directions. Some directions include target detection [73], mapping [70, 105] to distinguish between free and occupied spaces, and place recognition [46, 52, 60], SLAM [60, 101] and odometry estimation [32, 41, 106, 111].

The research topic of target detection, i.e. distinguishing true objects from background noise, dates back more than half a century. CA-CFAR was proposed in 1968 by Finn and Johnson [51], and since then, more than 25 variants have been proposed [73], including OS-CFAR and GO-CFAR [86]. Recent work where Constant False Alarm Rate (CFAR) has been applied to robotics localization has found that CA-CFAR is hard to tune and sensitive to different environments [40, 41].

Interpreting radar data for localization is generally considered challenging, largely because of the complex noise characteristics inherited from the radar modality in interaction with the environment. Interestingly, unfiltered radar data has – under some favorable conditions [65] – been successfully used for place recognition. On the contrary, radar odometry estimation has little to gain by leaving the filtering step out [9, 106] due to reliance on landmarks being consistent from different viewpoints. For that reason, previous work on radar odometry and pose tracking has largely focused on filtering and creating representations that are suitable for registration.

Representations used for pose tracking can be categorized as dense or sparse. Dense representations keep the original polar or Cartesian representation [15, 23, 81, 106]. This allows for modeling of both free and occupied space and retaining the detail level in the original data. Sparse methods [31, 40, 41, 59, 60] on the other hand attempt to extract lightweight point sets from the data for computationally efficient tracking. Some of these additionally compute per-point descriptors [21, 32], normals [5, 33, 101] or distributions [5, 69].

Previous methods for radar odometry extract features from regions with high intensity and low gradient [40, 41], from corner regions [60], or from regions with high intensity compared to fixed [69, 79] or adaptive estimation of noise level [33]. Some operate in two stages [33, 69, 101] by computing the surface normal or distribution after initial filtering. Feature extraction in two stages is adopted in this thesis.

Some methods circumvent the need for human expert knowledge to formulate the extraction or filtering criteria by instead relying on deep neural networks to learn the filtering of raw data [15, 23, 106], extraction of features [21, 32], or prediction of occupancy [70, 105]. Although recent work has demonstrated that extraction can be learned without the need for ground truth labeling [32], the learning-based approaches have required a significant amount of data [23, 32, 106]. The processing technique in state-of-the-art radar odometry, which is presented in **[Paper III]**, does neither require learning nor a significant amount of data, and generalizes to diverse environments without the need of retuning parameters. It is however unknown if, or to what extent, existing learning-based methods for filtering or extraction could provide further improvements to non-learning odometry pipelines such as the one presented in this thesis.

Finally, the discussion of noise suppression for pose tracking could benefit from a registration or mapping perspective, which may integrate further noise suppression techniques. For example, noise can be suppressed during registration with outlier rejection techniques, including graph theoretic approaches [40, 41, 59, 60], correspondence weighting [21, 40, 41] and robust loss functions. In the mapping stage, landmarks can be refined sequentially over a series of scans. Consistency of landmarks, that appear in multiple scans, has been utilized for online removal of false (inconsistent) landmarks [101], and for generating training data for learning of filtering [15].

2.2 Proposed methods for radar filtering

This section presents the radar filtering techniques k *strongest* and *BFAR*, proposed in **[Paper I-III]** and **[Paper IV]** respectively. These methods are intended as a first initial filtering step before creating representations that are more suitable for odometry. In addition to filtering, the purpose is to drop the dense raw representation – that additionally incorporates free space information as depicted in Fig. 1.1 – in favor of a sparser point cloud representation.

Both methods operate on azimuths independently for efficient processing. Additionally, these methods build on the insight that false landmarks generally appear at lower intensity levels, as pointed out by Kung et al. [69]. For that reason, the methods integrate fixed thresholding as part of their detection criteria, which is sufficient to filter a large portion of the noise present in the range data. Specifically, most speckle noise and weaker multipath reflections appear with lower intensity and can be removed accordingly.

2.2.1 Filtering with k strongest

k *strongest* was designed as a lightweight and learning-free method for target detection. The filter complements the thresholding of static intensity ($z_i : z_i > z_{min}$) by bounding the number of output detections to a maximum of k ranked according to intensity. This ensures the mitigation of stronger multipath reflections and saturated receiver returns, which appear with consistently high intensity. Despite the bounding, the filter generally provides an ample amount of detections for odometry estimation. However, challenging cases include observing walls with high incident angles. In such cases, multipath reflections appear with higher intensity compared to true landmarks, and additional context is required to correctly filter reflections.

This detection criterion can be implemented with computational efficiency and provides sufficient detections for odometry estimation. Additionally, bounding the number of detections speeds up later stages of the localization pipeline and improves overall efficiency.

2.2.2 Filtering with BFAR

The proposed filter BFAR was designed as a combination of fixed and adaptive CA-CFAR target detection. In contrast to k strongest, BFAR does not limit the size of the output detection set, and allows for the detection of additional landmarks. In contrast to k strongest, the intensity level that surrounds a target landmark, i.e. neighboring cells in the range-intensity signal, is used to adaptively estimate the background noise level in addition to the fixed. This results in a more fine-grained detection even at local maxima with slightly lower intensity levels, resulting in improved odometry. The implementation is however less

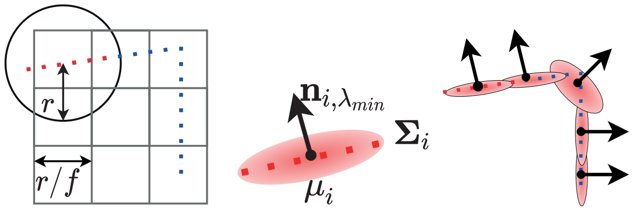

(a) (b) (c) (d)

Figure 2.1: The procedure of computing oriented surface points. (a) Points are inserted into a grid of cell size r/f. For each grid cell with points, the sample means and covariance are computed from all nearby points within a radius r of the cell centroid. (b) Sample covariance is used to estimate the direction of the surface normal. The final representation is shown in (c). Additional details are presented in (d).

efficient compared to k strongest, but real-time capable. Which filter to select depends on the application. More details are provided in [**Paper II-IV**].

2.2.3 Parameters and generalization

Both filtering methods depend on multiple parameters. Radar filtering parameters have previously been reported to be challenging to tune, particularly in different environments [40, 41]. For that reason, an ablation study was carried out (in [**Paper III**]) that investigates how parameters affect both drift and computational time. The study concluded that k strongest is largely insensitive to the selection of parameters, particularly with respect to CA-CFAR – which was included as a reference. BFAR, on the other hand, is more sensitive compared to k strongest, but less sensitive compared to CA-CFAR (as concluded in [**Paper IV**]).

For the application of odometry, which typically runs with limited prior knowledge of the target environment, the filtering needs to operate robustly without parameter changes. For that reason, a Spatial Cross-Validation (SCV) was carried out to observe how the parameter impacts drift within new unseen environments. In the study, the impact of numerous parameters was investigated, including the parameters of k strongest. Further details of this study are given in Sec. 2.5.4. To summarize the findings, the approximate same set of parameters was optimal for multiple environments and provided low drift. These results indicate the feasibility of robust filtering in a range of environments.

Figure 2.2: Radar data after filtering (green) and computing set of oriented surface points (black). Surface points are scaled according to planarity. Longer arrows indicate that the surfaces are more planar.

2.3 Radar representation - Oriented surface points

After having applied filtering as described in Sec. 2.2 to remove noise and keep significant points, the point cloud is further condensed into a sparser set of oriented surface points. Each oriented surface point contains a surface location, an estimated surface normal, and a landmark descriptor. This representation provides a compact summary of the filtered point set and can be utilized for efficient and accurate odometry. Specifically, a grid-based approach is used to compute intensity weighted distribution from points in nearby occupied grid cells. From the distribution, an oriented surface point is extracted from the surface location (mean) and surface normal (smallest Eigenvector). In addition, measures of planarity and sample size are kept with oriented surface points to aid in the association of correspondences during registration. An overview of the method is presented in Fig. 2.1 [1]. An example of an oriented surface point set computed within a city environment from the Oxford dataset [22] is shown in Fig. 2.2.

With an exception to the work by Kung et al. [69] – which was published in close time proximity to **[Paper II]** – surface normal or distribution was not explicitly modeled in contemporary sparse methods on radar odometry. Instead, these methods commonly attempt to detect landmarks that are stable in two directions [21, 40, 41]. Computing a sufficiently large set of such landmarks

[1] `https://tinyurl.com/RadarFiltering`

is challenging, although the need can be alleviated by a careful design of the registration cost function as discussed in [**Paper III**].

The benefit of modeling the surface normal or distribution, in terms of precision and robustness, was evaluated for the estimation of odometry in [**Paper III**], the alignment classification in [**Paper VI**], and for both the registration and the classification of the loop in [**Paper VI**]. In summary, the usage of the local point distribution allows for more accurate registration with less dense sets and for improved alignment classification. Furthermore, experiments carried out in [**Paper III**] demonstrate that considering surface normal is integral for robust odometry in feature-poor environments. The normals capture subtle but important geometrical variations in the scene and allow for later point cloud registration by minimizing distances in the direction in which the landmarks are stable. However, in addition to filtering and data representation, a range of additional factors play a critical role in robust odometry. Previous work – related to other aspects except filtering and feature extraction which were covered in Sec. 2.1 – on radar odometry is presented in the next section.

2.4 Related work – radar odometry estimation

The estimation of radar odometry has gained increasing interest in recent years [1], especially with improved sensor quality, the release of a series of public data sets [22, 34, 65, 95] and well-established benchmarking criteria [53].

Methods for spinning radar odometry can be categorized depending on whether sparse or dense representations are utilized in the registration. Note that the radar representations, filtering, and feature extraction were explicitly discussed in Sec. 2.

Sparse methods estimate odometry by registering feature sets [15–17, 21, 31, 32, 40, 41, 59, 60, 69, 72, 92, 101, 111], via RANSAC [31] or by minimizing the sum of (weighted) point-to-point, or point-to-distribution distances, over a correspondence set. The correspondences are found through the nearest neighbor search, by finding inliers through consistency graphs [59, 60], using motion constraints [17] or using RANSAC [31]. An important advantage of matching sparse representations is that efficiency can be retained at the original sensor resolution without the need for downsampling. However, performance is largely affected by the correspondence and minimization criteria. An extensive discussion on this topic is presented in [**Paper III**].

Dense methods, on the other hand, register scans – without the explicit need to find correspondences – by maximizing image correlation [23, 43, 81, 106] using the Fourier Mellin transform [85]. On one hand, the formulation is inherently well-suited for spinning radar as it utilizes information from the full sweeps. Unfortunately, the computational need scales poorly with sensor resolution and may require downsampling for sufficiently fast online usage. Although some work has been focused on improving the efficiency of dense registration [81, 106], current methods run at a significantly slower speed compared to sparse methods for point set registration. This is important as localization methods are designed with a trade-off between speed and performance, and computational resources may be scarce in autonomous systems.

An important source of recent advances has been the accounting of motion distortion [17, 31, 59, 60, 101, 111]. The effect is caused by the relatively low spinning rate of current spinning radars (commonly 4 Hz), especially compared to lidars which generally run at a faster rate(commonly 10-75 Hz). Finally, some methods perform estimate odometry by registering the latest scans to multiple [92], or to a local submap [58, 101, 111]. While automotive radar methods adopt this technique to overcome sparsity [58, 69], the advantages apply similarly to sparsely extracted feature sets. In summary, numerous aspects of radar odometry have been considered. The next section presents a method that integrates these techniques (presented in this section) for highly robust radar odometry. The impact of these techniques is thoroughly quantified in **[Paper III]**.

2.5 CFEAR radar odometry

Article **[Paper I-III]** investigates a pipeline for sparse and efficient radar odometry. An overview of the system is presented in Fig. 2.3. The pipeline can be separated into three parts:

1. Computing of CFEAR features, which filters range data (see Sec 2.2), computes oriented surface points (as previously discussed in Sec. 2.3), and compensates for motion distortion. The motion compensation uses a constant velocity model and timestamp of each measurement to *deskew* the filtered range data by applying a non-rigid transformation. The ablation study in **[Paper III]** confirmed the validity of the constant velocity model.

2. Local mapping, which stores historical keyframes for registration.

3. Registration, which estimates motion by aligning sets of oriented surface points. Registration minimizes a surface points distance metric. Residuals between surface points are scaled by a robust loss function and weighted according to correspondence quality. Correspondences are found between the latest scan and a set of previous keyframe scans. A starting point for registration is provided by a constant velocity model.

One of the key contributions is the registration cost function, which jointly registers the current scan with multiple keyframe scans. This technique improves drift, pose accuracy, reduces bias, and enhances robustness to changes in the environment.

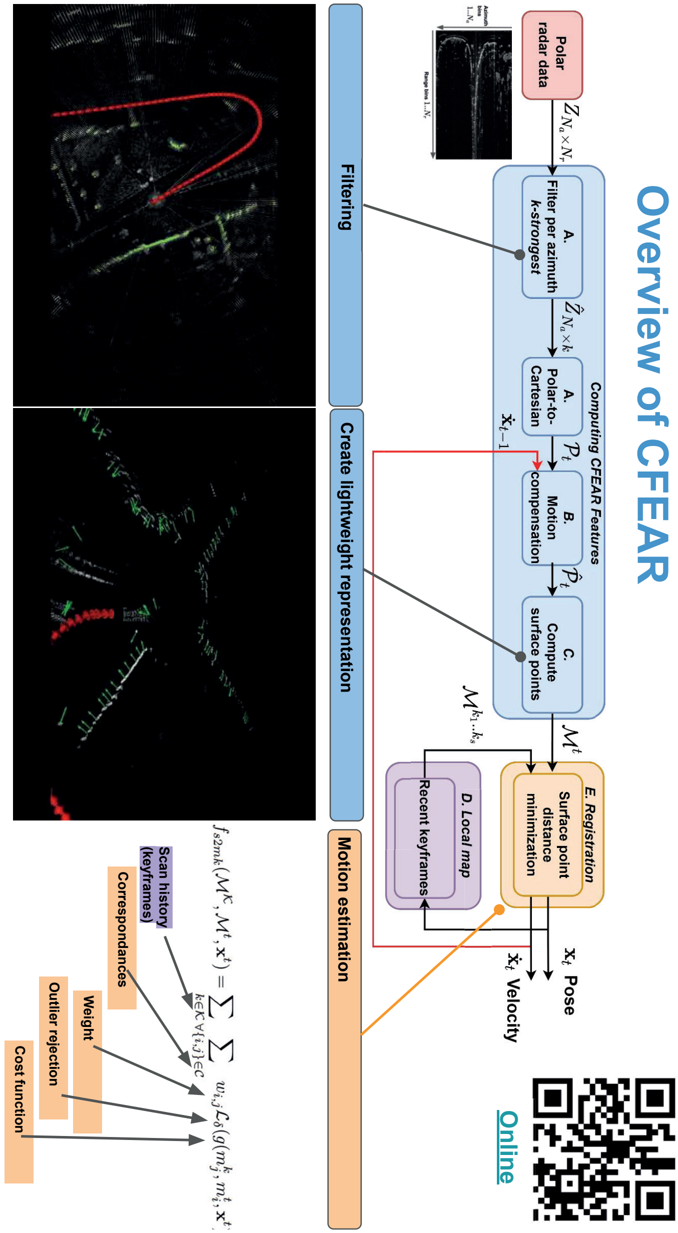

Figure 2.3: The CFEAR radar odometry pipeline.

The evaluation presented in **[Paper III]** quantifies performance in terms of drift and accuracy. Note that, pose accuracy (Relative Pose Error (RPE)) is measured between consecutive poses, while drift is measured over longer distances [53]. For that reason, drift is a preferred measure of robustness, as even infrequently occurring failures (i.e., large errors in position or orientation) have a large impact over longer distances. Therefore, our evaluation of robustness and generalization is based on measurements of drift rather than pose accuracy.

2.5.1 Cost function and keyframes

When **[Paper I-II]** were designed, previous sparse methods for 2D radar odometry relied on the minimization of point-to-point distances and did not attempt to utilize estimates of surface normal or distribution for registration. For that reason, the work took inspiration from lidar localization research, where minimization registration in the direction of the estimated surface normal [44, 88], or distribution [27, 74] has been a successful technique for localization and mapping [24, 89, 94, 110, 115]. For that reason, a point-to-line cost function for minimization was adopted in **[Paper I]** when registering the latest scan to a single previous keyframe scan. With a single keyframe scan, the point-to-normal distance minimization was found to be significantly more accurate and less sensitive to sparsity compared to the point-to-point distance. Later in **[Paper III]**, the ablation study included a characterization of the behavior of the cost function (point-to-point, point-to-line, point-to-distribution) then gradually increasing the number of keyframe scans, starting from a single. Surprisingly, as the number of keyframe scans was increased, the performance ranking between these cost functions was reversed, and point-to-point provided lower drift. The drift decreased strictly as additional keyframe scans were introduced, even up to 50 scans when a mean translation drift of 1.09% was reached. This interaction between multiple keyframes and cost function was a surprising discovery.

After the publication of **[Paper II]**, more recent work has moved toward omitting the point-to-point cost functon in favor of point-to-normal [101] or point-to-distribution [111], and toward using cost functions with multiple keyframe scans [92] as presented in **[Paper II-III]**.

2.5.2 Evaluation datasets

The proposed radar odometry method–as well as other methods for fault detection and SLAM presented in later sections–was evaluated on the datasets presented in Fig. 2.4. The datasets are: Oxford Radar RobotCar Dataset [22], MulRan [65], and the sequences VolvoCE, Kvarntorp and Orkla from the diverse ORU dataset [9]. These datasets include a spinning radar used throughout the work of this thesis. Fig. 2.4.(a-c) are part of the *Diverse ORU dataset* published together with **[Paper III]** [9]. In that dataset, three sequences from substantially different environments were recorded, corresponding to different real-world industrial scenarios. The first sequence (Fig. 2.4.a) was recorded on a test track for work machines and autonomous vehicles within a forest environment, driving a wheel loader, with a radar mounted on the roof, over a distance of 1.6 km. The second sequence was recorded in an underground mine, driving a car with a roof-mounted radar over 1.2 km through partly feature-poor environments. The third sequence was recorded in an indoor warehouse, navigating a forklift with radar over a distance of 200 m through aisles with shelves.

The Oxford and MulRan datasets are publicly available radar datasets with ground truth positions. Oxford contains repetitions of a single route within strictly urban settings. The MulRan dataset includes variations of sequences in mixed environment types, including less structured environments with open fields and bridge crossings.

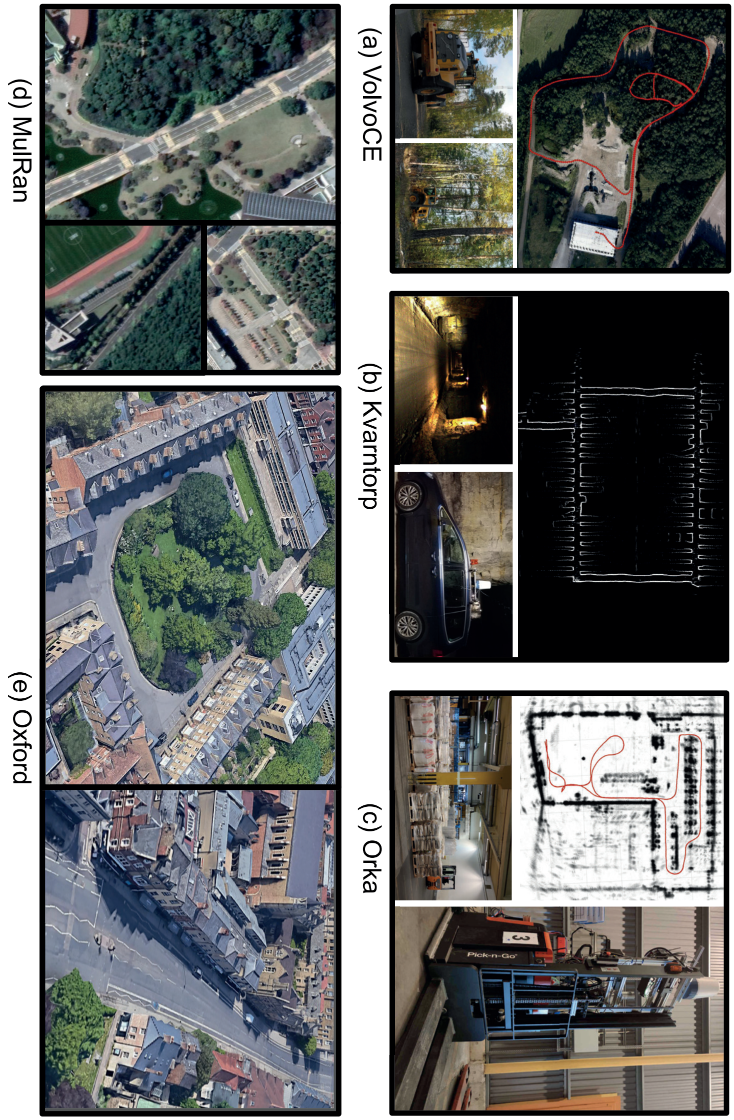

Figure 2.4: Radar datasets used for evaluation in this thesis. (a-c): Diverse ORU datasets [9], including sequences VolvoCE, Kvarntorp, and Orkla. (d): MulRan [65] with mixed structured/semi-structured sequences. (e) Oxord radar RobotCar dataset [22] includes repetitions of a route in a structured urban environment.

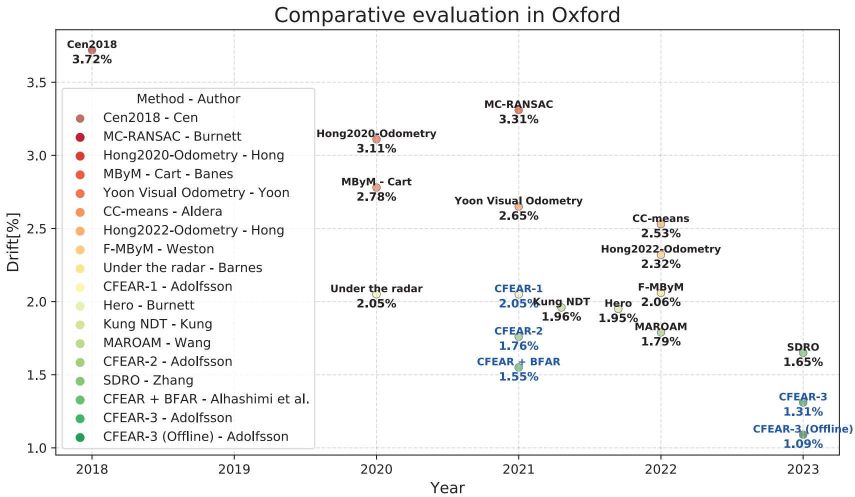

Figure 2.5: Comparative evaluation in the Oxford radar odometry benchmark [22]. Methods proposed in this thesis are marked in blue with performance as measured by the time of publication. Note that improved versions of CFEAR-1 and CFEAR-2 with lower drift and higher efficiency were presented in [9]. The most accurate configuration of CFEAR (CFEAR-3-S50 Offline) achieves the lowest drift (1.09%) among comparable methods for radar odometry.

2.5.3 Configurations and comparative evaluation

[Paper III] presents multiple configurations, optimized for efficiency, performance, or a balance between these properties. The most efficient configuration in terms of execution time (CFEAR-1) is similar to the initial workshop paper (**[Paper I]**). The balanced configuration (CFEAR-2) is similar to the conference paper (**[Paper II]**). The configuration tuned for performance (CFEAR-3) is presented in **[Paper III]**. In the Oxford dataset, these methods run between 44 and 160 Hz, with a drift between 1.31 and 1.71%. Finally, the configuration CFEAR-3-s50 achieves the lowest drift but is more suitable for offline processing because of the slower update rate. The performance of these configurations is compared to other methods in Fig. 2.5. The figure compares results in the public Oxford radar odometry benchmark since the release of the dataset.

2.5.4 Robustness to feature-poor environments

Improving robustness within feature-poor environments was addressed in [**Paper III**]. In the Kvarntorp scenario in Fig. 2.4.b, the radar is located within long, narrow, and feature-poor passages located in an underground mine. The article presents multiple techniques that enable effortless navigation within that scenario, including the point-to-line cost function, the weighting of surface correspondences according to matching quality, and the use of multiple keyframes. These techniques alleviate the issue that the cost function is ill-conditioned in the longitudinal direction where correspondence quality degrades.

An alternative method to enable navigation within feature-poor scenes is presented in [**Paper XII**]. An external sensor, e.g. wheel odometry measurements, provides complementary pose (with uncertainty) measurements in the longitudinal direction independently of the scene. Integrating such measurements into the cost function—as extra penalizing terms—resolves the ill-conditioned cost function, and enables navigation even in situations without features. The method seamlessly shifts the weighting to ego-motion constraints according to the decrease of features in the scene.

2.5.5 Generalization to new unseen environments

A key requirement of odometry is the property of maintaining low odometry drift, even during the exploration of previously unseen environments. Prior to [**Paper II**], the best-performing methods relied on learning to extract features [21, 32] or remove noise [23]. Additionally, previous experiments suggested that the performance of the best-performing method was obtained by overfitting to the geographical locations over which it was trained [23, 104]. However, it is desired that drift does not increase substantially during navigation within new unseen environments. For that reason, [**Paper III**] includes a qualitative study and an extensive Spatial Cross-Validation (SCV) generalization study. The SCV study aims to determine how optimal parameters for one environment affect performance within another environment. The MulRan and Oxford data sets, visualized in Fig. 2.4.(d-e), were used in the experiments which included a total of *98 million* pose estimates. The conclusion of the SCV study is that the optimal parameter set in one dataset is near-optimal in the other dataset. In other words, the benefit of tuning the parameters for the new target environment, rather than keeping the default parameters, is negligible.

The qualitative generalization study demonstrated that the parameters tuned in the Oxford environments could be used directly for odometry estimation in the sequences from the Diverse ORU dataset. Specifically, the odometry quality was high in the sequences VolvoCE (outdoor forest vehicle test track), Kvarntorp (underground mine, partly feature-poor), and Orkla (indoor logistics scenario). The consistently high performance, achieved in highly diverse environments and over an extended period of time, indicated a high level of robustness.

However, errors may eventually occur, from which safe recovery is critical to achieving resilient navigation. The next chapter focuses on the detection of errors and the integration of failure detection together with CFEAR for robust radar SLAM.

Chapter 3
Quality-aware localization and mapping

This chapter presents a literature overview and novel methods for localization fault detection and radar SLAM. Sec. 3.2 summarizes Papers V-VI, which propose a novel method for point cloud-based fault detection in lidar and radar. Sec. 3.4 discusses the integration of fault detection in radar SLAM for robust verification of loop closure constraints.

3.1 Related work - localization fault detection

Ascertaining robust localization by quantification of localization quality and by detecting localization errors has been studied from multiple perspectives. Some methods use multiple (redundant) positioning systems [78, 98] to detect failures or additional sensors to learn the detection of errors [16]. The aim of this thesis is the detection of localization errors, without the need for additional sensors. Other methods attempt to predict the risk of localization failure by analyzing scene geometries [80], or, in the event in which the risk is high, plan alternative navigation routes [39] or disable loop closure [60]. In contrast, the aim of this thesis is the detection of errors after the registration has been carried out.

One formulation aims to detect localization errors when navigating using previous maps of the environment [12–14]. The method of Akai et al. [13] detects localization errors from measurement residuals between the lidar scan and the prior 2D occupancy maps of the environment. The work was later extended to 3D using the distance field representation [14]. In their work, range measurement residuals are not treated independently, and the relationship between measurements is jointly considered using Markov random fields.

A family of methods computes the alignment quality between registered images [23] or point clouds [25, 29, 71, 82], with the quality represented as a normal distribution (covariance matrix). This representation has been used mainly for data fusion [61] and SLAM [57], and, to a limited extent, alignment classification [19]. In contrast, the work carried out in this thesis aims to represent the alignment quality by a single scalar for the purpose of detecting localization errors. Specifically, the *localization quality* is formulated as an *alignment quality* assessment (and classification) problem that operates on point cloud pairs. Alignment quality assessment between point cloud pairs is a broad formulation, hence enabling applicability to verify different stages in a localization system, for example, pose tracking, global localization estimates, topometric pose estimation, or loop closure.

A range of methods aims at assessing the quality between pair of point clouds. The most well-known measure is the Iterative Closest Point (ICP) fitness score – the sum (or averaged) squared Euclidean point-to-point distances [26, 88]. The measure is commonly used to verify potential loop closures in SLAM [94]. However, the ICP measure is sensitive to the environment, and previous studies have shown that the detection of small errors is challenging with the measure [19, 96]. For this reason, alternative quality measures have been proposed. Makadia et al. [75] proposed a measure of the consistency of normals. Almqvist et al [19] proposed measures based on a point-to-distribution between overlapping NDT cells. Yin. et al. [107] combined point cloud overlap, differences between 2D point cloud projections, and normal distribution parameters for point-to-point and point-to-plane distance. Similarly to Yin et al. the method presented in this thesis combines multiple complementary measures within a logistic regression model. Combining multiple measurements improves generalization as presented in [**Paper VI**], and reduces sensitivity to different error magnitudes as discussed in [**Paper VII**].

Until today, Almqvist et al. presented the most rigorous characterization of alignment classification techniques using 3D lidar data [19]. The article includes a study that investigates how methods robustly generalize to unseen environments. A similar characterization and generalization study is presented in [**Paper VI**] for spinning radar.

To this day, the detection of small errors remains challenging. Particularly, with generalization across diverse environments [19], and with tolerance to changing overlap conditions. These challenges motivate the novel work presented in the next section.

3.2 CorAl Introspection

CorAl introspection belongs to the class of alignment quality classification. CorAl was first proposed for lidar data in [**Paper V**] and later extended to operate on radar data in [**Paper VI**]. For use with spinning radar data, however, an intermediate point cloud extraction step is required. This step extracts peaks – computed from the initial filtering step presented in the previous section – that are more suitable for alignment classification.

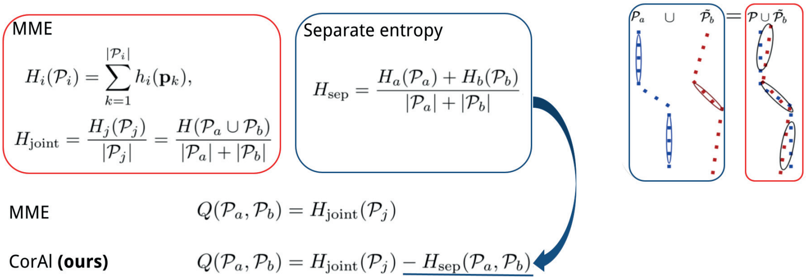

Figure 3.1: Principle of CorAl quality measure

The principle of the *CorAl quality measure* is demonstrated in Fig. 3.1. The method takes inspiration from the Mean Map Entropy (MME) quality measure [49, 84], which has been used to evaluate point cloud refinement. MME measures averaged *per-point differential entropy* – a measure of spread, or uncertainty, from points sampled within a radius. The metric, which can be seen as a quantification of the overall crispness in a point cloud, varies significantly across environments. CorAl additionally estimates the *separate entropy* The separate entropy can be seen as the crispness owed to the scene itself, before merging the point cloud pair. As depicted in Fig. 3.1, the separate entropy is subtracted from the MME. Thus, obtaining a new environment-insensitive quality measurement, which is instead sensitive to point cloud alignment quality. For classification, the subtraction of these two measures is instead replaced with a linear combination where coefficients are learned inside a linear regression classifier. This reduces the already high bias of the classification model.

3.2.1 Training phase

Training data for CorAl are generated by computing alignment quality measures before and after applying a rigid alignment offset between scan pairs. This process is repeated for a set of scan pairs and allows for generating training data from ground-truth pose estimates or odometry. The training process is visualized in Fig. 3.2 [1].

3.2.2 Generalization of detection

Part of the results of the generalization studies in **[Paper V-VI]** is presented in Fig. 3.3. The study learns the model parameters for classification in one environment and evaluates them in the other and vice versa. An interesting finding in the evaluation was the consistently high rate of CorAl classification, even

[1] `https://tinyurl.com/CorAlTraining`

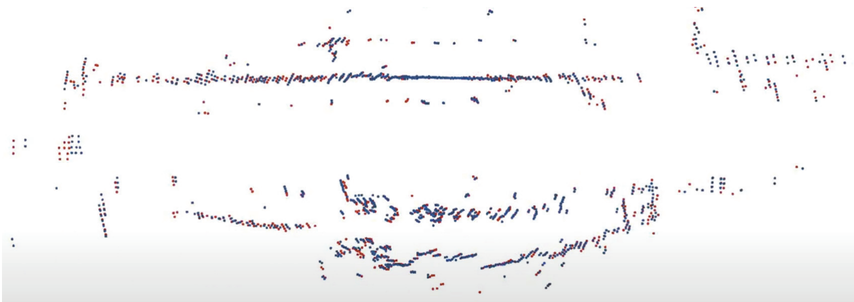

(a)

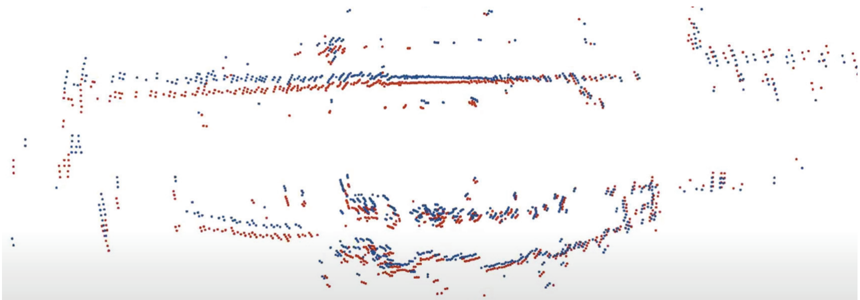

(b)

Figure 3.2: Training phase of CorAl which uses scan pairs with corresponding ground truth or estimated poses. The CorAl measures (and training labels) are obtained for both aligned scan pairs (a) and intentionally misaligned scan pairs (b). This process is repeated for misalignments with varied translations and rotations – to reflect that localization errors may have varied characteristics.

when operating in new environments outside of the training set. The generalization characteristics were similar for both lidar and radar throughout this study. Specifically, the level of generalization was higher when training in less structured areas and testing in more structured environments. On the contrary, generalization was lower when training in more structured environments and testing in less structured environments. By training on data from both structured and unstructured environments, it was possible to find a model that is jointly accurate in structured, semi-structured, and unstructured environments. Thus, a higher level of generalization can be ensured with less structured and more diverse data in the training set.

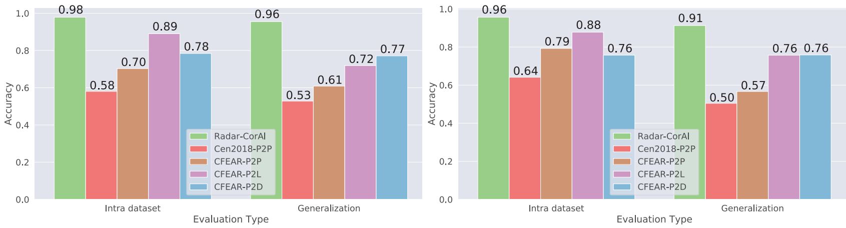

(a) Intra dataset:
Trained and evaluated on Oxford.
Generalization:
Trained on MulRan, evaluated on Oxford.

(b) Intra dataset:
trained and evaluated on MulRan.
Generalization:
trained on Oxford, evaluated on MulRan.

Figure 3.3: Classification of small localization errors (0.5 m) and generalization to new unseen environments.

3.2.3 Different error characteristics

The study in **[Paper VI]** shows that alignment classification remains challenging. In line with previous research, the study confirms that detection is challenging for small errors, and with lower overlap. Additionally, quantitative and qualitative experiments showed that it is difficult to determine which level of error is acceptable. For example, vibrations from uneven ground or wheels can introduce non-rigid distortions in the scans, which are reflected in the CorAl score and, accordingly, challenge alignment classification – even with the best possible rigid alignment between the scans. Rejecting distorted scans is clearly the desired behavior for visual reconstruction of environments (i.e. SLAM) where the resulting map should be cleanly created from non-distorted scans. However, for robot navigation, it can be sufficient to determine if the current pose is correct with a higher level of acceptance for distortions.

A weakness of the CorAl measure from **[Paper V-VI]** is the choice of radius required to compute per-point differential entropy. It is generally hard to find a radius that works well for very different error magnitudes, which is essential, as the error characteristics are not known in advance. When adapting CorAl to loop closure classification as presented in **[Paper VII]**, the detection of larger errors was important to reject false loop candidates. For that reason, the logistic regression model of CorAl was extended to input additional quality measures. These are computed from the CFEAR Point-to-Line measure, which is more suitable for quantifying larger alignment errors. This results in a more versatile classifier, suitable for the verification of loop constraints in SLAM.

Sec. 3.4 discusses the integration of CorAl for the verification of loop closure constraints in SLAM. Going even further, **[Paper VII]** shows how fine-grained alignment verification of CorAl can be unified with place recognition for improved place retrievals. Therefore, it allows for safer and more accurate localization in prior maps and SLAM as presented in the following sections.

3.3 Related work – radar SLAM

Radar-based localization and mapping has been a research topic for many years. Already in 1998, a beacon-based radar localization system was presented by Clark and Dlurrant-Whyte [45]. A few years later, Dissanayake et al. [48] presented the first radar-based method for SLAM, with landmark included within the state of an Extended Kalman Filter (EKF). Later methods proposed alternative techniques to estimate motion, more commonly referred to as registration. Checchin [43] employed the Fourier-Mellin Transform for dense registering of radar scans, while Chandran and Newman [42] formulated motion estimation as a map quality optimization problem. Other methods performed the matching of occupancy grid maps [87], or SIFT features [36]. The previously mentioned methods rely on additional sensors, such as a wheel encoder and gyroscope. In contrast, Callmer et al. [36] demonstrated the feasibility of SLAM with radar as a single sensor.

In 2013, Mark et al. [76] adapted the Rao Blackwellized particle filters by Grisetti et al. [54] for use with radar. After that time, most of the presented approaches [58, 60, 101] relied on pose graph optimization – separate from the task of odometry estimation and mapping – to correct the estimated trajectory. An exception, however, is the work of Schuster et al. [93] which integrates landmark measurements as additional constraints within pose graph optimization.

Recently, Holder et al. [58] presented a method based on automotive Doppler radar. The method estimated odometry using ICP with external sensors providing a starting point for ICP. Loops were computed via GLARE place descriptor matching and registering of local sub-maps. Hong et al. [59, 60] presented a highly capable method for radar-only SLAM. Loop constraints are obtained by computing and matching M2DP [56] descriptors from a point cloud, extracted from the highest intensity levels per azimuth. Wang et al. [101] adapted Lidar Odometry And Mapping (LOAM) [110] for radar, with Scan Context [64, 66] for loop detection.

The work presented within this thesis relies solely on Doppler-free spinning radar and follows the paradigm of clear separation between odometry, loop detection, and trajectory correction. In contrast to the methods listed in this section, the presented work focuses on achieving robustness in loop closure by careful verification of computed loop constraints. A literature survey on loop verification in visual and lidar SLAM is presented in **[Paper VII]**.

3.4 TBV Radar SLAM

TBV Radar SLAM integrates the work presented in the previous chapters of the thesis, including the odometry method *CFEAR* and misalignment detector *CorAl*, and reuses the previously presented radar data representations. The addition is the inclusion of a pose graph optimization back-end, robust loop detection, and verification model, with Scan Context as a framework for descriptor matching. On the basis of such integration, TBV presents a series of steps to substantially increase the rate of verified correct loop retrievals. An overview of the complete TBV Radar SLAM system is presented in Fig. 3.4, and visualized externally [2].

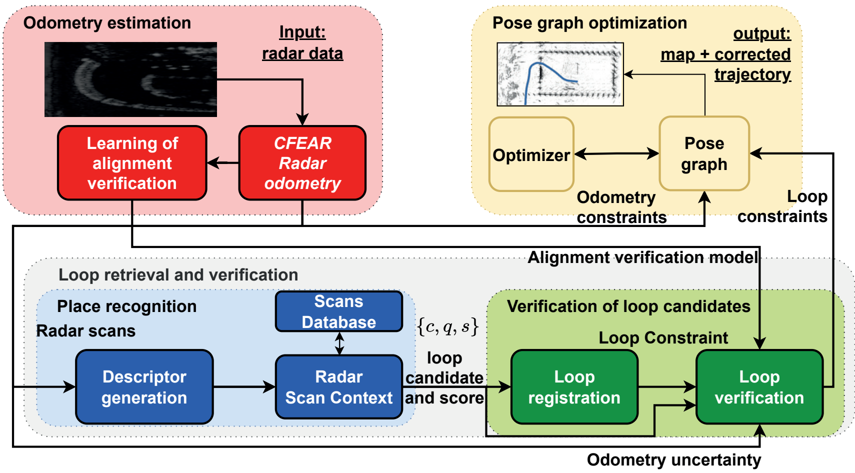

Figure 3.4: Overview of TBV Radar SLAM.

In summary, CFEAR is used to estimate odometry and compute scene representation for loop retrieval and verification. Once in an offline training phase, a model for alignment verification is automatically learned from odometry and scans within the CorAl framework. After initial offline training, alignment quality can be computed and utilized for loop closure classification. Alignment verification is combined with odometry uncertainty and place similarity score, resulting in an improved quality measure for more robust loop classification. Loop candidates are found by matching of place descriptors, computed from aggregated sparse sets of radar peak detections. For the best matching candidates, both loop constraints and quality measures are calculated. From these constraints, only one can be selected according to the highest loop quality.

[2] `https://tinyurl.com/TBVRadarSLAM`

3.4.1 Evaluation

TBV Radar SLAM was evaluated on localization and mapping performance, and on loop closure computation and verification. An overview of the loop closure ablation study is presented in Fig. 3.5. The figure presents the baseline method *radar Scan Context* [65, 66] (plot 1 in the figure), and proposed improvements (plot 2-6 and 8). In the figure, each point, i.e., pair of precision and recall, corresponds to a selection of a loop quality threshold $y_{\mathtt{th}} \in [0, 1]$. Note that both the predicted loop label and the computed loop constraint have to be correct (within a tolerance) for a loop to be considered a true positive. Retaining a high level of precision is critical for loop closure. False positives, which lower precision, may corrupt the resulting map and can be hazardous to navigation. Fortunately, the different versions presented in Fig. 3.5 can be tuned for close to 100% precision – with a varied level of recall – to ensure the rejection of incorrect loop constraints.

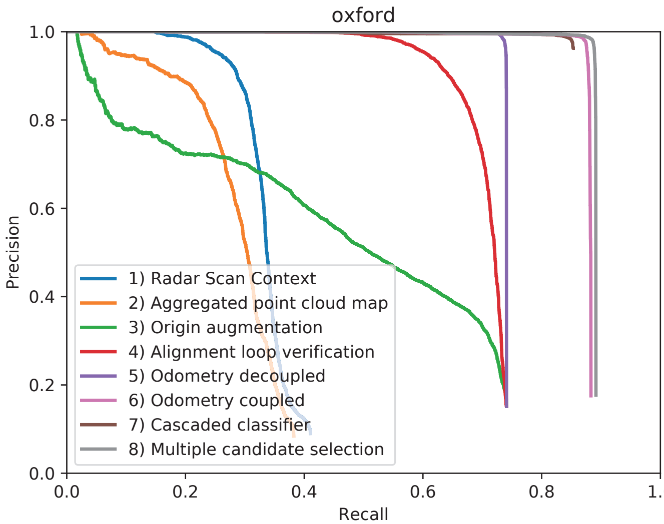

Figure 3.5: Ablation study of the introspective loop closure of TBV Radar SLAM. (Plot 1) is a plain integration of CFEAR and Radar Scan Context [65] without additional verification. (Plot 8) utilizes all techniques presented in **[Paper VII]** and can be tuned for both high precision and recall. Most of the improvement can be attributed to the verification which combines place similarity, odometry uncertainty and point cloud alignment.

3.4.2 Unified verification and retrieval

TBV Radar SLAM was initially designed to first retrieve and verify place retrievals, followed by loop registration and separate geometric analysis and verification using CorAl. However, loop retrieval could fail due to erroneously low measures of place similarity for correct candidates. Unfortunately, this would result in candidates being discarded before final alignment verification. For that reason, the verification of the place retrieval and the final geometric verification were integrated and carried out jointly. The advantage of this strategy is that place retrieval verification is informed and complemented by alignment verification carried out after registration.The former technique ("separate" or "cascaded" verification) is evaluated in Fig. 3.5 (plot 7), while the latter technique (joint verification) is presented in Fig. 3.5 (plot 6).

Rather than verifying a single top loop retrieval, this step was generalized to register, verify and select among multiple competing candidates. Among the candidates, only the candidate with the highest quality (Fig. 3.5 (plot 8)) could be selected. These improvements – combined with additional techniques for retrieval and verification – result in robust loop classification with more than 90% recall close to 100% precision. This makes the proposed method suitable for robust large-scale mapping, which requires only short revisited segments to effortlessly detect and close loops, even without the need for a robust back-end pose graph optimizer.

3.4.3 Filtering for place recognition

The radar intensity peaks (extracted as detailed in [**Paper VI**]) were used to calculate the place descriptors in (Fig. 3.5 (plot 2-8)). Specifically, neighboring scans are aggregated into a local map before computing the place descriptors. This representation is more sparse compared to Hong [60] – which extracts intensities one standard deviation above the intensity mean – and compared to the dense method [65, 101] where scans are merely downsampled. The strategy "raw unfiltered scan" (Fig. 3.5 (plot 1)) is more accurate compared to using intensity peaks (Fig. 3.5 (plot 2)), at least without further improvements (Fig. 3.5 (plot 3-8)). This raises the question of how much data should be filtered. It is likely that the descriptors computed from detailed local maps could improve loop retrieval, benefiting from the rich information inherent in current spinning radar.

3.4.4 SLAM performance

This section presents the results of the complete SLAM system, i.e. integration of the loop closure module and CFEAR radar odometry, into TBV Radar SLAM. The parameters of TBV Radar SLAM were tuned for a single sequence in the Oxford dataset [22], and evaluated on the Oxford, MulRan [65] and the

ORU Diverse radar dataset, which was released in **[Paper VI]** [9]. TBV Radar SLAM achieved an average Absolute Trajectory Error (ATE) of 3.90 m (over 10 km) in the Oxford dataset, and 2.9 m (over 5.9 km) in the MulRan dataset. With this level of accuracy, TBV Radar SLAM improves the previous state of the art [60] with 65% reduced trajectory error. Furthermore, TBV consistently performs well within substantially different environments without the need to change parameters. Examples of estimated trajectories are visualized from the Oxford dataset in Fig. 3.6, from MulRan in Fig. 3.7a [3], and from ORU Diverse radar dataset in Fig. 3.7c [4].

[3] `https://tinyurl.com/KAIST02TBV`
[4] `https://tinyurl.com/VOLVOTBV`

(a) Before loop closure

(b) After loop closure

Figure 3.6: Radar SLAM in the Oxford dataset before and after loop closure. Ground truth (red) and estimated trajectory (blue).

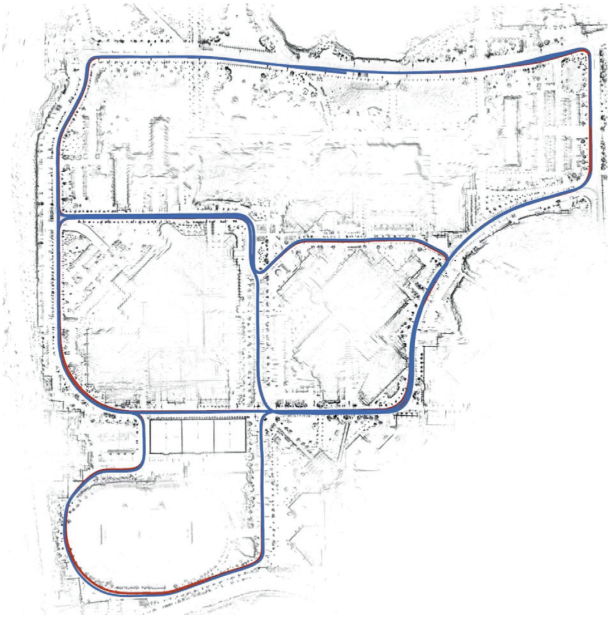

(a) MulRan - KAIST02: Ground truth (red) and estimated trajectory after loop closure (blue).

(b)

(c) ORU Diverse radar dataset-VolvoCE: Before (green) and after (blue) loop closure

(d)

Figure 3.7: Radar SLAM in the MulRan and ORU Diverse radar dataset.

3.4.5 Evaluation challenges

The Oxford and MulRan dataset sequences start and end in close proximity, with long overlapping trajectory segments. For that reason, a few loop closures at the end of the trajectory are sufficient to correct the drift accumulated over the trajectory. As a consequence, trajectory quality in these datasets is largely dependent on odometry performance rather than the loop closure potential. For that reason, the benefit of robust loop closure is not fully reflected in the evaluated trajectories. Therefore, to fully demonstrate the benefit of the proposed loop closure module, additional experiments are required in more challenging datasets.

Furthermore, during the evaluation, inconsistencies were found between the aligned radar data and the ground-truth positions provided in the Oxford dataset. Specifically, after computing and visually inspecting the correctness of the loop constraints, the estimated transformation deviated up to ≈ 3 meters from the ground-truth position. This error was found in the lateral direction of movement and therefore cannot be explained by the speed of the vehicle. Note that such a level of ground truth inaccuracy is similar to the level of accuracy of TBV Radar SLAM, so further improvement to SLAM may be hard to evaluate due to the ground truth not being accurate enough. As an additional consequence, the performance presented in the evaluation of TBV Radar SLAM may be pessimistic.

3.4.6 Learned place recognition

Radar Scan Context, which was adopted for place matching, uses hand-crafted descriptors for place matching, thus limited by the expertise of humans to design the system. An alternative method for place recognition, namely "NDT Transformers", is presented in [**Paper IX**]. In contrast, NDT Transformers learns the computing of global descriptors for place matching by the use of a deep neural network. The contributions of NDT Transformers are in the domain of 3D point cloud place recognition, and the method is not directly applicable to dense 2D radar without modifications.

The state-of-the-art radar place recognition [52] operates on Cartesian radar images – which are expensive to compute at high resolution – and learns the whole process of extracting features from raw data. However, some early steps of the problem, for example, filtering and motion compensation (and computing of the NDT representation), can be addressed without learning. Particularly given that these tasks are already addressed by the concurrently running odometry estimator. Therefore, delegating these tasks could potentially simplify the learning of the subsequent place descriptor. For that reason, NDT transformers could potentially be suitable for efficient and accurate radar place recognition.

Chapter 4
Rapid global localization in prior maps

This chapter presents methods for rapid localization in prior maps with high robustness and accuracy. The chapter starts with an introduction to topometric localization and presents a method for improving map and localization quality. A brief summary of related work on global localization is presented in Sec. 4.2. Sec. 4.3 presents the novel localization pipeline: *Localizing faster* for rapid global localization in prior maps. Localizing faster, from [**Paper VIII**], leverages predictions of localization quality – conceptually incorporating *self-aware* localization – which improves both localization speed and robustness.

4.1 Submap representation for localization

In the TBV Radar SLAM pipeline presented in the previous section, single scans are used as the atomic map representation when reconstructing the environment in pose graph SLAM. Once SLAM is finished, alternative atomic map representations, rather than single scans, may be more suitable for localization. Scans could be fused into smaller submaps or, more commonly, into a single monolithic map, before being utilized in the localization system. [**Paper XI**] discusses the choice of atomic map representation for localization in prior maps. Specifically, in a topometric localization framework, how should maps be represented? From numerous small submaps, each consisting of a few scans, to a single monolithic map that integrates all scans.

The article presents a method that groups scans of an observed scene from a similar perspective to simplify the map complexity within each group. Simplified map complexity improves localization quality in submaps created from these groups. The method operates as an intermediate partitioning step that runs between the mapping and the deployed localization. This method has a limited contribution to this thesis, and serves as a motivation for the map modality adopted in the work on global localization in Sec. 4.3.

4.1.1 Submaps for topometric localizaiton

Submaps have been widely used in localization research. Early methods adopted submaps to manage computational time and memory requirements and to enable real-time localization for filter-based approaches [28, 30]. Since then, submaps have been used extensively to improve localization quality for tasks including, but not limited to, lidar [110], radar [101, 111], and visual [37] odometry estimation. As such, fusing measurements into local maps is widely considered a sound strategy for building accurate scene representations, providing means to reduce sparsity, remove redundancies, or inconsistencies, and create more complete maps.

Localization in prior maps has made use of single monolithic maps [90] representations, submap representations [50], or similarly, topometric representations with vertices of submaps [33] or single scans. Topometric representations are typically motivated by their simplicity, enabling localization without the need of globally consistent maps [77]. The work presented in the next section argues for the usage of such topometric submap representations – even for consistent maps. The topometric submap representations have the aforementioned advantages of local maps, but with a lower map complexity compared to monolithic maps – modeling the scene from a single perspective only. Thus, the work in the next section proposes grouping globally consistent scans into local submaps for localization in these.

4.1.2 Topometric localization

After grouping of scans into submaps, the resulting submaps are sparser in location compared to the original pose graph but updated from multiple scans for better coverage and thus more suitable for localization. Following this step, the robot can perform metric localization (pose tracking) within the submaps. The robot pose is estimated with respect to the submap that was previously updated close to the robot's current location, thus ensuring that the most accurate, among nearby submaps, is being used to localize. [**Paper XI**] presents different strategies to partition scans into local maps, using online or offline techniques, and based on scan locations, point cloud geometry, or a combination of these two.

4.1.3 Findings on topometric submap localization

The main finding in [**Paper XI**] is that localization in small submaps is preferred to using monolithic, which integrates all measurements in a single map

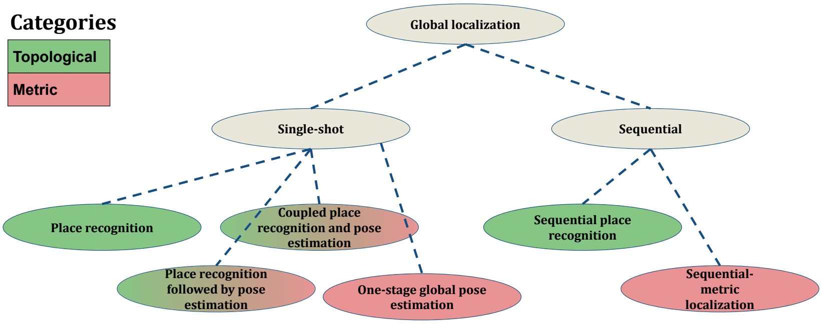

Figure 4.1: Categorizes of global localization in prior maps.

representation. Localization in submaps can reduce localization error by up to 46% compared to monolithic maps, and reach an accuracy of as low as 1.7 cm within a realistic working scenario of a robot navigating within a warehouse.

The preference for smaller submaps may seem contradictory to the findings of [**Paper III**], where the use of numerous scans is preferred for the estimation of odometry. There are, however, some key differences. First, scan-to-submap registration is adopted for topometric localization in [**Paper XI**], whereas the method for radar odometry in [**Paper III**] never fuses scans into submaps, but instead maintains scans separately and register the latest scan to multiple jointly. Second, the use of more scans in [**Paper III**] mainly reduces odometry drift, while the impact on pose accuracy is already minor after some have been used. This submap representation was adopted as part of the global localization system presented in the next section.

4.2 Related work - Global localization

Global localization from range data can be divided into *single-shot* and *sequential* approaches as depicted in Fig. 4.1. Single-shot approaches localize using a single measurement set, whereas sequential approaches operate on a continuous stream of measurements. Single-shot approaches can be further subdivided into four categories: place recognition only, place recognition followed by pose estimation, coupled place recognition and pose estimation, and one-stage global pose estimation [109].

Place recognition only localizes topologically by computing and matching place descriptors. NDT transformers [114] (presented in [**Paper IX**]) belong to this category. Other popular methods in this category operate on lidar point clouds [67, 99], radar images [52, 65, 68, 91], or both [108].

Place recognition followed by pose estimation localizes first topologically, followed by a metric refinement of the pose, e.g., using registration.

This approach is often adopted in topometric localization systems [46] (that integrate global localization) and in SLAM for loop closure [58, 60].

Coupled place recognition and pose estimation, address these steps in an integrated fashion. An advantage is that pose estimation together with a final geometric verification can be leveraged to filter incorrect place retrievals [109]. This idea motivates the unified radar loop retrieval and verification module in TBV Radar SLAM ([**Paper VII**]). Scan context [64, 66] (which was used within [**Paper VII**]) can be placed into this category as place matching retrieve and utilize the relative heading between descriptors.

One-stage global pose estimation estimates the current pose in a single step, contrary to the previous metric localization techniques. Deep pose regression is commonly adopted for one-shot global pose estimation [62, 63, 83, 100, 102, 103]. This approach to localization has a key advantage: The efficiency of the method depends only on the complexity of the neural network as it circumvents the need for descriptor matching, which is tedious in large-scale environments. Instead, the map is implicitly stored within the weights of the network. Motivated by its efficiency, this technique was included as part of [**Paper VIII**].

Sequential localization operates on a stream of measurements, such as a series of scans produced by a spinning range sensor. These measurements can be complemented by additional information sources, e.g., wheel odometry. Traditional approaches to sequential (metric) localization are MCL [47], and more recent versions for accurate 3D [90] and introspective localization [12].

Although MCL is robust and accurate, it becomes intractable for large-scale 6-DOF localization. Interestingly, MCL and one-shot deep pose regression provide complementary properties for global localization. Hence, it can be combined for robust and accurate, yet efficient, localization. The rest of this section discusses such an implementation proposed in [**Paper VIII**]. The method makes use of lidar scans and odometry for 6DOF localization without the need for additional sensors, such as IMU.

4.3 Localizing Faster: Efficient and precise lidar-based robot localization in large-scale environments

The method for global localization, proposed in [**Paper VIII**] achieves fast localization by combining one-shot deep pose regression with uncertainty, with MCL localization in local NDT maps. These two systems are referred to as System 1 and System 2, and provide complementary properties.

4.3.1 System 1 - Deep pose regression with uncertainty

System 1 is a version of one-shot global localization. Specifically, 6-DOF deep-pose regression with uncertainty. The model has been trained on a set of corrected poses and observations (range data) to operate within a specific environment.

First, 3D point clouds are aggregated via odometry and superimposed around the current pose. After removal of the ground plane, the point cloud is converted into a birds-view grayscale image, with pixel intensity according to point height over the ground plane. An example of a scene represented by a lidar point cloud with corresponding grayscale image representation is visualized in Fig. 4.2. Pose regression is then learned from the grayscale images, supervised by the corrected poses in the world frame. In this system, the map of the environment is *implicit* in the weights of the neural network.

4.3.2 System 2 - Submap NDT-MCL

System 2 is based on NDT-MCL [90], with the *explicit* topometric maps presented in the previous section. NDT-MCL is a particle filter for localization based on the NDT voxel representation. In summary, localization is carried out by applying a motion model based on odometry, a measurement model between scan and map, normalizing the particle weights, and resampling pose particles according to particle likelihood.

Rather than localizing in monolithic maps [90], nearby submaps are used in a scan-to-map measurement model. The submap selection criterion in **[Paper XI]** was adopted. Each particle selects the submap that was updated closest to the robot's current position. Note that the selection of the submap can be precomputed for each possible location. Hence, introducing submaps does not alter the online computational complexity for particle-based localization.

4.3.3 Hybrid one-shot / Monte Carlo localization

System 1 and System 2 are seamlessly integrated within the particle filter. Specifically, a set of N_{sys1} particles is sampled from each predicted pose distribution and merged with the MCL particles (before applying the motion and measurement model). In the resampling step, the total particles are reduced from $N_{sys1} + N_{sys2}$ back to N_{sys1}. These steps are repeated when new observations are available, without the need to switch modes between global localization and pose tracking.

Note that system 2 provides (i) prediction of a global localization pose estimate, and (ii) prediction of the quality of that pose. System 2 is integrated with system 1 by sampling pose particles according to the predicted pose distribution. This is motivated by that the accuracy of global pose estimates is correlated with the predicted quality as indicated by the pose covariance. Therefore, the predicted quality can be used to adjust the size of the search space. When the pose estimate is accurate and the correctly predicted quality is high (low uncertainty), the pose particles are confidently sampled from a smaller area. On the contrary, less accurate pose estimates with correctly predicted low quality adjust the sampling to a larger area which reflects the larger search space. If the quality is predicted as inadequate, that is, if the variances in the quality covariance matrix exceed a limit, the pose estimate is omitted. This reflects that the robot may be navigating outside of the boundaries of the given map.

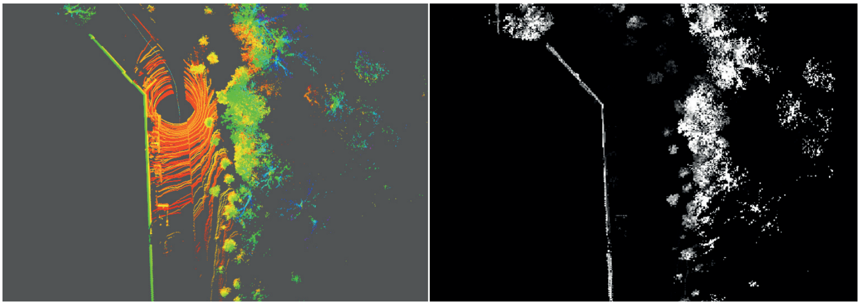

Figure 4.2: Bird's-eye view leveraged for global localization. Left: 3D Point cloud colored by height, right: corresponding gray-scale image with height over the ground plane encoded in gray-scale.

4.3.4 Evaluation

Results on localization performance are presented in Tab. 4.1. In the experimental setup, the robot is tasked to repeatedly localize within a 250×250 m large outdoor campus environment from the NLCT dataset [38]. The evaluation compares localization with predicted distribution (using self-awareness of localization quality), fixed distribution (no self-awareness), and an NDT-MCL baseline with uniform particle initialization in all possible locations.

The results show that self-awareness of localization quality improves robustness, compared to using a fixed distribution, with a localization success rate of 93.3% (up from 92.5%). Localization time is reduced by 16%, reaching a median localization time of 0.8 s (down from 0.95 s). Compared to the baseline MCL with uniform sampling, the localization time is reduced by 99.5%, and the failure rate is reduced by 86%. The improvement compared to baseline can be attributed to the significantly reduced search space which enables localization with fewer particles more densely distributed around the correct pose. Leveraging the localization quality measure adapts the sampling according to the pose estimate quality which further increases the likelihood that the particles are sampled in the vicinity of the actual robot location, and accordingly, the likelihood of convergence.

	Success rate (%)									Localization time [s]	
Method	Feb	April	May	June	Aug	Oct	Nov	Dec	Overall	Mean	Median
Hybrid - GP Cov (ours)	97.2	**100**	**94.4**	**94.9**	**95.8**	**94.7**	**81.2**	**88.7**	**93.3**	**1.94 ± 3.0**	**0.80**
Hybrid - Fixed Cov (ours)	**97.7**	99.0	93.0	93.5	94.8	94.0	79.7	88.3	92.5	2.32 ± 3.3	0.95
NDT-MCL uniform init.	62.0	70.6	57.3	59.6	52.7	51.7	37.6	40.1	54.0	154.29 ± 46.2	157.93

Table 4.1: Success rate of hybrid one-shot / MCL localization (presented in this thesis). The evaluation compares the proposed method with predicted localization quality (GP Cov) and without using predicted quality (Fixed Cov). The baseline is NDT-MCL with a uniform sampling of pose particles in the map.

4.3.5 Application to radar

The method presented here adopts image-based pose regression from grayscale images. No specific assumptions are tied to the usage of a specific modality to produce such images. Interestingly, this representation bears a close resemblance with Cartesian spinning radar images, yet with a lower level of noise. This representation has recently been adopted in state-of-the-art radar place recognition [52], and has even been applied to radar pose regression [102]. For that reason, the presented method is likely suitable for global localization with a 2D radar.

Chapter 5
Conclusions

This chapter answers the research questions, presents future work, and summarizes the contributions of this thesis.

5.1 Summary of research questions

5.1.1 *RQ1: How can radar data be efficiently filtered and represented for robust radar odometry?*

This thesis presents a strategy that condenses dense raw data into a sparse set of oriented surface points. Surface points are computed in two steps: first, dense range data is filtered using k strongest for rapid conservative filtering that provides an upper bound of detections or with BFAR for increased sensitivity and detection rate. Evaluation filtering within the CFEAR pipeline shows that k strongest outperforms CA-CFAR in terms of speed and accuracy, with lower parameter sensitivity. However, the experiments in [**Paper IV**] demonstrate that BFAR is more suitable where the highest accuracy is required, at the expense of computational time. In the Spatial Cross-Validation of [**Paper III**], the optimal parameter set of k strongest demonstrated a high level of generalization between environments.

Oriented surface points can be computed using a grid-based method. In occupied grid cells, the surface point and normal are estimated by computing the intensity weighted point distribution, followed by eigendecomposition. The use of surface normals enables accurate odometry, even from sparse sets, with a high level of efficiency. This produces a set of features that are stable in at least one direction, which are suitable features, in particular, within feature-poor environments.

5.1.2 *RQ2: How can accurate and robust odometry be achieved with radar?*

Accurate and robust radar odometry can be achieved by considering multiple aspects of the odometry estimation problem. The challenge of non-rigid distortions arising from *scanning while moving* is addressed in spinning radar using a constant velocity model. The relative timestamp of each measurement is used to compensate (deskew) the range data. This strategy works well in realistic scenarios and tolerates a fair amount of acceleration at a system level without substantially deteriorating odometry quality.

Radar noise is largely addressed during filtering using k strongest or BFAR. However, subsequent steps, including multiple keyframes, weighted correspondences, and robust loss function, similarly suppress remaining false radar detections. Thus, the impact that radar-specific false detections have on odometry is greatly alleviated.

More generally, registration utilizes oriented surface normals, multiple keyframes, and weighted correspondences, to improve robustness within feature-poor environments and to changes in environment scale.

The drift is reduced by increasing the number of keyframes with which the latest scan is registered, even up to a large amount. However, a larger number of scans negatively impact computational time and are more suitable for offline mapping.

As an optional strategy to improve odometry within feature-poor environments, registration can be made robust by tightly coupling range measurements with externally provided ego-motion uncertainty estimates. This strategy seamlessly adapts reliance between point-cloud alignment and the external ego-motion depending on the available features in the scene. Thus preventing deterioration of registration within feature-poor environments. This method was evaluated on lidar data, but the principle applies to any method to range sensors, in general.

When these techniques were integrated for radar odometry (without external sensors), an average translation drift of 1.09% was reached in the urban Oxford radar RobotCar dataset. This corresponds to an estimated error of 4.36 m over a distance of 400 m. The most accurate method, after those presented in this thesis, reaches a drift of 1.65% [111]. The drift is consistently low, even in substantially different datasets. Additionally, these results are close to what can be achieved using lidar, which has higher precision and less noise compared to radar. However, the best-performing methods (and their implementations) in lidar and radar odometry have not yet been compared under similar conditions; for example, lidar odometry is slightly simplified in the most established benchmark “KITTI” [9, 53]. In the KITTI dataset, lidar point clouds have already been compensated for motion, even though the problem is of considerable importance. The absence of this issue is probably a contributing factor to why odometry performance is significantly lower on the KITTI benchmark compared to MulRan [65]. Another reason for the large discrepancy in performance can be that the experimental setup in MulRan additionally favors Radar, with the lidar having a restricted field of view. Additional experiments in fair benchmarks are required to understand how localization compares between sensor modalities.

5.1.3 *RQ3: How can localization quality be assessed and leveraged for robust detection of localization failures?*

In this thesis, localization quality assessment is approached by measuring and classifying the alignment between point clouds. In general, the detection of small localization errors is challenging, particularly under lower spatial overlap and for a variety of error magnitudes.

Evaluating the ability to detect errors is difficult. The ground poses generally contain some level of noise which can challenge learning and evaluation, particularly for a sensitive quality measure. What level of errors can be considered acceptable, and how sensitive the quality measure should be, depends on the application. Therefore, even small, acceptable errors in alignment can be incorrectly classified as misaligned. Despite this, the proposed CorAl method has demonstrated the ability to detect small localization errors with high accuracy and the ability to generalize across environments.

Interestingly, evaluation in both lidar and radar consistently showed that generalization was highest when training within less structured environments and evaluating in a more structured environment. Classification for a variety of errors could be improved by combining multiple quality measures, each tailored for the detection of small or large errors respectively, within a single classifier.

5.1.4 RQ4: *How can self-awareness of localization quality be utilized to enhance the robustness of a localization system?*

For global localization, localization quality is represented using a 6-DOF pose distribution, e.g. mean and covariance in both position and orientation. The presented method for hybrid one-shot/Monte Carlo localization can take advantage of its quality assessment – by sampling accordingly – to improve the localization success rate. Additionally, the localization time was reduced by narrowing the search space according to the predicted pose distribution or even omitting predictions with high uncertainty. The full method enables efficient and rapid global localization within a large-scale environment, achieving subsecond median localization time.

For loop closure in SLAM, the loop quality assessment is performed by combining separate measures of place similarity, consistency odometry uncertainty, and CorAl alignment quality. Using the loop quality measure, loops can be robustly classified with close to 100% precision with recall exceeding 90%, without the need to consider sequential information for an additional loop consistency check. Based on the highly robust loop closure and odometry proposed in this thesis, pose-graph optimization was able to correct drift without the need for a more resilient optimization back-end. This robustness, together with a high level of generalization, enables accurate and reliable mapping of new unseen environments.

5.2 Future work

During the work of this thesis, interesting directions for future research were discovered. Some of these are mentioned here.

5.2.1 Fusing scans into radar submaps

The experiments in [**Paper III**] show that the lowest radar odometry drift was achieved when leveraging a larger number of keyframe scans for registration. There is, however, a trade-off between drift and computational time that restricts the amount that can be considered while retaining the efficiency of CFEAR. Thus, a promising research direction should focus on fusing multiple scans into local maps without significantly altering the run-time performance. Recent work by Zhang et al. [111] and Wang et al. [101] included a focus on fusing scans into local maps within their work. However, the benefit was not systematically evaluated.

As an additional aspect, building detailed local maps may have a positive benefit for place recognition. The technique presented in [**Paper VII**] relied merely on a sparse set of radar peaks, aggregated into local maps. Therefore, most radar measurements are discarded. It is likely that more detailed local maps, created from multiple scans, can provide a higher level of detail required to correctly recognize places.

5.2.2 Localization in prior maps

The work carried out in this thesis explored SLAM with radar as well as global localization on previous lidar maps. However, the presented method for global localization using a hybrid deep learning / MCL localization method is agnostic to the sensor type and could be adapted for usage with radar. For example, one-shot deep pose regression (system 1 in Sec. 4.3) could be learned from superimposed scans, fused into detailed maps as discussed in 5.2.1. On the other hand, an MCL measurement model (system 2 in Sec. 4.3) could be derived based on the sparse set of oriented surface points to achieve efficient sequential localization.

5.2.3 Localization fault detection

This thesis proposed a method for localization fault detection and showed how the detection of alignment errors could improve the robustness of loop closure. Note that the usage within loop closure is primarily motivated by its otherwise relatively high failure rate in relation to errors in odometry estimation. In a similar fashion, pose tracking could be more carefully verified, given sufficiently reliable fault detections.

In general, more work is required to advance detection performance to improve the safety of autonomous systems. This problem may be addressed from multiple perspectives, including certifiable registration, by combining multiple sensors, or misalignment detection after registration, as done within this thesis.

5.3 Thesis contributions and summary

In summary, the most important contributions of this thesis are:

1. A pipeline for efficient and accurate radar odometry that robustly generalizes across environments. A detailed ablation study quantifies the most important aspects of achieving high odometry quality. The proposed method achieves state-of-the-art performance in public benchmarks. An optional sensor fusion technique is presented that incorporates an external source of odometry with uncertainty, into registration. Thus, retaining robustness within highly feature-poor environments.

2. A method that evaluates localization quality and detects localization errors by analyzing alignment quality between point-cloud pairs.
   The method robustly generalizes across environments. Higher generalization can be achieved by training in less structured environments.

3. A pipeline for an introspective and accurate real-time method for state-of-the-art radar SLAM. The method robustly generalizes across environments and produces consistent radar maps. Robust loop retrieval and computation of loop constraints achieved with more than 90% recall at close to 100% precision. This is achieved by jointly verifying place retrieval and loop constraint, and by coupling loop retrieval and verification.

4. A topometric localization framework that improves pose tracking in prior maps by the use of specialized submaps. The framework is motivated by the insight that topometric submap representations can be preferred to monolithic maps, even for globally consistent maps – a set of smaller maps is preferred to larger submaps for accurate pose tracking.

5. A method for fast global localization in prior maps. The method combines the prediction of pose and quality using deep regression, together

with MCL in topometric submaps, thus achieving fast and accurate localization. Predictions of localization quality are leveraged for improved robustness and localization speed.

The research presented in this thesis has contributed to the knowledge of large-scale localization and mapping, addressing numerous remaining challenges in the endeavor to achieve robust and safe autonomous navigation. Covered research topics include radar interpretation and odometry estimation, sensor fusion, localization fault detection, SLAM, map representation, global localization, and pose tracking in prior maps. The proposed methods have been evaluated using publicly available benchmarks, presenting state-of-the-art results in comparative evaluations. They have been further evaluated on data collected from diverse underground, off-road environments, and industrial indoor scenarios. Evaluations carried out in this thesis demonstrate versatile applicability of the proposed methods, particularly for use in diverse and feature-poor environments. Ablation studies have been presented that demonstrate the impact of components and parameters. Generalization studies have shown that the methods consistently operate accurately even within previously unseen environments. Research material including code, collected data, and evaluations have been released for some articles to promote re-usability and comparability, to further advance the research field.

References

[1] Nader J. Abu-Alrub and Nathir A. Rawashdeh. Radar odometry for autonomous ground vehicles: A survey of methods and datasets, 2023.

[2] Daniel Adolfsson, Stephanie Lowry, and Henrik Andreasson. Improving localisation accuracy using submaps in warehouses. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Workshop on Robotics for Logistics in Warehouses and Environments Shared with Humans*, 2018.

[3] Daniel Adolfsson, Stephanie Lowry, Martin Magnusson, Achim Lilienthal, and Henrik Andreasson. A submap per perspective - selecting subsets for SuPer mapping that afford superior localization quality. In *2019 European Conference on Mobile Robots (ECMR)*, pages 1–7, 2019.

[4] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim Lilienthal, and Henrik Andreasson. Oriented surface points for efficient and accurate radar odometry. In *Radar Perception for All-Weather Autonomy, a Half-Day Workshop, ICRA*, 2021. URL `https://arxiv.org/abs/2109.09994`. arXiv.

[5] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. CFEAR radarodometry - conservative filtering for efficient and accurate radar odometry. In *IROS*, pages 5462–5469, 2021.

[6] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. CFEAR radarodometry - conservative filtering for efficient and accurate radar odometry. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5462–5469, 2021.

[7] Daniel Adolfsson, Martin Magnusson, Qianfang Liao, Achim J. Lilienthal, and Henrik Andreasson. CorAl – are the point clouds correctly

aligned? In *2021 European Conference on Mobile Robots (ECMR)*, pages 1–7, 2021.

[8] Daniel Adolfsson, Manuel Castellano-Quero, Martin Magnusson, Achim J. Lilienthal, and Henrik Andreasson. CorAl: Introspection for robust radar and lidar perception in diverse environments using differential entropy. *Robotics and Autonomous Systems*, 155:104136, 2022. ISSN 0921-8890. URL `https://www.sciencedirect.com/science/article/pii/S0921889022000768`.

[9] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Lidar-level localization with radar? The CFEAR approach to accurate, fast and robust large-scale radar odometry in diverse environments. *T-RO*, pages 1–20, 2022.

[10] Daniel Adolfsson, Mattias Karlsson, Vladimír Kubelka, Martin Magnusson, and Henrik Andreasson. TBV radar SLAM – trust but verify loop candidates. *IEEE Robotics and Automation Letters*, 8(6):3613–3620, 2023.

[11] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Lidar-level localization with radar? the CFEAR approach to accurate, fast, and robust large-scale radar odometry in diverse environments. *IEEE Transactions on Robotics*, 39 (2):1476–1495, 2023.

[12] Naoki Akai. Reliable monte carlo localization for mobile robots. *Journal of Field Robotics*, n/a(n/a). URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.22149`.

[13] Naoki Akai, Luis Yoichi Morales, Takatsugu Hirayama, and Hiroshi Murase. Misalignment Recognition Using Markov Random Fields With Fully Connected Latent Variables for Detecting Localization Failures. *IEEE Robotics and Automation Letters*, 4(4):3955–3962, October 2019. ISSN 2377-3774.

[14] Naoki Akai, Yasuhiro Akagi, Takatsugu Hirayama, Takayuki Morikawa, and Hiroshi Murase. Detection of localization failures using markov random fields with fully connected latent variables for safe lidar-based automated driving. *T-ITS*, 23(10):17130–17142, 2022.

[15] R. Aldera, D. D. Martini, M. Gadd, and P. Newman. Fast radar motion estimation with a learnt focus of attention using weak supervision. In *2019 (ICRA)*, pages 1190–1196, 2019.

[16] R. Aldera, D. D. Martini, M. Gadd, and P. Newman. What could go wrong? introspective radar odometry in challenging environments. In *2019 IEEE (ITSC)*, pages 2835–2842, 2019.

[17] Roberto Aldera, Matthew Gadd, Daniele De Martini, and Paul Newman. What goes around: Leveraging a constant-curvature motion constraint in radar odometry. *RAL*, 7(3):7865–7872, 2022.

[18] Anas Alhashimi, Daniel Adolfsson, Martin Magnusson, Henrik Andreasson, and Achim J. Lilienthal. BFAR – bounded false alarm rate detector for improved radar odometry estimation, 2021. URL `https://arxiv.org/abs/2109.09669`.

[19] Håkan Almqvist, Martin Magnusson, Tomasz P. Kucner, and Achim J. Lilienthal. Learning to detect misaligned point clouds. *JFR*, 35(5):662–677, 2018.

[20] Henrik Andreasson, Daniel Adolfsson, Todor Stoyanov, Martin Magnusson, and Achim J. Lilienthal. Incorporating ego-motion uncertainty estimates in range data registration. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1389–1395, 2017.

[21] D. Barnes and I. Posner. Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar. In *ICRA*, pages 9484–9490, 2020.

[22] Dan Barnes, Matthew Gadd, Paul Murcutt, Paul Newman, and Ingmar Posner. The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset. In *ICRA*, pages 6433–6438, 2020.

[23] Dan Barnes, Rob Weston, and Ingmar Posner. Masking by moving: Learning distraction-free radar odometry from pose information. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *CoRL*, volume 100 of *CoRL*, pages 303–316. PMLR, 30 Oct–01 Nov 2020.

[24] Jens Behley and Cyrill Stachniss. Efficient surfel-based slam using 3d laser range data in urban environments. In *RSS*, volume 2018, page 59, 2018.

[25] Ola Bengtsson and Albert-Jan Baerveldt. Robot localization based on scan-matching—estimating the covariance matrix for the idc algorithm. *RAS*, 44(1):29–40, 2003. ISSN 0921-8890. URL `https://www.sciencedirect.com/science/article/pii/S0921889003000083`. Best Papers of the Eurobot '01 Workshop.

[26] P.J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE PAMI*, 14(2):239–256, 1992.

[27] P. Biber and W. Strasser. The normal distributions transform: a new approach to laser scan matching. In *IROS*, volume 3, pages 2743–2748 vol.3, 2003.

[28] Jose-Luis Blanco, Juan-Antonio FernÁndez-Madrigal, and Javier GonzÁlez. Toward a unified bayesian approach to hybrid metric–topological slam. *IEEE Transactions on Robotics*, 24(2):259–270, 2008.

[29] Silvère Bonnabel, Martin Barczyk, and François Goulette. On the covariance of icp-based scan-matching techniques. In *2016 American Control Conference (ACC)*, pages 5498–5503, 2016.

[30] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller. An atlas framework for scalable mapping. In *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, volume 2, pages 1899–1906 vol.2, 2003.

[31] Keenan Burnett, Angela P. Schoellig, and Timothy D. Barfoot. Do we need to compensate for motion distortion and doppler effects in spinning radar navigation? *IEEE RAL*, 6(2):771–778, 2021.

[32] Keenan Burnett, David J. Yoon, Angela P. Schoellig, and Tim Barfoot. Radar odometry combining probabilistic estimation and unsupervised feature learning. In *RSS*, Virtual, July 2021.

[33] Keenan Burnett, Yuchen Wu, David J. Yoon, Angela P. Schoellig, and Timothy D. Barfoot. Are we ready for radar to replace lidar in all-weather mapping and localization? *IEEE RAL*, 7(4):10328–10335, 2022.

[34] Keenan Burnett, David J. Yoon, Yuchen Wu, Andrew Zou Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith Y. K. Leung, Angela P. Schoellig, and Timothy D. Barfoot. Boreas: A multi-season autonomous driving dataset, 2022. URL `https://arxiv.org/abs/2203.10168`.

[35] 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. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016.

[36] Jonas Callmer, David Törnqvist, Fredrik Gustafsson, Henrik Svensson, and Pelle Carlbom. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011(1), sep 2011.

[37] Carlos Campos, Richard Elvira, Juan J. Gómez Rodríguez, José M. M. Montiel, and Juan D. Tardós. Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021.

[38] Nicholas Carlevaris-Bianco, Arash K. Ushani, and Ryan M. Eustice. University of Michigan North Campus long-term vision and lidar dataset. *International Journal of Robotics Research*, 35(9):1023–1035, 2015.

[39] Manuel Castellano Quero, Tomasz Piotr Kucner, and Martin Magnusson. Alignability maps for ensuring high-precision localization. In *13th IROS Workshop on Planning, Perception, Navigation for Intelligent Vehicles*.

[40] S. H. Cen and P. Newman. Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions. In *ICRA*, pages 6045–6052, 2018.

[41] Sarah H. Cen and Paul Newman. Radar-only ego-motion estimation in difficult settings via graph matching. In *ICRA*, pages 298–304, 2019.

[42] Manjari Chandran and Paul Newman. Motion estimation from map quality with millimeter wave radar. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 808–813, 2006.

[43] Paul Checchin, Franck Gérossier, Christophe Blanc, Roland Chapuis, and Laurent Trassoudaine. Radar scan matching SLAM using the fourier-mellin transform. In *Springer Tracts in Advanced Robotics*, volume 62, pages 151–161, 01 2009. ISBN 978-3-642-13407-4.

[44] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, pages 2724–2729 vol.3, 1991.

[45] S. Clark and H. Durrant-Whyte. Autonomous land vehicle navigation using millimeter wave radar. In *ICRA*, volume 4, pages 3697–3702 vol.4, 1998.

[46] Daniele De Martini, Matthew Gadd, and Paul Newman. kRadar++: Coarse-to-fine FMCW scanning radar localisation. *Sensors*, 20(21), 2020. ISSN 1424-8220. URL `https://www.mdpi.com/1424-8220/20/21/6002`.

[47] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, volume 2, pages 1322–1328 vol.2, 1999.

[48] M.W.M.G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions on Robotics and Automation*, 17(3): 229–241, 2001.

[49] D. Droeschel and S. Behnke. Efficient continuous-time slam for 3d lidar-based online mapping. In *ICRA*, pages 1–9, May 2018.

[50] Philipp Egger, Paulo V K Borges, Gavin Catt, Andreas Pfrunder, Roland Siegwart, and Renaud Dubé. Posemap: Lifelong, multi-environment 3d lidar localization. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3430–3437, 2018.

[51] HM Finn. Adaptive detection mode with threshold control as a function of spatially sampled clutter-level estimates. *RCA Rev.*, 29:414–465, 1968.

[52] Matthew Gadd, Daniele De Martini, and Paul Newman. Contrastive learning for unsupervised radar place recognition. In *ICAR*, pages 344–349, 2021.

[53] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *(CVPR)*, 2012.

[54] Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE Transactions on Robotics*, 23(1):34–46, 2007.

[55] Kyle Harlow, Hyesu Jang, Timothy D. Barfoot, Ayoung Kim, and Christoffer Heckman. A new wave in robotics: Survey on recent mmwave radar applications in robotics, 2023.

[56] Li He, Xiaolong Wang, and Hong Zhang. M2dp: A novel 3d point cloud descriptor and its application in loop closure detection. In *IROS*, pages 231–237. IEEE, 2016.

[57] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar SLAM. In *ICRA*, pages 1271–1278, 2016.

[58] Martin Holder, Sven Hellwig, and Hermann Winner. Real-time pose graph slam based on radar. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1145–1151, 2019.

[59] Ziyang Hong, Yvan Petillot, and Sen Wang. Radarslam: Radar based large-scale slam in all weathers. In *IROS*, pages 5164–5170, Oct 2020.

[60] Ziyang Hong, Yvan Petillot, Andrew Wallace, and Sen Wang. RadarSLAM: A robust simultaneous localization and mapping system for all weather conditions. *IJRR*, 41(5):519–542, 2022. URL `https://doi.org/10.1177/02783649221080483`.

[61] Xiaoliang Ju, Donghao Xu, and Huijing Zhao. Scene-aware error modeling of lidar/visual odometry for fusion-based vehicle localization. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):6480–6494, 2022.

[62] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *2016 IEEE international conference on Robotics and Automation (ICRA)*, pages 4762–4769. IEEE, 2016.

[63] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015.

[64] Giseop Kim and Ayoung Kim. Scan context: Egocentric spatial descriptor for place recognition within 3D point cloud map. In *IROS*, Oct. 2018.

[65] Giseop Kim, Yeong Sang Park, Younghun Cho, Jinyong Jeong, and Ayoung Kim. Mulran: Multimodal range dataset for urban place recognition. In *ICRA*, pages 6246–6253, Paris, May 2020.

[66] Giseop Kim, Sunwook Choi, and Ayoung Kim. Scan context++: Structural place recognition robust to rotation and lateral variations in urban environments. *IEEE Transactions on Robotics*, 38(3):1856–1874, 2022.

[67] Jacek Komorowski. Minkloc3d: Point cloud based large-scale place recognition. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1789–1798, 2021.

[68] Jacek Komorowski, Monika Wysoczanska, and Tomasz Trzcinski. Large-scale topological radar localization using learned descriptors. In Teddy Mantoro, Minho Lee, Media Anugerah Ayu, Kok Wai Wong, and Achmad Nizar Hidayanto, editors, *Neural Information Processing*, pages 451–462, Cham, 2021. Springer International Publishing. ISBN 978-3-030-92270-2.

[69] Pou-Chun Kung, Chieh-Chih Wang, and Wen-Chieh Lin. A normal distribution transform-based radar odometry designed for scanning and automotive radars. In *ICRA*, pages 14417–14423, 2021.

[70] Pou-Chun Kung, Chieh-Chih Wang, and Wen-Chieh Lin. Radar occupancy prediction with lidar supervision while preserving long-range sensing and penetrating capabilities. *IEEE Robotics and Automation Letters*, 7(2):2637–2643, 2022.

[71] David Landry, Francois Pomerleau, and Philippe Giguère. Cello-3d: Estimating the covariance of icp in the real world. In *ICRA*, pages 8190–8196, 2019.

[72] Hyungtae Lim, Kawon Han, Gunhee Shin, Giseop Kim, Songcheol Hong, and Hyun Myung. Orora: Outlier-robust radar odometry, 2023.

[73] José Raúl Machado-Fernández, Norelys Mojena-Hernández, and Jesús de la Concepción Bacallao-Vidal. Evaluation of cfar detectors performance. *Iteckne*, 14(2):170–178, 2017.

[74] Martin Magnusson, Achim Lilienthal, and Tom Duckett. Scan registration for autonomous mining vehicles using 3D-NDT. *Journal of Field Robotics*, 24(10):803–827, 2007. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.20204`.

[75] A. Makadia, A. Patterson, and K. Daniilidis. Fully automatic registration of 3d point clouds. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 1, pages 1297–1304, June 2006.

[76] Jan Willem Marck, Ali Mohamoud, Eric vd Houwen, and Rob van Heijster. Indoor radar SLAM: A radar application for vision and GPS denied environments. In *2013 EuRAD*, pages 471–474. IEEE, 2013.

[77] Mladen Mazuran, Federico Boniardi, Wolfram Burgard, and Gian Diego Tipaldi. Relative topometric localization in globally inconsistent maps. In *International Symposium of Robotics Research*, 2015.

[78] Juan Pablo Mendoza, Manuela M. Veloso, and Reid Simmons. Mobile Robot Fault Detection based on Redundant Information Statistics. 11 2009. URL `https://kilthub.cmu.edu/articles/journal_contribution/Mobile_Robot_Fault_Detection_based_on_Redundant_Information_Statistics/6607376`.

[79] Malcolm Mielle, Martin Magnusson, and Achim J Lilienthal. A comparative analysis of radar and lidar sensing for localization and mapping. In *ecmr*, September 2019.

[80] Simona Nobili, Georgi Tinchev, and Maurice Fallon. Predicting alignment risk to prevent localization failure. In *ICRA*, pages 1003–1010, 2018.

[81] Y. S. Park, Y. S. Shin, and A. Kim. PhaRaO: Direct radar odometry using phase correlation. In *2020 IEEE (ICRA)*, pages 2617–2623, 2020.

[82] Sai Manoj Prakhya, Liu Bingbing, Yan Rui, and Weisi Lin. A closed-form estimate of 3d ICP covariance. In *2015 14th IAPR (MVA)*, pages 526–529, May 2015.

[83] Noha Radwan, Abhinav Valada, and Wolfram Burgard. Vlocnet++: Deep multitask learning for semantic visual localization and odometry. *IEEE Robotics and Automation Letters*, 3(4):4407–4414, 2018.

[84] Jan Razlaw, David Droeschel, Dirk Holz, and Sven Behnke. Evaluation of registration methods for sparse 3d laser scans. *2015 European Conference on Mobile Robots (ECMR)*, pages 1–7, 2015.

[85] B.S. Reddy and B.N. Chatterji. An fft-based technique for translation, rotation, and scale-invariant image registration. *IEEE Transactions on Image Processing*, 5(8):1266–1271, 1996.

[86] Hermann Rohling. Radar cfar thresholding in clutter and multiple target situations. *IEEE Transactions on Aerospace and Electronic Systems*, AES-19:608–621, 1983.

[87] R. Rouveure, M.O. Monod, and P. Faure. High resolution mapping of the environment with a ground-based radar imager. In *2009 International Radar Conference "Surveillance for a Safer World" (RADAR 2009)*, pages 1–6, 2009.

[88] Szymon Marek Rusinkiewicz. Efficient variants of the ICP algorithm. In *The Third International Conference on 3D Digital Imaging and Modeling*, pages 145–152, 2001.

[89] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, Juha Ala-Luhtala, and Achim J. Lilienthal. Normal distributions transform occupancy maps: Application to large-scale online 3d mapping. In *ICRA*, pages 2233–2238, 2013.

[90] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, and Achim J. Lilienthal. Normal distributions transform monte-carlo localization (NDT-MCL). In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 382–389, 2013.

[91] Stefan Saftescu, Matthew Gadd, Daniele De Martini, Dan Barnes, and Paul Newman. Kidnapped radar: Topological radar localisation using rotationally-invariant metric learning. In *ICRA*, pages 4358–4364, 2020.

[92] Carl H. Schiller, Bruno Arsenali, Deran Maas, and Stefano Maranó. Improving marine radar odometry by modeling radar resolution and exploiting additional temporal information. In *IROS*, pages 8436–8441, 2022.

[93] F. Schuster, C. G. Keller, M. Rapp, M. Haueis, and C. Curio. Landmark based radar SLAM using graph optimization. In *ITSC*, pages 2559–2564, 2016.

[94] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5135–5142, 2020.

[95] Marcel Sheeny, Emanuele De Pellegrin, Saptarshi Mukherjee, Alireza Ahrabian, Sen Wang, and Andrew Wallace. Radiate: A radar dataset for automotive perception. *arXiv preprint arXiv:2010.09076*, 2020.

[96] Luciano Silva, Olga R.P. Bellon, and Kim L. Boyer. Precision range image registration using a robust surface interpenetration measure and enhanced genetic algorithms. *TPAMI*, 27(5):762–776, May 2005.

[97] Li Sun, Daniel Adolfsson, Martin Magnusson, Henrik Andreasson, Ingmar Posner, and Tom Duckett. Localising faster: Efficient and precise lidar-based robot localisation in large-scale environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4386–4392, 2020.

[98] P. Sundvall and P. Jensfelt. Fault detection for mobile robots using redundant positioning systems. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 3781–3786, 2006.

[99] Mikaela Angelina Uy and Gim Hee Lee. Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4470–4479, 2018.

[100] Abhinav Valada, Noha Radwan, and Wolfram Burgard. Deep auxiliary learning for visual localization and odometry. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 6939–6946. IEEE, 2018.

[101] Dequan Wang, Yifan Duan, Xiaoran Fan, Chengzhen Meng, Jianmin Ji, and Yanyong Zhang. Maroam: Map-based radar slam through two-step feature selection, 2022. URL `https://arxiv.org/abs/2210.13797`.

[102] Wei Wang, Pedro P. B. de Gusmão, Bo Yang, Andrew Markham, and Niki Trigoni. Radarloc: Learning to relocalize in fmcw radar. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5809–5815, 2021.

[103] Wei Wang, Bing Wang, Peijun Zhao, Changhao Chen, Ronald Clark, Bo Yang, Andrew Markham, and Niki Trigoni. Pointloc: Deep pose regressor for lidar point cloud localization. *IEEE Sensors Journal*, 22 (1):959–968, 2022.

[104] R Weston. *Deep probabilistic methods for improved radar sensor modelling and pose estimation*. PhD thesis, University of Oxford, 2022.

[105] Rob Weston, Sarah Cen, Paul Newman, and Ingmar Posner. Probably unknown: Deep inverse sensor modelling radar. In *2019 (ICRA)*, pages 5446–5452. IEEE, 2019.

[106] Rob Weston, Matthew Gadd, Daniele De Martini, Paul Newman, and Ingmar Posner. Fast-mbym: Leveraging translational invariance of the fourier transform for efficient and accurate radar odometry. In *ICRA*, pages 2186–2192, 2022.

[107] Huan Yin, Li Tang, Xiaqing Ding, Yue Wang, and Rong Xiong. A failure detection method for 3d lidar based localization. In *2019 Chinese Automation Congress (CAC)*, pages 4559–4563, 2019.

[108] Huan Yin, Xuecheng Xu, Yue Wang, and Rong Xiong. Radar-to-lidar: Heterogeneous place recognition via joint learning. *Frontiers in Robotics and AI*, 8, 2021. ISSN 2296-9144. URL `https://www.frontiersin.org/article/10.3389/frobt.2021.661199`.

[109] Huan Yin, Xuecheng Xu, Sha Lu, Xieyuanli Chen, Rong Xiong, Shaojie Shen, Cyrill Stachniss, and Yue Wang. A survey on global lidar localization, 2023. URL `https://arxiv.org/abs/2302.07433`.

[110] Ji Zhang and Sanjiv Singh. Low-drift and real-time lidar odometry and mapping. *Autonomous Robots*, 41:401–416, 02 2017.

[111] Rongxi Zhang, Yuanhui Zhang, Duo Fu, and Kang Liu. Scan denoising and normal distribution transform for accurate radar odometry and positioning. *RAL*, 8(3):1199–1206, 2023.

[112] Tan Zhang, Wenjun Zhang, and Madan M. Gupta. Resilient robots: Concept, review, and future directions. *Robotics*, 6(4), 2017. ISSN 2218-6581. URL `https://www.mdpi.com/2218-6581/6/4/22`.

[113] Yi Zhou, Lulu Liu, Haocheng Zhao, Miguel López-Benítez, Limin Yu, and Yutao Yue. Towards deep radar perception for autonomous driving: Datasets, methods, and challenges. *Sensors*, 22(11), 2022. ISSN 1424-8220. URL `https://www.mdpi.com/1424-8220/22/11/4208`.

[114] Zhicheng Zhou, Cheng Zhao, Daniel Adolfsson, Songzhi Su, Yang Gao, Tom Duckett, and Li Sun. NDT-transformer: Large-scale 3d point cloud localisation using the normal distribution transform representation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5654–5660, 2021.

[115] Robert Zlot and Michael Bosse. *Efficient Large-Scale 3D Mobile Mapping and Surface Reconstruction of an Underground Mine*, pages 479–493. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-642-40686-7. URL `https://doi.org/10.1007/978-3-642-40686-7_32`.

Publications *in the series* Örebro Studies in Technology

1. Bergsten, Pontus (2001) *Observers and Controllers for Takagi – Sugeno Fuzzy Systems*. Doctoral Dissertation.
2. Iliev, Boyko (2002) *Minimum-time Sliding Mode Control of Robot Manipulators*. Licentiate Thesis.
3. Spännar, Jan (2002) *Grey box modelling for temperature estimation*. Licentiate Thesis.
4. Persson, Martin (2002) *A simulation environment for visual servoing*. Licentiate Thesis.
5. Boustedt, Katarina (2002) *Flip Chip for High Volume and Low Cost – Materials and Production Technology*. Licentiate Thesis.
6. Biel, Lena (2002) *Modeling of Perceptual Systems – A Sensor Fusion Model with Active Perception*. Licentiate Thesis.
7. Otterskog, Magnus (2002) *Produktionstest av mobiltelefonantenner i mod-växlande kammare*. Licentiate Thesis.
8. Tolt, Gustav (2003) *Fuzzy-Similarity-Based Low-level Image Processing*. Licentiate Thesis.
9. Loutfi, Amy (2003) *Communicating Perceptions: Grounding Symbols to Artificial Olfactory Signals*. Licentiate Thesis.
10. Iliev, Boyko (2004) *Minimum-time Sliding Mode Control of Robot Manipulators*. Doctoral Dissertation.
11. Pettersson, Ola (2004) *Model-Free Execution Monitoring in Behavior-Based Mobile Robotics*. Doctoral Dissertation.
12. Överstam, Henrik (2004) *The Interdependence of Plastic Behaviour and Final Properties of Steel Wire, Analysed by the Finite Element Metod*. Doctoral Dissertation.
13. Jennergren, Lars (2004) *Flexible Assembly of Ready-to-eat Meals*. Licentiate Thesis.
14. Jun, Li (2004) *Towards Online Learning of Reactive Behaviors in Mobile Robotics*. Licentiate Thesis.
15. Lindquist, Malin (2004) *Electronic Tongue for Water Quality Assessment*. Licentiate Thesis.
16. Wasik, Zbigniew (2005) *A Behavior-Based Control System for Mobile Manipulation*. Doctoral Dissertation.

17. Berntsson, Tomas (2005) *Replacement of Lead Baths with Environment Friendly Alternative Heat Treatment Processes in Steel Wire Production.* Licentiate Thesis.

18. Tolt, Gustav (2005) *Fuzzy Similarity-based Image Processing.* Doctoral Dissertation.

19. Munkevik, Per (2005) *"Artificial sensory evaluation – appearance-based analysis of ready meals".* Licentiate Thesis.

20. Buschka, Pär (2005) *An Investigation of Hybrid Maps for Mobile Robots.* Doctoral Dissertation.

21. Loutfi, Amy (2006) *Odour Recognition using Electronic Noses in Robotic and Intelligent Systems.* Doctoral Dissertation.

22. Gillström, Peter (2006) *Alternatives to Pickling; Preparation of Carbon and Low Alloyed Steel Wire Rod.* Doctoral Dissertation.

23. Li, Jun (2006) *Learning Reactive Behaviors with Constructive Neural Networks in Mobile Robotics.* Doctoral Dissertation.

24. Otterskog, Magnus (2006) *Propagation Environment Modeling Using Scattered Field Chamber.* Doctoral Dissertation.

25. Lindquist, Malin (2007) *Electronic Tongue for Water Quality Assessment.* Doctoral Dissertation.

26. Cielniak, Grzegorz (2007) *People Tracking by Mobile Robots using Thermal and Colour Vision.* Doctoral Dissertation.

27. Boustedt, Katarina (2007) *Flip Chip for High Frequency Applications – Materials Aspects.* Doctoral Dissertation.

28. Soron, Mikael (2007) *Robot System for Flexible 3D Friction Stir Welding.* Doctoral Dissertation.

29. Larsson, Sören (2008) *An industrial robot as carrier of a laser profile scanner: Motion control, data capturing and path planning.* Doctoral Dissertation.

30. Persson, Martin (2008) *Semantic Mapping Using Virtual Sensors and Fusion of Aerial Images with Sensor Data from a Ground Vehicle.* Doctoral Dissertation.

31. Andreasson, Henrik (2008) *Local Visual Feature based Localisation and Mapping by Mobile Robots.* Doctoral Dissertation.

32. Bouguerra, Abdelbaki (2008) *Robust Execution of Robot Task-Plans: A Knowledge-based Approach.* Doctoral Dissertation.

33. Lundh, Robert (2009) *Robots that Help Each Other: Self-Configuration of Distributed Robot Systems.* Doctoral Dissertation.

34. Skoglund, Alexander (2009) *Programming by Demonstration of Robot Manipulators*. Doctoral Dissertation.

35. Ranjbar, Parivash (2009) *Sensing the Environment: Development of Monitoring Aids for Persons with Profound Deafness or Deafblindness*. Doctoral Dissertation.

36. Magnusson, Martin (2009) *The Three-Dimensional Normal-Distributions Transform – an Efficient Representation for Registration, Surface Analysis, and Loop Detection.* Doctoral Dissertation.

37. Rahayem, Mohamed (2010) *Segmentation and fitting for Geometric Reverse Engineering. Processing data captured by a laser profile scanner mounted on an industrial robot.* Doctoral Dissertation.

38. Karlsson, Alexander (2010) *Evaluating Credal Set Theory as a Belief Framework in High-Level Information Fusion for Automated Decision-Making*. Doctoral Dissertation.

39. LeBlanc, Kevin (2010) *Cooperative Anchoring – Sharing Information About Objects in Multi-Robot Systems*. Doctoral Dissertation.

40. Johansson, Fredrik (2010) *Evaluating the Performance of TEWA Systems*. Doctoral Dissertation.

41. Trincavelli, Marco (2010) *Gas Discrimination for Mobile Robots.* Doctoral Dissertation.

42. Cirillo, Marcello (2010) *Planning in Inhabited Environments: Human-Aware Task Planning and Activity Recognition.* Doctoral Dissertation.

43. Nilsson, Maria (2010) *Capturing Semi-Automated Decision Making: The Methodology of CASADEMA*. Doctoral Dissertation.

44. Dahlbom, Anders (2011) *Petri nets for Situation Recognition*. Doctoral Dissertation.

45. Ahmed, Muhammad Rehan (2011) *Compliance Control of Robot Manipulator for Safe Physical Human Robot Interaction.* Doctoral Dissertation.

46. Riveiro, Maria (2011) *Visual Analytics for Maritime Anomaly Detection*. Doctoral Dissertation.

47. Rashid, Md. Jayedur (2011) *Extending a Networked Robot System to Include Humans, Tiny Devices, and Everyday Objects*. Doctoral Dissertation.

48. Zain-ul-Abdin (2011) *Programming of Coarse-Grained Reconfigurable Architectures*. Doctoral Dissertation.

49. Wang, Yan (2011) *A Domain-Specific Language for Protocol Stack Implementation in Embedded Systems*. Doctoral Dissertation.

50. Brax, Christoffer (2011) *Anomaly Detection in the Surveillance Domain*. Doctoral Dissertation.

51. Larsson, Johan (2011) *Unmanned Operation of Load-Haul-Dump Vehicles in Mining Environments*. Doctoral Dissertation.

52. Lidström, Kristoffer (2012) *Situation-Aware Vehicles: Supporting the Next Generation of Cooperative Traffic Systems*. Doctoral Dissertation.

53. Johansson, Daniel (2012) *Convergence in Mixed Reality-Virtuality Environments. Facilitating Natural User Behavior*. Doctoral Dissertation.

54. Stoyanov, Todor Dimitrov (2012) *Reliable Autonomous Navigation in Semi-Structured Environments using the Three-Dimensional Normal Distributions Transform (3D-NDT)*. Doctoral Dissertation.

55. Daoutis, Marios (2013) *Knowledge Based Perceptual Anchoring: Grounding percepts to concepts in cognitive robots*. Doctoral Dissertation.

56. Kristoffersson, Annica (2013) *Measuring the Quality of Interaction in Mobile Robotic Telepresence Systems using Presence, Spatial Formations and Sociometry*. Doctoral Dissertation.

57. Memedi, Mevludin (2014) *Mobile systems for monitoring Parkinson's disease*. Doctoral Dissertation.

58. König, Rikard (2014) *Enhancing Genetic Programming for Predictive Modeling*. Doctoral Dissertation.

59. Erlandsson, Tina (2014) *A Combat Survivability Model for Evaluating Air Mission Routes in Future Decision Support Systems*. Doctoral Dissertation.

60. Helldin, Tove (2014) *Transparency for Future Semi-Automated Systems. Effects of transparency on operator performance, workload and trust*. Doctoral Dissertation.

61. Krug, Robert (2014) *Optimization-based Robot Grasp Synthesis and Motion Control.* Doctoral Dissertation.
62. Reggente, Matteo (2014) *Statistical Gas Distribution Modelling for Mobile Robot Applications*. Doctoral Dissertation.
63. Längkvist, Martin (2014) *Modeling Time-Series with Deep Networks*. Doctoral Dissertation.
64. Hernández Bennetts, Víctor Manuel (2015) *Mobile Robots with In-Situ and Remote Sensors for Real World Gas Distribution Modelling*. Doctoral Dissertation.
65. Alirezaie, Marjan (2015) *Bridging the Semantic Gap between Sensor Data and Ontological Knowledge*. Doctoral Dissertation.
66. Pashami, Sepideh (2015) *Change Detection in Metal Oxide Gas Sensor Signals for Open Sampling Systems*. Doctoral Dissertation.
67. Lagriffoul, Fabien (2016) *Combining Task and Motion Planning*. Doctoral Dissertation.
68. Mosberger, Rafael (2016) *Vision-based Human Detection from Mobile Machinery in Industrial Environments*. Doctoral Dissertation.
69. Mansouri, Masoumeh (2016) *A Constraint-Based Approach for Hybrid Reasoning in Robotics*. Doctoral Dissertation.
70. Albitar, Houssam (2016) *Enabling a Robot for Underwater Surface Cleaning*. Doctoral Dissertation.
71. Mojtahedzadeh, Rasoul (2016) *Safe Robotic Manipulation to Extract Objects from Piles: From 3D Perception to Object Selection.* Doctoral Dissertation.
72. Köckemann, Uwe (2016) *Constraint-based Methods for Human-aware Planning*. Doctoral Dissertation.
73. Jansson, Anton (2016) *Only a Shadow. Industrial Computed Tomography Investigation, and Method Development, Concerning Complex Material Systems.* Licentiate Thesis.
74. Sebastian Hällgren (2017) *Some aspects on designing for metal Powder Bed Fusion.* Licentiate Thesis.
75. Junges, Robert (2017) *A Learning-driven Approach for Behavior Modeling in Agent-based Simulation.* Doctoral Dissertation.
76. Ricão Canelhas, Daniel (2017) *Truncated Signed Distance Fields Applied To Robotics*. Doctoral Dissertation.

77. Asadi, Sahar (2017) *Towards Dense Air Quality Monitoring: Time-Dependent Statistical Gas Distribution Modelling and Sensor Planning*. Doctoral Dissertation.

78. Banaee, Hadi (2018) *From Numerical Sensor Data to Semantic Representations: A Data-driven Approach for Generating Linguistic Descriptions*. Doctoral Dissertation.

79. Khaliq, Ali Abdul (2018) *From Ants to Service Robots: an Exploration in Stigmergy-Based Navigation Algorithms*. Doctoral Dissertation.

80. Kucner, Tomasz Piotr (2018) *Probabilistic Mapping of Spatial Motion Patterns for Mobile Robots*. Doctoral Dissertation.

81. Dandan, Kinan (2019) *Enabling Surface Cleaning Robot for Large Food Silo*. Doctoral Dissertation.

82. El Amine, Karim (2019) *Approaches to increased efficiency in cold drawing of steel wires*. Licentiate Thesis.

83. Persson, Andreas (2019) *Studies in Semantic Modeling of Real-World Objects using Perceptual Anchoring*. Doctoral Dissertation.

84. Jansson, Anton (2019) *More Than a Shadow. Computed Tomography Method Development and Applications Concerning Complex Material Systems*. Doctoral Dissertation.

85. Zekavat, Amir Reza (2019) *Application of X-ray Computed Tomography for Assessment of Additively Manufactured Products*. Doctoral Dissertation.

86. Mielle, Malcolm (2019) *Helping robots help us—Using prior information for localization, navigation, and human-robot interaction*. Doctoral Dissertation.

87. Grosinger, Jasmin (2019) *On Making Robots Proactive*. Doctoral Dissertation.

88. Arain, Muhammad Asif (2020) *Efficient Remote Gas Inspection with an Autonomous Mobile Robot*. Doctoral Dissertation.

89. Wiedemann, Thomas (2020) *Domain Knowledge Assisted Robotic Exploration and Source Localization*. Doctoral Dissertation.

90. Giaretta, Alberto (2021) *Securing the Internet of Things with Security-by-Contract*. Doctoral Dissertation.

91. Rudenko, Andrey (2021) *Context-aware Human Motion Prediction for Robots in Complex Dynamic Environments*. Doctoral Dissertation.

92. Eriksson, Daniel (2021) *Getting to grips with cartons: Interactions of cartonboard packages with an artificial finger.* Doctoral Dissertation.

93. Dinh-Cuong, Hoang (2021) *Vision-based Perception For Autonomous Robotic Manipulation*. Doctoral Dissertation.

94. Akalin, Neziha (2022) *Perceived Safety in Social Human-Robot Interaction*. Doctoral Dissertation.

95. Han Fan (2022) *Robot-aided Gas Sensing for Emergency Responses*. Doctoral Dissertation.

96. Tomic, Stevan (2022) *Human Norms for Robotic Minds*. Doctoral Dissertation.

97. Kondyli, Vasiliki (2023) *Behavioural Principles for the Design of Human-Centered Cognitive Technologies, The Case of Visuo-Locomotive Ewperience*. Doctoral Dissertation.

98. Morillo-Mendez, Lucas (2023) *SOCIAL ROBOTS / SOCIAL COGNITION. Robots' Gaze Effects in Older and Younger Adults*. Doctoral Dissertation.

99. Yang, Yuxuan (2023) *Advancing Modeling and Tracking of Deformable Linear Objects for Real-World Applications*. Doctoral Dissertation.

100. Adolfsson, Daniel (2023) *Robust large-scale mapping and localization*. Doctoral Dissertation.