

Safety-Aware Radar Data Processing in Automated Parking Functions

Zur Erlangung des akademischen Grades eines
DOKTORS DER INGENIEURWISSENSCHAFTEN (Dr.-Ing.)
von der KIT-Fakultät für
Elektrotechnik und Informationstechnik
des Karlsruher Instituts für Technologie (KIT)
angenommene

DISSERTATION

von
Tim Benedikt Brühl, M.Sc.
geb. in Neuwied

Tag der mündlichen Prüfung:	11.09.2026
Hauptreferent:	Prof. Dr.-Ing. Sören Hohmann
Korreferent:	Prof. Dr.-Ing. Steven Peters

Preface

First and foremost, I would like to express my heartfelt gratitude to Prof. Dr.-Ing. Sören Hohmann for supervising this work, for the constant support, and for the constructive feedback during our regular meetings. From you, I learned what scientific research truly means. It was reassuring to have an advisor who appreciated the relevance of this topic and continuously encouraged me to develop new concepts. My sincere thanks also go to Prof. Dr.-Ing. Steven Peters for the keen interest in my research and for taking on the role of co-referee.

My time as a doctoral researcher at Dr. Ing. h.c. F. Porsche AG has been an extraordinary experience. I thank Martin Boß for the trust placed in me and Dr.-Ing. Sebastian Söhner for the insightful guidance and strategic direction in our research areas. Furthermore, I thank Dr.-Ing. Simon Isele for his supervision and the critical review of this thesis that helped make my explanations clearer and more comprehensible in so many respects.

My special thanks go to my Ph.D. peers – Dr.-Ing. Stefan Roos, Dr.-Ing. Lukas Ewecker, Dr.-Ing. Robin Schwager, Tin Stribor Sohn, and Tim Eberhardt. You made the past three years unforgettable through your ideas, professional discussions, humor, and friendship. I feel very fortunate to have been part of a team where I could always rely on my colleagues. I am also grateful to the entire campus group in ADAS pre-development at Porsche for their stimulating discussions and to the Porsche doctoral network for its fruitful exchange of ideas.

I would like to thank all my students and thesis candidates for their motivation and the invaluable support that made my work so much easier: Alex, who is no longer with us – far too soon and beyond comprehension. Manuel, Aravind, Felix Blahak, Julia, Max, Jenny, Jan, Benedikt, Felix Otterbach, and Vico. Your energy and confidence in me inspired me to give my best every day, and you delivered in return. This thesis is also your achievement.

I am indebted to my friends for their encouragement and patience whenever travels or weekend plans had to be postponed or canceled. My heartfelt thanks go to Arne and Antonia for always making me laugh and for standing by me. You have been incredibly important to me throughout my journey. Finally, I would like to express my deepest gratitude to my parents, Anja and Herbert, for their unwavering support and unfaltering belief in me.

Abstract

The development of automated parking functions is reaching a level of automation that fully relieves the user from the driving task. As the human driver will no longer serve as a fallback, the safety requirements for highly automated parking functions increase significantly. Any potential hazard of other road users by the function must be eliminated to the greatest extent possible. Due to its advantages in distance and velocity estimation, radar-based environment perception plays a central role in this context.

This thesis presents a method for radar data processing that is tailored to the increasing safety demands of highly automated parking functions and is applicable at multiple abstraction levels. Existing approaches reduce noise inherent to the measurement principle by applying filtering operations, but in doing so may also remove potentially relevant information. In contrast, the method proposed here evaluates the criticality of all radar data within each processing step based on a driving tube that represents the space the ego vehicle will occupy in the immediate future. Radar data identified as critical are bypassed around the signal processing chain and thus remain guaranteed available to the planner. The work considers the complete radar data processing chain, which derives a spectrum from the raw radar data, then generates a point cloud, radar clusters, and ultimately object hypotheses. The proposed criticality assessment is designed to be independent of the respective processing stage. For each stage, specific strategies are discussed that enable the extraction of further indicators of critical information. These strategies also extend into the temporal context of previously processed data in order to recover information that would otherwise have been discarded. In addition, a systematic procedure is introduced to determine the resulting error rate. The proposed approaches are evaluated using existing datasets and in a real-world experiment. Moreover, the idea that discarding information may lead to critical situations is transferred from object detection to motion estimation. For this domain as well, a concept for recognizing and preventing critical actions is developed and verified.

By performing data processing based on the criticality assessment defined by the vehicle's driving tube, the proposed method connects the perception with the planning function. Together with high-performance radar sensors, it establishes the foundation for deploying radar sensors in highly automated parking functions with guaranteed error rates.

Kurzfassung

Die Entwicklung automatisierter Parkfunktionen befindet sich an der Schwelle eines Automatisierungsgrades, der den Nutzer vollständig von der Fahraufgabe entbindet. Da damit der menschliche Fahrer künftig als Rückfallebene entfällt, steigen die Sicherheitsanforderungen an hochautomatisierte Parkfunktionen erheblich. Eine Gefährdung anderer durch die Funktion muss weitestgehend eliminiert werden. Durch die Vorteile bei der Distanz- und Geschwindigkeitsmessung kommt der radarbasierten Wahrnehmung der Umgebung hierbei eine zentrale Rolle zu.

Diese Dissertation präsentiert eine Methode zur Verarbeitung von Radardaten, die auf die steigenden Sicherheitsanforderungen hochautomatisierter Parkfunktionen zugeschnitten ist und auf mehreren Abstraktionsebenen ansetzt. Bestehende Verfahren reduzieren das messprinzipbedingte Rauschen durch Filteroperationen, entfernen dabei jedoch potenziell relevante Informationen. Im Gegensatz dazu bewertet die hier vorgeschlagene Methode die Kritikalität sämtlicher Radardaten innerhalb jedes Verarbeitungsschritts anhand eines Fahrschlauchs, der den Raum beschreibt, den das Ego-Fahrzeug in naher Zukunft einnehmen wird. Als kritisch identifizierte Radardaten umgehen die Filteroperationen über Bypässe und bleiben somit garantiert für den Planer verfügbar. Die Arbeit betrachtet die vollständige Radarverarbeitungskette, die ausgehend von den Radar-Rohdaten zunächst ein Spektrum, dann eine Punktwolke, Radar-Cluster und schließlich Objekthypothesen bildet. Die beschriebene Kritikalitätsbewertung ist unabhängig von der jeweiligen Verarbeitungsstufe anwendbar. Für jede Stufe werden spezifische Strategien diskutiert, die das Extrahieren weiterer Hinweise auf kritische Informationen ermöglichen. Diese Strategien beziehen auch den temporalen Kontext bereits verarbeiteter Datenmengen mit ein, um zuvor verworfene Informationen wiederherstellen zu können. Zudem wird eine systematische Vorgehensweise zur Bestimmung der resultierenden Fehlerrate vorgestellt. Die Ansätze werden sowohl anhand bestehender Datensätze als auch in einem Realexperiment evaluiert. Darüber hinaus wird die Annahme, dass das Verwerfen von Informationen zu kritischen Situationen führen kann, von der Objekterkennung auf die Bewegungsschätzung übertragen. Auch hierfür wird ein Konzept zur Erkennung und Verhinderung kritischer Aktionen entwickelt und verifiziert.

Indem die Datenverarbeitung auf Basis der durch den Fahrschlauch des Fahrzeugs vorgegebenen Kritikalitätsbewertung vorgenommen wird, verbindet die vorgestellte Methode Wahrnehmungs- und Planungsfunktion. Zusammen mit leistungsfähigen Radarsensoren schafft sie damit die Grundlage für den Einsatz von Radarsensoren in hochautomatisierten Parkfunktionen mit garantierten Fehlerraten.

Contents

Preface	i
Abstract	iii
Kurzfassung	v
1 Introduction and Motivation	1
1.1 Motivation	1
1.2 Safety Guarantees on Radar-Detected Objects	4
1.3 Structure	5
2 State of the Art	9
2.1 Fundamentals	9
2.2 Perception Sensors	15
2.3 Radar Sensor Data Processing	16
2.4 Vehicular Safety	28
2.5 State of Research	32
2.6 Discussion and Research Gaps	37
3 Area Criticality Estimation	41
3.1 Introduction	41
3.2 Driving Tube Construction	44
3.3 Criticality Evaluation	49
3.4 Discussion	57
4 Processing Adaptations in Radar Spectra	61
4.1 Information Representation in Radar Spectra	61
4.2 Subsequent Radar Peak Search	65
4.3 Evaluation of Safety-Aware Peak Detection	78
4.4 Discussion	88
5 Processing Adaptations on Point Clouds	93
5.1 Radar Point Clouds in Object Detection	93
5.2 Successive Point Cloud Treatment	96
5.3 Failure Rate Estimation	101
5.4 Evaluation on RadarScenes Dataset	108

5.5	Discussion	118
6	Processing Adaptations at the Object Level	123
6.1	Tracking of Object Hypotheses	123
6.2	Clustering and Tracking Modifications	125
6.3	Evaluation on Real World Data	128
6.4	Discussion	140
7	Application in Radar Odometry Estimation	145
7.1	Motion Estimation for Automated Parking Functions	145
7.2	Safe Radar Motion Estimation	150
7.3	Evaluation on RadarScenes Dataset	157
7.4	Discussion	164
8	Conclusion	169
A	Appendices to Chapter 2	173
B	Appendices to Chapter 4	179
	Abbreviations and Symbols	181
	List of Figures	191
	List of Tables	193
	References	195

1 Introduction and Motivation

This chapter presents major challenges in the area of perception systems for parking functions with a high level of automation¹. Section 1.1 highlights the current situation: although mature functions exist, their functional safety is still insufficient for market release as highly automated parking functions. The importance of the perception component for ensuring functional safety is a key motivation of this thesis. The research problem arises from the critical role of the perception component and is formulated in Section 1.2. Finally, Section 1.3 provides an outline of the content of this thesis.

1.1 Motivation

Today's society strives for safe, fast, and convenient individual transportation. Vehicle passengers should experience a comfortable and, ideally, brief travel time. Parking is often an unpleasant part of the journey, first because the passenger's destination has already been reached, and second due to the difficulty of the parking task, which may involve maneuvering through narrow and poorly illuminated environments. The potential to enhance user satisfaction by relieving passengers of the parking task is substantial.

These needs are addressed by offering functions that can assume the parking task through lateral and longitudinal vehicle control. Such systems typically comprise three major components, following the sense-plan-act paradigm [Bro91]:

- A perception system that continuously monitors the vehicle's environment to enable adaptation to other traffic participants.
- A planning system that generates policies for core driving tasks, such as navigation, guidance, and control [Don82].
- An actuation system that extends existing actuators, such as the drivetrain, braking system, and steering, with control interfaces.

Many partially automated parking functions are already available on the market [Gee21; Mer24]. The associated safety risk induced by these partially automated functions is low, as a human driver continuously monitors the environment. If the planning system makes decisions that could lead to collisions or endanger other traffic participants, the human driver can intervene and override the system's control actions. Moreover, if the

¹ Formal definition of automation levels given in Subsection 2.1.3.

function reaches its operational limits, it can request the human driver to take over the parking task. However, a major hurdle for highly automated parking functions that no longer require a driver is that all such situations must be managed autonomously by the system.

Current approaches demonstrate that assisted functions in both the parking and the driving domain are already capable of confidently handling a variety of scenarios. At the same point, these functions are expected to exceed the reliability of a human driver, even in unforeseen situations. A recent survey on the confidence in fully automated driving functions revealed that people believe vehicles with such functions will increase safety in road traffic (mean: 6.24/11 on Thurstone scale)[Wei+21]. Furthermore, they would rather trust a fully automated vehicle than today's driver in road traffic (6.09/11). While the release of functions with high automation levels remains restricted in Europe, companies such as Waymo are already conducting large-scale fleet experiments in urban areas of the United States [Way25]. A recent study compared failure rates of the "Waymo Driver" with human driver error rates in the same region. The Waymo system achieved 0.6 incidents per million miles, whereas the incident rate of human drivers was determined to be 2.1 incidents per million miles [Kus+24]. Assuming an average speed of 15 mph, Waymo achieves an approximate incident rate of $1 \times 10^5 h^{-1}$. However, automotive safety standards [Int18] demand failure rates as low as $1 \times 10^9 h^{-1}$. This shifts the challenge in the development of automated driving functions from handling typical scenarios within the intended Operational Design Domain (ODD)² to ensuring safe operation in all circumstances, including potentially unknown ones.

These expectations impose new requirements on the safety and accuracy of the perception, planning, and actuation systems. As the risks involved in the parking domain are generally lower than those in urban or highway driving scenarios, automated parking functions may be introduced earlier than highly automated driving functions and pave their way for production readiness. Nevertheless, it is essential to recognize that particularly Vulnerable Road User (VRU)s³ may still be severely injured during incidents in parking environments. According to the National Highway Traffic Safety Administration (NHTSA), 60,577 pedestrians were injured and 7,388 pedestrians were killed in traffic crashes in 2021 [Nat23]. A study conducted in the state Colorado reported that 10.4 % of all pedestrian crashes occurred in parking lots [GO 12]. In North Carolina, the share was 9.4 % [TLF14]. The city of Toronto reported that "pedestrian hit in parking lot" was the cause of 11 % of all pedestrian crashes [Bas+15]. Although the statistics do not report the number of injured or killed VRUs in parking garages, the existence of a relevant hazard can be reasonably assumed.

The number of collisions involving pedestrians in the context of automated parking functions highlights the critical importance of robust VRU detection for the planning system to generate safe trajectories. In the perception systems of highly automated functions, the integration of various sensor modalities is standard practice. The most

² Formal definition given in Subsection 2.1.3.

³ Formal definition given in Section 2.2.

important technologies include camera, lidar, ultrasonic, and radar sensors [WWN20]. Due to their use of ultrasonic wave propagation through air and the resulting limited range [Kin+00], ultrasonic sensors are primarily employed for near-range obstacle detection. Consequently, they are unsuitable for object detection during transfer maneuvers. During transfer maneuvers, the vehicle travels from the entrance of a parking garage to the parking space assigned to the vehicle by the backend management system. The remaining sensors — camera, lidar, and radar — are all capable of object detection in higher distances. However, camera performance degrades under poor illumination or low contrasts, such as when a pedestrian wears dark clothing in shadowed environments [Gha+24]. Lidar signals, while accurate under ideal conditions, suffer significant attenuation in the presence of rain or fog [Dre+23]. Although this is less common in indoor parking situations, it poses a substantial drawback in adverse weather conditions that affect other highly automated driving domains.

Radar sensors, by contrast, are more robust under adverse weather conditions, being less affected by fog and rain [Nor+20]. Moreover, unlike camera and most lidar systems, radar systems can directly measure the Doppler velocity of the reflection, its velocity component with respect to the sensor [Ric22]. As a result, radar sensors have become widely adopted in automotive perception systems. While radar can detect the presence of objects independently, its limited object classification capability [Sch+20a] necessitates fusion with other sensor modalities, such as cameras [Yao+24].

The reflection behavior of radar waves depends heavily on the material composition of the object. The lower the energy absorption and the higher the relative permittivity ϵ_r , the better the object reflects radar waves [KC20]. For example, materials with a high water content and metallic surfaces are strong reflectors. While the former facilitates VRU detection, the latter assists in recognizing bicycles or vehicles. However, metal structures may also elevate the overall reflected signal power and increase noise levels in the received signal. Signal processing filters can mitigate these effects by suppressing unwanted noise. Up to now, radar information that is likely to represent noise has often been removed. The probability of noise is higher for signals measured with low reflected power without supporting evidence in their spatial neighborhood. However, this approach may result in weak reflectors being absent from the final output of the radar processing. If such reflectors correspond to objects that require a response from the automated function, this filtering can create a safety hazard.

Automated parking functions operate in environments such as parking garages, where radar sensors must contend with metal-rich infrastructure including pillars, ceiling and fencing. Effective radar data processing is thus essential to manage these challenges. When radar sensors serve as the primary sensing modality for highly automated parking functions, they must not only achieve high detection performance but also ensure the safety of the object detection. Quantifying the rate of missed detections is crucial to meet stringent safety targets associated with high levels of autonomy. Accurate and safe perception protects the lives of VRUs in parking facilities, fosters public trust in automated systems, and enhances the customer's confidence in automated vehicles.

1.2 Safety Guarantees on Radar-Detected Objects

This thesis addresses the problem of safe radar-based object detection, which can be interpreted as hypothesis testing problem. Hypothesis testing introduces two types of errors, both of which are central concepts in this work. Incorrect decisions in the hypothesis testing process can lead to unsafe behavior of the automated function.

The notation of [NP33], who first described these error types, is adopted in the following. A null hypothesis is defined as follows: "The observation area is free; i.e., no object or structure is present at the position (x_{obj}, y_{obj}) ". The radar measurement may either support or contradict this null hypothesis. Depending on the truth value of the null hypothesis and the system's decision to accept or reject it, errors of type I or type II may occur.

Definition 1.1 (False Positive (FP) detection (Type I Error))

The null hypothesis is rejected, meaning that an object is suspected to be present at the position (x_{obj}, y_{obj}) , although the null hypothesis is actually true.

In this case, an object is erroneously detected. For the automated parking function, this means that the planner's intended trajectory is unnecessarily inhibited. While this may lead to deadlocks or reduced system robustness, it typically does not result in harm of other traffic participants.

Definition 1.2 (False Negative (FN) detection (Type II Error))

The null hypothesis is accepted, meaning that no object is suspected to be present at the position (x_{obj}, y_{obj}) , although the null hypothesis is actually false.

A type II error results in a hazardous situation if a present object is not detected and thus not incorporated into the planning task of the automated parking function. Type II errors in radar perception can propagate through the planning and acting stack, potentially leading to a violation of a safety goal under International Organization for Standardization (ISO) 26262, thereby necessitating guarantees at the sensor level.

Definition 1.3 (Hazardous Situation)

A hazardous situation is defined as a scenario in which an unsolved conflict arises between the policy of an automated vehicle and the intention of a traffic participant, and where injury of this traffic participant is likely if no action is taken.

Type I errors may trigger unnecessary braking or deadlocks, contributing to user dissatisfaction and reducing trust in the system. Excessive occurrence of this error type can lead to degraded user experience and may trigger inappropriate fallback behavior. However, in contrast to type II errors, type I errors can be tolerated to a certain extent because they do not create immediate safety risks. In partially automated functions, radar measurements are typically filtered to reduce the occurrence of type I errors to an acceptable level. In doing so, a limited number of type II errors may be tolerated as a trade-off.

This thesis proposes a radar data processing approach designed to eliminate type II errors that originate from the removal of valid information early in the processing pipeline. A core hypothesis of this work, subject to formal validation, is that object detection can be rendered safe by incorporating safety awareness into radar data processing. Safety-relevant information could be preserved throughout the radar processing pipeline. By doing so, the type II error rate can be reduced sufficiently to meet safety requirements. Moreover, safe radar perception depends on the presence of information about relevant objects in the raw radar data. If such information is absent at the sensor level, it cannot be recovered later in the signal processing pipeline. Consequently, radar sensors must provide a baseline level of fundamental detection performance. Type II errors caused by missing raw signal information fall under the domain of sensor hardware development and are outside the scope of this thesis.

1.3 Structure

An overview of the content of this work, divided into eight chapters, is provided below to facilitate comprehension of its structure. An illustrated version is provided Figure 1.1.

Chapter 2 introduces and clarifies the terminology used in this work. The elements of environmental perception formulated by the six-layer model are outlined. As this thesis specifically addresses radar data processing, another section describes the processing steps from the received radar echo to an object hypothesis. Moreover, mechanisms leading to erroneous radar detections and noise points are discussed. A review of the recent state of research is also provided, which is fundamental for deriving the research gap and resulting research questions.

In Chapter 3, a method is developed to estimate the criticality of radar measurement data based on the current motion state of the vehicle and the trajectory provided by the planner. The trajectory can be interpreted as a driving tube, representing the area that the vehicle will occupy in the immediate future. The criticality of radar measurement information is derived from its position within the driving tube. The drive-tube-based




Figure 1.1: Structure of the thesis, showing the first- and second-level headings. The pillars represent the research contributions of this work.

criticality evaluation can be applied at different processing levels of the radar system and is capable of incorporating additional attributes such as velocity measures.

One of the earliest steps at which radar data processing can be adapted is the peak detection in the radar spectrum. Chapter 4 introduces a method that intervenes at this step by modifying an established algorithm for peak detection, thereby providing a safety-aware process. This detection approach is intended to ensure that critical but weak peaks are detected while avoiding noise detections, particularly in regions with minimal impact on the automated function. To achieve this, previous radar measurements are revisited to search for evidence on object presence. In addition, this chapter presents an experiment evaluating the effectiveness of this methodology.

Chapter 5 discusses the criticality-aware filtering of radar point clouds. Each point in the cloud receives a criticality score. The filter criteria of default filters are reduced or suspended in regions where critical points are identified. This enables the reintroduction of points into subsequent processing, even if they were filtered out by default. An experiment demonstrates the effect of the approach. Furthermore, the chapter discusses how safety of the perception function can be quantified by counting perception malfunctions that contribute to hazardous situations.

Chapter 6 proposes concepts for utilizing the criticality evaluation of points to adapt the downstream algorithms that generate objects based on the point cloud. These objects represent the final stage of radar data processing and serve as input to the automated function. All traffic participants requiring a response from the automated function must be included in this object list. Two modifications are introduced: one adapts the clustering method, and the other accelerates the generation of a hypothesis if a point is sufficiently critical. Since critical scenes with pedestrians and corresponding radar data are not yet covered by existing datasets, a dedicated dataset is recorded to enable conclusions regarding the performance of these approaches.

In Chapter 7, the concept that some radar points in the cloud are more crucial for safe behavior than others is transferred to radar odometry. The algorithm for estimating the ego vehicle's longitudinal velocity is first introduced. With this setup, erroneous estimates are identified in an existing dataset, particularly in environments that challenge the radar sensor, such as tunnels or construction sites with close metal guard rails. Methods are presented to identify and mitigate failures in radar odometry estimation that could have hazardous consequences. This chapter demonstrates that the problem arises across all domains in which sensor data are used to guide the vehicle, extending beyond object detection, which is arguably the most intuitive application of criticality-aware radar data processing.

Chapter 8 concludes this work by recapitulating its most important findings and summarizing its contributions to the field. Furthermore, it outlines future research directions that build upon these contributions.

2 State of the Art

This chapter establishes the theoretical and technical foundation of the present work. Since this thesis addresses an automotive context, the system under consideration is assumed to be integrated into a vehicle, hereafter referred to as the *ego vehicle*.

Definition 2.1 (Ego Vehicle)

The ego vehicle is the vehicle equipped with radar sensors that generate radar measurement data. It serves as the center of the environmental model constructed from these sensors. This environmental model is a central component that enables automated functions within the ego vehicle.

Section 2.1 introduces fundamental mathematical concepts, such as coordinate systems and motion models. Since the requirements and limitations of the perception system are derived from the functionality of the targeted application, this section also outlines the relevant terminology. The perception system observes the environment surrounding the ego vehicle, which is structured according to the six-layer model. The role of radar sensors in fulfilling this perception task is described in Section 2.2. Following this, Section 2.3 explains the operational principle of radar sensors, including the signal processing chain that transforms raw measurements into intermediate representations such as radar spectra, point clouds and an object list. This object list serves as an input to the planning subsystem. Given that the objective of this work is to integrate safety considerations into the perception task, which contributes to an end-to-end safe automated function, Section 2.4 introduces the key concepts of vehicular safety and the applicable international norms. The literature review in Section 2.5 provides insights on the current state-of-research in the fields that are relevant for the aspects covered in this work. Finally, Section 2.6 discusses the strengths and limitations of existing work, identifying research gaps, such as linking planner's intentions to the radar data processing. The research questions of this work are derived from the insights gained in this chapter.

2.1 Fundamentals

This section outlines the foundational assumptions underpinning this work. Subsection 2.1.1 defines the coordinate systems used throughout the methodology, including a

description of the various coordinate frames and their transformations. Subsection 2.1.2 presents the motion models applied to estimate the ego vehicle's trajectory based on discrete measurements from onboard sensors. In particular, the Constant Velocity (CV), Constant Turn Rate and Velocity (CTRV), and Constant Turn Rate and Acceleration (CTRA) models are discussed. Finally, Subsection 2.1.3 introduces the automated parking function including its ODD. This subsection bridges the gap between signal processing in perception and the downstream modules for planning and control.

2.1.1 Coordinate Systems

The interaction between sensor data that perceive objects in the environment and data measured by the ego vehicle introduces multiple coordinate systems. The *vehicle* coordinate system is the most relevant coordinate frame applied in this work. Its origin is located in or on the ego vehicle. All *sensor* coordinate systems are dependent on the *vehicle* coordinate system. In the literature, two widespread vehicle coordinate systems are used: the one specified in ISO 8855 [Int11], and the one specified in Society of Automotive Engineers (SAE) International (SAE) J670 [SAE22]. The convention adopted in this work is ISO 8855. This standard defines the x-axis as pointing forward, the y-axis as pointing to the left, and the z-axis as pointing upward. The rotation around the x-axis is referred to as the roll angle, the rotation around the y-axis as the pitch angle, and the rotation around the z-axis as the yaw angle. ISO 8855 defines the reference point to be located somewhere within the vehicle's sprung mass. In contrast, a slightly modified variant is applied in this work, which locates the reference point at the center of the rear axis on the ground ($z = 0$). The advantage of this definition is that the z-positions of objects external to the vehicle can be determined independently of the vehicle's current damper compression. This simplifies the intuitive understanding of spatial relations in the environment. However, if the z-position becomes relevant, the deviation from ISO 8855 must be taken into account. For the purpose of this work, in which environmental data is mostly transformed into a 2D plane using either (r, φ) or (x, y) coordinates, this definition proves to be effective.

Additionally, the *odom* coordinate system is employed as a stationary reference frame. Its origin coincides with that of the ISO 8855 coordinate system but is fixed at the start point $t = t_0$. Positional shifts resulting from ego vehicle motion, estimated using wheel-speed sensors and the yaw rate sensor, are added to this initial point applying the motion models outlined in the following Subsection 2.1.2. The *world* coordinate system is also stationary. It defines the position of the *vehicle* coordinate system within the environment. Due to inaccuracies in the localization process, the *odom* frame drifts from the *world* frame over time. However, for the investigated area, this deviation is negligible. The *odom* frame would coincide with the *world* frame in the hypothetical case of perfect odometry without measurement errors.

The *sensor* coordinate system uses the sensor's center as its origin. This coordinate

system is defined as a static transformation relative to the *vehicle* coordinate system. For radar sensors, measurements are typically acquired in polar coordinates, which can be converted to Cartesian coordinates.

The ISO 8855 *vehicle* coordinate system, the *odom* and *world* coordinate systems, and the *sensor* coordinate system (exemplary for the front left radar sensor) are illustrated in Figure 2.1.




Figure 2.1: Overview of the coordinate systems used in this work: *odom* (black lines), *world* (grey lines), *vehicle* (blue lines), and *sensor* (red lines).

2.1.2 Vehicle Model

To determine the position of the ego vehicle in the *odom* coordinate frame, the CTRV vehicle model is employed. This model discretizes the vehicle's motion states $\mathcal{S}$ and assumes constant velocity v_k and constant yaw rate ω_k between two time steps. Each vehicle state is defined as $\mathcal{S} = [x_{\text{odom}}, y_{\text{odom}}, v, \varphi, \omega]$. By applying the time difference Δt between two time steps k and $k + 1$, the consecutive state $\mathcal{S}[k + 1]$ of the vehicle can be computed using the following formula:

$$\mathcal{S}_{\text{CTRV}}[k + 1] = \begin{bmatrix} x_{\text{odom}}[k + 1] \\ y_{\text{odom}}[k + 1] \\ v[k + 1] \\ \varphi[k + 1] \\ \omega[k + 1] \end{bmatrix} = \begin{bmatrix} x[k] + \cos \varphi \cdot v[k + 1] \cdot \Delta t \\ y[k] + \sin \varphi \cdot v[k + 1] \cdot \Delta t \\ v[k + 1] \\ \varphi[k] + \omega[k + 1] \Delta t \\ \omega[k + 1] \end{bmatrix}. \quad (2.1)$$

This simplified formula assumes linear motion rather than circular motion, which is a valid approximation for sufficiently low yaw rates ω . The CV model is a special case of the CTRV model where $\omega = 0$.

In some vehicles, acceleration is directly measured using simple and low-cost accelerometers. In such cases, the CTRA motion model may be applied. Since spatial position in this model is derived via double integration, the impact of measurement noise is more pronounced than in the CTRV model. The corresponding state equations are:

$$\mathcal{S}_{\text{CTRA}}[k+1] = \begin{bmatrix} x_{\text{odom}}[k+1] \\ y_{\text{odom}}[k+1] \\ v[k+1] \\ a[k+1] \\ \varphi[k+1] \\ \omega[k+1] \end{bmatrix} = \begin{bmatrix} x[k] + \frac{1}{2} \cos \varphi \cdot v[k] \cdot \Delta t \\ y[k] + \sin \varphi \cdot v[k] \cdot \Delta t \\ v[k] + a[k+1] \cdot \Delta t \\ a[k+1] \\ \varphi[k] + \omega[k+1] \Delta t \\ \omega[k+1] \end{bmatrix}. \quad (2.2)$$

In this thesis, these motion models are also used as state models in radar object tracking, as described in Subsection 2.3.5. This step in the processing of radar data produces motion estimates for each object hypothesis, based on previous state estimates and the latest radar measurement. The motion models introduced in this section are well-suited for the movement of VRUs [Bor+25]. More advanced concepts of VRU motion prediction incorporate motion patterns learned via Convolutional Neural Network (CNN)s, head orientation, the presence of other road users, and various environmental aspects [Gho+22]. While these methods may provide more accurate trajectory estimation than the CTRV model, most of the required input features are not available from radar data, limiting the applicability of such methods in this context.

2.1.3 Automated Valet Parking

This work presents an approach for safety-aware processing in automated parking functions. This subsection defines the terminology and scope of the function for which the considerations of this work are valid.

Definition 2.2 (Automated Driving Function [Win+24])

An automated driving function in the automotive context takes over tasks such as the longitudinal and/or lateral control of the vehicle, thereby partially or fully disengaging the driver from the driving task. It operates vehicle actuators, such as steering or braking system. The control outputs are determined by the automated driving function using information obtained, for example, from maps or perception sensors. The driver can override the automated driving function at any time.

Automated functions can be categorized into various levels of automation. A widely adopted classification is defined by the SAE J3016 standard [SAE21]. If the automation

level does not correspond to the highest level (SAE Level 5), the functionality is limited to operating within its function-specific ODD.

Definition 2.3 (Operational Design Domain (ODD) [SAE21])

Operating conditions under which a given driving automation system or feature thereof is specifically designed to function, including, but not limited to, environmental, geographical, and time-of-day restrictions, and/or the requisite presence or absence of certain traffic or roadway characteristics.

Further distinctions can be made based on the operational environment. *Automated driving functions* typically perform driving tasks on public roads, such as in urban, rural or highway scenarios, while *automated parking functions* are responsible for maneuvering the vehicle into a parking space after it has reached its destination. The primary differences between automated driving and automated parking functions are the expected speeds of surrounding traffic participants, which are generally lower in parking environments, and the significantly higher demand for positional accuracy in parking maneuvers.

In most automated functions currently on the market, the driver is still required to continuously monitor the system and remains responsible for its decisions. This corresponds to SAE Level 2 of the automation classification [SAE21]. However, a setup in which the driver must remain present during the parking procedure is only conditionally useful in parking situations. In contrast, the introduction of highly automated SAE Level 4 functions for parking use cases appears particularly beneficial.

Definition 2.4 (SAE Level 4 Function)

An SAE Level 4 function (hereafter called Level 4 function) is capable of operating the vehicle without the requiring human driver intervention [SAE21]. Consequently, the function can operate without a driver within a defined ODD. If the Level 4 system encounters boundaries beyond its operational scope, it transitions to a predefined safe state.

The Automated Valet Parking (AVP) system considered in this work is specified by ISO 23374 and is classified as a Level 4 function [Int23]. It enables passengers to exit the vehicle at a designated drop-off location prior to the commencement of the parking towards a parking lot distributed by the parking facility and identified by a parking lot Identifier (ID). Similarly, upon request, the system autonomously exits the parking lot, navigates through the parking garage, and allows passengers to board the vehicle comfortably at the pick-up location.

AVP systems can be categorized into two distinct architectural approaches, which differ in the tasks taken by the Original Equipment Manufacturer (OEM) and the facility operator. These are contrasted in Table 2.1. However, hybrid or mixed implementations combining characteristics of both types are also possible [Ban+17].

Table 2.1: Comparison of AVP Type 1 and AVP Type 2 (adapted from [Kra22]).

	AVP Type 1	AVP Type 2
Perception Sensors	Mounted at vehicle	Mounted in parking facility
Map required by vehicle	Yes (can be achieved by crowd data)	No (routing handled by facility)
Route / Trajectory Planning	By vehicle	By facility backend
Object Detection	Mainly by vehicle	By facility backend
Localization	By vehicle	By facility backend
Data exchange	Low, one-time (Parking Lot ID)	High, constantly (Trajectory)
Effort for OEM	High	Low
Effort for facility operator	Low	High
Liability in case of incidents	OEM	Facility operator

The maximum admissible speed $v_{\max, \text{domain}}$ of the system is decisive for the hazard posed by the AVP function, especially if operation in mixed traffic is considered. Such mixed traffic is found, e.g., in a parking garage used by both vehicles controlled by an automated function and manually driven vehicles [Ger23]. In such cases, VRUs such as pedestrians approaching and leaving their parked vehicle are a permanently present object type that needs to be handled by the function. In the section on basic performance requirements, the ISO 23374 standard [Int23] specifies a maximum speed of $v_{\max, \text{domain}} = 30 \text{ km/h}$. This limit is also adopted in the position paper [Ger23] and is well suited to the characteristics of present SAE Level 2 automated parking systems for the Chinese market. In contrast, European systems often operate at a lower maximum speed to ensure user safety in the event of a malfunction.

An AVP system typically comprises several functional components: an environmental model, consisting of detected objects and an ego vehicle location, a mission control, also called routing, behavior and motion planning, and vehicle control [Ban+17]. This thesis concentrates on safety risks that arise from the perception function. Therefore, the components and stages related to perception elements are described in detail in the following sections.

2.2 Perception Sensors

The foundation for all actions performed by an automated vehicle's function is an environmental model, which typically consists of various types of information. All data required for safe planning decisions can be assigned to one of the layers within the six-layer model [Sch+21a]. Table 2.2 lists all layers along with exemplary components and facilities in the vehicle providing information about the respective layer.

Table 2.2: Distribution of environmental information across the six-layer model.

	Layer 1	Layer 2	Layer 3	Layer 4	Layer 5	Layer 6
Description	Road	Traffic Infrastructure	Temporal Modifications of Street and Infrastructure	Moveable Objects	Environmental Conditions	Data and Communication
Examples	Road width, road surface	Traffic signs, speed bumps	Construction side fences, obstacles	Cars, Trucks	Fog, Sun dazzling	Car-to-Infrastructure
Additional information source	-	-	-	-	Temperature sensors, humidity sensors, light intensity sensors	Wireless networks
Map	X	X	-	-	-	-
Camera Sensor	-	X	X	X	X	-
Radar Sensor	-	-	X	X	(X)	-

As shown, the road level in layer 1 and traffic infrastructure level in layer 2 are typically part of an High-dimension (HD) map. In current research, there are investigations into automated functions that can operate without an HD map [Ort+20]. Nevertheless, since the components of the first two layers are static, they seldom fall within the domain covered by perception sensors. In contrast, layer 3 contains temporary modifications of layers 1 and 2, under which are obstructions in the drivable space, such as parking vehicles, poles, or guardrails. The elements of this layer, referred to here as *geometries*, must be measured by perception sensors. However, a failure to detect geometries is generally not safety-critical in the parking domain, as it would only result in material damage, even in the case of a significant misinterpretation of sensor data. Instead, hazardous situations primarily arise from VRUs.

Definition 2.5 (Vulnerable Road User (VRU) [Shi17])

A VRU is a traffic participant who is not protected by seat belts or the structural shield of the vehicle, unlike a passenger in a car or truck. VRUs are primarily riders of two-wheeled motorized vehicles such as mopeds, motorcycles, and scooters, as well as bicyclists and pedestrians. The risk of injury for a VRU in the event of a collision with a vehicle is significantly higher compared to other traffic participants.

VRUs are a subclass of layer 4 (moveable objects). In most automated functions, the contents of layer 4 are represented as a list of *objects*.

Definition 2.6 (Object [Win+24; Sch+21a])

The final product of a perception system that represents elements of the fourth layer is called object. Typically, an object includes a position, a velocity, and class estimate. It may also contain additional attributes such as dimensions and an existence probability. An object originates from an object hypothesis derived from measurement data gathered by perception sensors.

An object constitutes the final stage of signal processing within the perception system, which is then used by the planner to determine its policy. The accurate and comprehensive use of data acquired by perception sensors is essential to ensure the integrity of all layer 4 elements.

Attributes of layer 5 should also be estimated by perception sensors but are mainly used for holistic interpretation of the driving scene. They do not immediately entail safety-relevant tasks. In contrast, layer 6 attributes are obtained by the vehicle via telecommunication systems such as cellular networks.

In summary, the six-layer model applies the concept of describing the environment in sets of instances. For the dynamic parts of the environment, these sets are also known as agent sets [Jia+23] or *object lists*. The approach is also a major element of architectures for automated driving [Ras+22] and part of standards, e.g., OpenSCENARIO®[ASA22].

2.3 Radar Sensor Data Processing

This section describes the individual steps of radar data processing, beginning with the received radar echoes and concluding with an object list representation. In this work, the term radar data refers to information generated at different stages of the radar

processing chain, that are introduced below. Radar signals are considered a specific subset of radar data corresponding to early-stage signal representations.

Definition 2.7 (Radar Signal)

A signal is a function of time that represent the evolution of one or more independent variables [Bou06; Sun24], typically time. In the context of radar systems, a radar signal denotes the electrical or digital representation of electromagnetic waves received by the radar sensor, particularly in the early stages of the radar processing chain [Ric22].

As automated parking functions generally operate using high-level representations of the environment, a multi-stage processing of the radar data is required. Various references describe radar data processing pipelines for object detection [Eng+21; Mar22; Win+24]. Maintaining an object list requires associating newly detected object candidates with previously identified objects and continuing their observations using tracking algorithms [SM23]. The referenced works differ in the granularity of signal processing and the definition of the endpoint of the processing pipeline. For the purposes of this work, a tailored overview is synthesized from the aforementioned literature. The complete radar data processing pipeline is illustrated in Figure 2.2. The processing chain considered begins with the wave arriving at the receiver antenna and ends with an object list. The intermediate outputs of each processing step are indicated along the transition arrow.

This section starts with a definition of radar signal noise in Subsection 2.3.1. Subsequently, Subsection 2.3.2 provides an overview about general radar signal processing, including the required Digital Signal Processing (DSP) and radar characteristics. An important intermediate product of radar data processing is the radar point cloud. The clustering and tracking algorithms described in Subsection 2.3.4 and Subsection 2.3.5 are based on this point cloud and are used to generate cluster centroids and object hypotheses, respectively. The planner receives these hypotheses in the form of an object list.

2.3.1 Radar Noise

Noise in signal processing is a component of the received signal which corrupts the wanted signal. Noise operates mostly as an addition or multiplication to the wanted signal. Details about noise in signal processing can be found in [Tuz02]. Consult [Wil90] for an in-depth discussion about signal recovery, which is — for radar signals — usually performed during the early signal processing.

Radar signals are subject to various noise effects. Two sources of radar noise can be distinguished: noise received through the antenna and noise generated in the radar




Figure 2.2: Radar data processing pipeline, illustrating the steps from the reflected wave to the object list.

receiver itself [Ric22]. Both external and internal noise lead to a higher noise level, causing FN detections. The complex mechanisms of wave propagation make it difficult to process and remove these disturbances. Research has addressed the impact of phase noise, which causes the oscillator to produce an imperfect sinusoidal wave, resulting in a widened central peak and distorted sidebands [Sid+19]. Recently, a comprehensive analysis of multiple noise factors affecting 4D radars was published, categorizing the reasons for the occurrence of noise into distinct types [Cha+25]:

1. Piece to Piece: Sensor-specific variance arising from the manufacturing process.
2. Change over Time: Aging effects in electronics and mechanical components, such as the radome.
3. Usage: Noise caused by integration into the vehicle, including imprecise calibration, mechanical vibration, and contaminants.
4. Environment: Noise induced by the adverse weather conditions, ground clutter, and multipath reflections, potentially resulting in ghost objects.
5. System Interaction: Disturbance caused by interaction with other systems, such as instable power supply, interference with signals from other electronic devices in the vehicle, or reflections from the vehicle's own bodywork.

Noise can lead to both FP and FN errors. FN errors occur if the noise level is increased, thus, weak reflections may be suppressed, as elaborated in Section 4.1, and the signal does not result in a detection point. FP points can, e.g., result from clutter effects or high-frequency disturbances. While some noise effects cannot be directly controlled, others can be mitigated to varying degrees. For instance, interference can be reduced using tensor decomposition techniques applied to the received signals [Wan+24b], and multipath propagation effects [Zhe+24] can also be addressed through advanced processing methods.

Interference is becoming increasingly significant due to the growing number of radar sensors deployed in automotive applications. It is caused by external sources of radio waves, primarily other radar sensors operating at similar frequencies, which affect the received signal and raise the overall noise level. Multipath propagation is a natural property of electromagnetic waves. While it can sometimes be exploited advantageously, such as in the detection of concealed objects, it can also lead to ambiguity. For example, reflections originating from ground clutter or mirror effects may generate false detections if the actual source of the signal is indeterminable.

Figure 2.3 illustrates these effects. On the left hand side, a True Positive (TP) pedestrian produces a FP, so-called ghost detection, due to a mirroring effect. Clutter can similarly cause additional FPs when the wave propagates directly or via indirect paths. A strong reflector such as a car can absorb the weak reflections of a pedestrian. As a consequence, this pedestrian is not detected (FN). On the right hand side, an example shows how signal processing can benefit from multipath reflections. Even occluded objects can be detected via reflections between the ground and the undercarriage of the occluding vehicle.




Figure 2.3: Environmental effects leading to erroneous detections [Cha+25], causing FP and FN detections (red). Direct and indirect multipath reflections, clutter, and absorption contribute to these effects. Multipath propagation, however, can enable TP detections (example in blue).

A substantial body of research is dedicated to processing these complex signals to better associate ghost targets and TPs [LS22; Zhe+24]. However, safety-aware frameworks

processing points that lead potentially to a hazardous situation, even if these points could not be assigned to ghost or multipath reflections, remain underdeveloped.

2.3.2 Radar Raw Signal Processing

During the development of automated driving functions such as Adaptive Cruise Control (ACC) or Automated Emergency Braking (AEB) systems, radar sensors were identified as a suitable solution for object detection. Originally used primarily in military applications, radar sensors were particularly valued for their long range and ability to measure relative velocities. These features are also required for highway driving scenarios.

The working principle of radar sensors is to transmit radio waves that are reflected by obstacles and objects in the environment. The reflected waves are evaluated by the DSP and processed to generate an environmental representation containing information about geometries and objects. In addition to enabling collision-free path planning, the resulting data can also be used for the ego vehicle's localization and motion estimation. The received signal power P_{rec} of the reflected wave can be calculated using the radar equation [Gam20], Equation 2.3. It depends on the transmission power P_{transm} , the antenna gains of the transmitting and receiving antennas G_{transm} and G_{rec} , respectively, the wavelength λ , the Radar Cross Section (RCS) of the reflector σ , and the range r between the sensor and the reflector.

$$P_{\text{rec}} = \frac{P_{\text{transm}} \cdot G_{\text{transm}} \cdot G_{\text{rec}} \cdot \lambda^2 \cdot \sigma}{(4\pi)^3 \cdot r^4}. \quad (2.3)$$

The RCS is a key parameter in this work for distinguishing radar points, and is therefore defined below.

Definition 2.8 (Radar Cross Section (RCS))

The RCS of a target is the fictional area intercepting that amount of power which, when scattered equally in all directions, produces an echo at the radar equal to that from the target [Sko+80].

The RCS is characteristic for each instance reflecting the radar wave. Its value relates the reflection properties of the target to those of an isotropic reference reflector, typically a sphere with perfect conductive properties. As described later, the RCS is derived from the peak amplitude of the reflected signal in the frequency spectrum.

Historically, two frequency bands have been used for radar perception sensors: 24 GHz and 77 GHz. The advantages of 24 GHz sensors include moderate costs and low signal absorption. However, due to the limited maximum bandwidth of only 200 MHz and the resulting poor range resolution, these sensors are no longer suitable for perception tasks in complex, highly-automated Level 3 and 4 functions. Moreover, the European Union has mandated a gradual phase-out of the 24.25 – 26.65 GHz frequency band since 2022 [ETS20].

The higher frequency of 77 GHz systems results in smaller antennas, allowing more antennas to be integrated into a single sensor. This increased antenna density improves angular resolution. Currently, Frequency-Modulated Continuous-Wave (FMCW) radars dominate the market. These systems transmit multiple so-called frequency chirps to compute the range, velocity, and angle of reflections. Details regarding the operation of FMCW radar sensors are provided by [Jan18; Gam20]. Additionally, a brief comprehension is given in Section A.1.

2.3.3 Radar Peak Detection

The product of raw signal processing of the received signal is a discrete radar spectrum, also referred to as a Radio Frequency (RF) image.

Definition 2.9 (Radar Spectrum [Ric22])

A radar spectrum can be described as a grid of intensity values extending over discrete bins in range, angle, and Doppler velocity. These intensity values represent the power of the received signal and are indicative of the presence and strength of reflectors in the scene.

Detection candidates are extracted from the measurement by identifying peaks in the spectrum.

Datasets used for the development of radar signal processing algorithms differ significantly in the spectral representations they provide [Reb+22]. Range-Azimuth (RA) maps are divided into range and angular cells. Range-Doppler (RD) views are two-dimensional maps that represent both range and Doppler velocity information. Finally, also three-dimensional Range-Azimuth-Doppler (RAD) maps exist, incorporating all three attributes range, Doppler velocity, and azimuth angle. The RA map is commonly used for object detection, as it delivers distinct spatial information about the present reflectors. For this reason, the further explanations focus on RA spectra.

Each chirp in an FMCW measurement cycle produces an individual spectrum. Within this spectrum, amplitude values in the bins are discretely distributed across so-called cells. Using the cell's amplitude, the goal of peak detection is to identify all signal peaks

while maintaining a constant FP rate, that is, the rate of detections not associated with actual reflectors.

Fixed and Adaptive Threshold Detection

A fixed detector, such as the Neyman-Pearson detector [Ric22], applies a single, fixed threshold value T . This approach leads to a high number of FP points in radar spectra with a high noise level, while weak detections are frequently suppressed. Therefore, fixed detectors are generally unsuitable for radar application across diverse environmental conditions [Gam20].

An alternative suited for variable environments is the use of an adaptive threshold T_{adapt} . In such methods, the spatial and situational noise level $\mathcal{N}$ is estimated, and the threshold is typically set as a multiple of this noise level: $T_{\text{adapt}} = k \cdot \mathcal{N}$. This strategy ensures that the number of FP detections remains consistent regardless of the noise level in the current detection scenario. Hence, such methods for threshold computation are referred to as Constant False Alarm Rate (CFAR) [Lud13].

Figure 2.4 illustrates the advantage of noise-estimating algorithms. Two detection scenarios are depicted: one at a parking space with a relatively low noise level on the left hand side, and another in an urban environment with elevated noise levels on the right hand side. The detected peaks from both the Neyman-Pearson detector and an adaptive-threshold detector are compared. The blue line illustrates the received signal S . The green line shows the static threshold T_{fix} applied by the Neyman-Pearson detector, and the red line shows the dynamic threshold T_{adapt} that is adaptive to the noise in the respective cells. The static threshold T_{fix} is selected to detect moderately reflective objects such as pedestrians in this environment. In the environment with low noise level, the fixed threshold fails to detect even peaks with a high Signal-to-Noise Ratio (SNR). In the environment with high noise level, while the higher peak in range cell $i = 50$ is detected, the lower peak at range cell $i = 30$ does still not exceed the threshold.

Cell Averaging (CA)-CFAR

Although alternative approaches are currently under investigation, most practical applications still rely on CFAR and its specific variants for peak detection. The noise level within a single chirp sequence is generally stable, however, reflectors may naturally influence the noise floor. Most CFAR algorithms are based on CA, in which the power level of the cell under test C is compared to the average power level attributed to background noise. To estimate the background noise, the power levels of unaffected cells $\tilde{C}$ within a kernel $\mathcal{K}$ surrounding the current cell are averaged. The number of cells included in this kernel should be sufficiently large to ensure a reliable average, yet small enough to avoid contamination from neighboring peaks. The kernel is formally defined as $\mathcal{K} = (\tilde{C}_1, \tilde{C}_2, \dots, \tilde{C}_{n_{\mathcal{K}}})$, and hence contains $n_{\mathcal{K}}$ cells. The cell under test is




Figure 2.4: Schematic comparison of peak detection with a fixed threshold and a noise-adaptive threshold. The fixed threshold (green line) fails to detect peaks under low-noise conditions, whereas the adaptive threshold adjusts dynamically with the noise level.

excluded from the kernel, just as a predefined number of adjacent guard cells n_{guard} , which are assumed to be influenced by the signal of the cell under test, particularly if this cell contains a peak. The estimated noise floor is then multiplied by a threshold factor k_{thresh} to establish a margin between the signal and the detection threshold.

Given an accepted false alarm probability p_{FA} and the kernel size n_K , the threshold factor can be analytically computed. Subsequently, the detection probability P_d may be derived based on the SNR [GK88]:

$$k_{\text{thresh}} = (p_{FA})^{-\frac{1}{N}} - 1 \quad (2.4)$$

$$P_d = \left(1 + \frac{k_{\text{thresh}}}{1 + \text{SNR}}\right)^{-N} \quad (2.5)$$

The threshold value for the cell under test is computed by:

$$T_{\text{CA-CFAR}} = k_{\text{thresh}} \cdot \frac{\sum \mathcal{K}}{n_K}. \quad (2.6)$$

In some algorithmic variants, both a factor $k_{\text{thresh},\times}$ for multiplication as well as a factor $k_{\text{thresh},+}$ for addition are applied.

$$T_{\text{CA-CFAR}} = k_{\text{thresh},\times} \cdot \frac{\sum \mathcal{K}}{n_K} + k_{\text{thresh},+}. \quad (2.7)$$

At the edges of the spectrum, the bounds of the kernel are limited by the first cell r_0 , φ_0 and the last cell r_{max} , φ_{max} respectively. This impairs the quality of the noise level estimation in these regions. The principle of CA-CFAR is illustrated in Figure 2.5. In this example, the cell $C_{i,j}$ is evaluated. The number of range averaging cells, in some literature referred to as training cells, is $b_{av,r} = 3$, with one guard cell $b_{guard,r} = 1$. The size of the kernel need not necessarily be square. In this instance, the number of angular averaging cells is $b_{av,\varphi} = 5$, and the number of angular guard cells is $b_{av,\varphi} = 2$.

In this structure, the size n_K of the kernel K can be determined with Equation 2.8.

$$n_K = b_{av,r} \cdot b_{av,\varphi} - b_{guard,r} \cdot b_{av,\varphi}. \quad (2.8)$$




Figure 2.5: Concept overview of CA-CFAR, including the cell-under-test (CuT) within a radar spectrum and averaging cells bounded by guard cells.

The configuration of the averaging and guard cell parameters should be selected based on the characteristics of the radar spectrum. Modifications of the standard CA-CFAR involve alternative operations in place of averaging [Reb+22].

- *CA-Greatest-Of CFAR* computes more than one threshold value and selects the maximum among all sub-thresholds. Typically, in the one-dimensional case, the first threshold value is calculated as the average of all cells on the left-hand side of the cell under test, while the second threshold value is the average of all cells on the right-hand side. This method is particularly suitable when FPs should generally be suppressed.
- *CA-Smallest-Of CFAR* operates similarly to the Greatest-Of CFAR variant but selects the smallest sub-threshold. In this configuration, a greater number of detections are typically raised at the cost of an increased FP rate.

- *Ordered Statistic CFAR* sorts all cells within the kernel and selects a threshold based on a predefined rank k . The signal value of the cell under test is then compared to the cell at this ranked position. The rank is typically set such that $\frac{k}{n_K} \approx 0.75$, assuming the values are sorted in ascending order [Hab+08].

A detection is registered if the signal value S of the tested cell exceeds the threshold computed specifically for that cell, see Equation 2.7:

$$\text{detection} = \begin{cases} 1, & (S - T) > 0 \\ 0, & (S - T) \leq 0 \end{cases} \quad (2.9)$$

2.3.4 Radar Clustering

While pedestrian detections often result in only a few reflection points, vehicles such as cars and trucks may generate a large number of radar returns, particularly when they are close to the radar sensor. Since it is not feasible to formulate a separate hypothesis for each individual radar detection, clustering algorithms are employed to group radar points into clusters. These clusters are a preliminary stage of object hypotheses that can be used in planning components of automated functions. Clustering involves identifying radar points that exhibit shared similarities. These similarities are often spatial (i.e., the points are located near each other), but clusters can also form based on other shared attributes, such as similar Doppler velocities.

A widely used method for this purpose is the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm [Est+96]. DBSCAN is computationally efficient and requires only two input parameters ϵ and min_{pts} . DBSCAN categorizes points into core points and border points. For each radar point q , the ϵ -neighborhood is determined, which includes all points within a distance ϵ or less from q . A point is classified as a core point if it has at least min_{pts} neighbors within its ϵ -neighborhood. A border point has fewer than min_{pts} neighbors but lies within the ϵ -neighborhood of at least one core point.

A point p in this ϵ -neighborhood is said to be directly density-reachable from point q if p lies within the ϵ -neighborhood of q and q is a core point, i.e. it has a minimum number of $\text{min}_{\text{pts}} \in \text{neighbors}$. A point p_1 is density-reachable from q if there exists a chain of directly density-reachable points from q to p_1 . If there exists a point o from which both p_2 and q are density-reachable, then p_2 and q are density-connected. All points in a cluster must be density-connected to at least one other point in the same cluster. Points that do not belong to any cluster are classified as noise points. Figure 2.6 illustrates this concept with two clusters, one of size five, and one of size three. Noise points are located between the two clusters. This example uses the minimum number of neighbors $\text{min}_{\text{pts}} = 2$.




Figure 2.6: DBSCAN clustering concept showing core and border points. Opaque points represent core points, while semi-transparent points indicate border points. The number within each point denotes the number of its ϵ -neighbors.

The parameters $\min_{pts}$ and ϵ can be selected using the k -dist method, which involves plotting each point's distance to its k -nearest neighbor in the set. The threshold point in this graph indicates the maximum acceptable k -distance for a point to be considered part of a very thin cluster. Points with larger k -neighbor distances are likely to be noise, while points with lower k -distances belong to clusters. The original authors recommend setting $k = 4$, as the curve changes only marginally for $k > 4$. Once an acceptable noise percentage is determined, the threshold should be chosen to separate noise points from clustered points [Est+96].

Since its introduction, the standard DBSCAN algorithm has been extended in several ways. While DBSCAN is relatively rigid in its cluster definition, it can be advantageous to incorporate additional information about the radar points into the clustering process, when available. One such extension is Hierarchical-DBSCAN [CMS13], which provides greater flexibility. Hierarchical-DBSCAN organizes points into a cluster hierarchy by constructing a minimum spanning tree. The structure of this tree is governed by the selected so-called mutual reachability distance between two points p and q , which determines when new clusters are formed. This algorithm can also be integrated with evaluation metrics based on labeled data, such as the B^3F measure, increasing the likelihood that points with the same label are grouped into the same cluster. The applicability of this method for introducing camera information in radar clustering, thereby achieving a safer radar free space detection, is covered in [Brü+24e].

2.3.5 Radar Tracking

The final stage of radar data processing produces objects that serve as an input to the planner, representing layer 4 in the environmental perception model (see Section 2.2). Object hypotheses are initiated by clusters that are aggregations of spatio-temporally

associated radar points. These hypotheses must be maintained over multiple measurement cycles, during which new observations are integrated with existing hypothesis data. The Kalman filter offers a method for combining both the predicted evolution of a system state and a new observation of that state [Kal60]. It has been proven to deliver the optimal state estimate, i.e., an unbiased estimate with the minimum achievable covariance, under the assumption that the process and observation errors are Gaussian-distributed and the noise is white with known covariances. Figure 2.7 illustrates the concept of Kalman-based tracking.




Figure 2.7: Kalman-filter-based tracking scheme, showing the prediction and innovation steps, including the association of new observations.

The system is modeled using a system matrix A , and an input $u(t)$ affects the state via the control matrix B . The system states are observable through measurements, and this observation process is described by the observation matrix C . The influence of process noise is represented by the process noise matrix G . In the first step of the tracking process, new observations z_k are associated with existing hypotheses $\hat{x} \in \mathcal{H}_k$. If the association fails, a new object hypothesis is initiated. For each successfully associated hypothesis, its state $\hat{x}_k$ is updated in the *innovation step*. Subsequently, the updated state of the hypotheses is propagated forward in time to account for the state dynamics between two consecutive time steps t_k and t_{k+1} in the *prediction step*, applying their state model. In both states, prediction and update, not only the state vector $\hat{x}$ but also the error covariance matrix P is estimated. This covariance matrix represents the uncertainty associated with the current state estimate.

2.4 Vehicular Safety

System-wide safety can only be ensured if every stage of the system operates reliably, including perception, planning, and acting. Perception errors, in particular, cannot be compensated for by subsequent planning and acting stages. Recent research has investigated safe planning [KLD21], safe electrical actuation systems such as brake-by-wire systems [Fan+24b], and architectures that integrate both [Dor+20]. However, the transfer of these considerations to perception systems has received little attention. This section introduces safety concepts intended to ensure safety in planning and acting systems. The objective is to apply these concepts in subsequent chapters of this work to derive a similar procedure for perception systems.

A central and pivotal metric for assessing the general safety of a function is the Failure-in-Time (FIT) rate.

Definition 2.10 (FIT rate)

The FIT rate quantifies the number of failures over a specific time horizon when the system is used as intended within its ODD or exposed to foreseeable misuse. In widely adopted standards such as ISO 26262 [Int18], this time horizon is defined as 1×10^9 hours.

No system can be entirely fault-free. Nevertheless, the FIT rate must be reduced to a negligible level to ensure safe usage by customers. In the case of automated driving functions, this implies that the system's FIT rate must be significantly lower than that of human drivers. Discussing the factor $\frac{\text{FIT}_{\text{system}}}{\text{FIT}_{\text{human}}}$ leads away from technical more into ethical considerations. Furthermore, it requires extremely large fleets and long periods to conclusively demonstrate that a system outperforms human drivers [Nee19]. In this context, a system is considered functional safe if it complies within the FIT thresholds defined in the applicable standards, as discussed in the following Subsection 2.4.1 and Subsection 2.4.2.

Amongst others, uncertainty originating from the interaction of the ego vehicle with other traffic participants is a challenge for providing safety evidence for automated driving systems [GCT25]. To tackle this challenge, perception systems must reliably detect presence and motion state of these traffic participants. Miss-detection inevitably leads to an increased FIT rate.

The low speeds involved in the ODD of automated parking functions significantly reduce the consequences of functional errors. It is, therefore, reasonable to introduce higher automation levels in such low-risk scenarios before addressing more complex, high-speed situations.

The safety of automated driving functions can be dissected into two complementary aspects: functional safety, as defined by ISO 26262 [Int18], and Safety Of The Intended Functionality (SOTIF), defined by ISO 21448 [Int22]. ISO 26262 governs safety-critical electrical and electronic components by assigning safety goals and requiring components to meet specific FIT targets. ISO 21448 provides design guidelines for preventing hazards arising from functional inadequacies or foreseeable misuse during normal operation. These two standards are elaborated upon the following subsections.

2.4.1 Functional Safety

The ISO 26262 standard ensures functional safety in electrical systems by prescribing development processes and mitigation mechanisms for incidences deemed to carry a certain level of risk [Int18]. Identification of such incidents is conducted through a Hazard Analysis and Risk Assessment (HARA). This process identifies potentially hazardous events that may result in harm due to system malfunctions under operational conditions. Each hazardous event is assigned a risk classification, from which an Automotive Safety Integrity Level (ASIL) is derived. The ASIL is linked to a corresponding safety goal that must be achieved to prevent the hazardous outcome.

The risk classification entails a systematic evaluation of various aspects of driving scenarios. Since this thesis focuses on preventing hazardous situations resulting from perception system malfunctions, the HARA methodology is described in detail below.

Three key attributes define the risk associated with a given situation: to what extent can the driver or operator *control* the consequences of the system's malfunction? How frequently are traffic participants *exposed* to the situation in which the malfunction causes the hazard? How *severe* are the potential consequences if the situation is not avoided?

- **Severity:** The severity is divided into four categories: S0 = no injuries, S1 = moderate injuries, S2 = severe injuries with probable survival, S3 = severe injuries with uncertain survival.
- **Probability of Exposure:** Five exposure categories are defined: E0 = incredibly unlikely to happen, E1 = very low probability ($\leq 1 \text{ year}^{-1}$), E2 = low probability ($\leq 10 \text{ year}^{-1}$), E3 = medium probability ($\leq 20 \text{ year}^{-1}$), E4 = high probability ($\geq 100 \text{ year}^{-1}$).
- **Controllability:** The more difficult it is for individuals to control a hazardous event to avoid harm, the higher the controllability class. The classification is divided into four categories: C0 = controllable by everyone, C1 = controllable by $> 99\%$ of individuals, C2 = controllable by $> 90\%$, C3 = controllable by $< 90\%$.

These three attributes are combined to determine the ASIL. The highest level, ASIL D, is assigned when all attributes fall into the most critical categories. An automated

parking function such as AVP can reach this level for certain safety goals. For instance, during a transfer ride to a parking lot at a speed of $v_{\text{cur}} = 12 \text{ km/h}$ (E4), a pedestrian enters the ego vehicle's driving tube, but the ego vehicle fails to respond due to a brake malfunction. Since the vehicle operates driverlessly, there is no way to stop it in the short time available (C3). The pedestrian may be trapped beneath the vehicle and suffer life-threatening injuries (S3) [Int18]. According to [Int18], the target value for random hardware failures, in this case the brake malfunction, for such a safety goal is $< 10^{-8} / \text{h} = 10 \text{ FIT}$.

Even if that the brake fulfills this safety goal, the same risk still persists if the perception system operates incorrectly. The pedestrian may be insufficiently detected and the planner generates a path that collides with the pedestrian. For this reason, a similar failure rate can be assigned to errors in the perception system. Based on the considerations in ISO 26262, several aspects relevant to the criticality assessment of individual situations can be derived. These aspects form the basis of the criticality assessment developed in this work.

- *Severity* primarily depends on the current velocity v_{cur} of the automated vehicle. Other impact factors such as vehicle type, front bumper design, and roadway properties are difficult to model due to their complexity and the limited number of empirical studies investigating them [Hus+19].
- *Exposure* must be assumed to be consistently high, as in parking situations, crossing VRUs are assumed to occur continuously during the routine operation of the automated parking function.
- *Controllability* does not imply that the function must prevent all dangerous actions by other road users. Most traffic participants can avoid a collision with an automated vehicle through their natural foresight and behavior in traffic. This supports the argumentation that provoked accidents, i.e., those initiated by the conscious actions of other road users, must be considered to a certain extent in the safety assessment.

2.4.2 Safety of the Intended Functionality (SOTIF)

The standard ISO 26262, introduced in the previous section, is complemented by ISO 21448, which extends the scope of functional safety to more complex systems and addresses functionalities that are typically implemented in software. ISO 21448 ensures that SOTIF is achieved for an automated driving function.

Definition 2.11 (SOTIF [Int22])

SOTIF is the absence of unreasonable risk due to hazards resulting from functional insufficiencies of the intended functionality or its implementation.

The central principle of SOTIF is that scenarios can be categorized into four distinct areas: known and non-hazardous scenarios (area 1), known and hazardous scenarios (area 2), unknown and hazardous scenarios (area 3), and unknown and non-hazardous scenarios (area 4). The goal of safety-aware development for all components involved in the automated function is to reduce area 2 through functional enhancements, and reduce area 3 through comprehensive verification and validation processes [Int22].

This approach is shown in Figure 2.8. The primary focus of this work is on reducing area 2, while techniques are also applied in Subsection 5.3.1 to minimize area 3 through systematic scenario selection.




Figure 2.8: SOTIF-based approach to reducing both unknown and hazardous areas of the ODD to achieve SOTIF. For this work: Improvement on the sensor set and perception data processing reduces the number of hazardous scenarios (arrow I), and a higher number of situations in training and test data sets reduces the number of unknown scenarios (arrow II). Adapted from [Int22].

These concepts are directly applicable to this thesis, which aims to design a perception system that operates safely within the ODD, providing all objects to the planner for a safe operation. To keep risk at an acceptable level, the principles of functional development according to SOTIF are respected. Area 2 is mitigated by implementing safety-aware sensor data processing. Area 3 is addressed by developing an agnostic approach that considers a wide range of driving situations, including potentially unknown scenarios. Additionally, the validation dataset should include as many relevant scenarios as possible from the ODD.

To quantify the FIT rate, a HARA is required. This analysis can be conducted in several ways [Abd17]. A key finding in system-theoretical accident analysis is that inadequate control can lead to losses. The System-Theoretic Process Analysis (STPA) analysis [Lev16] aims to identify hazards that arise from system states under conditions of

inadequate control. In a subsequent step, the potential of each identified hazard to cause a loss is evaluated. According to ISO 21448, STPA can be applied as a hazard analysis approach. It is capable of identifying triggering conditions for hazardous behavior. These triggering conditions are essential for generating a test scenario catalogue [HSE22], as they initiate system responses that may lead to hazardous events [Int22]. For each identified hazard, a corresponding SOTIF requirement can be derived. In Section A.2, an STPA is performed for the automated function targeted in this work. This analysis serves as the foundation for developing test scenarios and for identifying potential losses that may occur during the operation of the AVP function.

2.5 State of Research⁴

Current research explores methods to utilize radar data more effectively by introducing advanced radar signal processing concepts, such as those based on CNNs, as outlined in Subsection 2.5.1. However, the majority of these studies focus on improving general object detection metrics, which do not account for the relevance of detected objects in relation to the driving task. To incorporate this aspect, it is necessary to apply methods that evaluate detected objects and measurements with respect to their impact on the driving situation. For this evaluation, criticality assessment metrics can be used. An overview of criticality assessment metrics is provided in Subsection 2.5.2. Due to the high noise level inherent in radar data, filtering techniques are often employed, as discussed in Subsection 2.5.3. However, filtering introduces a certain level of risk, as it may remove potentially important information. To mitigate this risk, some filtering algorithms incorporate safety-aware mechanisms that aim to preserve crucial information. Recent literature, summarized in Subsection 2.5.4, introduces safety-aware object detection metrics derived from criticality assessment methods (Subsection 2.5.2). Additionally, approaches have been proposed that adapt the processing algorithm based on the current task of the automated function. Nevertheless, detailed investigations specifically addressing radar sensors remain limited in current literature. Advancing this field by developing novel safety-aware radar data processing methods requires suitable datasets. An overview of such datasets is provided in Subsection 2.5.5.

2.5.1 Advanced Radar Signal Processing

While this thesis emphasizes the explainability of safety-relevant decisions, processing radar point clouds using deep learning algorithms has shown promising results. Two major areas of relevance are identified reviewing current publications. The first is signal processing in emerging four-dimensional radar systems, which are capable of measuring distance, Doppler velocity, azimuth angle, and elevation angle of a reflection

⁴ Section builds on [Brü+24c].

(see e.g. [Xu+21; Din+23; Fan+24a]). These systems often operate in combination with increased bandwidth, which enables finer range resolution. The second area concerns the sensor fusion of radar with features from camera and lidar systems (see e.g. [Wu+23; Gu+25]). Existing work has demonstrated fusion of 4D radar data with camera perception. In such approaches, a three-part feature extraction pipeline is often applied. Query points are projected into three-dimensional space, where they intersect with camera features. These features are subsequently processed into bounding box proposals using a combinations of classification and regression heads [FPC24].

CNNs are well known for their high performance in tasks such as object detection and semantic segmentation in images. However, applying CNNs to point clouds requires converting the point cloud into a suitable format. One common approach is to use a so-called range image, which is a spherical projection of the point cloud. Several studies propose applying a CNN to this range, before re-projecting it back into the three-dimensional space [Mil+19]. However, this method was originally developed for per-point segmentation of lidar data and requires specific adaptations to be effective for radar sensors.

CNNs can also be applied directly to radar spectra, as these resemble images in their structure and appearance. One study proposes using three parallel CNN-based autoencoders, each processing a two-dimensional spectrum constructed from two of the three radar tensors: range, angle, and Doppler velocity. A fusion module then combines the features extracted by each autoencoder. The work further suggests a training pipeline augmented with camera data to enhance learning [Gao+21].

Low-level features extracted from radar spectra can be combined with target-level features from point clouds using a CNN composed of three stages: encoding the data from a three-dimensional block that fuses point and spectral values, extracting class information, and computing class scores. Object hypotheses are subsequently obtained by clustering the labeled targets [Pal+20].

Three-dimensional bounding boxes can also be derived from radar spectra using a graph-based representation of the radar tensor. A graph neural network processes this representation, and its output is further processed using a pyramid ResNet [MKT21]. This approach outperforms traditional CNNs applied to optical⁵ representations of the radar tensor [Maj+19].

Since three-dimensional convolutions applied to single radar measurements do not incorporate causality, another work proposes a hybrid approach that combines two-dimensional CNNs with Recurrent Neural Networks [Dec+24]. The method operates on radar spectra and supports various radar image formats, such as Range-Azimuth or Range-Doppler representations. Training was conducted on the ROD2021 dataset [Wan+21a]. The framework demonstrates flexibility in handling a range of detection

⁵ Optical means in this context that the representation is based on two-dimensional images.

tasks, including semantic segmentation. Its superior performance compared to conventional CNN-based detection methods underscores the importance of incorporating temporal features and suggests the use of transformer networks, which are particularly well-suited for modeling temporal dependencies.

Transformer networks, which introduce positional encoding and self-attention mechanisms, were originally developed to overcome limitations in Recurrent Neural Networks [Vas+17]. Foundational work has extended transformers to process three-dimensional point clouds [Zha+21]. In this context, positional encoding is applied to both the attention vectors and the feature transformations. These are aggregated with a point transformer layer, which encodes local neighbors of each point. The point transformer is capable of performing both semantic and part segmentation on point clouds and can retrieve point clouds with similar shapes to a given query. Recent literature applies these transformer models to radar point clouds for semantic segmentation [Zel+23]. Given the sparse nature of radar point clouds, the standard *softmax* function in the original point transformer [Zha+21] is replaced with a Gaussian function, allowing for effective segmentation even from single radar measurements. This adaptation outperforms both the original point transformer [Zha+21] and a radar-optimized network based on multi-layer perceptrons [Sch+21b].

2.5.2 Criticality Estimation

Methods to estimate the criticality of an object for a driving task are extensively compared in [Wes+23]. The study shows that various metrics are based on the well-known Time-to-Collision (TTC), including potential TTC, time-integrated TTC, and time-to-maneuver. All these metrics measure the remaining time until a specific event, such as reaction, maneuver, stop, or collision. Some metrics partially account for the assumption that actions such as braking or evasion are taken immediately. Other metrics determine the required longitudinal or lateral acceleration necessary to avoid a collision [Jan05]. These accelerations can optionally be combined with the minimum distance to the object during the evasive action [Jun+18]. The Object Criticality Model [CM23] combines three weighted criteria: the current distance to the object, the closest distance that the vehicle and the object would reach, and the remaining time until this distance is reached. Hence, this metric also rates situations as more critical if less reaction time remains.

An online safety evaluation is proposed based on the required acceleration. Furthermore, it incorporates both the longitudinal distance required to stop the vehicle and the acceleration needed for longitudinal or lateral maneuvers. The model is capable of detecting errors in motion prediction of potential collision objects resulting from perception uncertainties [WPW22]. While these metrics correlate with the probability of a collision occurring, they often fail to account for its severity. The Pedestrian Risk Index (PRI) [Caf+11] combines the TTC of both the vehicle and the pedestrian, as well as the vehicle's time to stopping, t_{stop} , with the severity, which is primarily correlated with

the impact velocity v_{coll} . The impact velocity is computed by subtracting the velocity reduction achieved through immediate braking from the current vehicle velocity v_{cur} :

$$\text{PRI} = v_{\text{coll}}^2 \cdot (\text{TTC}_{\text{veh}} - t_{\text{stop}}). \quad (2.10)$$

Other work compares a velocity-based collision severity metric to the established Traffic Conflict Technique metric introduced in [Hyd87]. The authors argue that their critical event severity metric reflects the dangerousness of conflicts more realistically [Bag13]. The conflict severity is computed using

$$\text{Conflict Severity} = \underbrace{(v_{\text{veh}} - v_{\text{VRU}})}_{\Delta v} - \underbrace{\frac{m_{\text{VRU}}}{m_{\text{VRU}} + m_{\text{veh}}}}_{\text{weight factor}} \cdot t_{\text{stop}} \cdot a_{\text{lat}}. \quad (2.11)$$

It applies the velocity difference Δv , the available braking time t_{stop} which correlates with the TTC, the masses m_{veh} and m_{VRU} of both colliding partners, and their deceleration capabilities. Higher severity for VRUs is accounted for via a weight factor. In the case of an unequal mass distribution between the collision partners, this factor tends toward zero, and the conflict severity remains high. The success rate of this metric surpasses both the TTC-based metric and the Traffic Control Technique metric.

2.5.3 Filtering Mechanisms

Several works point out that object existence probability correlates with signal measures such as signal power, SNR, and RCS, weighting the occupancy probability in a grid based on the RCS feature of the signal [Li+20]. Further literature shows that the false alarm probability increases for targets with low SNR values [LBR22]. Hence, many approaches control the noise in radar point clouds by removing points with low signal power or related measures.

Recent research suggests removing objects from the object list based on their criticality before passing the list on to the planning task [CM24] in order to improve the planner's performance. Object lists are evaluated using the Object Criticality Model [CM23] from one of the authors' previous works, as explained in Subsection 2.5.2. Since the authors refer to literature indicating that qualitative planning depends mainly on the perception task [PKF20], the reference metric for their safety-aware filtering is the Planning Kullback-Leibler (KL)-divergence, which measures differences between two trajectories determined by the same planner. The work is evaluated on all objects in the dataset, and is thus not limited to VRUs. Results show that the Planning KL-divergence improved in most cases. However, the criticality filtering tends to remove too many objects in some challenging scenarios, resulting in worse scores.

2.5.4 Safety-aware Perception Algorithms

The literature describes the perception-control linkage, that is, the impact of perception errors on automated functions, and highlights the necessity of a holistic consideration of the function [SAC19]. As a first step, perception functions should be evaluated according to their ability to detect critical objects. This can be done by utilizing criticality assessment of objects, as described in Subsection 2.5.2, to weight each object's relevance. Perception failures concerning objects that are crucial for a safe driving task are penalized [WDE21].

A perception failure can be categorized as either task-relevant or non-task-relevant, independently of the error magnitude. Assuming the availability of a Ground Truth (GT) representation of the environment, plausible scenes are generated and each scene is evaluated with respect to its risk to the driving task. This enables the determination of the risk costs associated with perception failures. By storing scenarios with high risk costs, the concept can support the identification of scenes prone to perception errors [Ant+23].

Introducing safety considerations in perception systems is not limited to radar systems. In a safety-aware image processing pipeline incorporating a lidar sensor, external and internal cueing are applied to identify candidate regions. External cueing uses object locations provided by the lidar sensor, which are projected onto the image plane and sorted by distance and recency. Internal cues are derived from optical flow in the camera image. The ranking of detection candidates is used to reduce processing time while maintaining the safety of the perception function [LSA23]. Further literature presents an approach to determine critical pedestrian objects based on simple criticality estimation methods such as TTC and distance-based criticality, without consideration of specific planners. The object is evaluated both in terms of its threat metric and the performance metric of an image-based object detector. This dual evaluation supports the selection of relevant scenes in measurement data [Lys+22].

Other research describes a SOTIF-oriented analysis of the automated functions, identifying worst-case operational constraints. These constraints allow for deriving a minimum required perception area [SDW19]. This perception area is then adapted according to the characteristics of the sensing space and sensor capabilities, enabling the computation of position-dependent requirements on perception error. In a case study, these requirements on perception performance are illustrated for a specific operational scenario [Chu+23]. Similar work introduces the task-specific required error tolerance of an object detector, where the task is described by a simplified set of scenarios [Phi+22].

Recent research [Mei+24] introduces an algorithm for object relevance in radar detection, which is combined with a hybrid point- and grid-based CNN for radar object detection [Ulr+22] to enable safety-aware object detection. A fuzzy set is used to combine four different metrics for object criticality evaluation: distance, relative velocity, TTC, and collision distance. The model incorporates the relative motion between the potentially

critical object and the ego vehicle. The relevance of the object is determined using a fuzzy set. The experiments indicate an improvement in both the maximum Average Precision and the Intersection-over-Union metric of the object detector for objects evaluated as more relevant.

2.5.5 Datasets

Datasets are essential for testing alternative signal processing approaches and quantifying error rates. The requirements for a dataset are derived from the objectives of the experiments.

- *Radar data coverage*: There are various intermediates in radar data processing, as summarized in Subsection 2.3.2: radar spectra, radar point clouds, and radar objects.
- *Location coverage*: The data should be recorded preferably within the ODD of the AVP function, including parking lots and indoor parking garages, which exhibit unique characteristics relevant to radar applications.
- *Object coverage*: Since hazardous situations mainly arise from interactions with VRUs, datasets should preferably include these objects.
- *Situation coverage*: Scenarios should include conflicts between the VRU's intention and the ego vehicle's trajectory. Ideally, these situations require sudden changes in the driving command, such as evasive maneuvers or braking.

Multiple datasets widely used in the research community have been evaluated against these criteria. However, some radar datasets, such as the Oxford Radar Robotcar [Bar+20] and RADIATE [She+21], use radar sensors that distinguish significantly from standard automotive systems. While the standard in the automotive industry is a combination of four to six corner radar sensors mounted behind the vehicle's bumper, these datasets employ a Navtech CTS350-X [Nav25] radar system, which performs 360 degree sweeps and is mounted on the vehicle's roof. Due to the resulting discrepancies in radar data characteristics, datasets with such configurations were not considered in more detail. Other datasets using industry-standard sensor systems are listed in Table 2.3.

2.6 Discussion and Research Gaps

Current research on the processing of radar spectra and point clouds focuses on accurate per-point class-labeling [Zha+21; Zel+23]. Semantic point labeling significantly contributes to scene understanding and more precise object definitions. Additionally, semantically labeled point clouds support localization algorithms, as dynamic objects that are not consistently present across measurements can be removed, and class labels

Table 2.3: Overview of state-of-the-art radar datasets for radar-based automated driving functions.

	Radar Data	Location	Objects
CARRADA [Oua+21]	Spectra	Static, Parking Scene	Cars, VRUs
CRUW [Wan+21b]	Spectra	Static and dynamic, Parking lots, Urban, Highways	Cars, VRUs
RADDet [ZNL21]	Spectra	Static, Urban	Cars, VRUs
nuScenes [Cae+20]	3D Sparse Point Cloud	Dynamic, Parking lots, Urban, Rural	Trucks, Cars, VRUs
RadarScenes [Sch+21b]	3D Dense Point Cloud	Urban, Rural	Trucks, Cars, VRUs
View-of-Delft [Pal+22]	4D Dense Point Cloud	Urban, Motorway	Cars, VRUs

serve as useful attributes during point cloud registration [IHFZ21]. However, recent deep-learning-based algorithms introduce a black-box problem in decision-making [Hua+17]. Removing radar points solely based on their class labels — if the label suggests a radar noise class — can pose safety risks if this label is incorrect. In a first step, radar points should be processed based on their criticality instead, rather than relying on attributes such as estimated class labels. As rule-based algorithms offer guaranteed and interpretable behavior — fundamental characteristics for safety-oriented systems — this work focuses on such approaches and excludes deep learning. If deep learning is to be incorporated into safety-relevant perception tasks, the research focus must shift more towards explainability. This field is beyond the scope of this thesis.

Researchers emphasize that ensuring the safety of automated functions requires a comprehensive evaluation of perception, planning, and actuation, as well as the interfaces among these components [SAC19]. However, current studies show that criticality assessment is primarily based on object lists [Mei+24], while early sensor data, such as point clouds or image pixels, are rarely considered. Many existing approaches evaluate criticality based on the ego vehicle’s current driving state, including its velocity v_{cur} and steering angle δ [Wes+23]. These approaches cannot incorporate the planner’s intent. Hence, the procedure of evaluating criticality based on the current driving state may be less appropriate for systems driven by automated functions. Another notable gap in current research is that most criticality assessment metrics equate criticality with the likelihood of avoiding a conflict. The second dimension, namely, the potential severity of the threatening collision, is often neglected. This leads to the following research question:

(RQ1) How can the criticality of radar data be evaluated with respect to the current state of an automated parking function?

In contrast to lidar point clouds, with which they are often compared [Bil23; Mus+24], radar point clouds appear sparse and noisy. Consequently, many approaches in research and industry rely on combined mechanisms that aggregate point clouds from multiple measurements to densify the data and filter out noise points [Zha+23a; HB25]. The filtering process aims to remove points that are likely FPs, meaning the existence probability of an object at their position is low. However, the current state of research does not provide a concept that incorporates the criticality attribute into the point removal process. Some studies emphasize that perceived information is not equally significant across all situations; its relevance depends on the current driving context [WDE21; Lys+22]. These studies, however, are typically limited to later processing stages, such as the object-level representation.

The concept of identifying critical information as early as possible in the signal processing pipeline has not been widely discussed. It would help ensure that critical radar data is not inadvertently removed or disregarded by the planning function. Differentiating the treatment of relevant and irrelevant data at each step of radar data processing offers a way to balance functional safety and efficient data utilization. Current research does not sufficiently address which modifications can be induced by criticality and applied at the specific stages of the radar data processing chain. Instead, their proposed safety-aware processing methods focus on the characteristics of downstream automated functions [SDW19] or link perception failures to the current driving context [Ant+23] — with the objective of identifying challenging situations for the perception system in order to generate test scenarios for the validation process. The corresponding research question is thus formulated as:

(RQ2) How can radar data processing be altered according to the criticality assessment of the radar data to improve functional safety?

Any proposed modifications must be validated within the vehicle safety assurance process to quantify their overall impact on functional safety, including derivation of the FIT rate. While verification and validation require extensive datasets and substantial labeling effort, there may be potential for deriving the necessary information through methodical approaches. The derivation of total failure rates for perception functions and methodologies for quantifying perception failures is not sufficiently addressed in current literature. This leads to an additional research question:

(RQ3) How can the safety of radar perception functions be quantified based on methodical considerations?

This question also encompasses evaluating the approaches in (RQ2) with respect to their failure rates, thereby demonstrating their effectiveness.

The literature review has shown that radar sensor measurements are primarily used for object detection. The concept of criticality is particularly intuitive in this context, as failure to detect potential collision partners can lead to hazardous situations. However, the effects of using radar data, especially when it contains flaws, for other subtasks

within the automated function have not yet been thoroughly investigated. In addition to localization, radar-based motion estimation is an important area of application due to the high precision of radar measurements [Kel+13; Win+15]. This leads to the research question:

(RQ4) Can erroneous radar sensor data cause critical situations if it is used for motion estimation of the ego vehicle?

In Subsection 2.5.5, a literature review was conducted to provide an overview of publicly available datasets. It was found that datasets exist for many of the intermediate stages of radar data processing. Some datasets, such as RadarScenes [Sch+21b], are closer to functional application scenarios, whereas others, such as CARRADA [Oua+21], focus on general feature extraction from the radar signal. Datasets may be dynamic, recording from a moving vehicle equipped with radar sensors in real traffic, or static, recorded from a stationary platform. All datasets have share the limitation that they contain only non-critical situations, meaning that VRUs are not endangered in any recorded scene. Moreover, existing datasets lack scenarios inside parking garages, which pose unique challenges for radar sensors due to environmental effects such as heavy metal structures. These findings lead to the conclusion that existing datasets must be extended or modified to support the evaluation of safety-critical perception algorithms. Alternatively, a custom dataset could be recorded to more rigorously address the specific requirements of the intended research.

3 Area Criticality Estimation⁶

As outlined in Section 2.2, radar data provides information about objects that may potentially interfere with the policy of an automated vehicle. The goal of this chapter is to provide a basis for analyzing the incident risk associated with perception data, assuming that this data represents an object, referred to as criticality. The criticality estimation method evaluates the impact of the information on the planner's driving policy regardless of the format depending on the current processing stage. Criticality, along with other fundamental terms, is defined in Section 3.1. In the following Section 3.2, a method is proposed for constructing a spatial representation of the near-term positional allocation of an automated vehicle following its planner's policy. Finally, the calculation of a criticality score is outlined in Section 3.3. The chapter concludes with Section 3.4, which discusses potential limitations of the method. Furthermore, the premises for adapting the method to perception sensor data that slightly deviate from radars are elaborated.

3.1 Introduction

During the maneuvering process controlled by an automated parking function, the ego vehicle moves in a space shared with other traffic participants, specifically cars, bicyclists, pedestrians, and scooter riders. Compared to automated driving, maneuvering policies are significantly less strict in parking spaces and garages. To clarify, the function needs to be prepared for a free and extended use of the available drivable spaces by other traffic participants. Drivable space is defined by road markings and may often be used bidirectionally — a contrast to highway traffic, where no crossing traffic needs to be considered. Additionally, in Germany, the principle of mutual consideration applies. This increases the variety of potentially occurring scenarios. Shared use of the available space can lead to conflicts when interests of two participants overlap. The preferred approach to resolving such conflicts is to identify them timely and to negotiate a strategy to solve the conflict. This can happen in three different ways:

- *Vehicle-to-Vehicle Communication*: This method requires that both participants are vehicles equipped with a vehicle-to-vehicle communication interface and both are agents controlled by automated functions [ESE06].

⁶ Chapter builds on [Brü+24d].

- *Human-to-Vehicle Communication Mechanisms:* This method applies when one participant is an automated vehicle and the other participant is either a human or a vehicle operated by a human. Communication methods for the automated vehicle with the human may include visual or auditory signals, such as lights and sounds [Car+21].
- *Infrastructure Control:* In this case, the infrastructure manages traffic and informs both participants about right-of-way. It may communicate with the vehicle via vehicle-to-infrastructure services and with humans (e.g., by traffic lights or displays). In this scenario, the infrastructure assumes responsibility for ensuring a safe resolution to the situation [Mac+18].

Due to spontaneous behavior of VRUs, not all conflicts can be anticipated or negotiated in advance. This may lead to hazardous situations. In the event of such a situation, the automated function must⁷ recognize the conflict using its perception sensors and take immediate action to avoid a collision. An action, such as braking or evasive maneuvers, typically deviates from the primary objective of the automated function. Hazardous situations must be resolved within a short time frame.

Current research on perception methods often relies on evaluation metrics based on absolute counts of detected objects [PNS20; MKT21]. However, the failure to detect individual objects that are in conflict with the policy of the automated function can result in severe issues for the downstream function, particularly if the automated function fails to act upon these objects. Taking up the idea of safety-aware object processing, it is proposed that critical and non-critical information should not be weighted equally. This concept is largely neglected in current literature [PNS20], which does not differentiate between detection failures involving safety-relevant objects and those involving safety-irrelevant objects.

A key conclusion from such considerations is that the processing of sensor information must be adapted depending on the hazardousness of the region. Each indication for the potential presence of an object must be evaluated if its nonobservance has the potential to cause critical consequences. This evaluation needs to happen under constant consideration that the indication does not represent an object with a certain probability and may degrade the function's robustness. This object indication is called *object evidence* below.

⁷ According to ISO 26262, a very limited error rate of 10 FIT is acceptable for SAE Level 4 functions, see Section 2.4.

Definition 3.1 (Object Evidence)

Object evidence refers to any kind of information provided by a perception sensor that suggests the presence of an object. Depending on the processing stage, this information may take the form of a signal, point, cluster, hypothesis or bounding box. Object evidence may carry a probabilistic assessment regarding the likelihood of the object's existence.

To determine the relevance of object evidence, it is proposed to introduce a continuous function capable of quantifying the impact of sensor information on the current driving task of an automated function. Conventional object detection performance metrics are adequate for areas not affected by the actions of the ego vehicle. In contrast, a more sophisticated approach is required to interpret information in critical areas that may lead to hazardous situations. The key metric for this information is referred to as *criticality*.

Definition 3.2 (Criticality)

In a hazardous situation, the criticality measure assesses the potential hazard associated with a piece of object evidence. Independently of the existence probability of the object, it evaluates both the likelihood of avoidance as well as the severity of a possible incidence. The likelihood of avoidance depends significantly on the remaining time available for an action. The incident's severity increases with the expected severity of potential injury.

The criticality evaluation algorithm described in Section 3.3 computes a total criticality $crit_{total}$ by integrating these different aspects.

The preliminary work presented in this chapter is crucial to ensure the information flow throughout the perception system. Standard perception approaches typically estimate probabilities for environmental states and discard improbable hypotheses. In doing so, they neglect the role of the current policy for the perception system. Perception data may be discarded if it is irrelevant for the current driving task or driving intention. However, even low-probability data that contributes to a relevant or potentially critical hypothesis must be considered in a system designed to operate comprehensively safely. Figure 3.1 compares a conventional perception pipeline with the proposed pipeline incorporating area criticality estimation. The components related to the area criticality estimation are highlighted in blue.




Figure 3.1: Comparison of a standard architecture (top) with an architecture including area criticality estimation (bottom). The bottom architecture includes safety considerations in the perception stage.

3.2 Driving Tube Construction

Criticality of an object evidence, as defined in the definition above, depends strongly on temporal and spatial constellation as well as on the current and future motion state of both the ego vehicle and the object in the vehicle's surrounding. Traffic participants at risk of colliding with the ego vehicle must either occupy or intersect the space that the ego vehicle is expected to use. For this reason, it is reasonable to apply the ego vehicle's driving tube to determine the relevant region for criticality assessment.

Definition 3.3 (Driving Tube)

The driving tube of the ego vehicle is a spatiotemporal model representing the area that will be occupied or swept by the ego vehicle over a short time horizon in the near future. The driving tube begins from its current position, and ends with the position planned by the automated function at the end of the planning time t_{plan} . The time horizon is sufficiently long to bring the ego vehicle in a safe state, i.e. usually to stop it ($v_{ego} = 0$ m/s).

One of the simplest representations of a driving tube is a circular segment defined by the vehicle's width and its turn radius, which is derived from the current steering angle δ . The length of the circular segment depends on the current velocity v_0 . A basic assumption to apply this model is that the vehicle continues to move at a constant

velocity and with a constant steering angle. However, this concept is inherently limited and unsuitable when accelerations, decelerations, or complex trajectories such as S-curves are involved. This also implicates that this concept cannot incorporate planned trajectories.

As outlined in Subsection 2.5.4, research has examined the spatial zone necessary to ensure the safe operation of an automated function and has applied this idea to the AVP system [SDW19]. This zone, referred to as the minimum required perception zone, is derived from the full range of scenarios that the function is expected to handle. It defines the area that must be covered by the vehicle's perception sensors to ensure the detection of all objects that could pose a hazard to the automated function. While this approach helps to establish the concept of an AVP system and highlights potentially critical areas, it does not provide a continuous measure of criticality that varies with the position of object evidence and the specific situation.

Instead, as long as the vehicle is under the control of the automated parking function, it is preferable to use the output of the trajectory planner. This approach has a crucial advantage: it explicitly reflects the current intention or policy that the function is about to follow within the specified time horizon. Any object evidence item which overlays with this policy is considered in the criticality assessment. The trajectory planner operates within the time interval t_{plan} and produces a discrete set of N vehicle states $\mathcal{T} = [\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_{N-1}]$ which the vehicle should pass through within the next moments, referred to as the vehicle's trajectory $\mathcal{T}$. Each vehicle state is expressed in the *vehicle* coordinate system and is defined as $\mathcal{S} = [x, y, \gamma, v, a, \delta]^\top$, where (x, y) denote the position, γ is the alignment angle of the ego vehicle in yaw axis, v is the ego vehicle velocity, a is the ego vehicle acceleration, and δ is the steering angle⁸. Figure 3.2 gives an overview of all measures defining the ego vehicle's state.




Figure 3.2: Components of the state vector $\mathcal{S} = [x, y, \gamma, v, a, \delta]^\top$.

⁸ The positional component z_k and the alignment angle of the ego vehicle in pitch axis are excluded from the considerations since the radar perception systems in this work are not capable of measuring vertical components. Analogously, the planner computes two-dimensional states only. The valid approximation are flat surfaces for driving, as present in most parking garage levels.

All vehicle states are defined over a uniform time interval $\Delta t = \frac{N}{t_{\text{plan}}}$. This means that the planning time frame t_{plan} defines the maximum horizon for the criticality assessment. Research on controllers for automated parking functions mentions varying time horizons depending on the objective of the function, e.g. $N = 50$ samples with a sample time $\Delta t = 0.1$ s if the maximum velocity is $v_{\text{max}} = 2.9$ m/s [Kne+21]. As described in Subsection 2.1.3, this work assumes a maximum velocity of $v_{\text{max}} = 30$ km/h, driven during the transfer phase, to offer a flexible parameterization of the AVP function for worldwide markets. The planning time horizon can be computed by

$$t_{\text{plan}} = t_{\text{react}} + \frac{-v_{\text{max}}}{a}. \quad (3.1)$$

Assuming a deceleration of $a = -3.0$ m/s² and a reaction time of $t_{\text{react}} = 0.25$ s, a planning time of approximately $t_{\text{plan}} = 3.0$ s can be derived. The deceleration of modern vehicles can reach up to $a_{\text{max}} = -11.0$ m/s² under ideal conditions. However, decelerations exceeding $a = 2.00$ m/s² are perceived as uncomfortable [Int97]. Although the vehicle operates without passengers during automated parking and comfort is therefore less relevant, such decelerations may not be expected by other drivers under normal driving conditions, i.e., outside hazardous situations. The reaction time depends on the brake system and is discussed in detail in Section 3.3.

In this planning time, the vehicle can either stop or sufficiently adapt to new situations. The maximum relevant time frame for criticality assessment corresponds to the time required for both the vehicle and the object to react and decelerate to a standstill. Accordingly, the planning horizon t_{plan} is generally sufficient for the criticality assessment to identify hazardous situations and bring the automated vehicle to a stop in order to prevent an incident.

For each state S_k , a rectangle is constructed at the vehicle's position (x_k, y_k) . The rectangle is rotated by the current orientation angle γ_k . This rectangle represents the position of the ego vehicle at the time step k . It has dimensions $(l_{\text{veh}}, w_{\text{veh}})$ according to the length l_{veh} and the width w_{veh} of the vehicle which is under automated control. In addition, the vehicle's footprint is extended by a safety boundary d_{safe} . The value for the size of this boundary is chosen to account for perception uncertainties and to ensure sufficient clearance for other traffic participants.

The set of corner points from all such rectangles forms the driving tube. If the trajectory consists of N states, the driving tube comprises $4 \cdot N$ points. For computational efficiency, the driving tube is approximated using the convex hull of the points. The convex hull is the smallest convex set that contains all edge points of the driving tube. By definition, a convex set completely contains all line segments connecting any pair of two points of the driving tube. This ensures that the entire area swept by the vehicle is enclosed in the driving tube. If the line segment connecting an arbitrary point to the vehicle's current position (x_0, y_0) lies partially outside the convex hull, then that point

cannot be a member of the convex hull. This approach enables efficient point-in-region determination. The convex hull provides an estimate for the space utilized by the vehicle during its maneuvers. This convex estimate is conservative, meaning that it typically encompasses slightly more space than strictly necessary, especially in tight bends where the driving tube is considerably concave at the inner side of the bend. The concept is illustrated in Figure 3.3.




Figure 3.3: Driving tube construction using multiple rectangles and forming a convex hull around all rectangles.

This method, which represents the trajectory by the convex hull of the ego vehicle's future corner points, is particularly suitable when the driving tube calculation must be performed on limited hardware, as it is computationally efficient. A slightly modified version uses the trajectory $\mathcal{T}$ with its states $S_1, \dots, S_N$. It approximates the trajectory using spline interpolation between all states and transforms the space into a Frenet coordinate system. This approach is advantageous for determining the lateral and longitudinal displacements between the individual radar points and the vehicle's trajectory. Therefore, this method is applied in the subsequent Section 3.3.

The area criticality estimation is especially relevant if the vehicle is controlled by the automated function. However, it can also be applied in assisted driving scenarios, where the vehicle is operated by a human driver. In this case, the trajectory computed by the planner and serving as an input for the driving tube calculation is replaced by a trajectory which is estimated from the vehicle's current motion state. The ego vehicle provides its current velocity v_{veh} as well as information about the steering state. This can be either the steering wheel angle δ , the steering rack position Ψ or the front-left and front-right wheel angles δ_L and δ_R , respectively. This work employs this option whenever motion data of the ego vehicle is available and trajectory data from planners are unavailable, e.g., due to dataset constraints.

The following considerations apply to a four-wheeled, two-track vehicle with two axles and a steering mechanism on the front axle. For further information on the application of the driving tube principle specifically to four-wheel-steered vehicles, which became widespread in recent years, see, e.g., [HP16]. A single-track model is applied. The relevant quantities are depicted in Figure 3.4.




Figure 3.4: Single-track model with turn radius, illustrating the derivation of the turn radius from the single-track model (adapted from [BF18]).

Applying the Ackermann angle δ_A and the vehicle's wheelbase d_{wb} , the turn radius

$$r = \frac{d_{wb}}{\tan \delta_A} \quad (3.2)$$

can be determined. The Ackermann angle

$$\tan \delta_A = \frac{d_{wb}}{r} \quad (3.3)$$

can now be determined for both front wheel angles. Steer-by-wire systems may provide individual wheel angles δ for each wheel. If the steering system is constructed conventionally, the front wheel angles δ_L and δ_R are functions of the steering rack position: $\delta_L = f(\Psi)$ and $\delta_R = \tilde{f}(\Psi)$. Recent dynamic steering systems [BF18] feature superimposition gears that can modify the steering ratio and steering force dependent on velocity, drive mode, etc.. However, assuming a constant steering ratio, the steering rack position is a function of the steering wheel angle: $\Psi = g(\delta)$. Furthermore, the vehicle's yaw rate ω combined with the current velocity v_{veh} provides sufficient information to compute the curvature of the driving tube:

$$r = \frac{v_{\text{veh}}}{\omega}. \quad (3.4)$$

With this radius, the trajectory interface of the criticality assessment function can be filled. The time frame remains the same as for the planned trajectory. Discrete states $\mathcal{S}_k$ are computed by applying

$$\varphi_k = \frac{v \cdot \Delta t_k}{r}, \quad (3.5)$$

$$x_k = r \cdot \sin \varphi_k, \quad (3.6)$$

$$y_k = r \cdot (1 - \cos \varphi_k). \quad (3.7)$$

The velocity v_k and acceleration a_k in the individual states $\mathcal{S}_k$ are assumed to remain constant.

To summarize the preceding considerations, two methods exist for constructing the driving tube. Both rely on data typically provided by the vehicle's bus system:

1. Using data from the Electronic Power Steering (EPS), e.g. the steering wheel angle δ or the steering rack position Ψ .
2. Using the vehicle's yaw rate ω , which is often provided by the Electronic Stability Control (ESC).

3.3 Criticality Evaluation

The objective of the criticality assessment is to assign a continuous value between $[0, 1]$ to object evidence generated by perception sensors, quantifying the likelihood that the associated situation could develop into a hazardous one as a result of this evidence. In the radar application case, object evidence may include spatial attributes such as the position in Cartesian coordinates (x, y) or polar coordinates (r, φ) , as well as motion information, including velocity components (v_x, v_y) or a measured Doppler velocity v_{rel} . Additional attributes, such as the signal power or signal shape, may also be considered. These features can support classification of the detected object. The continuity of the criticality value allows for graduated responses based on the severity of the situation. As introduced in Section 3.1, two principal factors should be reflected in the criticality metric: the severity of a potential collision and the chances of avoiding the collision. The collision severity depends mainly on the released kinetic energy during the collision. Given the ego vehicle's velocity v_{veh} and its mass m_{veh} , and the object's velocity v_{obj} and its mass m_{obj} , the total kinetic energy W_{kin} in the system can be approximated by

$$W_{\text{kin}} = \frac{1}{2} m_{\text{veh}} v_{\text{veh}}^2 + \frac{1}{2} m_{\text{obj}} v_{\text{obj}}^2. \quad (3.8)$$

However, estimating the mass m_{obj} of the potential collision object from signal-level or point-level data is both complicated and unreliable. Moreover, this work addresses critical collisions with VRUs, in particular pedestrians. Collisions between two cars, or between a car and a truck during parking maneuvers, typically do not result in serious injuries but primarily in material damage. Although such collisions are relevant at higher velocities, e.g. on rural roads or highways, they are considered non-critical in this work. For collisions between the ego vehicle and a VRU, such as a pedestrian, the mass of the collision counterpart can be assumed to be significantly smaller than that of the ego vehicle, i.e. $m_{\text{obj}} \ll m_{\text{veh}}$. Therefore, the object's kinetic contribution represented by the second term of the equation is omitted. Since the ego vehicle's mass m_{veh} is constant, the severity is directly proportional to the square of the ego vehicle's velocity v_{coll} at the point of impact. For this purpose, a velocity criticality factor crit_{vel} is introduced. Instead of using the current velocity measure v_{cur} , the intended velocity of the trajectory v_{coll} at the point of the collision with the potential object is used. The square of this velocity is normalized with the square of the maximum velocity $v_{\text{max, domain}}^2$ within the functional domain. The specific value of $v_{\text{max, domain}}$ depends on functional design and is discussed in Subsection 2.1.3.

The velocity criticality factor is computed by

$$c_{\text{vel}} = \frac{v_{\text{pred}}^2}{v_{\text{max, domain}}^2}. \quad (3.9)$$

The chances of avoidance depend on the remaining decision time and the velocity reduction until the potential collision occurs. To prevent a collision, the TTC must be sufficient to confirm the object presence ($t_{\text{obs, min}}$), prepare to apply the braking system (t_{react}), and decelerate the vehicle accordingly (t_{brake}). The TTC is the sum of the required times for these three actions:

$$\text{TTC} \geq t_{\text{obs, min}} + t_{\text{react}} + t_{\text{brake}}. \quad (3.10)$$

TTC correlates with the longitudinal distance d_{long} that the vehicle travels along its planned path until it potentially collides with the object evidence. For the subsequent analysis, the relative position between the object evidence and the vehicle is described using a Frenet frame in the two-dimensional Euclidean space $\mathbb{R}^2$ with longitudinal and lateral displacement as the coordinate dimensions. The unit tangent vector s and the unit normal vector d of the trajectory $\mathcal{T}$ define the base of the Frenet frame [RC24]. The position of a point p in this Frenet frame is given as (s, d) . In this frame, the scalar s is the arc length along the trajectory segment to q , and corresponds to the *longitudinal distance* to the object evidence, denoted d_{long} . The scalar d is the length of the normal vector from the point q on the trajectory to p , and corresponds to the *lateral distance* to

the path, denoted d_{lat} . The relative velocity in the Frenet frame is given as $(\dot{s}, \dot{d})$. It is composed of the vehicle's velocity and the velocity of the object evidence. For example, if the vehicle is moving at a constant velocity v_{veh} toward a stationary point within the driving tube, the Frenet velocity components are $\dot{s} = v_{\text{veh}}$ and $\dot{d} = 0$. The concept is illustrated in Figure 3.5.




Figure 3.5: Vehicle trajectory in Frenet coordinates, showing trajectory $\mathcal{T}$ in the Frenet frame with longitudinal distance s and lateral distance d_{lat} .

The spatial attributes of the object evidence, originally expressed in either polar or Cartesian coordinates, can be transformed into the Frenet frame using the Frenet transformation. The Frenet coordinates of a point p are computed as described by the algorithm in Figure 3.6. The condition $\mathbf{p} + \kappa \mathbf{d}_i = \mathbf{S}_i + \lambda \mathbf{s}_i$ defines a linear system of equations, which is solvable if $\kappa \geq 0$ and $0 \leq \lambda \leq 1$ simultaneously. The longitudinal distance d_{long} corresponds to the first coordinate s of the Frenet frame. Analogously, the lateral distance d_{lat} from the object evidence to the trajectory equates to the second Frenet coordinate d . This method can be applied when the angle between the trajectory segments is sufficiently obtuse. Specifically, it holds when the angular difference between the tangential vector of the current segment $\mathbf{s}_i$ and the vector of the subsequent segment $\mathbf{s}_{i+1}$ is less than 10° . To improve numerical stability while searching for the intersection point, the parameter λ is constrained to the interval $[-0.05, 1.05]$.

The TTC is computed by solving the integral over the intended velocity $v(t)$ of the trajectory for the travel distance s on the trajectory to the object evidence.

$$s = \int_{t_0}^{t_0 + \text{TTC}} v(t) dt. \quad (3.11)$$

To solve this equation numerically, the continuous velocity $v(t)$ is discretized to $v[k]$ and linear interpolation is assumed between the discrete states $\mathcal{S}$ in the trajectory $\mathcal{T}$. The discrete integral, starting from the state k_0 , is now

$$s = \sum_{i=k_0}^{k_0+k-1} v[i]\Delta t + \tilde{s}_{k_0+k}v[k_0+k]\Delta t, \quad (3.12)$$

$$TTC = \frac{s_0}{v[k_0]} + \frac{s_1}{v[k_0+1]} + \cdots + \frac{s_{k-1}}{v[k_0+k-1]} + \frac{\tilde{s}_k}{v[k_0+k]}. \quad (3.13)$$

A longitudinal distance criticality factor $crit_{long}$ is introduced based on this calculated TTC. This factor reflects the remaining time for a collision-preventing action. If the evidence is so close that no evasive action can be initiated, it assigns $crit_{long} = 0.0$. Although the evidence is highly critical, a collision cannot be prevented in any case. The evidence is hence evaluated as insignificant for the signal processing. However, if the time is sufficient to observe and start to decelerate the vehicle, then the evidence is very useful as it may avert the hazard. Thus, $crit_{long}$ is set to 1.0. The criticality then decreases following a quadratic function to reflect the diminishing energy in the system. The criticality receives the value $crit_{long} = crit_{stop}$ if the time is sufficient to bring the vehicle to a complete stop, given that immediate action is taken. Beyond this point, the criticality continues to decrease linearly, identifying future-relevant regions where early object observation should begin. Figure 3.7 illustrates this behavior by plotting the criticality metric $crit_{long}$ as a function of the TTC to the evidence.

Figure 3.6: Frenet coordinate calculation from discrete trajectory states $\mathcal{S}$, illustrating the algorithm to compute the Frenet coordinates of a point p in Cartesian space using the discrete states $\mathcal{S}$ of the vehicle's trajectory $\mathcal{T}$.

Algorithm 1 Parameter calculation for per-point criticality.

```

1: Input: Trajectory  $\mathcal{T} = N$  States  $[\mathcal{S}]$ , Cartesian Point  $p = [x, y]$ 
2: Output: Point in Frenet coordinates  $p_{frenet} = (s, d)$ 
3: function CART2FRENET(Trajectory  $\mathcal{T}$ , Point  $p$ )
4:    $s \leftarrow 0$ 
5:    $d \leftarrow 0$ 
6:   for  $i \leftarrow 0$  to  $N$  do
7:     Tangential Vector  $\mathbf{s}_i = [s_x, s_y] \leftarrow \mathcal{S}_{i+1} - \mathcal{S}_i$ 
8:     Normal Vector  $\mathbf{d}_i \leftarrow [-s_y, s_x]$ 
9:     if  $p + \kappa \mathbf{d}_i == \mathcal{S}_i + \lambda \mathbf{s}_i$  then
10:        $s \leftarrow s + \|\lambda \mathbf{s}_i\|$ 
11:        $d \leftarrow \|\kappa \mathbf{d}_i\|$ 
12:       break
13:    $s \leftarrow \|\mathbf{s}_i\|$ 
14: return  $s, d$ 

```




Figure 3.7: Longitudinal criticality $crit_{long}$ as a function of the TTC of an object evidence.

The reaction time t_{react} is the elapsed time between the decision of the automated function or the driver, respectively, to initiate the vehicle's deceleration, and the moment when the brake force applies. Literature sources report times between $t_{react} = 0.15$ s [Rob25] for electromechanical brake booster systems that are commonly used in automated vehicles, and $t_{react} = 0.26$ s [Gro+08] for emergency braking systems that rely on mechanical brake boosters. During this time, the vehicle continues moving at a constant velocity.

The observation time frames $t_{obs, min}$ and $t_{obs, reg}$ depend on the applied perception function. Object evidence must be observed either in at least two varying time frames or by two independent sensor modalities. Alternatively, a single observation may suffice if the information content of the evidence combined with its criticality justifies an immediate action. However, monitoring the evidence during the braking process, with the option to withdraw the action if necessary, is also acceptable. The maximum achievable deceleration a_{max} can be calculated based on several factors: the friction coefficient μ , the gravitational constant g , the vehicle's mass m , the aerodynamic lift F_{aero} , and the retarding force F_{ret} [Pau+21]:

$$a_{max} = \mu g + \frac{\mu F_{aero}}{m} + \frac{F_{ret}}{m}. \quad (3.14)$$

Modern cars typically achieve a maximum deceleration a_{max} of about 10.0 m/s^2 [Pau+21]. Given that the maximum deceleration a_{max} is negative during braking, the braking time t_{brake} can be derived as

$$t_{brake} = \frac{v_{cur}}{a_{max}}. \quad (3.15)$$

As a third component, the lateral criticality $crit_{lat}$ is determined. The lateral criticality assesses whether the object evidence is inside the vehicle's driving tube or likely to intersect it laterally as the vehicle progresses. Object evidence within the driving tube ($|d_y| \leq \frac{w_{veh}}{2} + d_{safe}$) is assigned the maximum lateral criticality value $crit_{lat} = 1.0$. To account for uncertainties in the currently measured position and the future estimated position of the object evidence, a lateral safety margin with a distance of $d_{safe} = 2.0$ m is applied around the vehicle's driving tube. If the object evidence lies outside this boundary but near the driving tube, the lateral criticality decreases according to a continuous cubic polynomial. Given the lateral distance d_y between the trajectory $\mathcal{T}$ and the object evidence, the lateral criticality $crit_{lat}$ is computed with the following formula:

$$crit_{lat}(d_y) = \begin{cases} 1.0, & \text{if } |d_y| \leq \frac{w_{veh}}{2}, \\ 2\left(\frac{d_y - \frac{w_{veh}}{2}}{d_{safe}}\right)^3 - 3\left(\frac{d_y - \frac{w_{veh}}{2}}{d_{safe}}\right)^2 + 1, & \text{if } \frac{w_{veh}}{2} \leq |d_y| \leq \frac{w_{veh}}{2} + d_{safe}, \\ 0.0, & \text{otherwise.} \end{cases} \quad (3.16)$$

Figure 3.8 shows the lateral criticality $crit_{lat}$ as a function of the lateral distance d_y .




Figure 3.8: Lateral criticality $crit_{lat}$ as a function of the lateral distance d_y of an object evidence from the driving tube.

An optional class criticality $crit_{class}$ may be introduced when the object evidence allows conclusions to be drawn about the object class. However, in this work, class criticality is excluded from the criticality assessment, as such classification is not feasible at the radar point level or signal level.

Eventually, with each component reflecting a critical aspect of the object evidence's hazard, the total criticality $crit_{total}$ is determined by

$$crit_{total} = crit_{vel} \cdot crit_{long} \cdot crit_{lat}. \quad (3.17)$$

It can reach values in the range $[0, 1]$. Figure 3.9 illustrates the total criticality $crit_{total}$ for a scenario where a VRU crosses laterally through the vehicle's driving tube while the vehicle travels at a constant velocity $v_{veh} = 12.0$ km/h.




Figure 3.9: Simulated criticality in a VRU crossing scenario, showing the criticality evolution for a VRU crossing the ego vehicle's driving tube.

As shown, the total criticality reaches a peak of $crit_{total, max} \approx 0.16$ just before the VRU enters the driving tube. However, briefly after entering the driving tube, it falls to $crit_{total} = 0.00$, as there is no longer sufficient time to initiate an emergency action. The reason for a maximum criticality that is significantly below 1.0 is the normalization to the maximum domain velocity. In this example, the maximum domain velocity is set to $v_{max, domain} = 30.0$ km/h. The normalization with the maximum domain velocity provides a metric that relates each radar point to a point with the maximum expected criticality. Individual normalization velocities can be defined, e.g., as the velocity at which the probability of severe injury for the VRU collision counterpart exceeds a specified threshold.

A collision at 12.0 km/h releases only 16 % of the energy released in a collision at 30.0 km/h, which is reflected in the scaled criticality score. For this reason, the maximum criticality score appears to be comparatively low with $crit_{max} = 0.16$.

In addition to the criticality score, further metrics may be determined that influence the reaction to obstacles in its path. The calculations of these metrics are introduced in the following paragraph:

- Stopping Distance s_{stop} , the distance required to stop the vehicle, including the reaction time t_{react} :

$$s_{\text{stop}} = v_{\text{cur}} \cdot t_{\text{react}} + \frac{v_{\text{cur}}^2}{2 \cdot a_{\text{max}}}. \quad (3.18)$$

- Stopping Time t_{stop} , the time to stop the vehicle, including the reaction time t_{react} :

$$t_{\text{stop}} = t_{\text{react}} + \frac{v_{\text{cur}}}{a_{\text{max}}}. \quad (3.19)$$

- Reduced velocity Δv to the collision, if the travel distance to the object evidence is s and the reaction time is t_{react} :

$$\Delta v = \sqrt{2 \cdot a_{\text{max}}(s - v_{\text{cur}} \cdot t_{\text{react}})}. \quad (3.20)$$

If motion information is available within the object evidence, it should be integrated into the criticality assessment. Objects that are moving into the vehicle's driving tube present a higher risk than those that are stationary or even moving away from the driving tube. To account for this, available velocity information is used to predict the object's future position by incorporating its contribution into the current position of the object evidence. In the *vehicle* coordinate system, the position of the object evidence at the time $k \cdot \Delta t$ can be computed as

$$x_k = x_0 + v_x \cdot k \Delta t, \quad (3.21)$$

$$y_k = y_0 + v_y \cdot k \Delta t. \quad (3.22)$$

Using this equation, the future position of the object evidence can be estimated discretely for each step k . If such motion prediction is incorporated, each estimated positional value must be transformed in the Frenet coordinate system. The criticality assessment is then performed for each updated Frenet position (s_k, d_k) . The maximum criticality value across all steps k is used as the final result. This approach enables the forecasting of object evidence which is likely to become more relevant in the near future.

The concept is illustrated in Figure 3.10.

Radar sensors provide a measurement of the Doppler velocity v_{rel} . This corresponds to the relative radial velocity between the reflective object and the radar sensor. It is important to emphasize that the object may also possess a hidden component, the tangential velocity v_{tang} . This component is orthogonal to the radial velocity component




Figure 3.10: Integration of VRU motion into the criticality metric by propagating the object evidence's position p_k using motion information (v_x, v_y) . The states $\mathcal{S}_k$ represent the ego vehicle's trajectory.

and can therefore not directly be observed. It can only be inferred by observing the object over multiple measurement cycles. If this observation is not done and criticality is determined based on a single frame, this tangential velocity is assumed to be $v_{\text{tang}} = 0$ and following simplified equations apply:

$$\varphi_0 = \tan^{-1} \frac{y_0 - y_{\text{sensor}}}{x_0 - x_{\text{sensor}}}, \quad (3.23)$$

$$x_k = x_0 + \cos \varphi_0 \cdot v_{\text{rel}} \cdot \Delta t, \quad (3.24)$$

$$y_k = y_0 + \sin \varphi_0 \cdot v_{\text{rel}} \cdot \Delta t. \quad (3.25)$$

By applying these equations, the position of the evidence for future states $k = 1, \dots, N$ can be determined. This concept is applied within the criticality assessment to account for dynamic changes of ego vehicle and VRUs in the environment. The advantage of this approach is that the dynamic model is derived from only a single radar frame.

3.4 Discussion

This chapter presents a concept for assessing the criticality of so-called object evidence in relation to the currently intended driving maneuver. It incorporates both the current motion state and the future driving intention of the automated function in the shape of the trajectory output to determine whether a potential object interferes with the immediate driving policy. It models the occupied space of the vehicle as a driving tube. The Frenet coordinate system enables efficient computation of the remaining distance to the object evidence along the trajectory and facilitates the determination of whether the

object lies inside or outside of the driving tube. Future positional changes of the object evidence can also be incorporated by propagating its estimated motion and determining and evaluating the maximum total criticality value across all predicted positions. As the driving tube is considered as a convex shape, the sheering behavior of the vehicle is incorporated adequately. The convex shape may lead to a consideration of object evidence that is close to the driving tube, but not inside of it. This design decision is made consciously to ensure that the considered set of object evidence is complete and does not omit FNs.

The criticality assessment is designed to make use of all available information to evaluate the evidence's relevance as accurately as possible. It can be applied to any type of sensor data that contains at least positional attributes in the *vehicle* coordinate frame. This is a major difference to conventional approaches, which typically assess relevance based on an object list. As outlined in Subsection 2.3.2, an object list is the final stage of the perception system's processing pipeline. It thus contains only a reduced, processed amount of information compared to the information content of the sensor's raw data. As a result, the criticality assessment could only change the behavior of the planning and actor system, but is not able to adapt processes in the perception system. By contrast, the proposed method allows criticality to be evaluated at earlier stages of signal processing. This opens the possibility of incorporating additional sensor systems as information sources, particularly when the existence of an object evidence is uncertain or disputed by the primary sensor.

This uncertainty is usually incorporated into the filtering and processing pipeline of perception sensors. Uncertainty is interpreted as the likelihood that an object actually exists where object evidence is observed. To enhance their robustness, many existing approaches rely on the estimated existence probability of objects, such as cars or VRUs, as a primary indicator of object relevance and remove object evidence that is unlikely to correspond to an actual object. However, such methods do not distinguish between objects of different criticalities. As a result, a critical object whose existence is considered less probable may be disregarded. This proceeding neglects the principle that functional safety is defined as the frequency of injury-causing events to passengers or other road users in a given timeframe. If a collision with an object is likely to result in injury, even objects with a low existence probability must be considered to satisfy functional safety requirements. Similarly, incorrect classifications can also lead to hazardous situations if filtering is performed based on the semantic label of information. In this case, object evidence representing a VRU could erroneously be classified as clutter and deleted. The proposed framework addresses this shortcoming by incorporating criticality into decision-making at early stages.

The proposed area criticality estimation utilizes the policy and motion state of the automated vehicle as inputs to evaluate the relevance of arbitrary sensor information. This methodology addresses (RQ1), which concerns the relationship between perception data relevance and the downstream planner. Since the criticality evaluation function employs measurement parameters such as position (x, y) and Doppler velocity v_{Dopp} ,

which are already available in radar raw data, it can assess the criticality of radar data at various processing stages. Besides, a class-level criticality measure $crit_{class}$ can incorporate information about an object's class derived from other sensor modalities, such as cameras. However, the use of cameras for criticality assessment is not investigated in this work, as the described methodology is applied only when the presence of an object is uncertain, that is, when no other sensor modality provides confirmed information about the object's presence. This condition implies that no other object information than the object evidence of the radar is available.

The criticality assessment may exhibit FN scores, meaning that certain information is classified as non-critical even though it may require continued observation to derive a potentially urgent action. A common case of FN criticality is imprecise lateral positioning, particularly in a large lateral distance to the vehicle. Another source of FN results is incorrect motion information, which may cause incorrect propagation of the object's future positions. Each occurring FN object evidence leads to the case that this information is lost prematurely. A safety hazard arises if no additional object evidence appears in subsequent sensor cycles. Although the likelihood of an object being present only during one frame is low, the FN case needs to be prevented proactively. A further risk is the "one cancels another" characteristic of the individual criticality components. The total criticality becomes zero if any of the following conditions hold:

1. The intended velocity of the vehicle at the collision point with the object evidence is $v_{coll} = 0$.
2. The object evidence appears too close or too far away from the vehicle for an action to be deemed reasonable.
3. The object evidence seems to not appear to enter the vehicle's driving tube.

If even one of these conditions is met, the object evidence is classified as non-critical. Incorrect conclusions can be partly prevented by storing information with at least two non-zero criticality components. Conversely, each noisy signal or point that falls inside of the driving tube is treated as critical, as it is indistinguishable from actual object evidence originated from real reflections. This means that FP instances make up the largest proportion of the identified object evidence instances. Consequently, the precision of the criticality assessment may be extremely low, potentially approaching zero. If a high number of object evidence instances trigger emergency actions, the system's robustness is significantly compromised. The task of the downstream processing steps in the perception system is to handle the additional information of the criticality assessment method.

However, latent contextual information, which may be crucial for determining the relevance of the information, is currently not integrated in the assessment. For example, it is often difficult to determine from radar data alone whether an object is traversable. This uncertainty arises because the reflection may originate from the ground, a pole,

or another object outside the class of VRUs. As a result, every piece of object evidence must be handled with equal caution.

With minor adaptations, the fundamental concept of criticality assessment can be applied to various stages of radar data processing. In the following chapters, these processing stages and the possible actions that may be triggered by critically assessed object evidence are formally defined and discussed.

4 Processing Adaptations in Radar Spectra⁹

This chapter discusses a framework for safety-aware processing of radar spectra, incorporating criticality assessment of RA cells into a modified variant of CA-CFAR. This framework, described in Section 4.2, fosters the detection of weak peaks in critical areas of the radar spectrum. This enables the detection and confirmation of weakly reflecting objects that may be endangered by the trajectory planned by the automated function. As outlined in Subsection 4.2.4, object candidates detected in images can serve as region proposals for safety-aware peak detection in RF images. Section 4.3 evaluates this framework using the ROD2021 dataset, which features spectra from various types of VRUs, e.g., pedestrians and bicyclists. Section 4.4 concludes this chapter by discussing its main findings.

4.1 Information Representation in Radar Spectra

As can be seen in the exemplary processing pipeline presented in Subsection 2.3.2, the RF image constitutes one of the earliest steps in radar processing providing information about object presence. Conventional radar data processing mechanisms extract points from the RF image, however, a significant proportion of the information provided by the radar sensor is already lost in this step. While Analog-to-Digital Converter (ADC) data represents the pre-stage of the radar spectrum, apart from losses due to discretization, the spectra themselves include all information about the received radar signal. The algorithms applied to the spectrum extract points corresponding to signal peaks in the spectrum, thereby generating a point cloud for downstream algorithms. Each point in this point cloud is characterized by three fundamental attributes: range r , azimuth angle φ , and Doppler velocity v_{rel} .

The key indicators for the performance of peak detection are the number of type I errors (FP) and type II errors (FN).

Definition 4.1 (FP and FN detections in radar spectra)

A FP detection in radar spectra is a peak that is not associated with an actual reflector.

A FN detection is a present reflector that lacks a corresponding identified peak in the radar spectrum.

⁹ Chapter builds on [Brü+26].

Precision and *recall* are two fundamental performance metrics for object detection directly derived from the number of FP and FN reflectors counted across a set of measurement frames.

Definition 4.2 (Precision and Recall)

The precision metric quantifies the proportion of detections in the detection set that correspond to actual (i.e., existing) objects.

$$\text{Precision} = \frac{TP}{TP + FP}. \quad (4.1)$$

The recall metric quantifies the proportion of actual objects that have been correctly detected.

$$\text{Recall} = \frac{TP}{TP + FN}. \quad (4.2)$$

Both metrics are essential for the parameterization of a robust object detection algorithm: low precision results in reduced robustness, and low recall results in an unreliable detection function. Hence, both metrics are commonly combined into the so-called F_1 -score, computed by

$$F_1\text{-score} = \frac{2 \cdot TP}{2 \cdot TP + FP + FN}. \quad (4.3)$$

A perception function can be parameterized towards increasing the recall, which means to accept a higher number of FP, or towards increasing the precision, which means to accept a higher number of FN. The maximum F_1 -score typically offers a suitable compromise between the precision and the recall value.

Figure 4.1 shows a qualitative ratio between recall and precision, including the maximum F_1 -score.

Moving to the right of the curve increases recall at the expense of lower precision, i.e., robustness. Conversely, robustness can be increased by moving to the left and accepting more FNs, which reduces safety. Better data quality, achieved for example through performant sensors or the application of sensor fusion, can shift the curve significantly toward the upper-right corner. Improved data-processing methods, such as the adapted treatment of information according to its classification [Lys+22], can also yield positive effects. The figure compares performance under weak sensor data (red line) with performance under average data quality (blue line). However, the curve still requires




Figure 4.1: Relationship between recall and precision, maximum F_1 -score of each curve, and possible improvement of the perception function through a safety-aware system.

balancing between recall and precision. The objective of the methods proposed in this chapter and the subsequent chapters is to improve the recall rate, particularly in critical areas, while maintaining robustness in non-critical areas. The advantage is that recall can be enhanced in regions where FNs have a significant impact on functional safety without compromising the function's overall robustness. A qualitative representation of such a safety-aware function is depicted as the dotted blue line in Figure 4.1.

Moreover, the ability to separate peaks is relevant. In discrete cells, two peaks at cells a and b can only be distinguished if there exists a cell c between a and b such that the amplitude satisfies $S(c) < S(a)$ and $S(c) < S(b)$. However, if the differences between $S(a)$, $S(b)$, and $S(c)$ are small, the two peaks may merge into one, and the corresponding reflectors are not separated.

The fundamentals of CA-CFAR, which is a common peak detection algorithm for radar spectra, are outlined in Subsection 2.3.3. It estimates the general noise floor of the radar spectrum by averaging the discrete measurement cells. A multiple of this estimated noise floor is then used as a threshold for peak detection. The peaks detected in these spectra are transferred to a point cloud.

In this work, the regular CA-CFAR threshold is extended by a merging strategy for detections from single RA cells. If a peak spreads across multiple cells, this often results in an accumulation of detections. To reduce the redundancy, Connected Component Labeling (CCL) is applied, collapsing groups of connected detections into single representative points. CCL is primarily known from image processing and also called blob extraction. It identifies connected components such as groups of neighboring cells with *detection* status. The von Neumann neighborhood is used, meaning that only the cells with a neighboring edge are considered neighbors. For a cell C with the range index r_i and the angular index φ_j , the four cells with following range and angular indices

are neighbors of C : (r_{i-1}, φ_{j-1}) , (r_{i-1}, φ_{j+1}) , (r_{i+1}, φ_{j-1}) , (r_{i+1}, φ_{j+1}) [KR96]. Multiple cells forming a neighborhood correlation are grouped into a single blob. During iteration over all cells with a *detection* status, the following labeling rules apply:

- If a cell is not a neighbor of any of all previously visited cells, it initializes a new blob.
- If a cell is a neighbor of an already labeled cell, it is assigned to the same blob.

To determine the center point or cell of a blob $B = (C_1, \dots, C_{n_{\text{cells}}})$, one method involves calculating the mean spatial center r_B, φ_B of the blob B :

$$r_B = \frac{1}{n_{\text{cells}}} \sum_{i=0}^{n_{\text{cells}}} r(C_i), \quad (4.4)$$

$$\varphi_B = \frac{1}{n_{\text{cells}}} \sum_{i=0}^{n_{\text{cells}}} \varphi(C_i). \quad (4.5)$$

However, this method may fail to locate the actual signal peak, especially if the object is spatially spread within the spectrum. This can directly lead to a safety hazard if a critical object such as a VRU is closer than determined by the peak detection algorithm. For this reason, an alternative strategy is proposed: selecting the cell with the maximum signal per blob over all chirps. In this case, the range attribute of the blob point is $r_B = r(\arg \max_{C \in B} S(C))$ and the angular attribute is $\varphi_B = \varphi(\arg \max_{C \in B} S(C))$.

The peak detection threshold depends solely on the noise level in a single time frame and does not consider previous detections. Furthermore, the threshold does not adapt to critical regions. As a result, weak reflectors may go undetected by the CFAR threshold algorithm because the amplitude of their reflected signal remains below the threshold. This information cannot be recovered in downstream clustering and tracking stages, as these steps rely on the point cloud to generate object hypotheses. Therefore, it is crucial to ensure that potential information is not lost during the peak detection step.

Three distinct types of errors can occur, which are illustrated in Figure 4.2. The first type is a FP, meaning that a detection point is raised despite the absence of an actual object. This error reduces the robustness of the downstream perception function. A FP may cause the ego vehicle to brake unexpectedly, which can lead to collisions with other vehicles. However, the injury risk during such collisions is significantly lower since humans are vehicle passengers and thus protected by passive safety. The second type is a FN, meaning that no detection is raised for an existing object. Minimizing this error type as much as possible is essential, as it constitutes a root cause for hazardous situations. If the objective is to achieve a high functional safety, the reduction of FN errors has a higher priority than the reduction of FP errors. The third error type is the merging of two objects into a single point. In an unfavorable situation, a strongly

reflecting object may absorb weakly reflecting objects in its vicinity. When the two objects are in close proximity, the impact of this merger on the downstream automated parking function is limited. However, such cases should still be avoided whenever possible.




Figure 4.2: Error types in CA-CFAR threshold detection.

4.2 Subsequent Radar Peak Search

The standard CFAR algorithm estimates the noise level and detects peaks that exceed this level in the radar spectrum. Consequently, a peak must exhibit a minimum multiple of the noise power to be detected, which must exceed the threshold factor k_{thresh} . As outlined in Subsection 2.3.2, the power of the received signal depends on the object reflectivity or RCS. This section discusses required threshold design to identify VRU reflections in the spectrum and the application of the cell criticality assessment.

4.2.1 Concept of Threshold Adaptation

The SNR of a pedestrian is quantified in the following to motivate the discussion. According to existing studies, the RCS value σ of a pedestrian for 79 GHz radars ranges between $[-6 \text{ dB m}^2, 9 \text{ dB m}^2]$ with a mean of approximately 3 dB m^2 [MNK12], and in the range between $[-5 \text{ dB m}^2, 0 \text{ dB m}^2]$ depending on the look angle, respectively [CC14].

The SNR value can be determined by applying Equation 2.3 and the operating figures of the respective radar sensor. For the Texas Instruments TI AWR1843 used in the experiments, these operating figures are listed in Table 4.1 of Subsection 4.3.1. A pedestrian target ($\text{RCS} = 0 \text{ dB m}^2$) at a distance of $R = 10 \text{ m}$ produces a peak with an SNR of 4.550 dB, corresponding to a linear factor of 2.851. If the threshold factor is set higher than this value, an average pedestrian peak will fall below the CFAR threshold and will not be detected. On the other hand, the threshold shall not be overly sensitive to the noise standard deviation σ_N . For clarification, Figure 4.3 illustrates the noise floor and a received signal containing two peak. Peaks stand out from the noise by a

specific margin relative to the noise floor. When signal power is plotted on a logarithmic scale, the SNR can be read directly. In the figure, the green arrows indicate the distance between the signal peaks and the noise floor.




Figure 4.3: Signal with two peaks that significantly exceed the noise floor.

Conventional CFAR algorithms are designed to configure the detection threshold in a way that balances both FP and FN rates. However, this standard approach overlooks a crucial aspect: it does not account for the relevance of a cell with respect to the downstream automated parking task. Every cell's threshold is determined in the same way, although some cells belong to the action space of the parking function, whereas others are too far away from the vehicle to create any significant impact. Overlooking objects may create hazardous situations in the former cells, while the consequences may be negligible in the latter cells. Hence, for cells belonging to the driving tube of the parking function, it may be justified to tolerate a higher number of FP cells in exchange for reducing the FN rate.

Secondly, due to the highly dynamic nature of both the noise and the reflection processes, the amplitude of a signal peak shows a high variance. This can result in an unstable detection scenario, in which the peak of the reflection signal of one object sporadically crosses above and falls below the threshold in consecutive radar spectra. As object evidence can be interpreted as a peak in the spectrum, it can systematically be identified by lowering the threshold around cells where a peak was detected in previous frames. In this way, more consistent information can be accumulated, and the validity of a peak can be confirmed more reliably.

The following subsection adapts criticality estimation to the specific use case of peak

detection in radar spectra. Based on this criticality evaluation and the principles of standard CFAR peak detection, an adaptive detection algorithm is proposed. This algorithm has the potential to identify weak reflections which barely exceed the noise level with a greater reliability.

Figure 4.4 compares this architecture with a standard peak detection pipeline. In critical areas, potential candidates are identified using a weak peak detection approach. These peak candidates are then verified based on temporal consistency across preceding radar spectra. In the final step, the verified peak candidates are combined with the peaks from the standard peak detection.




Figure 4.4: Comparison of the standard peak detection algorithm with the criticality-aware peak detection algorithm that integrates another peak detection cycle.

4.2.2 Criticality Evaluation on Single Cells

The evaluation of area criticality can be applied effectively to cells in radar spectra, provided that both the range r and the azimuth angle φ are known for each cell. Under this condition, the longitudinal criticality factor $crit_{lon}$ and the lateral criticality factor $crit_{lat}$ can be derived. The velocity criticality component $crit_{vel}$ defined in Section 3.3 depends solely on the spatial position of the cell and not on its Doppler velocity. Therefore, the method is applicable for RA spectra. The polar coordinates of a cell $C_{i,j}$ can be determined based on its range cell index i and the angle cell index j . The absolute range r_i is computed by multiplying the range resolution Δr with the cell index i . The absolute range r_i of the cell i can be derived by rearranging the equation for the mixing frequency f_M in terms of the frequency slope $\dot{f}_{transm}$:

$$f_M = 2 \cdot \Delta r \cdot i \cdot \frac{\dot{f}_{\text{transm}}}{c}, \quad (4.6)$$

$$r_i = \Delta f \cdot i \cdot \frac{c}{2 \cdot \dot{f}_{\text{transm}}}. \quad (4.7)$$

The angle φ_j in radians can be computed with

$$\varphi_j = -\frac{\pi}{2} + \frac{j}{n_{\text{cells}}} \pi. \quad (4.8)$$

The cartesian coordinates are computed from the polar coordinates with

$$x = r_i \cdot \cos(\varphi_j), \quad (4.9)$$

$$y = r_i \cdot \sin(\varphi_j). \quad (4.10)$$

Instead of calculating the criticality for each object evidence, the property of the discrete RA grid is utilized. For each cell $C_{i,j}$ with $i \in \mathbb{N}; 1 \leq i \leq n_{\text{rangebins}}$ and $j \in \mathbb{N}; 1 \leq j \leq n_{\text{anglebins}}$, a criticality score is determined. Figure 4.5 illustrates an exemplary case. The number of range bins and angular bins is equal: $n_{\text{rangebins}} = n_{\text{anglebins}} = 128$. A straight trajectory at a constant vehicle velocity of $v_{\text{veh}} = 4.0 \text{ m/s}$ is assumed. The velocity criticality component is set to $\text{crit}_{\text{vel}} = 1$.




Figure 4.5: Exemplary criticality assessment of range and angle cells for a trajectory moving straight ahead at a velocity $v_{\text{veh}} = 4.0 \text{ m/s}$.

This criticality assessment enables context-dependent processing in the next step by allowing specific measures to be taken based on the criticality of the respective cell. Standard CA-CFAR is applied to all cells not deemed critical in the current situation (i.e., when $\text{crit}_{\text{total}} = 0$).

4.2.3 Threshold Reduction in Critical Cells

Region(s) of Interest (RoI) are defined as spatially constrained areas in which the presence of an object should be examined in greater detail. In these RoI, the criteria for interpreting available information as evidence of object presence are relaxed, meaning that the signal power of the evidence needs to be lower to still be considered. The peak detection step extracts a set $\mathcal{D}_{\text{safe}}$ of radar points, also referred to as detections, from a three-dimensional RA spectrum. In this context, the detection threshold can be reduced within RoI to enhance the visibility of weak detections with low SNR.

The threshold reduction shall be applicable at the smallest available spatial resolution, which in this case is the RA cell. Secondly, it must support superimposability, that is, the ability to apply multiple reduction terms to a number of RoI concurrently, resulting in a total differentiable threshold function T^* .

In the absence of motion information, object evidence is assumed to persist in the same spatial cell. However, radar specific effects can cause that the peak variates in the cell location. The range $r(t+1)$ and angle cell $\varphi(t+1)$ in the next time step are modeled as being normally distributed around the original cell $C_{i,j}$, where the peak was previously detected, with the standard deviations σ_r for the range cells and σ_φ for the angular cells. If a radar system observes a radar reflector that is located at the constant position (r, φ) , the central limit theorem accounts for the measured range r_{meas} and azimuthal angle φ_{meas} . It is

$$r_{\text{meas}} \sim \mathcal{N}(\mu_r, \sigma_r), \quad (4.11)$$

$$\varphi_{\text{meas}} \sim \mathcal{N}(\mu_\varphi, \sigma_\varphi). \quad (4.12)$$

To conclude, if information was measured in a cell $C_{i,j}$ at time t , it is assumed that the position of the affected cells in the next time step can be modeled by a Gaussian distribution centered at the cell $C_{i,j}$. Hence, to find consecutive peaks, the threshold reduction must follow the shape of a reversed Gaussian function with its expected value located at cell $C_{i,j}$. This term reducing the threshold function is referred to as the *Gaussian Dimple* Δ .

Definition 4.3 (Gaussian Dimple Δ)

A Gaussian dimple is a three-dimensional Gaussian function that acts as a subtractive term, reducing the CA-CFAR threshold locally. The position of the dimple is defined by the cell at which its peak $(i_{\text{max}}, j_{\text{max}})$ occurs. The width of the Gaussian dimple is specified by w_r in the range dimension and w_φ in the angular dimension. Its depth is controlled via the parameter d .

$$\Delta(i, j) = d \cdot \exp\left(-\frac{(i - i_{\max})^2}{w_r} - \frac{(j - j_{\max})^2}{w_\varphi}\right). \quad (4.13)$$

The depth d of the Gaussian dimple can enhance the detection of weak peaks. This depth should be increased in proportion to the criticality of the cell in which the dimple is placed and the certainty that an object inducing a peak is indeed present at that position. The range width parameter w_r of a Gaussian dimple should be chosen according to the estimated variance in range when measuring a point target. Similarly, the angular width parameter w_φ corresponds to the expected variance in the angular dimension. A separate dimple function is generated for each RoI. This yields a set of dimple functions $\Delta_1, \dots, \Delta_{N_{\text{RoI}}}$, corresponding to the N_{RoI} RoIs. For each cell, the combined reduction threshold function Δ^* is computed as the maximum reduction over all individual dimple functions in the set.

$$\Delta^*(i, j) = \max_k \{\Delta_k(i, j)\}. \quad (4.14)$$

The reduced threshold T^* is then calculated by subtracting this combined reduction function from the original CA-CFAR threshold: $T^* = T - \Delta^*$.

Figure 4.6 illustrates an example based on a RA spectrum. A signal peak S is present at the cell $C_{i=51, j=51}$. However, as can be seen, this peak does not exceed the threshold T determined by the standard CA-CFAR method, which is shown as the red dotted line in the right-hand figure. When a Gaussian dimple is applied at the peak position, the standard threshold T is reduced there, resulting in T^* , shown as the red, solid line. The position of the Gaussian dimple can be derived by prior object evidence obtained from earlier measurements. In this example, the signal S , represented by the blue solid line in the right-hand figure, now exceeds the reduced threshold T^* at the peak, thus confirming the detection of the object.

A constructive method for applying this strategy is outlined below. Every peak identified by the CA-CFAR detection algorithm serves as evidence of a potential object in this area. However, depending on the configuration of averaging cells, guard cells, and threshold level, this peak may represent a FP. Such detections should be disproved. Conversely, TPs may be confirmed if the object evidence intensifies across additional frames or channels. To approve or disprove object evidence, it is advisable to incorporate additional sources of information that can help mitigate detection errors. In the context of radar peak detection, such information may either originate from additional radar chirps collected during the same measuring step, or spectra data from subsequent measurement steps. Both sources contribute individually to investigating object evidence more comprehensively.




Figure 4.6: Effect of reduction terms in the form of dimples for detecting weak signal peaks.

Individual radar chirps may contain random disturbances, provoked by stochastic processes, which are not necessarily present in the same form across other chirps within the same measurement. Examples of such disturbances include thermal noise, quantization noise [Ali+24], or random dropouts.

In contrast, certain effects may consistently impair all radar chirps within a single measurement frame. In such cases, multiple measurement frames can be leveraged to confirm the object's presence based on accumulated evidence. The impairments on detection may not persist in the next measurement frame if they are temporary but deterministic in nature, such as clutter effects from the ground, interference from other radio emitters [All+19], or multipath reflections.

As stated before, it is sometimes beneficial to aggregate information from multiple measurement frames to decide reliably regarding the presence of an object. However, this approach incurs the drawback of increased latency, as additional time is needed to collect enough frames for confirmation. This delay can be costly, particularly when timely action is required to avoid a potential collision with the detected object. A solution to this trade-off is to recapture past radar spectra in light of newly detected object evidence. This information from past spectra can either prove or falsify the object evidence. If such newly detected object evidence in the form of a peak in a spectrum is observed for the first time, it may indicate either a FP detection, or that the reflection was too weak to exceed the threshold in the preceding spectra before.

The framework includes maintaining a history of measurement frames for subsequent, backward peak search in the past spectra. Specifically, the past spectra within a defined time horizon t_{hist} are utilized. Given a measurement frequency f_{meas} , the number of measurement frames considered is $N_{\text{hist}} = t_{\text{hist}} \cdot f_{\text{meas}}$. The current frame is denoted by the index τ , and the preceding frames are indexed as $\tau - 1, \dots, \tau - N_{\text{hist}}$ accordingly.

Initially, any detected peak is treated as non-confirmed evidence. It generates a RoI around the peak cell and its adjacent cells. The CA-CFAR-thresholds from preceding measurement frames are then recaptured. Gaussian dimples are applied at each RoI, and the threshold detection procedure is repeated. Since these measurement frames span a time interval t_{hist} , it is necessary to consider potential object motion over time. Accordingly, the RoI is enlarged proportionally to the time difference Δt , accounting for increasing uncertainty about the object's position in older frames. This time difference is calculated for each preceding frame k with $1 \leq k \leq N_{\text{hist}}$: $\Delta t = \frac{k}{f_{\text{meas}}}$. Simultaneously, the corresponding dimple becomes shallower for higher Δt , reflecting the decreasing probability of an object presence at one spot of the investigated region.

As described previously, Gaussian dimples are shaped using the parameters d , w_r , and w_φ . To summarize, the following relationships between the time difference Δt (the difference of the preceding measurement frame k and the current measurement frame), also called the age of the frame, and the dimple configuration, can be derived:

- $\Delta t \uparrow \Rightarrow w_r \uparrow, w_\varphi \uparrow$: Consideration of position uncertainties with higher measurement age
- $\Delta t \uparrow \Rightarrow d \downarrow$: Probability to find object evidence in the formerly observed area decreases with higher measurement age

Figure 4.7 illustrates exemplary how the reduced threshold evolves over time. The reduction was applied to five preceding measurement frames ($N_{\text{hist}} = 5$). It can be observed that the target is only detected in measurement frames with lower age at the frame indices $\tau - 1, \tau - 2, \tau - 3$. Nevertheless, the threshold reduction area was expanded to remain receptive to dimples in more distant cells. In this figure, only one received signal is plotted due to better overview. The proposed method always compares the corresponding received signal $S_{\tau-k}$ with the reduced threshold for this signal $T_{\tau-k}^*$.

To combine the robustness of standard CA-CFAR detection with the sensitivity of threshold reduction for identifying weak detections, it is proposed to perform a two-stage peak detection process.

1. The first peak detection applies the standard factor k_{thresh} . The goal of the first peak detection is to detect strong peaks which can be considered confirmed detections due to their high SNR. These detections are referred to as *trusted detections* and are added to the set $\mathcal{D}_{\text{trust}}$.
2. The second peak detection applies a reduced factor $k_{\text{thresh},-} < k_{\text{thresh}}$. It identifies weaker peaks that are more likely to arise from noise and are less likely to correspond to real objects. Since these detections are not considered as trusted detections yet, they are termed *detection candidates* and collected in the set $\mathcal{D}_{\text{cand}}$.




Figure 4.7: Gaussian dimple for threshold reduction, illustrating the creation of Gaussian dimples at a RoI to reduce the threshold. The dimple becomes shallower and wider for thresholds T^* belonging to older measurement frames.

The detection candidates are subject to a voting process. In this process, the threshold from each preceding spectrum is lowered via Gaussian dimples in regions corresponding to detection candidates. Gaussian dimples are only applied in areas that have previously been assessed as critical, and their shape may be modulated according to the criticality score of the corresponding detection cell causing the dimple.

Definition 4.4 (Detection in the RA spectrum)

A detection $\mathcal{D}$ is a peak in a two-dimensional or three-dimensional RA signal spectrum that exceeds a threshold function T . It is characterized spatially by the index (i, j) of the cell $C_{i,j}$ in which the signal peak forms and the assessed criticality of this cell. If the detection candidate occurs in the context of safety-aware processing, this candidate has the number of votes in the preceding spectra as an attribute.

$$\mathcal{D} = \{i, j, c_C, n_{\text{votes}}\}$$

For each preceding spectrum $\tau - k, k \in [1, \dots, N_{\text{hist}}]$, another peak detection is performed using the reduced threshold $T_{\tau-k}^*$. The detections obtained from this step are counted as confirmations if they can be matched to the respective detection candidate. Each detection is characterized by its RA cell $C_{i,j}$ where the peak value is located. The

range index i and the angular index j of this cell are converted to cartesian coordinates (x_C, y_C) by applying Equation 4.10. The cartesian coordinates $(x_{C'}, y_{C'})$ of the detection candidate $C'_{r=i', \varphi=j'}$ are determined analogously. The detection is considered matching to the detection candidate if the Euclidean distance between the detection candidate and the evaluated detection is less than the predefined matching distance d_{match} . In this case, the following condition applies:

$$\text{matched} \leftrightarrow \sqrt{(x_{C'} - x_C)^2 + (y_{C'} - y_C)^2} \leq d_{\text{match}}. \quad (4.15)$$

A matching detection candidate in an evaluated preceding measurement spectrum contributes a vote for this candidate and increments its vote count by one. The number of votes required to confirm a candidate can be selected based on the characteristics of the spectrum and the CA-CFAR threshold factors used in both stages of peak detection. Figure 4.8 illustrates the voting mechanism with representative examples: a trusted detection exceeding the standard threshold k_{thresh} is always approved. A detection candidate is classified as an approved detection if its accumulated vote count exceeds a specified threshold N_{votes} after the evaluation of the preceding spectra.




Figure 4.8: Voting mechanism to approve or disprove detection candidates. In this example, at least three votes ($N_{\text{votes}} = 3$) are required to confirm a detection.

The set of approved detection candidates, denoted $D_{\text{cand, appr}}$, includes all detection candidates in the set D_{cand} whose vote count attribute is at least the number of required votes N_{votes} : $D_{\text{cand, appr}} = \{\text{det} \in D_{\text{cand}} | \text{det.votes} \geq N_{\text{votes}}\}$.

The final detection set $\mathcal{D}$, returned by the criticality-aware peak detection framework, comprises all trusted detections $\mathcal{D}_{\text{trust}}$ as well as all approved detection candidates with a sufficient vote count $\mathcal{D}_{\text{cand, appr}}$: $\mathcal{D} = \mathcal{D}_{\text{trust}} \cup \mathcal{D}_{\text{cand, appr}}$.

Accordingly, the method consists of three cascaded stages of peak detection: the first stage identifies strong peaks, the second stage finds peak candidates and the third stage evaluates these candidates in preceding spectra, applying an adaptive threshold based on Gaussian dimples to validate the presence of weak peaks. This framework is summarized in Figure 4.9. The corresponding pseudocode is provided in Section B.2.




Figure 4.9: Scheme for threshold reduction in critical cells of a radar spectrum.

4.2.4 Camera Information Fusion

In the previous section, it was demonstrated how a criticality assessment of the radar perception field, combined with a reduced CA-CFAR threshold, can produce detection candidates that are subsequently examined in greater detail to confirm or refute the presence of object evidence. While this method provides one approach to identifying relevant but weak reflections that may otherwise be missed by standard peak detection, another means of identifying potential regions of interest is through the incorporation of data from other perception sensors. The fusion of radar and camera data is a widely discussed topic in research, see, e.g., [BF18; Kim+23]. The concept of this subsection is inspired by these approaches and combines them with the radar cell criticality evaluation proposed earlier.

In an early-fusion strategy, camera information can support radar detection by proposing potential object locations, even when the object's radar reflection is too weak to trigger a detection on its own. To integrate camera data into the processing pipeline presented in this work, region proposals for potential object evidence must be provided by the camera. In a simplified shape, these evidences can originate from a image object detection algorithm. A camera image $\mathcal{I}$ from a monocular camera represents a two-dimensional projection of the real world. Hence, it lacks direct depth information. This missing dimension can either be estimated or calculated geometrically under certain assumptions.

In this case, the following assumptions are made to estimate the three-dimensional position of the image contents:

1. The measurement area is completely flat ($z = 0$), which is a valid assumption for indoor parking garages and a reasonable approximation for most parking

spaces. Referring to Chapter 3, the driving tube for criticality estimation is also constructed in two-dimensional coordinates.

2. The image processing algorithm is capable of identifying the coordinates $(u_{Z=0}, v_{Z=0})$ corresponding to the ground position of the object evidence.
3. The camera matrix K and the mounting height h_{cam} of the camera are known.

Figure 4.10 illustrates the constellation of the projection with all dimensions.




Figure 4.10: Camera projection scheme applying the vehicle coordinate system and assuming that projected points lie on the ground ($Z_{\text{veh}} = 0$), corresponding to $Z_{\text{cam}} = -h_{\text{cam}}$ in the camera coordinate frame.

Using the focal distances f_x and f_y , and the optical centers c_x, c_y , the image coordinates (u, v) can be transferred into world coordinates $(X_{\text{world}}, Y_{\text{world}}, Z_{\text{world}})$. The camera projection equation is:

$$X_{\text{world}} \cdot \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} -f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} Y_{\text{world}} \\ Z_{\text{world}} \\ X_{\text{world}} \end{bmatrix}. \quad (4.16)$$

Given these assumptions and taking $Z_{\text{world}} = -h_{\text{cam}}$ in the camera coordinate system, the longitudinal position (X_{world}) can be computed using the following formula:

$$v \cdot X_{\text{world}} = f_y \cdot Z_{\text{world}} + c_y \cdot X_{\text{world}}, \quad (4.17)$$

$$(v - c_y) \cdot X_{\text{world}} = -f_y \cdot h_{\text{cam}}, \quad (4.18)$$

$$X_{\text{world}} = \frac{f_y \cdot h_{\text{cam}}}{c_y - v}. \quad (4.19)$$

By inserting this value into the equation derived from Equation 4.16, the lateral position Y_{world} can be calculated:

$$Y_{\text{world}} = X_{\text{world}} \cdot \frac{c_x - u}{f_x}. \quad (4.20)$$

If an image-based object detector is used that performs semantic segmentation of pixels belonging to potential objects, the lowest pixel row v_{ground} is used to determine the world position $(X_{\text{world}}, Y_{\text{world}})$ of the object evidence. By contrast, image-based object detectors that return bounding boxes [Gir+14; Liu+16] typically output these box in the form (u, v, w, h) , where (u, v) are the pixel coordinates of the upper-left corner point of the box, and w and h represent the width and height of the bounding box, respectively. For these detectors, the lowest pixel in the image is consequently determined with $v_{\text{ground}} = v + h$.

To demonstrate this concept, the widely used *YOLOv10* model [Wan+24a] is employed in this work, using the pre-trained weight set *yolov10m.pt*. Other image-based region proposals, which could employ features on pixel level, can be used in a similar manner. *YOLOv10* is an end-to-end object detector with a latency of only a few milliseconds. Trained by default on the dataset *MS-COCO* [Lin+14], it is optimized for general-purpose object detection. Although retraining on domain-specific datasets for automated driving such as the *Waymo Open Dataset* [Sun+20] or *Cityscapes* [Cor+16] may yield slight performance improvements, the default weights are considered sufficient for a proof-of-concept implementation.

Figure 4.11 shows an example image from a street scenery. The *YOLOv10* model has been applied to the image, and the resulting bounding boxes are visualized. The projection is calculated exemplary for this scenario and the calculation is given in Section B.1. The border color of each bounding box indicates the detected object class: blue for cars, red for pedestrians, and green for bicyclists. For each bounding box, the variable v_{ground} is included to enhance interpretability. The bounding box corresponding to the man on the bicycle is detected by *YOLOv10*, but not shown in the figure for better clarity.

The likelihood of observing a bicyclist or car as a single weak radar reflection peak that is insufficient in its altitude to exceed the CA-CFAR threshold is low compared to pedestrians. Bicycles and vehicles have significantly better radar reflection properties due to their metal structures. Hence, they typically generate multiple radar detections and are not modeled as single-point targets. Since this work focuses primarily on pedestrian safety, only detection bounding boxes labeled as *pedestrians* are considered for the threshold reduction method.

Each bounding box that is returned by *YOLOv10* and labeled with a pedestrian tag is treated as object evidence. In a second step, the criticality of each detection is evaluated using the same methodology described in Subsection 4.2.3. Based on this evaluation,




Figure 4.11: Camera image of a street scenario included in the dataset used for the experiments of this section [Wan+21a], showing vehicles, pedestrians, and a bicyclist. The pixel row corresponding to the ground is highlighted with a text box.

Gaussian dimples are applied at the bounding box positions to reduce the detection threshold in these regions accordingly. This facilitates the detection of flat or weak peaks in regions where the camera contributes supporting evidence for object presence.

This methodology describes how camera systems can be employed to apply the CA-CFAR threshold reduction which was described in Subsection 4.2.3 before. It is compared to the subsequent peak search approach in the following section.

4.3 Evaluation of Safety-Aware Peak Detection

In the next step, the capability of the proposed framework — namely, reducing the standard CA-CFAR threshold applied during peak detection in the RA spectrum — is investigated. The evaluation requires a dataset that includes radar spectra collected in scenarios representative of the ODD of an automated parking function. The used dataset is introduced in Subsection 4.3.1, which describes its structure and the scenarios it covers. Based on this dataset, various experiments are conducted to examine the impact of different parameters on detection performance. The experimental setup and the performance metrics used are outlined in Subsection 4.3.2. Finally, Subsection 4.3.3 presents the experimental results.

4.3.1 Dataset

The number of publicly available datasets that provide RA or RAD radar spectra with pedestrian data is limited. Some datasets, such as [Dim+20], are recorded in urban driving scenarios where the recording vehicle moves too quickly to reflect a parking-related ODD. Others, like the CARRADA dataset [Oua+21], contain limited scenario diversity and are collected in static environments such as a single fixed parking lot location.

To evaluate the proposed framework, a dataset was selected that offers a balanced compromise: it includes some scenario variability while maintaining moderate complexity to support explainability and systematic evaluation. For this reason, the *ROD2021* dataset [Wan+21a] was used. *ROD2021* is a subset of the *CRUW* dataset [Wan+21b], and it includes RA radar spectra, images captured by a calibrated stereo camera, and GT object annotations in the form (x, y, class) of mostly urban traffic scenarios. The composition of scenes inside of the *ROD2021* dataset is illustrated in Figure 4.12.

	Parking Lot	Public Road near Campus	Urban Traffic	Highway Traffic
50 Sequences	25 val, 5 test	9 val, 1 test	3 val*, 2 test	3 val, 5 test

Green font: used in experiment, red font: not used in experiment
 * (fast urban traffic scenario was excluded from experiments)

Figure 4.12: Scenario structure of the *ROD2021* dataset, with scenarios used in the experiment for validation shown in green, and unused scenarios shown in red.

For evaluation in this work, the three highway sequences and one high-speed urban sequence were excluded because they do not align with the ODD of an automated parking function. Moreover, the test sequences were not used, as they lack object GT annotations and camera images.

ROD2021 was recorded using the TI AWR1843 FMCW radar sensor, whose operating figures are listed in Table 4.1.

The dataset includes RA spectra that are generated by applying a range Fast-Fourier-Transformation (FFT) followed by an angular FFT. Each spectrum comprises 128 discrete range bins and 128 angular bins. Both the real part and the imaginary part of the signal values are stored separately. To compensate for the amplitude's dependency on the reflector's range, min-max scaling is applied. Each measurement frame includes 255 chirps, but the dataset only provides spectra for the four chirp sequences with indices $[0, 64, 128, 192]$. The sensor covers a range of 25 m and an azimuth Field of View (FoV) of $\pm 90^\circ$. The magnitude of each chirp signal $|C|$ is computed from its real part $\text{Re}(C)$ and imaginary part $\text{Im}(C)$ as follows:

Table 4.1: Operating figures of the TI AWR 1843 radar sensor.

Description	Symbol	Value
Transmitter Power	P_t	12 dB m
Total Antenna Gain	$G_{\text{transm}} \cdot G_{\text{rec}}$	42 dB m
Radar Frequency	f	76.5 GHz
Temperature	T	20 °C
Bandwidth	B	650 MHz
Noise Factor	F	25.119
System Loss	L	6 dB
MIMO Configuration	<i>MIMO</i>	2 Tx, 4 Rx

$$|C| = \sqrt{\text{Re}(C)^2 + \text{Im}(C)^2}. \quad (4.21)$$


A plot of the magnitude heatmap for a single measurement is shown in the upper Figure 4.13a. High magnitudes are visualized in yellow colors, and low magnitudes in violet. The corresponding camera image is shown in the lower Figure 4.13b.

4.3.2 Experimental Setup

The experimental evaluation is divided into three components:

1. A baseline experiment.
2. Multiple variations of the safety-aware peak detection algorithm involving the use of preceding radar spectra.
3. Experiments that include camera image fusion, where the YOLOv10 object detector provides potential peak candidates based on its bounding box outputs.

The object GT labels in the dataset consist solely of a two-dimensional point (x, y) marking its center and a class label. Thus, they do not account for the physical extent of each object or the spatial imprecision of radar measurements. To address this spatial imprecision in the world coordinates, an Euclidean distance d_{accept} is introduced during the matching process from GT objects to detected radar peaks. This acceptance distance d_{accept} defines whether a detection can be considered a match to a GT object or vice versa. Due to varying dimensions of each object class, the value of d_{accept} is specific to the class label of the respective GT object. Table 4.2 lists the class labels along with their corresponding acceptance distances d_{accept} . The distances roughly reflect the average physical extent of objects in each class.



(a) Magnitude $|C|$ heatmap collected by the radar sensor (Chirp 0)



(b) Camera Image

Figure 4.13: Exemplary frame from the ROD2021 dataset: sequence *PMS2000*, frame 208.

The two principle metrics for assessing perception performance are FP peaks and FN detections.

Table 4.2: Acceptance distance for various GT annotations, showing the values of the acceptance distance d_{accept} for the different GT classes.

	Pedestrian	Bicyclist	Car
Distance	1.5 m	2.2 m	4.0 m

Definition 4.5 (FP Peak Detection)

A peak exceeding the CFAR threshold and registered as a detection is called a *FP peak detection* if it cannot be matched to any GT object. That is, the distance d_{NN} to the nearest GT annotation exceeds the object-specific acceptance distance d_{accept} .

Definition 4.6 (FN Non-Detected Object)

A GT object annotation is classified as a *FN non-detected object* if there is no corresponding detection within the acceptance distance. Formally, the distance d_{NN} from the annotation to the closest detection exceeds the object-specific acceptance distance d_{accept} .

Both an exemplary FP and a FN detection are illustrated in Figure 4.14. Grey dashed circles represent cases where detections and annotations lie within the acceptance distance d_{accept} and are thus matched. Red dashed circles indicate instances where the acceptance distance is exceeded. The detection is thus considered as a FP, and the annotation is evaluated as a FN.




Figure 4.14: Detection metrics derived from GT annotations, illustrating detections and annotations classified as TPs, FPs and FNs.

In addition to these detection counts, the precision and recall metrics are computed using Equation 4.1 and Equation 4.2, respectively. All performance metrics are calculated

separately once for the entire perception area and once for the driving tube which has been assessed as critical by the framework.

Baseline Experiment

The goal of the baseline experiment is to quantify the performance of standard CA-CFAR-based peak detection. The results serve as reference for evaluating the safety-aware enhancements proposed in this work. The baseline experiment consists of two parts: the cell parameter evaluation for averaging cells and guard cells, and the evaluation of the threshold factor k_{thresh} . In the first part of the experiment, peak detection is performed using a fixed threshold factor $k_{\text{thresh}} = 5.0$ in a reduced set of sequences altering the parameters $b_{\text{av}, r}$, $b_{\text{av}, \varphi}$, $b_{\text{guard}, r}$ and $b_{\text{guard}, \varphi}$. As the reduced set, four sequences¹⁰ are selected for their variety in scene complexity and object types: one sequence with a bicyclist, one sequence featuring only one pedestrian, one sequence featuring two pedestrians and one urban sequence. In the second part, the threshold factor k_{thresh} is varied in steps of 0.2 from 2.4 to 7.0 in the same sequence set to identify the optimum threshold factor $k_{\text{thresh}, \text{opt}}$ for peak detection.

For each combination of cell parameters and threshold factor, the core evaluation metrics are computed: FP, FN, recall, precision, and F_1 -score.

Experiments on Subsequent Radar Peak Search

The effectiveness of the subsequent, zone-aware radar peak search is investigated through five distinct experiments, labeled A to E. Each experiment evaluates a different configuration of the threshold reduction and use of information from preceding radar frames. The specifications for each experiment are detailed below.

- A. Subsequent peak search without additional dimples. After applying CA-CFAR with the standard threshold factor $k_{\text{thresh}, \text{opt}}$, a reduced threshold factor $k_{\text{thresh}, -} < k_{\text{thresh}, \text{opt}}$ is used to identify weaker peaks. A voting process is applied, and a weak candidate is confirmed if it receives more than two votes ($N_{\text{vote}} > 2$).
- B. Threshold reduction in critically assessed cells. In cells that are not assessed as critical, the standard optimum threshold factor $k_{\text{thresh}, \text{opt}}$ is used. In contrast, for cells $C_{i,j}$ deemed critical ($\text{crit}_{\text{total}}(C_{i,j}) > 0.0$), the same CA-CFAR algorithm is applied but with the same reduced threshold factor $k_{\text{thresh}, -}$ as in variant A.
- C. Subsequent radar peak search with additional dimples, using a low number of preceding spectra $N_{\text{hist}} = 3$ for backward analysis.
- D. Subsequent radar peak search with additional dimples, using a medium number of preceding spectra $N_{\text{hist}} = 5$ for backward analysis.

¹⁰ [2019_04_09_BMS1000, 2019_04_09_PMS1000, 2019_04_09_PMS2000, 2019_04_30_MLMS000]

- E. Subsequent radar peak search with additional dimples, using a high number of preceding spectra $N_{\text{hist}} = 8$ for backward analysis.

The parameters used for the experiment variants C to E, including the number of visited preceding spectra for backward peak search n_{hist} , the number of required votes for a weak peak candidate in the preceding frames N_{votes} , and the dimensions of the dimples w_r and d , are summarized in Table 4.3. The dimple width in the angular direction is fixed at $w_\varphi = 8.0$ for all experiments.

Table 4.3: Variation of the experimental setup to investigate the subsequent radar peak search framework.

Variant	N_{hist}	N_{votes}	$w_r(\tau - N_{\text{hist}}, \dots, \tau - 1)$	$d(\tau - N_{\text{hist}}, \dots, \tau - 1)$
C	3	2	[5.0, 3.0, 1.0]	[0.133, 0.267, 0.400]
D	5	3	[7.0, 5.5, 4.0, 2.5, 1.0]	[0.08, 0.16, 0.24, 0.32, 0.40]
E	8	5	[11.00, 9.57, 8.14, ..., 1.00]	[0.05, 0.1, 0.15, ..., 0.40]

The spatial spread of each dimple is derived based on the expected movement of VRUs during the observation time window. The frame rate in the ROD2021 dataset is $f_{\text{frames}} = 30$ Hz for both RA measurements and camera images. The observation time for backward analysis is calculated as $t_{\text{observ}} = \frac{N_{\text{hist}}}{f_{\text{frames}}}$. Given the length of one range bin $d_{\text{bin},r} = 0.195$ m and the maximum expected VRU velocity v_{VRU} , the maximum number of crossed range bins $N_{\text{bins, crossed}}$ during the observation time can be estimated as:

$$N_{\text{bins, crossed}} = \frac{v_{\text{VRU}} \cdot N_{\text{hist}}}{f_{\text{frames}} \cdot d_{\text{bin},r}}. \quad (4.22)$$

As an example, if the observation length is five frames ($N_{\text{hist}} = 5$) and the maximum expected VRU velocity is $v_{\text{VRU}} = 8.0$ km/h, the number of potentially crossed range bins is $N_{\text{bins, crossed}} = 6.84$, which is rounded up to determine the maximum dimple width in the range dimension, yielding $w_r(\tau - 5) = 7$.

Experiments with Image-Based Cell Proposal

In contrast to the backward radar peak search, which operates on preceding radar spectra, the image-based cell proposal approach uses only the current radar spectrum. However, it relies on objects detected in the corresponding camera image by the YOLOv10 object detector. A CA-CFAR standard threshold is computed using the standard threshold factor k_{thresh} . In a parallel processing step, YOLOv10 identifies objects in the camera image. All object bounding boxes labeled as *pedestrians* are projected onto the radar cell grid. Gaussian dimples are applied at those positions to locally reduce the threshold. Each Gaussian dimple has the same configuration, defined by its depth $d = 0.8$, the range width $w_r = 7.0$ and the angular width $w_\varphi = 7.0$.

4.3.3 Experimental Results

Results of the Baseline Experiment

The recall and precision values for various parameter combinations is presented in Table 4.4. This table includes all combinations of a kernel's shape:

- Widely separated guard cells and averaging cells, $\frac{b_{av}}{b_{av, \max}} \ll 1, b_{\text{guard}} = b_{\text{guard}, \max}$.
- Both guard and averaging cells close to the cell under test $\frac{b_{av}}{b_{av, \max}} \ll 1, \frac{b_{\text{guard}}}{b_{\text{guard}, \max}} \ll 1$.
- Both guard and averaging cells distant from the cell under test $b_{av} = b_{av, \max}, b_{\text{guard}} = b_{\text{guard}, \max}$.

Additionally, the best-performing configuration is shown in the final row.

Table 4.4: Detection performance for different CA-CFAR parameter combinations.

Param. $b_{av, r}, b_{av, \varphi}, b_{\text{guard}, r}, b_{\text{guard}, \varphi}$	FPS	FNs	F_1 -score	Recall	Precision
[4, 12, 2, 2]	725	5323	0.019	0.011	0.073
[4, 4, 2, 2]	1320	4536	0.222	0.157	0.380
[4, 12, 2, 10]	16143	1047	0.350	0.805	0.224
[12, 12, 2, 10]	13309	970	0.402	0.820	0.267
[12, 12, 10, 10]	10030	926	0.479	0.828	0.337
[12, 4, 2, 2]	7947	1095	0.505	0.797	0.370
[12, 12, 2, 2]	5214	1110	0.588	0.794	0.467
[12, 12, 10, 2]	4489	1139	0.614	0.788	0.502
[12, 4, 10, 2]	4020	1235	0.623	0.771	0.523
[12, 8, 10, 2]	4000	1226	0.624	0.772	0.524

It was observed that if the number of guard cells is too low, the receptive field of the kernel used to estimate the noise level is influenced by the peak itself, causing most detections to be missed. Conversely, large guard cell regions detect more objects but at the cost of a significantly higher number of FP detections. Interestingly, a kernel stretched in the range dimension and compressed in the angular dimension achieved the best trade-off between FPs and FNs, even though the range peak is typically narrower than the angular peak. The objective of the parameter sweep performed in this experiment was to identify a configuration with high recall while maintaining acceptable precision. Therefore, the parameter combination with the highest F_1 -score was selected for further evaluation. This parameter set is $b_{av, r} = 12, b_{av, \varphi} = 8, b_{\text{guard}, r} = 10, b_{\text{guard}, \varphi} = 2$.

Using this parameter set, the impact of the threshold factor k_{thresh} on detection performance was analyzed in more detail. Particularly, the balance between recall and precision was investigated. Figure 4.15 illustrates the effect of varying the threshold factor k_{thresh} .




Figure 4.15: Recall, precision, and F_1 -score as functions of the threshold factor k_{thresh} for the CFAR peak detection.

As shown, the optimum F_1 -score of 0.647 is achieved at a threshold factor of $k_{\text{thresh, opt}} = 5.6$. This threshold factor is used as the standard threshold to identify trusted detections. For the experiments, a reduced threshold factor of $k_{\text{thresh, -}} = 4.0$ is chosen, as this value differs significantly enough from the standard threshold and thus facilitates a comparison between the standard park detection and the safety-aware peak detection.

In the final step, peak detection was performed across all sequences in ROD2021 using the identified parameter set and the optimal threshold factor. As results for the baseline experiment, the number of FN is 15,145, and the number of FPs is 39,350. This results in a precision of 0.483, a recall of 0.666 and an F_1 -score of 0.560. Compared to the results on the reduced set of sequences, the performance on the full dataset is limited, with a drop of eleven percentage points in recall and four percentage points in precision. The remaining sequences in the dataset contain a larger number of objects and exhibit greater scene complexity. These factors generally lead to increased noise levels, which hinder accurate object detection. Moreover, the full dataset includes more structural elements, such as guard rails and lampposts, that may be detected by the radar sensor, but mistakenly counted as FPs in the evaluation.

Results of the Subsequent Radar Peak Search Experiment

Table 4.5 presents the evaluation results for all experiment variants related to subsequent radar peak search, as described in Subsection 4.3.2. The table reports the number of FPs, FNs, as well as recall, precision, and F_1 -score. The results are presented for the overall FoV and for the area identified as critical, as illustrated in Figure 4.5. Moreover, the table compares the performance metrics to the baseline values established in the previous Section 4.3.3.

Table 4.5: Comparison of detection performance between the baseline CA-CFAR and variants of the subsequent radar peak search.

	Baseline	A	B	C	D	E
FP	39,350	40,180	41,194	41,681	42,160	42,505
FP assessed crit.	1,959	2,781	3,855	4,286	4,770	5,118
FN	15,145	13,026	13,058	12,783	12,282	12,126
FN assessed crit.	4,521	2,532	2,801	2,284	1,844	1,713
Recall	0.666	0.712	0.711	0.718	0.729	0.732
Recall assessed crit.	0.823	0.901	0.890	0.911	0.928	0.933
Precision	0.483	0.499	0.503	0.493	0.496	0.497
Precision assessed crit.	0.925	0.908	0.884	0.867	0.859	0.852
F_1 -score	0.560	0.587	0.590	0.585	0.591	0.592
F_1 -score assessed crit.	0.871	0.905	0.887	0.888	0.892	0.891

All methods incorporating criticality-aware peak detection in the RA spectrum demonstrate a notable reduction in FN detections. A large effect was already achieved by investigating the preceding spectra without utilizing Gaussian dimples (experiment variant A). The number of FPs increases only slightly by 2.1 %, and the number of FNs is reduced by 14.0 %. Variant B, reducing the threshold in critical cells to achieve more sensitive peak detection there yields a comparable reduction of FNs. However, the number of FPs is considerably higher compared to the other experiment variants and nearly double the FP count of the baseline.

Introducing Gaussian dimples, which is investigated in the variants C to E, provides additional improvements: the number of FNs reduces about 1.9 % to 6.9 % relative to variant A where no dimples were applied. This improvement comes at the cost of an increased number of FPs, which is around 3.7 % to 5.8 % higher in the variant with dimples compared to variant A without dimples. It is important to note that the FP count may be partially inflated due to detections of objects that are not labeled in the GT. In such cases, TP detections could be misclassified as FPs.

The benefits of the proposed method are more pronounced when the cells of the driving tube — the region assessed as critical — are analyzed separately. The recall increases by up to eleven percentage points compared to the baseline, while the precision decreases by approximately seven percentage points in this area.

Variant C to E confirm the trend that a longer observation horizon results in fewer FNs, but also leads to a progressive increase in FPs.

Results of the Image-Based Cell Proposal Experiment

The results of the image-based cell proposal experiment are shown in Table 4.6.

Table 4.6: Comparison of detection performance between the baseline CA-CFAR and the image-based cell proposal.

	Baseline	Image-based cell proposal
FP	39,350	48,082
FP assessed crit.	1,959	7,266
FN	15,145	14,809
FN assessed crit.	4,521	4,284
Recall	0.666	0.673
Recall assessed crit.	0.823	0.832
Precision	0.483	0.443
Precision assessed crit.	0.925	0.777
F_1 -score	0.560	0.534
F_1 -score assessed crit.	0.871	0.804

The method leads to a negligible decrease of FN objects that are not detected, but introduces a significant increase in FP detections. Consequently, the overall F_1 -score degrades in comparison to the baseline. Most of the additional FPs issued by the image-based cell proposal method occur in the critical region, which has a negative impact on the robustness of the perception function.

4.4 Discussion

This chapter introduced several strategies for evaluating radar spectra with an emphasis on enhancing the safety of the overall perception function. Peak detection in radar spectra is a crucial processing step that reduces a continuous signal, containing values for each RA cell, into a small set of radar detections. As such, most of the information captured by the radar sensor is discarded in this process stage. If a peak in the signal is not detected by the peak detection during this stage, the corresponding information cannot be recovered later in the processing pipeline. Thus, it is essential to integrate safety-aware concepts already at the peak detection level.

The CA-CFAR algorithm remains the prevailing method for peak detection in radar signal processing and is well-suited for integration of additional context information. The core idea of safety awareness in radar signal processing is that the acceptable risk of missed detections should depend on their spatial location. Specifically, missing a detection in a safety-critical region, where it could lead to harm to VRUs, evokes greater consequences than in other regions. For this reason, a criticality assessment function is applied to the RA spectrum, assigning an individual criticality value to each cell depending on its relevance for the current driving task. In the experiments, this criticality function is assumed to be static, reflecting the nature of the ROD2021 dataset,

which contains only scenes that are representative of the ODD of an automated parking function. However, a concrete context that includes the intentions of the function is not given in the dataset. Therefore, the intention that is assumed in the experiment is a straight, linear movement of the sensor toward the scenery. However, this assumed motion does not necessarily result in a meaningful increase in criticality for all objects. Consequently, the experiments focused primarily on whether detections for critical GT objects were more likely to be identified or recovered using the proposed method.

It is important to acknowledge that the ROD2021 dataset does not include situations from indoor parking garages, which are a key part of the ODD for automated parking functions. Radar waves in such environments exhibit distinct propagation behaviors in parking garages due to structural elements such as metal ceilings, exposed pillars, and fences. These structures act as a strong reflector for radar waves that can obscure weaker reflectors from nearby VRUs. The limitations of the ROD2021 dataset prevent detailed investigation of these complex effects.

The proposed method of subsequent radar peak search, applied to preceding RA spectra, preserves the standard concept of peak detection, which extracts a point cloud from peaks in the radar signal for further processing. The novelty lies in additional mechanisms that enhances the treatment of weak but potentially relevant peaks. The core idea is that detections are likely to appear in more than one radar spectrum, even if the peak has a low height. A two-stage peak detection is applied: in the first stage, a standard CA-CFAR algorithm is used to identify peaks with a high SNR that are directly proven to be trusted detections of objects. In a second stage, a peak detector with a lower threshold identifies detection candidates, which may represent weaker reflectors. For each detection candidate, the peak detection threshold is reduced around the cells where detection candidates occurred. The approach leverages the idea that valid object detections are likely to reappear across multiple frames, even if their signal strength is weak. Backward peak detection in RoI increases the likelihood of capturing such weak signals and improving recall of detections.

The voting concept, which searches for peaks in preceding radar spectra and associates these peaks with detection candidates, was evaluated in experiment variant A. This approach resulted in overall recall increase of 4.6 percentage points compared to the baseline experiment. When considering only critical FN objects in the recall metric, the recall improved by nearly eight percentage points. These results provide evidence of temporal consistency in radar reflections from real objects, that is, a peak generated by an object in one frame is likely to be generated similarly in adjacent frames. The ratio of FP to FN detections remained acceptable, as indicated by the improved F_1 -score. Compared to the baseline, the modified approach accepts to generate more FP than FN detections, which is a sign for increased safety of the perception function. Additional FPs may still arise in some cases. One reason is an object moving rapidly while motion estimation is not accounted for due to missing Doppler information in the dataset. Another reason is an object just entering the radar sensor's FoV, making it impossible

to identify corresponding peaks in preceding frames since these earlier frames do not contain information about the object.

Lowering the threshold in cells assessed as critical produced a similar improvement in recall, but at the cost of reduced precision. This is likely because no information is available to validate whether a peak existed in preceding spectra already. It also proves the effectiveness of analyzing preceding spectra. Incorporating this information and facilitating the detection of weak peaks by the introduction of Gaussian dimples further enhanced recall. Variants C through E showed a one to two percentage point gain in recall compared to experiment variants A and B. Although the number of FPs increased relatively to the baseline, the precision remained stable. This is because the number of TP detections on existing objects also rose with the number of FPs, keeping precision effectively constant despite the absolute increase in FP counts.

Figure 4.16 illustrates the effect of voting and Gaussian dimples in detecting a low-intensity peak located near a high-intensity peak. In a standard CA-CFAR scenario, separating these two peaks across adjacent range bins is challenging. The lower peak is often masked by the noise floor estimation in this region, which is heavily biased by the adjacent high peak. Figure 4.16 contains four figures that each represent one measurement frame [529, 530, 531, 532] of the exemplary situation *2019_04_09_PMS2000*. In frame 532, the weaker peak fails to exceed the standard threshold T . Although similar peaks exist in previous frames 529, 530, and 531, they also go undetected by the standard threshold. However, when a reduced threshold T_* including Gaussian dimples is applied, the peak is successfully identified in all earlier frames. These earlier detections vote for the candidate peak in frame 532, thereby confirming its presence and preventing a FN.

Experiment variant C considered only the three most recent radar spectra, resulting in the shortest observation period among all evaluated configurations. In contrast, variant E included an observation period of eight preceding spectra, offering a significantly longer temporal horizon. The results suggest that a longer observation period is beneficial when higher safety guarantees are desired. Specifically, the probability of detecting even a TP with a weak or vanishing peak increases with a longer horizon, albeit at the cost of more FP detections. In addition to the observation period, the required number of votes to confirm a detection candidate serves as another parameter for controlling the robustness and sensitivity of the peak detection process. Requiring fewer votes leads to faster confirmation of the object, which can help reduce the number of FNs, but may also increase susceptibility to noise.

Gaussian dimples, once introduced, can also be applied to facilitate peak detection if other information about the presence of objects is available. In perception systems, this information may also be obtained from other sensor modalities, such as cameras. Accordingly, an image-aided approach was investigated, in which Gaussian dimples were introduced at locations where objects were detected by the object detector YOLOv10. However, this method did not yield promising results in the experiments. A very high

number of FPs was observed, accompanied by only marginal improvements in FN reduction. One reason for this might be a partially suboptimal quality of camera images, which manifests itself in blurred images. This affects the performance of YOLOv10 in these images. In other images, pedestrians were partially occluded or located in very shadowy environments, leading to missed or inaccurate detections by YOLOv10. Additionally, the projection of bounding boxes from two-dimensional image space to the three-dimensional radar coordinates may introduce localization errors. Even minor pixel-level inaccuracies in the bounding box location can propagate into an error of several meters in spatial positioning, especially if the object is far from the camera. Consequently, the Gaussian dimple may be applied to an incorrect region, failing to support the detection of the actual peak.

A central question of this thesis is how the processing of radar data can be adapted according to the criticality evaluation delivered by the method proposed in Chapter 3. Two approaches, subsequent radar peak search and image-based cell proposal, address (RQ2) at the radar spectrum data stage. Peaks that are too weak to exceed the standard CA-CFAR threshold, and therefore cannot progress to the point cloud stage, are exploited if they are located in a sufficiently critical region of the spectrum. The following chapters, Chapter 5 and Chapter 6, address the same question for subsequent processing stages, such as point clouds and object hypotheses.




Figure 4.16: Low peak detection using Gaussian dimples, showing a voting example for a weak peak detected and proved with threshold functions T^* reduced by Gaussian dimples.

5 Processing Adaptations on Point Clouds¹¹

Safety-aware processing on radar spectra is the best option for preserving as much information as possible about object presence. However, the amount of processed data in radar spectra is enormous, and the contextual information can be lower compared to later steps in the signal processing chain. Hence, the safety-aware peak detection method introduced in the previous chapter may not always be applicable. This chapter presents a method to incorporate safety awareness into radar data structured as a point cloud. First, the role of radar point clouds and their typical properties are outlined in Section 5.1. Subsequently, in Section 5.2, a framework is introduced that evaluates the potential impact of each radar point on the parking task for a given radar point cloud. For automated driving and parking functions, an essential challenge is to derive the quantitative failure rate of the perception system, which directly influences the occurrence of erroneous behavior in the overall function. Section 5.3 presents a methodological concept to determine this failure rate based on missed detections on the point cloud level. The evaluation of the safety-aware point cloud processing is described in Section 5.4, using an urban dataset of radar point clouds. Finally, Section 5.5 discusses the outcomes of the experiments and contextualizes them within the methodological considerations.

5.1 Radar Point Clouds in Object Detection

Radar point clouds are an intermediate representation in radar data processing between the digital signal and an object hypothesis. Although radar waves propagate continuously, the reflections can be expressed as a point cloud $\mathcal{P}$. This point cloud consists of $N_{\mathcal{P}}$ points, $\mathcal{P} = \{p_1, \dots, p_{N_{\mathcal{P}}}\}$. Each point is — at its most basic — characterized by a range r from the sensor, a relative velocity v_{rel} determined by the Doppler frequency shift, and the azimuth angle φ . These three components are direct outcomes of the three Fourier transforms, one for each dimension. Moreover, radar points often have additional attributes, such as the point's signal power, the SNR, which is closely related to the signal power, the RCS value, which also depends on the signal power, the elevation angle, if the antenna arrangement of the radar sensor is capable of measuring this dimension, and information from early preprocessing steps, such as a flag indicating whether the point belongs to a moving or stationary object. While VRUs typically generate only a single radar point, larger reflectors, such as cars or trucks, produce several reflections that support the classification of the object.

¹¹ Chapter builds on [Brü+24d].

Radar point clouds are popular because they reduce the amount of data significantly compared to radar spectra. However, the number of FP detections in a radar point cloud is still high. Consequently, the point cloud is noisy and requires filtering strategies to achieve robust object detection. On the other hand, the number of TP points associated with relevant objects in one radar measurement frame is sparse. This sparsity motivates the strategy of point cloud accumulation to densify the data. A simple filtering strategy is to remove radar points based on thresholds for particular point criteria. In the following, this approach is referred to as *rigid filtering*.

Definition 5.1 (Rigid Filter)

Let a point cloud $\mathcal{P}$ be given, containing a set of points p . Assume that each point p has an attribute a . A rigid filter removes all points from this point cloud $\mathcal{P}$ if the attribute a lies outside the value range $[a_1, a_2]$ specified by the filter.

$$\mathcal{P}_{\text{filt}} = \{p \in \mathcal{P} | a_1 \leq p.a \leq a_2\} \quad (5.1)$$

The criteria of a rigid filter result from the measurable attributes of a radar point. Rigid filters can be combined using Boolean algebra to include multiple conditions simultaneously.

Below, a suitable cascade of rigid filters for the object detection task is designed. Radar points at larger distances are relevant for tasks such as localization or free space estimation in parking scenarios. Moreover, they are predominantly used for automated driving tasks in rural or highway ODDs. For object detection in parking scenarios, however, these radar points are irrelevant because their distance is too large to cover areas prone to collisions with the ego vehicle. Therefore, the first rigid filter is applied to the x -position and y -position of each point. Considering a planning time frame of $t_{\text{plan}} = 3.0 \text{ s}$ for automated parking maneuvers, and assuming that the velocity of an object of interest is $v_{\text{obj}} \leq 4.0 \text{ m/s}$, points with $x \geq 20.0 \text{ m}$ or $y \geq 10.0 \text{ m}$ can be removed to reduce the size of the point cloud.

The measured Doppler velocity v_{rel} also provides an opportunity to remove potentially invalid radar points. Points caused by interference from other radars in the ego vehicle's radar belt system (intrinsic faults) or from foreign sensors (extrinsic faults) often exhibit unrealistic Doppler velocities. Considering a safety margin v_{margin} for objects exceeding the ODD's maximum speed, and the ego vehicle velocity v_{ego} , the maximum expected Doppler velocity $v_{\text{rel, max}}$ can be computed:

$$v_{\text{rel, max}} \leq v_{\text{ODD}} + v_{\text{ego}} + v_{\text{margin}} \quad (5.2)$$

Hence, the filter criterion for Doppler velocity depends on the current motion state of the ego vehicle. Moreover, the measured Doppler velocity can be compared with this motion state to remove static points. This strategy is useful for eliminating reflection points from geometries such as walls or fences, which would otherwise be processed into object hypotheses. The expected Doppler velocity $v_{\text{rel, exp}}$ for a static point at a measured azimuth angle φ can be determined using the azimuth alignment θ_S of the sensor relative to the driving direction and the current sensor velocity, which usually equals the vehicle velocity v_{veh} :

$$v_{\text{rel, exp}} = v_{\text{veh}} \cdot \cos(\varphi - \theta_S) \quad (5.3)$$

If the difference between the expected Doppler velocity of a static radar point and the measured Doppler velocity exceeds a tolerance v_{tol} , the point is classified as moving. This tolerance depends on the radar velocity resolution and is usually set to approximately $v_{\text{tol}} = 0.5 \text{ m/s}$.

In addition to removing points based on position and velocity, another effective strategy is to filter out points with low signal power values. The probability that a radar point is a FP increases if its signal power or reflectivity is low. Depending on the radar sensor, the measured value for this may be represented in different forms:

- *Signal Power*: The signal power is the absolute received power P_r of the reflected radar signal. For a reflector with a fixed RCS, the signal power increases as the reflector moves closer to the radar antennas.
- *SNR value*: The SNR value is the ratio between the received peak power P_r of the signal and the noise floor $\mathcal{N}$. For a reflector with a fixed RCS, the SNR decreases with increasing distance until the signal is no longer distinguishable from the noise floor, at which point the object becomes undetectable.
- *RCS value*: The RCS value is part of the radar equation (Equation 2.3) and can be estimated if the other parameters of the radar equation are known. Unlike signal power and SNR, the RCS normalizes for the distance between object and the sensor. Consequently, a point with the same RCS yields a weaker SNR when the distance to the radar sensor increases.

Applying a filter on one of these attributes is also a strategy to reduce FP points. However, this filtering algorithm also eliminates information about the environment. Specifically, the removed points may correspond to FN detections, such as points lost due to clutter, ghost detections, or reflections from objects. A point that is removed during the filtering process to avoid FNs can prevent an object hypothesis from being formed. A hazard arises if the removed point belongs to an *object of interest*.

Definition 5.2 (Object of Interest)

An object of interest is any element in the environment that the automated parking function must consider to ensure safe behavior. In general, an object of interest can be either a static geometry, such as a wall or a fence, or a dynamic object, such as a human, an animal, or a vehicle. In the parking ODD, collisions between the automated vehicle and static geometries are equally relevant as collisions between the automated vehicle and another car. However, a collision of the automated vehicle with a VRU can result in injury. Consequently, in this work, the objects of interest are limited to pedestrians, scooter riders and bicyclists.

Another hazard, distinct from the erroneous removal of FP points, occurs if the object of interest is not represented by a radar point at all. Objects missed in this way are referred to as *hardware-induced FNs*.

Definition 5.3 (Hardware-Induced False Negatives)

Hardware-induced FNs are objects without corresponding evidence returned by the radar sensor. This effect leads to hazards because information about the object's presence is entirely absent. The missing evidence originates at the level of the radar spectrum itself or at the raw signal level. This circumstance must be addressed during sensor development and cannot be mitigated by improved processing of the points contained in the raw point cloud.

In this chapter, the focus is on the subsequent enhancement of the filtered point cloud with critical points. This includes the criticality assessment of each point, the adaptation of the filtering mechanism in regions where points are expected, and the accumulation of the filtered point cloud together with the additional critical points. However, the mitigation of hardware-induced FNs is not part of this thesis.

5.2 Successive Point Cloud Treatment

The identification of critical points in the point cloud requires adapting the drive-tube-based criticality assessment. Each point p in the radar point cloud $\mathcal{P}$ is assigned a criticality value as an additional attribute. Critical points induce regions with reduced filtering criteria. In this way, points that would normally be removed by rigid filter criteria are reintroduced into the point cloud for further processing. The method thus offers a means of filtering a radar point cloud while retaining points that may be relevant.

In Chapter 3, a method to assess the impact of object evidence on the automated parking task was elaborated. This method can be directly applied to radar point clouds. Each radar point receives a criticality attribute $p.crit_{total}$, which consists of the longitudinal criticality $crit_{lon}$, the lateral criticality $crit_{lat}$, and the velocity criticality $crit_{vel}$.

5.2.1 Region Expansion Approach

The idea of the successive point cloud treatment is to utilize information from radar points relevant to the driving task for subsequent point clouds. As discussed in Chapter 4, additional object evidence is likely to appear in areas close to the position where the original object evidence was measured. Such evidence must not be removed by subsequent filtering cascades. It can be preserved by spatially loosening the filter policy around points that have been evaluated as critical. Section 4.2 has already demonstrated how information from earlier measurements can be recaptured in radar spectra. Here, a method is presented for treating successive radar point clouds.

The overall goal is to detect objects that would otherwise be removed by a standard detection algorithm employing filtering methods. Figure 5.1 illustrates an example of objects whose radar reflection points differ in terms of signal power. Rigid filtering leads to different constellations in the radar signatures of the objects. The first object $\mathcal{O}_1$ exhibits a high density of strongly reflecting radar points, sufficient to induce an object hypothesis. The second object $\mathcal{O}_2$ yields fewer radar points, and, moreover, the signal power is lower. Consequently, the remaining point density after rigid filtering is insufficient to induce an object hypothesis. The concept of subsequent point cloud treatment is to recover points with weak signal power, transfer as much information as possible to the subsequent processing step, and thereby provide the basis for forming an object hypothesis.

Radar points in the pre-filtered point cloud are assessed for criticality. A radar point is considered critical if it exceeds a criticality threshold $crit_{thresh}$. Each critical radar point in the point cloud $\mathcal{P}_\tau$ defines a *criticality region* in the N_{hist} subsequent point clouds $\mathcal{P}_{\tau+1}, \dots, \mathcal{P}_{\tau+N_{hist}}$. The size and shape of these criticality regions represent the uncertainties inherent in radar-based object measurements. Determining an object's exact position is challenging, since radar reflection points may shift along the object surface [Ber+17], and azimuth angle measurements are subject to relatively large errors [BJT21]. Furthermore, if a VRU is in motion, the object may shift between point cloud $\mathcal{P}_\tau$ and the subsequent point cloud $\mathcal{P}_{\tau+1}$. The region's position in the next point cloud with index $\tau + 1$, given the current position (x_τ, y_τ) , can be modeled by a two-dimensional multivariate normal distribution. The unknown variable is the relative motion vector $\mathbf{v} = \mathbf{v}_{rad} + \mathbf{v}_{tang}$. The radial velocity $\mathbf{v}_{rad}$ can be measured as the Doppler velocity by the radar sensor and is therefore observable. However, the tangential velocity $\mathbf{v}_{tang}$, which is perpendicular to the radial velocity $\mathbf{v}_{tang} \perp \mathbf{v}_{rad}$, is unobservable and cannot be directly measured. The relative motion vector $\mathbf{v}$ is a composition of the VRU's velocity




Figure 5.1: Comparison of conventional filtering and subsequent point cloud treatment. For object $\mathcal{O}_2$, points with low received signal power P_r are removed by default but can be recaptured through the proposed subsequent point cloud treatment.

vector $\mathbf{v}_{VRU}$ and the automated vehicle's velocity vector $\mathbf{v}_{veh}$. These factors influence the shape of the multivariate distribution, which is a function of the radial position r and the tangential position $tang$. The expected values can be computed using

$$\mu_{rad} = \mathbf{v}_{rad} \cdot dt, \quad (5.4)$$

$$\mu_{tang} = \mathbf{v}_{tang} \cdot dt, \quad (5.5)$$

which depend on the tangential measurement standard deviation σ_{tang} and the radial measurement standard deviation σ_{rad} . Given the Pearson correlation coefficient ρ of r and $tang$, the probability density function $f(r, tang)$ can be determined by

$$f(r, tang) = \exp \left\{ -\frac{1}{2[1 - \rho^2]} \left[\left(\frac{r - \mu_{rad}}{\sigma_{rad}} \right)^2 - 2\rho \left(\frac{r - \mu_{rad}}{\sigma_{rad}} \right) \left(\frac{tang - \mu_{tang}}{\sigma_{tang}} \right) + \left(\frac{tang - \mu_{tang}}{\sigma_{tang}} \right)^2 \right] \right\}. \quad (5.6)$$

Graphically, the multivariate normal distribution can be illustrated by an ellipsoid with varying probability density inside. Figure 5.2 illustrates a simplified version, where a point moves toward the radar sensor plane.

This consideration is only valid if the position p_t is assumed to be measured without uncertainties. In practice, positional errors $\epsilon(p_t)$ must also be taken into account, namely the radial error $\epsilon(r)$ and the angular error $\epsilon(\varphi)$.




Figure 5.2: Moving criticality region modeled by a multivariate normal distribution. The critical region of a point is interpreted as a moving distribution at time step t , defined by a multivariate normal function.

An ellipsoid is therefore applied, characterized by two parameters, e_1 and e_2 . Parameter e_1 is correlated to the radial standard deviation σ_{rad} , which is subject to the expected measurement error of the Doppler velocity $\epsilon_{v_{\text{rel}}}$. Parameter e_2 is correlated with the tangential standard deviation σ_{tang} subject to the expected agility of the VRU. It may reach its maximum value if the VRU moves solely in a tangential direction, without a radial velocity component v_{rad} . However, the parameter value is reduced if a radial component exists.

It holds

$$v_{\text{max}} = v_{\text{tang}} + v_{\text{rad}}, \quad (5.7)$$

$$v_{\text{tang}} = v_{\text{max}} - v_{\text{rad}}, \text{ and} \quad (5.8)$$

$$v_{\text{tang}} = \sqrt{v_{\text{max}}^2 - v_{\text{rad}}^2}. \quad (5.9)$$

Hence, the size and shape of the regions may vary from one piece of object evidence to another, depending on the information associated with each observation. The purpose of the criticality regions is to adapt the filter threshold. One possible strategy is to remove the filter entirely, another is to reduce the filter values in proportion to the

criticality and the abundance of object evidence within the region. The first method ensures that every available point is used in the subsequent object detection pipeline. The second method conducts a less exhaustive search for points but improves the robustness of object detection.

In Figure 5.3, the concept of successive point cloud treatment is illustrated. Each point p in the current radar point cloud $\mathcal{P}_\tau$ is assessed for criticality with respect to the currently intended driving tube. Critical points generate regions for each successive point cloud $\mathcal{P}_{\tau+1}, \dots, \mathcal{P}_{\tau+3}$. These regions are added to a region list and expand with their age. Regions that reach an age of N_{hist} are removed from the list. Each region represents an area where previous observations warrant closer examination for potential objects. The strategies for threshold reduction are outlined in the next subsection.




Figure 5.3: Concept of successive point cloud treatment. Critical points generate criticality regions that expand over time to represent increasing uncertainty.

5.2.2 Adapted Filtering

The radar filter policies can be relaxed in different ways so that points normally removed by standard rigid filtering are included in the safety-aware filtering approach. The safety-aware point cloud consists of the following sets:

- Radar points retained by *rigid filtering*, i.e., points considered relevant due to their properties, e.g., high signal power.

- Radar points *currently assessed as critical* in this radar measurement.
- Radar points located in a *critical region* where the filtering is suspended.

The latter two sets are reintroduced into the processing pipeline due to their significance for avoiding the loss of a critical object. In the case of criticality regions, even radar points not individually evaluated as critical are retained for further processing, as they may support the formation of an object hypothesis.

5.2.3 Point Cloud Clustering

The proposed method of criticality-aware filter adaptations for successive point cloud measurements affects the density of the processed point cloud. The effectiveness of the proposed approach is evaluated by analyzing the number of additionally generated cluster centroids. Clustering performance typically depends directly on the density of the point cloud. A standard DBSCAN algorithm, as described in detail in Subsection 2.3.4, is employed. A nearest-neighbor search using an ϵ -distance metric is applied to identify core points and boundary points. Furthermore, a minimum number of points $N_{\min, \text{Pts}}$ is required to form a cluster centroid. Both parameters must be adapted according to the density of the point cloud and the representation of objects within it. These parameters strongly depend on the applied sensor and the characteristics of the recorded environment.

The cluster centroids constitute an essential step in forming an object hypothesis in the downstream tracking task and are, therefore, a suitable metric for evaluation. Nevertheless, in downstream processing, information loss may still occur — for example, during the tracking step.

5.3 Failure Rate Estimation

Deriving and statistically validating low failure rates is a complex task. This section presents an approach to determine the failure rate for radar perception starting on the point cloud level in a systematic manner. Subsection 5.3.1 explains how scenarios representing the ODD can be identified methodically. Radar point clouds from these selected scenarios are recorded, and each recorded radar point — the smallest possible unit of object information — is assigned to particular quantification sets (Subsection 5.3.2). This quantification method forms part of a causal chain that may ultimately result in an incident, including injury to a person. The chain consists of several links, for each of which a transition probability can be estimated. These links are discussed in Subsection 5.3.3.

5.3.1 Methodical Scenario Search

As outlined in Section 2.4, the failure rate of a perception system must be validated prior to the release of any automated function that relies on this perception system. This requires determining all erroneous behaviors that may occur throughout its lifetime and across the entire fleet of vehicles. The full set of scenarios encountered by the perception system cannot be exhaustively investigated due to time and resource constraints.

Figure 5.4 shows the coverage of the ODD by recorded scenarios as a function of the total recording length qualitatively. As can be seen, a relatively large proportion of the ODD can already be covered using a small dataset. However, a vast number of rare situations are generally not covered completely within a scenario subset. These scenarios are more difficult to identify, and both the expected variance and the degree of deviation from standard scenarios are higher.




Figure 5.4: ODD coverage over time during statistical recording.

Therefore, subsets of scenarios are employed during validation to represent the requirements placed on the perception system by the automated function. The goal of these subsets is to cover as many ODD scenarios as possible while minimizing the effort required for scenario recording.

The main challenge in using a scenario subset is that system behavior in response to known unknowns and unknown unknowns is neither tested nor predictable. At the same time, the ODD contains a large number of rare situations that cannot be fully included in the testing set.

To maximize ODD coverage while keeping the set of testing scenarios manageable, a methodical scenario search is applied. As [Zha+23b] discusses, methods to identify critical concrete scenarios include a logical scenario model, an exploration method, and a criticality assessment. The scenario model determines parameters and combines them to construct varying scenarios. Exploration treats the scenario search as an optimization

problem, investigating the ODD and identifying optimized (i.e., critical or failing) scenarios. Both methods require criticality assessment as an essential element in defining the quality of the identified scenarios. Furthermore, inductive reasoning is mentioned [Zha+23b], which can be divided into two types of data sources for identifying critical scenarios: first, scenarios derived from accident data, and second, scenarios drawn from broader sources.

Based on these insights, the proposed methodical scenario search rests on multiple pillars:

- *Expert knowledge:* The ODD of the automated parking function is specified in official technical regulations, such as the ISO 23374 standard [Int23]. These regulations define, for example, the function's objectives, feature scope, activation conditions, and potential situations. The ODD derives directly from these definitions and frames the search space for critical scenarios. This search space can often be curtailed by considering only scenarios that have been identified as critical by experts [BMM18]. The role of experts is to derive testing scenarios from the ODD and their own domain knowledge, thereby contributing to the definition of a minimal set of testing scenarios.
- *Accident reports:* Numerous works discuss how accident reports can be used to generate critical scenarios by learning from past failures [Zha+23b; Guo+24]. While accident reports for automated parking systems are not yet available, it can be assumed that incidents involving human drivers may represent hazardous situations within the ODD that also pose challenges for automated systems. However, the underlying causes differ: humans, who rely mostly on vision [Xie+22], struggle with multi-tasking and with estimating velocities of other traffic participants, whereas radar systems are constrained by wave propagation effects and partially by low object reflectivity. As a result, radar perception failures may occur even in situations that humans can handle with relative ease.
- *Real-world datasets:* Datasets provide an arbitrary sample of typical situations occurring within the ODD, which can be used to enrich the scenario database [Rie+20]. Although these datasets often lack detailed descriptions, they usually contain odometry and sensor data. Current research investigates how this data can be structured and expressed in natural language, which facilitates efficient search and retrieval. A key advantage of such datasets is their relatively balanced representation of objects and locations, which may encompass critical or challenging scenarios.

The situation sets derived from these three sources are combined to form the SOTIF-driven minimal testing set. Observing the erroneous behavior of the perception function within this set enables a systematic estimation of the overall failure rate for the complete parking function. The following section provides an exemplary analysis of the ODD of the AVP function according to the three sources presented above.

Expert Knowledge

As described above, expert knowledge is a generic term that refers to information about the ODD not necessarily derived from data but from expert conclusions. In this work, the catalog of technical requirements for the autonomous driving function AVP, published by the Kraftfahrt-Bundesamt, Federal Motor Transport Authority (KBA), has been applied [Kra22]. This catalog defines details regarding the state machine for activating the AVP function, system monitoring, communication pathways, and the testing process. Alternatively, the ISO 23374 standard lists a scenario proposal for the performance evaluation of the AVP function [Int23].

The KBA catalog [Kra22] lists all scenarios that the AVP function is expected to handle. The relevant scenarios primarily concern collision avoidance, where radar sensors may serve as components of the perception system. All relevant test scenarios are extracted and aligned with scenarios identified from the other sources.

Accident Reports

As an example of an accident report, a study by the Gesamtverband der Deutschen Versicherungswirtschaft, German Insurance Association (GDV) was analyzed, focusing on accidents in parking situations involving pedestrians or bicyclists [Sch+20b]. This study provides detailed analyses of accident classes, causes, and involved traffic participants. A drawback of this study is that most documented accidents occurred outside parking garages. Moreover, it includes dooring accidents and incidents that occurred during loading or unloading. Since this work focuses on safety-oriented evaluation during transfer rides and the parking process itself, such accidents are excluded. The study distinguishes between direct and indirect accidents. Direct accidents involve collisions between a VRU and a vehicle during the actual parking maneuver. Indirect accidents involve parking vehicles that are causative factors in the incident. Both accident types are relevant for this work: the automated function must address challenges arising from parking vehicles (indirect accidents) during transfer rides to the parking lot as well as challenges during the maneuver into the parking space (direct accidents).

Real World Datasets

Real world datasets may contain scenarios not identified by the first two methods. This often occurs when hazardous situations arise from a combination of factors. Such datasets provide an opportunity to assess the completeness of the test catalog, turning known unknowns into known knowns. If a dataset is still unstructured, applying a methodology to evaluate the criticality of situations reduces the manual effort required to screen the entire dataset. For instance, a method may evaluate situations in which objects are located inside the driving tube with a low TTC. This approach is conceptually

similar to the area criticality estimation introduced in Chapter 3. Ideally, the dataset should contain scenarios consistent with the ODD, and its recorded sensor modalities should be comparable with the perception system under evaluation. In this work, the RadarScenes dataset, described in Subsection 5.4.1, is used as an example. This dataset was also applied to investigate the successive point cloud treatment introduced in Section 5.2.

Summary of the Testcases

Table 5.1 aligns scenarios from accident reports (GDV) with expert knowledge (KBA) and adds a generalized type description summarizing each relevant scenario. When environmental attributes — such as interference from other sensors or weather effects — are considered in the variation of these scenarios, a high coverage of the AVP function’s ODD can be achieved. The dataset includes these scenarios and supports the aforementioned variation of the test cases.

Table 5.1: Condensed, relevant scenarios used to derive perception-system failure rates for the AVP function, integrating scenarios from multiple sources.

GDV ID	KBA ID	Generalized Type Description
D1.1	2.3.2 (1.5)	VRU in parking lot and vehicle enters parking lot backwards
D1.2	2.3.2 (1.5)	VRU in parking lot and vehicle enters parking lot forwards
D1.4	2.3.1 (1.1)	VRU in front of parking lot and vehicle leaves parking lot forwards
I1	2.3.4 (1.1/1.2)	VRU moving laterally, covered by vehicle
I4	2.3.4 (1.1/1.2)	VRU moving laterally at crossing
I7	2.3.1 (1.2)	VRU in front of parking lot and vehicle leaves parking lot backwards
I10		VRU moving longitudinally, covered by vehicle
	2.3.1 (1.3/1.4)	VRU laterally next to vehicle
	2.3.3 (1.1)	VRU in the driving tube violating the minimum distance requirement
	2.3.2 (1.1/1.2)	VRU in driving tube, not moving
	2.3.2 (1.3/1.4/1.6/1.7)	VRU behind ramp
	2.3.3 (2.1)	VRU moving longitudinally parallel to vehicle
	2.3.4 (1.3)	VRU approaches vehicle longitudinally

This methodical scenario search approach integrates information from three distinct sources into a comprehensive scenario catalogue. This strategy is applied to achieve

the objectives of the SOTIF considerations introduced in Subsection 2.4.2, namely by reducing the number of unknown scenarios, as illustrated in area 3 of Figure 2.8.

5.3.2 Point Set Quantification

To ensure that the information available from radar sensors about present objects is fully exploited, any loss of information must be accounted for. Information leaks can emerge within each stage of the the processing pipeline. Especially if safety-aware methods are not applied during signal processing, critical points may remain undetected. At the same time, the information provided by the point cloud is sometimes insufficient to form an object hypothesis. The objective of the approach described in this section is to quantify the occurrence of such situations, which may still arise if safety-aware processing, as proposed in Section 5.2, is applied.

These eventualities are investigated by defining four quantities A , B , C , and F . Quantity A comprises all points that received specific treatment because they were assessed as sufficiently critical by the drive-tube-based criticality evaluation method. Points on critical objects that could lead to a hazardous situation if not processed belong to quantity B . Using these two quantities, the numbers of FPs and FNs can be derived. Here, the terms *False Positive* and *False Negative* refer to the point's criticality assessment and not to the detection or lacking detection of objects of interest.

A hazardous situation can occur if TP critical points are removed from the point cloud after safety-aware processing. For this reason, quantity F is defined, containing all points that remain after applying the filtering cascade. The points in F serve as the input for the clustering algorithm.

The most important measure of the approach is how many relevant radar points in B are processed into a cluster. These clusters form the foundation for the next step in radar data processing: the generation of object hypotheses. To capture this, quantity C is defined, containing radar points that have been successfully added to a cluster. This occurs if the density of radar points reflected by an object of interest is high enough to enable clustering. For this chapter, the processing chain is considered up to the formation of radar clusters.

Figure 5.5 visualizes the quantities A , B , and C .

Each subset of points can be interpreted in terms of its relevance to the safe perception function:

- $A \cap \bar{B}$: FP — high criticality assessed erroneously, which may degrade the robustness of the perception function.
- $\bar{A} \cap B \cap C$: FN — criticality of the radar point not identified, but the point was added to a cluster and thus successfully exploited.




Figure 5.5: Quantification of critical radar-point processing. The method distinguishes between points assessed as critical (set A), points with a critical impact (set B), and successfully clustered points (set C).

- $\bar{A} \cap B \cap \bar{C}$: FN — criticality of the radar point not identified, and the point density is too low for clustering. This may lead to a hazardous situation.
- $A \cap B \cap C$: TP — radar point is critical and successfully clustered.
- $A \cap B \cap \bar{C}$: TP — radar point is critical, but density is insufficient for clustering despite successful recapture.

Each point in $\bar{F}$ is not considered for clustering. Thus, investigating quantity $\bar{F}$ allows for identifying potentially missed hazardous objects. If a radar point belongs to the set $B \cap \bar{F}$, it has been erroneously removed from the point cloud, thereby degrading the safety of the perception function. A FN radar point whose criticality is not correctly assessed (set $\bar{A} \cap B$) will also fall into this set if it is not located in any criticality regions and if it is too weak to pass the rigid filters. Nevertheless, an object of interest can still be detected without this point, provided that enough points in F remain to form a cluster.

5.3.3 Failure Rate Derivation

The methods outlined in the preceding subsections contribute to the overall derivation of the failure rate. For the release of the perception function for SAE Level 4 automated parking functions, the probability of an incident resulting in personal injury is particularly relevant. This subsection estimates how likely such an incident is to occur as a consequence of a removed radar point.

Figure 5.6 illustrates all links from the use of the automated parking function within the ODD to the occurrence of an incident. Each link has an associated transition probability, derived methodically.




Figure 5.6: Illustration of the causal chain from applying the function inside the ODD to an incident resulting in personal injury.

The first transition is the scaling from a general function application in the ODD to specific critical scenarios, as described in Subsection 5.3.1. The second transition describes the conditions that lead to an object not being detected for at least one cycle, quantified methodically in Subsection 5.3.2. This transition is the transition discussed in this thesis.

The third transition considers the case in which a non-detected object causes a collision. This occurs if an object remains undetected for multiple frames. The relevant observation time begins at the moment when the object is about to enter the driving tube and ends when no further action for collision avoidance or mitigation can be taken anymore. The closer an object, the higher is the probability for an undetected object to cause a collision. However, the detection of closer objects should be easier at the same time.

The fourth transition derives the probability of an injury in the event of a collision. This probability corresponds to the severity classes defined in ISO 26262. The classes divide into four categories: no injuries, light injuries, severe injuries (survival probable), and fatal injuries (survival improbable). Both the likelihood and severity of injury increase with higher collision speed.

5.4 Evaluation on RadarScenes Dataset

This section evaluates the potential to reduce FN points through the proposed successive point cloud treatment method. To enable large-scale conclusions and support further research in this field, a public dataset is used that is described in Subsection 5.4.1. As the dataset lacks sufficient critical scenarios for a comprehensive analysis of the method, a crucial part of the experiment involves modifying the dataset. These modifications enable the investigations of critical situations in data that otherwise consist primarily of driving scenes recorded in typical traffic conditions, which naturally contain only a few critical situations.

5.4.1 Dataset

For this evaluation, the objective was to obtain a radar point cloud representative of the state of the art, meaning that the cloud has both high point density and frame rate. For this reason, the RadarScenes dataset [Sch+21b], recorded in German cities, was selected. RadarScenes features point cloud data from four radars, which cover the front and parts of the vehicle's sides with minor overlaps in their FoVs. The mounting angles of the sensors in the vehicle coordinate system are $[-85^\circ, -25^\circ, 25^\circ, 85^\circ]$. Each radar sensor has a FoV covering an azimuth angle of $\varphi = 120^\circ$ and a range of $r = 100$ m. The sensor cycle time is 60 ms. Each point in the point cloud includes attributes such as timestamp, ID of the sensor recording the point, measured range r , measured azimuth angle θ (converted to Cartesian coordinates x and y in both the *vehicle* and *odom* coordinate system), Doppler velocity v_{rel} , RCS value, a unique radar point ID, a GT object hypothesis ID, and a class label.

A key requirement for evaluating the proposed methodology, fulfilled by the RadarScenes dataset, is the availability of odometry data. For each point cloud, the current position and orientation of the recording vehicle in the odometry coordinate system $(x_{\text{veh}}, y_{\text{veh}}, \gamma)$, its current velocity v_{cur} , and its yaw rate ω are stored. This enables the generation of a drive-tube-based on odometry data, which in turn allows for per-point criticality assessment. In addition, the dataset contains camera data intended for scene overview by human observers. However, the objects in the images are anonymized, and the camera is not calibrated. RadarScenes covers a total length of 92.01 km with 4.45 h of recordings split into 158 sequences. Each individual segment of a sequence is called a *scene*. As the dataset was designed to provide a balanced representation of typical urban traffic, it contains a relatively high number of pedestrians and pedestrian groups.

5.4.2 Experimental Setup

This section introduces the required modifications to the dataset as a preparatory step for the experiments. It then outlines the setup of three experiments, which include investigation of the criticality estimation and criticality-aware filtering, as well as a comparison between the proposed criticality metric and an alternative one.

Reachability Analysis

As a first step, the dataset is examined for a sufficient number of critical scenarios. For this purpose, the number of critically assessed radar points is analyzed using the criticality assessment outlined in Chapter 3. The results show that critical radar points in the driving tube occur far too rarely in the unmodified dataset. However, the shape of the driving tube can be altered to artificially induce critical situations. Since the recording vehicle does not follow this altered driving tube, the concept cannot

be applied to analyze multiple frames in sequence. Nevertheless, it produces a large number of critical scenes.

To generate additional critical points, the per-point labels provided in RadarScenes are utilized. Many VRUs move along pedestrian walkways adjacent to the recording vehicle, while the vehicle trajectory is not directed toward them. For each frame, each point labeled as pedestrian, pedestrian group, or bicyclist is analyzed for its reachability under the current motion state of the recording vehicle.

Definition 5.4 (Point Reachability)

A radar point is defined as reachable and can potentially be evaluated as a critical point if a circular arc with radius r and angle α between the recording vehicle and the radar point can be constructed such that the following conditions hold:

- $r \geq r_{\min}$: the radius must be at least equal to the minimum drivable radius, assumed to be $r_{\min} = 6.0$ m.
- The fictitious TTC t_{TTC} must be below 3.0 s to consider the point reachable and sufficiently critical.
- The longitudinal acceleration of the vehicle along the arc must not exceed $a_{\text{long, max}} = 10.0 \text{ m/s}^2$ and the lateral acceleration must be less than or equal to $a_{\text{lat, max}} = 8.0 \text{ m/s}^2$.

Figure 5.7 illustrates the variables used in this definition.

The radius can be derived from the right-angled triangle:

$$r^2 = x_{\text{VRU}}^2 + (r - y_{\text{VRU}})^2, \quad (5.10)$$

$$r = \frac{x_{\text{VRU}}^2 + x_{\text{VRU}}^2}{2 \cdot y_{\text{VRU}}}. \quad (5.11)$$

If the radar point is generally reachable, a critical trajectory toward the point can be computed. To maximize the number of reachable points, a constant acceleration of at most $a_{\text{long, max}}$ toward the point is assumed. Consequently, the velocity increases linearly from v_{cur} to v_{coll} while moving along the circular arc. The fictitious TTC can then be computed with:

$$t_{\text{TTC}} = \frac{2 \cdot r \cdot \arcsin \frac{x_{\text{VRU}}}{r}}{v_0 + v_{\text{coll}}} - t_0. \quad (5.12)$$




Figure 5.7: Method to determine radar point reachability. Illustration of the reachability of a radar point located at position (x_{VRU}, y_{VRU}) .

In the RadarScenes dataset, two main modes of operation are considered: fictitious acceleration toward pedestrians crossing the road in front of the recording vehicle, and steering toward pedestrians on the sidewalk.

Experiment for Criticality Estimation

This experiment utilizes the artificial critical points induced by the method described above, with the objective of determining whether they are correctly evaluated as critical when applying the modified area-based criticality assessment method. The challenge lies in identifying GT critical points and employing a metric that is independent of the assessment procedure itself. In the experiments, all reachable points are considered GT critical points. Additionally, each radar point with the same *track ID* attribute as a reachable point is also considered critical and should receive a non-zero criticality score. To ensure comparability, only frames in which the recording vehicle's velocity does not exceed $v_{\max, \text{domain}} = 8.5 \text{ m/s}$ are included. Situations in which the vehicle exceeded this velocity were excluded in the experiment. Due to this approach, no continuous value for GT criticality is available. Instead, each GT critical point is compared with the evaluated metric.


- A GT critical point with an assessed criticality greater than zero is counted as a TP: $p \in (A \cap B)$.
- All points with an assessed criticality $p_{\text{crit}} > 0$ but not corresponding to GT critical points are counted as FPs: $p \in (A \cap \bar{B})$.

- Points with an assessed criticality of zero, despite being reachable and labeled as a VRU reflection, are counted as FNs: $p \in (\bar{A} \cap B)$.

A scene in a sequence from RadarScenes containing both TP and FP assessments is depicted in Figure 5.8. Figure 5.8a shows the corresponding camera image, and Figure 5.8b the radar point cloud with critically assessed points highlighted in orange. The modified path toward the reachable point corresponding to the VRU is shown as a red solid line. Since the vehicle is stationary in this scene, no original odometry path is visible in the plot.



(a) Camera Image [Sch+21b]



(b) Radar Image

Figure 5.8: Exemplary radar point reachability. Reachability of a radar point at position (x_{VRU}, y_{VRU}) for sequence 63 at timestamp 20727984889861.

Contrary to the safety-aware treatment of the radar spectrum, the actions determined in this experiment are not dependent on the actual criticality value of a point. Therefore, a threshold $crit_{thresh}$ must be chosen to classify a point as critical. If the threshold is set

too low, all noise points within the driving tube may be retained, even if they are only marginally critical. Conversely, an excessively high threshold may pose a safety hazard by causing critical points to be erroneously excluded. The resulting sets A , containing points with $p.crit \geq crit_{thresh}$, and B , including the reachable TP critical points of VRUs, are compared across different threshold values to identify the setting of $crit_{thresh}$ that best balances the performance metrics.

Experiment for Criticality-Aware Filtering

After evaluating the criticality assessment, the next step investigates the impact of adapting the filtering process to be safety-aware on the number of correctly processed critical points. In this experiment, the individual components — driving tube calculation, radar point criticality assessment, rigid filtering, clustering, and the recapturing mechanism — are combined into a single framework, illustrated in Figure 5.9.




Figure 5.9: Framework for safety-aware radar point cloud processing. The proposed framework comprises elements for conventional perception (blue) and components for recovering points potentially critical to the perception task (red).

Points that pass the rigid filtering stage are concatenated with points evaluated as critical and subsequently passed to the clustering algorithm. The artificial critical trajectory derived from the reachability analysis is used as the trajectory reference. As a baseline, the experiment first evaluates how many radar points are removed by standard rigid filtering. Four different RCS thresholds are applied, while filtering by position and velocity remains unchanged. The tested RCS thresholds are -15.0 dB m^2 , -10.0 dB m^2 , -5.0 dB m^2 , and 0.0 dB m^2 . A standard DBSCAN configuration is applied for the clustering, with an ϵ -distance of $\epsilon = 1.0$ and the minimum number of points for each cluster $\min_{pts} = 4$. These parameters were determined empirically and correspond well to the characteristics of a radar point composed of measurements from multiple sensors, in which at most two sensors overlap in their FoV. For each case, the number of points

in sets C and F is computed. This enables a direct comparison between the baseline and the clustering results obtained with the safety-aware filtering method, that includes the recapturing mechanism. Besides of the evaluation in the RadarScenes dataset, this mechanism is also evaluated in an experiment with real-world data, outlined in Section 6.3.

Experiment for Comparison with Doppler-Based Criticality Metric

The third experiment is designed to compare the proposed criticality metric with an alternative metric for point-level criticality assessment. Unlike the proposed method, which is based on the position of points within the driving tube, this alternative relies on the velocity of points relative to the recording vehicle. The underlying assumption is that radar points corresponding to non-static objects have a greater impact on potential policy changes of the underlying automated parking function than points corresponding to static objects. Accordingly, non-static points should receive a high criticality, while static points should not be assigned criticality. This reflects the intuition that moving VRUs are more likely to induce hazardous situations. The purpose of this experiment is to evaluate the performance of the proposed drive-tube-based criticality metric against a metric that does not incorporate the ego vehicle's driving state or its predicted occupied space at the next time steps, but instead relies solely on relative velocity as a criticality measure. The results of this experiment validate the proposed approach. Using the current velocity of the recording vehicle v_{cur} , the measured azimuth angle φ in the sensor coordinate system, and the measured Doppler velocity v_{rel} of the point, the motion-compensated relative velocity $v_{\text{rel, comp}}$ can be computed as

$$v_{\text{rel, comp}} = v_{\text{rel}} - v_{\text{cur}} \cdot \cos \varphi. \quad (5.13)$$

Due to inaccuracies in angle and Doppler velocity measurements, the compensated velocity $v_{\text{rel, comp}}$ of static objects may fluctuate around zero. To account for this, the velocity-based criticality value $\tilde{crit}_{\text{vel}}$ ramps up linearly for $0 \leq v_{\text{rel, comp}} \leq 1.5$ m/s. The criticality value is determined as

$$\tilde{crit}_{\text{vel}} = \begin{cases} 0.0 & \text{if } |v_{\text{rel, comp}}| = 0.0 \text{ m/s,} \\ \frac{|v_{\text{rel, comp}}| - 0.4 \text{ m/s}}{1.1} & \text{if } 0.0 \text{ m/s} < |v_{\text{rel, comp}}| \leq 1.5 \text{ m/s,} \\ 1.0 & \text{if } |v_{\text{rel, comp}}| > 1.5 \text{ m/s.} \end{cases} \quad (5.14)$$

In this experiment, the module that recaptures radar points based on criticality derived from the ego vehicle's driving tube is replaced by a module that recaptures points according to their velocity. The rigid filtering process remains the same as described in Subsection 2.5.3, with the additional condition that all radar points with an RCS value less than or equal to -10.0 dB m² are removed.

5.4.3 Experimental Results

This section presents the results of the conducted experiments that were described in the previous section.

Results of Criticality Estimation Experiment

Figure 5.10 shows recall, precision, and F_1 -score as functions of the criticality threshold value.




Figure 5.10: Detection metrics score as a function of the criticality threshold. Illustration of recall, precision, and F_1 -score across varying criticality thresholds.

Recall decreases approximately linearly with increasing criticality threshold, while precision reaches its maximum at thresholds in the range $0.3 < crit_{\text{thresh}} < 0.5$, although the overall precision remains relatively low. The reasons for this limited precision are discussed in Section 5.5. Two thresholding strategies can be identified. If high recall of 95 % or more is required, a low criticality threshold is advisable; in this case, $crit_{\text{thresh}} = 0.10$ is selected. If both recall and precision are to be balanced, a higher threshold of $crit_{\text{thresh}} = 0.35$ is preferable, as this threshold features the optimal F_1 -score. Table 5.2 summarizes the differences between the safety-focused and balanced criticality thresholds.

The safety-focused threshold includes 50.4 % more radar points in set A compared to the balanced threshold. Consequently, recall is higher, which is particularly relevant for criticality-oriented processing of subsequent point clouds. For this reason, the subsequent evaluation of safety-aware point cloud filtering was conducted using the lower criticality threshold, $crit_{\text{thresh}} = 0.10$. Under this configuration, all radar points

Table 5.2: Comparison of detection metrics for two criticality thresholds. Evaluation of safety-focused and balanced criticality thresholds $crit_{\text{thresh}}$.

	$crit_{\text{thresh}} = 0.10$	$crit_{\text{thresh}} = 0.35$
$\sum$ Points in set A	4,086,920	2,716,981
Recall	95.02 %	74.81 %
Precision	3.69 %	4.37 %
F_1 -score	7.10 %	8.26 %

with a criticality greater than this threshold are considered critical points and included in set A .

Results of the Baseline Filtering Experiment

In this experiment, the effect of applying a rigid filter to set B , consisting of TP reachable and critical points, is investigated. A total number of 187,106 radar points are assigned to set B . Table 5.3 lists the quantities of remaining points $\sum p_{\text{rem}}$ and clustered TP radar points $\sum B \cap C$ after applying standard rigid filters with varying RCS thresholds. These values are contrasted with the number of incorrectly removed TP points $\sum B \cap \bar{F}$ and the number of TP points that are not clustered $\sum B \cap \bar{C}$.

Table 5.3: Detection performance without safety-aware filtering. Baseline experiment showing the performance of filtering and clustering with rigid filters only.

RCS Threshold	$\sum p_{\text{rem}}$	Filter Rate	$\sum C$	$\sum B \cap C$	$\sum B \cap \bar{F}$	$\sum B \cap \bar{C}$
0.0 dB m ²	1,448,442	1.44 %	163,018	77	185,203	1,085
−5.0 dB m ²	2,840,936	2.83 %	595,928	783	176,752	8,930
−10.0 dB m ²	5,160,073	5.14 %	1,721,085	13,039	146,314	27,040
−15.0 dB m ²	9,319,639	9.28 %	4,281,094	60,165	98,530	27,651

The filter rate denotes the proportion of points that remain after rigid filtering. While most points are removed due to positional limits, the number of remaining points p_{rem} decreases substantially when stronger filters, such as 0.0 dB m², are applied. With an increasing number of neighboring points, the number of clustered points in set C grows disproportionately, provided that more points remain for clustering. Points in set B can follow three outcomes: they can be added to a cluster, which is the desired behavior ($B \cap C$), removed by the rigid filter ($B \cap \bar{F}$), or remain as individual radar points not belonging to any cluster ($B \cap \bar{C}$). Lower RCS thresholds result in higher numbers of correctly clustered points. Nevertheless, even under less stringent RCS filtering, a substantial share of TP points is removed by rigid filtering ($\frac{\sum B \cap \bar{F}}{\sum B}$): 78.5 % at a threshold of −10.0 dB m² and 52.9 % at −15.0 dB m². This finding suggests that

the filter threshold must either be lowered further — at the expense of introducing additional FPs — or that complementary strategies, such as recapturing points evaluated as critical, must be employed. Some points, however, appear to be located in regions of insufficient density to form clusters and therefore cannot be converted into clusters.

Results of the Criticality-Aware Filtering Experiment

Table 5.4 presents the results of successive point cloud treatment with recapture of critical points. As before, four different RCS thresholds were applied, and the percentage improvement over the standard processing pipeline without recapture of critical points is reported.

Table 5.4: Detection performance with safety-aware filtering. Criticality-aware filtering experiment showing the performance of filtering and clustering with safety-aware mechanisms applied.

RCS Threshold	Filter Rate	$\sum C$	$\sum B \cap C$	$\sum B \cap \bar{F}$
0.0 dB m ²	6.83 %	3,961,374 (+2,330 %)	142,699	31,392
−5.0 dB m ²	8.01 %	4,337,161 (+628 %)	142,862	30,929
−10.0 dB m ²	10.04 %	5,339,292 (+210 %)	143,312	30,004
−15.0 dB m ²	13.66 %	7,601,955 (+78 %)	144,934	28,383

The results indicate that the safety-aware filtering approach consistently produces a higher number of clustered points than the standard method, even though the filter rates are approximately equal. Moreover, the number of correctly treated points remains largely independent of the RCS threshold, stabilizing at approximately 76 %.

The proportion of TP points in set B that are erroneously removed ($\frac{\sum B \cap \bar{F}}{\sum B}$) decreases significantly when criticality-aware filtering is applied. For the weakest rigid filter (RCS threshold = −15.0 dB m²), the share is reduced from 52.9 % to 15.2 %. For the strongest filter (RCS threshold = 0.0 dB m²), 16.8 % of TP points are erroneously removed, compared with 99.4 % in the baseline case. Despite these improvements, further reduction of erroneously removed points is desirable. The causes of these remaining cases, which may pose potential safety hazards, are discussed in Section 5.5.

Results of the Comparison with Doppler-Based Criticality Metric

Table 5.5 compares the drive-tube-based recapture method with the velocity-based method for criticality evaluation.

The velocity-based evaluation method recaptures more than twice as many points as the drive-tube-based method. Consequently, the filter rate is higher. This approach also increases the number of FNs ($B \cap \bar{C}$), as more points are identified as critical

Table 5.5: Comparison of two criticality assessment metrics, comparing the perception performance metrics between the drive-tube-based criticality metric and the velocity-based criticality metric.

	Drive-tube-based metric	Doppler-velocity-based metric
Filter Rate	10.04 %	14.11 %
FP Points ($\sum A \cap \bar{B}$)	3,884,178	8,626,278
FN Points ($\sum \bar{A} \cap B$)	36,243	42,193
Points in set A ($\sum A$)	4,086,920	9,608,622
Correctly treated points ($\sum B \cap C$)	76.60 %	82.04 %
Points in set C ($\sum C$)	5,339,292	5,098,049

($\sum A$). The share of points incorrectly assessed as critical increases from 19.4 % in the drive-tube-based metric to 22.5 % in the Doppler-velocity-based metric. The increase in FPs — points erroneously identified as critical ($\sum A \cap \bar{B}$) — is even more pronounced. The number of FPs is 122 % higher than in the drive-tube-based method, indicating substantially lower robustness. Interestingly, the total number of clustered points is nearly the same across both methods. The velocity-based method benefits from the higher filter rate, resulting in a slightly higher proportion of correctly clustered points $B \cap C$ (82.04 % compared to 76.60 %).

5.5 Discussion

Filtering point clouds by removing points that do not meet certain criteria is a common method for reducing noise and clutter. The previous Chapter 4 highlighted that applying such filters independently of the perception and driving context may lead to critical behavior in the automated function. Relevant points may be removed erroneously and therefore not incorporated by the planner when generating its policy. This chapter presented a method for introducing safety aspects into radar-based perception at the point cloud level. The experiments show that the proposed approach significantly reduces the number of filtered critical points and is even able to keep this number at a constant, low level, regardless of the general filter limits. Nevertheless, the number of unprocessed critical points (set B) could still not be reduced to zero. At the same time, the precision rate remains very low. The reasons for this can be identified in various aspects.

First, the presented point cloud assessment method operates exclusively at the point cloud level. It is therefore not capable of distinguishing between radar reflections caused by clutter or noise, reflections from static objects such as pillars or street lamps, and reflections from VRUs. According to the definition used in this work, only ignoring the latter group of objects leads to a safety hazard. Although the representation of a pillar


and a pedestrian may be similar within a radar point cloud, the intended processing differs significantly. To address this issue, downstream point cloud processing must exploit additional information to maintain the number of FPs at an acceptable level. This can be achieved by associating successive point clouds, fusing radar points with camera images, and applying neural networks for object detection in point clouds, which are currently advancing rapidly. However, no classification at the point cloud level was conducted in this experiment, since such classification inevitably introduces safety risks if it fails for individual critical radar points. In this experiment, this limitation results in a significantly reduced precision.

Second, in this chapter, the RadarScenes dataset is modified to induce critical situations for evaluating the pipeline using reachability analysis. Nevertheless, it remains imperfectly suited for a detailed investigation of the underlying mechanisms. One reason is that the trajectory of the recording vehicle is altered artificially, while time-dependent methods are difficult to analyze due to the interdependence between the perceived environment and the vehicle's motion. The approach used in this experiment classifies radar points as critical as soon as they are reflections of a VRU within the altered driving tube. However, this approach may omit other critical points and include points that are not critical due to their current motion. Consequently, the composition of TP points in set B may not be fully optimal. These considerations highlight the importance of investigating the proposed approaches, such as successive point cloud processing, on real-world datasets containing genuine critical situations. Such an experiment is part of the subsequent Chapter 6. Furthermore, the dataset does not contain labels for static objects or noise points. While retaining reflections from static objects is desirable, another relevant objective of the processing pipeline is to remove as many noise points as possible. This limitation in the dataset prevents a detailed comparison between removed static-object points and noise points using the proposed algorithm.

The experiments show that processing successive point clouds does not prevent a significant number of points in set B from being removed by the filter. Consequently, these points are not included in set F , which is used for clustering. The recall determined in this experiment does not approach a value of 1. This is because the criticality of points decreases rapidly once they lie outside the driving tube. During the experiment, all radar points sharing the same track ID as a reachable point were stored as TP critical points. However, this approach can lead to cases in which points in set B are located close to the driving tube but are already outside it, particularly for large objects. Another situation in which points belonging to the same cluster receive a criticality value of zero occurs when their distance to the recording vehicle is too small to allow for a reaction. Figure 5.11 illustrates such a scenario: one point of a pedestrian group is identified as the reachable point. The trajectory, shown as a red solid line, is modified toward this radar point. Blue circles around neighboring points indicate that the target point, marked by a red circle, belongs to a larger group of reflection points representing a pedestrian. Although the target point p receives a high criticality value ($p.crit > 0.6$), other points with the same track ID are assigned a criticality value of zero.



(a) Camera Image (Image is about 1.0 s ahead of point cloud for better visibility of VRUs) [Sch+21b]



(b) Radar Point Cloud

Figure 5.11: Variation in criticality values among neighboring points. Radar points belonging to the same GT track ID receive different criticality values due to positional differences.

A key characteristic of the area-based criticality assessment is that the calculated criticality still depends strongly on the spatial position of the point, even though measurement uncertainties are incorporated. By generating criticality regions, this limitation is partially compensated for by recapturing newly occurring radar points in nearby locations, while tolerating that they may not be measured at exactly the expected position.

In addition to these aspects, which affect the overall experimental setup and motivate further investigation in a dedicated real-world experiment, another observation concerns the clustering approach. A qualitative assessment conducted within the experiments revealed that requiring a minimum of four radar points to form a cluster ($\min_{pts} = 4$) can pose a safety risk. Although measurements from four sensors are fused into a single point cloud, some VRUs do not generate a sufficient number of reflections to satisfy this requirement. In particular, under high noise conditions, pedestrians in the RadarScenes dataset may be represented by only two or three points. This conflicts with

the objective of maintaining a low number of FP tracks when only a few noise points are present. One possible mitigation strategy is to allow clusters to form with fewer radar points, provided that these points exhibit sufficiently high criticality values. However, a thorough evaluation of this approach requires analysis at the object-hypothesis level. Since the tracking step is not included in the framework presented in this chapter, this discussion is deferred to the subsequent Chapter 6.

In summary, this chapter presents a methodology for improving the filtering procedure in safety-relevant situations by avoiding the unintended omission of critical radar points during processing. This is achieved by incorporating contextual aspects such as the planned driving tube and the recapturing of radar points that are potentially relevant for safe functional behavior. The proposed approach addresses the next part of (RQ2) at the radar point level by recapturing points that have received a high criticality assessment. The availability and spatial density of these points at the clustering and tracking stages play a major role in detecting and classifying the corresponding objects. Consequently, the proposed method increases the likelihood that critical points are recognized as objects, thereby promoting higher functional safety.

Furthermore, this chapter presents a method to quantify the FIT rate of radar point-cloud-based perception. The core idea of this method is to count the radar points that lead to hazardous behavior within a representative set of scenarios that approximates the ODD of the function under test as accurately as possible. Section 5.3 describes a procedure for deriving the FIT rate attributable to the perception sensor based on points that are erroneously excluded during processing. This procedure accounts for the consequences of FN object detections within the ODD by evaluating the probability of a hazardous situation when the investigated perception pipeline is applied. In this way, (RQ3) is addressed. The method focuses on the impact of the perception function itself, spanning the processing chain from raw sensor data to the object list, while excluding effects caused by sensor underperformance or limitations of the downstream planner. Applying this method to the proposed safety-aware processing algorithms may demonstrate that these algorithms provide stronger guarantees for safe perception.

6 Processing Adaptations at the Object Level¹²

This chapter provides insights into utilizing the estimated criticality in later stages of radar data processing, which occur downstream of the radar point cloud. While Chapter 5 discusses the recovery of radar points through regional awareness during the filtering process, the evaluated criticality of the radar points can also serve as a means to generate object hypotheses with criticality awareness. Section 6.1 motivates the need for investigating these methods and discusses their potential. In Section 6.2, three approaches are developed, based on the radar data processing pipeline illustrated in Subsection 2.3.2. The evaluation of these approaches is performed on a dataset that was explicitly recorded for this purpose. The content of this dataset and the evaluation methods are outlined in Section 6.3. This section also presents a comparison of the three approaches with a baseline object detection pipeline that lacks criticality awareness mechanisms. The chapter concludes with Section 6.4, which reviews the potential of the proposed algorithms and the established dataset, as well as general opportunities and drawbacks of modified object detection in the later stages of radar data processing.

6.1 Tracking of Object Hypotheses

For safe radar data processing, it is crucial to preserve object evidence throughout the initial processing steps. However, the point cloud resulting from DSP and peak detection via CA-CFAR still contains too many points to incorporate each of them into the planning process. The automated function reacts to objects in an object list that represent the dynamic elements of the current scenario, such as VRUs moving close to the ego vehicle. Before an object hypothesis is evaluated by the planner, the object evidence such as a set of radar points undergoes clustering and tracking, which serve as additional filtering steps that merge information from several radar measurement frames. As Figure 6.1 illustrates, the number of instances decreases from the point level to the object level, while the computational effort per instance increases. Tracked object hypotheses include an estimated motion vector and are associated with existing hypotheses, possibly continuing the believe. They can also bridge gaps for the hypothesis in cases where no radar point is observed during single cycles. Consequently, an object hypothesis may persist in a location even if no radar point has recently been observed, and a single measurement point does not directly guarantee the creation of an object

¹² Chapter builds on [Brü+25b].

hypothesis. Moreover, it is important to emphasize that conventional tracking algorithms create object hypotheses with uniform sensitivity across the sensor's perception field, based solely on probabilistic models, and independently of the relevance of the observation to the driving task.




Figure 6.1: Schematic comparison of levels in late radar data processing. Comparing of the number of instances and the computational effort per instance at each level of late radar data processing.

This raises the question of whether object hypotheses should be created even when the existence probability for an object is low, but the potential consequences of ignoring it could be severe. In Chapter 3, a methodology is proposed to evaluate the criticality of radar measurement data at the point level. The information can then be utilized to modify the downstream processing of radar points to introduce safety-aware behavior. In particular, these modifications are applied to the clustering and tracking algorithms. Adapting the behavior of clustering and tracking fosters the creation of object hypotheses when points are assessed as critical. Existing literature applies such criticality evaluations at the object list level [WDE21; Lys+22; CM24], but pays little attention to the fact that critical objects may not appear on the object list at all if their hypotheses were removed by the clustering or tracking algorithm. While research on the algorithmic side has concentrated on advanced methods, such as the application of deep learning, as described in Subsection 2.5.1, modifications to the core algorithmic concept itself have not been sufficiently considered.

Object evidence that remains in the radar point cloud can disappear during the clustering stage due to the high sparsity of reflections on the object or because reflection points are spread over a large distance. In radar point clouds, such dispersion is often caused by imprecise angular measurements, which spread the points far enough apart that algorithms such as DBSCAN cannot group them. At the tracking stage, object evidence may vanish if the association between the prediction and the new observation fails, often due to imprecise measurements in one or both steps. In such cases, the weighting between the motion model and the recent observation may be unbalanced.

A detailed analysis of the mechanisms in modified clustering and tracking algorithms requires a dataset that contains critical scenarios. Since the complete AVP function is under consideration, it is important that this dataset should be recorded in the target

domain of this function, ideally a parking garage with a high density of metal structures. The literature review in the preceding chapters revealed that such datasets, particularly those featuring critical scenarios with VRUs, are not yet available. For this reason, a custom dataset must be systematically recorded to evaluate the methods developed in this chapter.

6.2 Clustering and Tracking Modifications

This section proposes two modifications that facilitate the creation of objects based on points assessed as critical. Both modifications rely on per-point criticality evaluation. Subsection 6.2.1 presents a modification of the DBSCAN algorithm, which allows clusters to form more readily if the points in the cluster are evaluated as critical. The second variant, described in Subsection 6.2.2, enables critical points to be directly tracked as object hypotheses.

6.2.1 Clustering Modification Approach

Radar points are grouped together via clustering algorithms such as DBSCAN to reduce the sensitivity of the parking function to individual FP points. The DBSCAN algorithm is explained in detail in Subsection 2.3.4. It is based on a distance metric that identifies points located close to each other. In the simplest case, this distance metric is Euclidean. Literature further proposes incorporating Doppler velocity [LLK18], which has the advantage that points with similar velocity, likely reflections from the same object, are more easily grouped into a cluster. The differences $(\Delta x, \Delta y, \Delta v_{\text{rel}})$ between two points $p(x_p, y_p, v_{\text{rel},p})$ and $q(x_q, y_q, v_{\text{rel},q})$ are computed by subtracting the individual values. With these values, the distance metric ϵ is defined as

$$\epsilon(\Delta x, \Delta y, \Delta v_{\text{rel}}) = \sqrt{\Delta x^2 + \Delta y^2 + \Delta v_{\text{rel}}^2}. \quad (6.1)$$

This conventional distance metric still does not incorporate the criticality of the point, meaning that the conditions for points to form a cluster are identical regardless of their relevance. The idea is to establish a distance metric that depends on both the Euclidean distance and the assessed criticality, denoted as $d_{\text{crit}}(\Delta x, \Delta y, \Delta v_{\text{rel}}, \text{crit}_p, \text{crit}_q)$. In contrast to other adaptations of DBSCAN aimed at increasing the number of clusters, such as decreasing the min_{pts} parameter or increasing the ϵ -distance, introducing a criticality-aware distance metric results in a differentiated clustering behavior. The more critical two points p and q are, the smaller their effective distance d_{crit} becomes. In this way, two points that are spatially too far apart to form a cluster may be moved closer together. Consider the situation in which two points are the only reflections from a VRU.

If one of these points lies beyond the ϵ -distance required by the DBSCAN algorithm, neither can generate a cluster. Consequently, the VRU remains undetected, and the planner will not react to it, creating a genuine safety hazard. To mitigate this issue, the criticality-aware ϵ -metric reduces the effective distance between two points by up to 20 %, using a factor k_{crit} as a multiplier ranging between 0.8 and 1.0. This factor can be parameterized; however, a stronger reduction diminishes the informative value of the spatial relationship between points. Furthermore, the factor should saturate for sufficiently large criticality values crit_{max} . The saturation is controlled by the parameter $a = \frac{1 - k_{\text{crit}, \text{min}}}{\text{crit}_{\text{max}}}$. For example, if the criticality distance should be reduced by up to 20 %, which occurs if one of the points p and q has a criticality $\text{crit} \geq 0.05$, the parameter is set to $a = 4$. With this variable, the reduction factor is computed as

$$k_{\text{crit}} = 1 - \min(a \cdot \max(\text{crit}_p, \text{crit}_q), 0.2). \quad (6.2)$$

The criticality-aware distance is then obtained using the factor k_{crit} as

$$d_{\text{crit}} = k_{\text{crit}} \cdot \sqrt{(x_p - x_q)^2 + (y_p - y_q)^2}. \quad (6.3)$$

The resulting distance contraction, which depends on the point with the highest criticality value, is illustrated in Figure 6.2.




Figure 6.2: Distance criticality d_{crit} shown as dependent on the evaluated criticality of the points to be clustered.

This procedure can also be interpreted as a spatial distortion that contracts distances within the ego vehicle's driving tube, while leaving areas outside of the driving tube undistorted.

6.2.2 Minimalistic Tracking

Many critical radar reflection points do not progress to a sufficient stage in radar data processing to be considered in an object hypothesis, as they are either filtered out or cannot be matched to a cluster. Therefore, a second approach for criticality-aware signal processing is to introduce a bypass for critical points. This bypass generates an object hypothesis directly from single radar points if their criticality exceeds a defined threshold. All object hypotheses (tracked objects) created in this way are considered by the planner of the automated function. The bypass procedure ensures that critical points are directly observed and that the time delay between measurement and the AVP function's reaction is eliminated. This prioritized treatment of single critical points makes it possible to detect many pedestrians, who would otherwise remain undetected, as they are often represented by only a single radar point at their torso [VM+24], especially when moving laterally to the radar sensor [BFR12]. However, the drawback is that most points assessed as critical are in fact FP points. For this reason, advanced hypothesis handling must be applied. This hypothesis handling manages the set of hypotheses $\mathcal{H} = h_1, \dots, h_n$ by associating the most recent cluster set with the hypotheses in this set $\mathcal{H}$.

Definition 6.1 (Associated Object)

An object is associated if the distance between the predicted position of its hypothesis $(\hat{x}_h, \hat{y}_h)$ and a cluster generated from the current point cloud (x_c, y_c) does not exceed the association distance d_{assoc} .

If the ratio between criticality and existence certainty is sufficiently high, a hypothesis is labeled as *confirmed*.

Definition 6.2 (Confirmed Object)

An object hypothesis is usually observed over several observation cycles n_{confirm} before its presence is confirmed. The tracking algorithm applies a confirmation flag for each hypothesis. The planner of the automated function considers only confirmed objects when generating its policy.

The longer an object hypothesis is usually observed, the greater is the certainty in predicting whether the object is present or absent. On the other hand, longer observation times also result in delayed reaction times. Each hypothesis is characterized by two parameters: its age n_{hist} in cycles and the number of cycles since the object was last observed n_{observ} . A hypothesis is removed if it is not observed for $n_{\text{observ, max}}$ consecutive cycles. For newly initiated tracks, there is the additional requirement that the hypothesis

must be observed in at least $n_{\text{init}, \text{min}}$ out of n_{init} measurement cycles. Figure 6.3 illustrates the conditions under which a hypothesis is confirmed and under which a confirmed track is continued. In this figure, the parameters are chosen as follows: $n_{\text{init}, \text{min}} = 2$, $n_{\text{init}} = 3$, and $n_{\text{observ}, \text{max}} = 4$.




Figure 6.3: Mechanisms of track management in radar object tracking. Illustration of track refusal, track continuation, and object hypothesis confirmation, showing the number of cycles during which a radar cluster can be associated with a track.

Typically, a hypothesis must be observed for at least $n_{\text{confirm}} \geq 3$ cycles before it is confirmed and returned as a relevant object. This means the planner can only react after three cycles. At an update rate of 16 Hz and a maximum domain speed of 8.5 m/s [Int23], the automated vehicle travels a distance of 1.53 m in that time. This distance can be essential for collision avoidance. However, in most cases, the confirmation cycles increase the stability of objects and keep the FP rate low enough to employ the perception system in an automated function. For this reason, minimalistic tracking should only be applied in situations where a decision must be made quickly, namely, when the distance between the critical point and the vehicle is small and the object is either inside the driving tube or about to enter it. Non-critical points, in contrast, should continue to be subject to longer confirmation times.

6.3 Evaluation on Real World Data

Both approaches, adapted clustering and adapted tracking, are evaluated in a realistic setup. In the preceding chapters, the methodology was developed and validated on the datasets ROD2021 [Wan+21a] and RadarScenes [Sch+21b]. However, these datasets do not fully align with the ODD and the research focus on safety-aware radar data processing. This thesis contributes to research by evaluating the methodology

using a custom dataset. The creation and recording of this dataset are outlined in Subsection 6.3.1. Subsection 6.3.2 presents the evaluation of the two methods described in Section 6.2. The successive point cloud treatment method, discussed in Section 5.2, is evaluated again on this dataset. Finally, Subsection 6.3.4 lists the results of the investigations.

6.3.1 Custom Dataset

To evaluate the methods, a dataset was recorded that fulfills preconditions unmet by existing publicly available datasets. These preconditions are described in Subsection 2.5.5 in detail. The custom dataset meets the preconditions in the following way:

- *Radar data coverage:* The radar data of the recorded dataset should be processed as little as possible. The test vehicle provides radar data in both point cloud and tracked object form. Since tracked objects in the test vehicle are generated using a black-box algorithm from the point cloud, they offer no significant additional value for considerations of the point cloud processing. Therefore, only the radar point cloud is recorded.
- *Location coverage:* The dataset was collected entirely in a parking garage with numerous metal structures, which may lead to higher noise levels compared to investigations in free space. The parking garage was in a typical occupation state, with both free and occupied parking lots. Consequently, diffuse reflections could occur at fences, supports, rails, and other parked vehicles.
- *Object coverage:* In a parking garage, the presence of VRUs must be expected. The most endangered VRU types are children and toddlers. Accordingly, tests were conducted with three different types of pedestrians, which are introduced in detail below. No tests with bicyclists or scooter riders were performed, as these VRUs are less typical for the ODD of the AVP function.
- *Situation coverage:* Based on the methodical approach developed in Subsection 5.3.1, a scenario with a high injury risk to VRUs was selected. This scenario is described in detail below. It can be classified as a transfer-drive scenario, since the speeds of the ego vehicle in these cases are higher than during the actual parking maneuver, thereby increasing injury risk in case of a collision.

The scenario in the dataset aims to reproduce a worst-case situation with maximum criticality, in which the standard object detection pipeline may still be prone to failure. Thus, a modification of scene 33 of the catalogue provided in the ISO 23374 standard is applied for the investigations. In this scenario, the ego vehicle moves at a constant velocity v_{ego} along the garage driveway. A pedestrian emerges perpendicularly from behind a parked vehicle, occluded until the crossing movement, and proceeds across the driveway. If no braking action occurs, the collision point with the ego vehicle is located

at the center of the vehicle's front. Figure 6.4 illustrates a two-dimensional bird's-eye view of the scenario.




Figure 6.4: Overview of the scenario used for the real-world experiment. Bird's eye view of a parking garage scenario involving the ego vehicle and a VRU, representing a critical case for the investigation of adapted point cloud processing and object tracking.

Data from four 77 GHz corner radar sensors are captured: two mounted behind the front bumper and two mounted behind the side skirts beneath the B-pillar. The approximate sensor positions are shown in Figure 6.4. Together, they provide consistent coverage of the ego vehicle's front, with the exception of a blind spot directly in the front of the vehicle. This blind spot is not relevant for the selected scenario, since the critical observation for the radar perception system is whether a VRU enters the ego vehicle's driving tube. If this moment is successfully detected, a braking action can be taken and the collision is prevented. The four radar sensors transmit their data asynchronously. However, due to the short measurement period, the point clouds from all four sensors are merged into a single point cloud that is processed every 60 ms. The technical specifications of the radar sensors are summarized in Table 6.1.

To support scene interpretation during the analysis, video data from the front parking camera is also recorded. In addition, the ego vehicle's longitudinal velocity v_x and its yaw rate ω are captured. This allows the current motion state to be used either as a substitute for the planning policy or to generate suitable trajectories, on which the driving tube is based.

Table 6.1: Overview of the main specifications of the corner radar sensor applied in the experiment.

Band Width	$b_1 = 600 \text{ MHz}, b_2 = 750 \text{ MHz}$
Measurement Range	$r_{1,\max} = 92.6 \text{ m}, r_{2,\max} = 24.0 \text{ m}$
Range Accuracy	0.2 m
Azimuthal FoV	150°
Azimuth Angle Accuracy	0.15°
Azimuth Angle Separability	7°
Measurement Cycle	60 ms

All reflection points from the pedestrian target are manually labeled. For this purpose, reflections captured by the radar sensors are transformed into one common *odom* coordinate system, compensating for sensor movement during the sequence. The reflections of the target are then labeled by drawing a bounding box around the trace of radar reflection points. This procedure ensures that most of the points labeled as TPs correspond to actual reflections from the target. Nevertheless, a small portion of the labeled points may represent noise (FPs), which is technically difficult to distinguish from true reflections.

6.3.2 Experimental Setup

The parameters of the test case are varied to investigate their influence. The distance to the occluding parked car d_{park} determines whether parts of the VRU are occluded and at what point the VRU becomes fully visible to the radar sensors. Nevertheless, even during occlusion, the VRU may remain detectable due to multipath propagation effects. The VRU's velocity v_{VRU} affects the available observation time, and the same applies to the ego vehicle's velocity v_{ego} . Furthermore, the severity of a potential collision and thus the criticality of radar reflections from the VRU increases at higher ego vehicle speeds. The discrete parameter values are listed in Table 6.2. For comparison purposes, the parameter specifications of Scene 33, as defined in the ISO 23374 standard [Int23], are also listed in this table. The variation of the test case in this experiment closely aligns with this standard. Three measurement runs are recorded for each parameter combination.

Three different types of pedestrian types are used. For safety reasons and to ensure reproducibility, the VRUs are represented by pedestrian targets with articulated legs that are driven by a belt-based propulsion system. Their reflectance properties are comparable to those of humans.

1. *Adult pedestrian*: The first target, representing an adult pedestrian, is expected to be easiest for the radar perception system to detect. For this VRU type, higher

Table 6.2: Scenario parameters modified for the experiment.

Symbol	Parameter Description	ISO 23374 Specifications	Parameter Values
d_{park}	Distance to parking lots	[0.6 m; 2.0 m]	[1.0 m; 2.0 m]
v_{ego}	Velocity of vehicle-under-test	0.5 m/s to 8.3 m/s	[4.0 km/h; 8.0 km/h; 12.0 km/h]
v_{VRU}	Velocity of VRU	5 km/h	[3.0 km/h; 5.0 km/h; 8.0 km/h]

walking velocities of 5.0 km/h and 8.0 km/h are tested.

2. *Child pedestrian*: The second target type represents a seven-year-old child, with walking velocities of 3.0 km/h and 5.0 km/h.
3. *Toddler on toy car*: The final target is a two-year-old toddler on a toy car. This case is particularly challenging, since no Doppler effects from moving legs are present. Moreover, the probability of a toddler suffering serious injuries in a potential collision is higher than for the other two targets as it can be overrun by the vehicle more easily. The assumed velocity for this target is $v_{\text{VRU}} = 3.0 \text{ km/h}$.

To evaluate the effectiveness of both the adaptation in clustering and the adaptation in tracking, a complete perception pipeline is implemented. By default, this pipeline consists of a rigid filtering of the captured point cloud, a DBSCAN-based clustering step, and tracking using a Kalman filter. The basis for criticality-aware processing is provided by the drive-tube-based criticality evaluation of the measured point cloud. No automated parking function such as AVP is used in the experiment; instead, the vehicle is manually controlled by a human driver during the recordings. For this reason, the ego vehicle's current motion state is used to approximate the predicted future trajectory. The motion signals required for this procedure are read from the ego vehicle's communication bus.

Each stage of the standard processing is then extended. In the filtering stage, critical points are reintroduced into the processing chain. The clustering and tracking stages are modified according to the methods described in Section 6.2. To maximize their effect, all radar points evaluated as critical bypass the rigid filtering module and are used directly as inputs to the adapted clustering and tracking algorithms. Figure 6.5 illustrates the complete framework. The blue boxes on the left represent the stages of the standard perception pipeline, while the green boxes highlight the safety-aware extensions.

As a measure of reflectivity of an object, the radar sensors in the test vehicle provide two values: the signal power and the SNR. Contrary to the signal power, the SNR value adapts to environments at various noise levels. For this reason, the SNR value is used in the rigid filter. To determine the filtering threshold, the noise situation of the scenario is analyzed using a test measurement. The objective is to remove as many noise points as




Figure 6.5: Framework for radar data processing with criticality-aware extensions. Object detection framework used in the experiments, comprising the successive point cloud treatment, clustering modification, and minimalistic tracking approach.

possible while retaining most of the TP reflections from the VRU. Figure 6.6 shows the distribution of SNR values for both noise points and TPs, based on a random sample of 1,000 points within the ego vehicle’s driving tube.

From this analysis, the SNR threshold of the rigid filter is set to 32.0 dB m². Points with a lower SNR are removed from the point cloud before further processing, while points with a higher SNR are passed to the clustering algorithm. In addition, another rigid filtering criterion applies to Doppler velocity. All points with a relative Doppler velocity of $v_{\text{rel}} \geq 14.0$ m/s are discarded. This threshold is derived from the assumed maximum velocity of 10.0 m/s for other traffic participants in the parking ODD, combined with the ego vehicle’s maximum velocity of 12.0 km/h ≈ 4.0 m/s in this experiment. A spatial rigid filter is also applied in the *vehicle* coordinate system, removing points which are irrelevant for the driving context located beyond a longitudinal distance of $|x| \geq 20.0$ m or a lateral distance of $|y| \geq 10.0$ m. To enable the detection of objects with very few reflection points, the minimum cluster size parameter is set to $\min_{\text{pts}} = 2$. This ensures that if at least the front-right and middle-right radar sensors each detect a single point of the VRU, a cluster is formed. The neighborhood radius parameter is set to $\epsilon = 0.70$, which provides a suitable compromise for detecting cars as well as VRUs such as pedestrians. For vehicle-only detection tasks, a larger value is recommended to reduce FP clusters.

Since the target motion in this scenario is linear, without turns or accelerations, the CV motion model is chosen. An extended Kalman filter is unnecessary, as no nonlinearities are present in this scenario. The estimated state vector for each object hypothesis is




Figure 6.6: Comparison of the SNR distribution for TP reflections from VRUs and FP noise points, demonstrating the rationale for rigid filtering.

$\hat{\mathbf{x}} = [x, y, v_x, v_y]$, including its two-dimensional position (x, y) and velocity (v_x, v_y) relative to the ego vehicle. The tracking parameters are determined heuristically to optimize both the initial confirmation and temporal stability of object hypotheses:

- The maximum association distance between a hypothesis and a cluster is set to $d_{\text{assoc}} = 1.0$ m, using the Euclidean distance metric.
- The number of cycles to confirm the object is set to $n_{\text{init}, \text{min}} = 3$, which corresponds to a confirmation time of $t_{\text{init}} = 180$ ms. The new object must be associated in at least two out of this three cycles.
- An object hypothesis is removed if it is not observed for $n_{\text{observ}, \text{max}} = 10$ cycles, corresponding to an observation gap of $t_{\text{observ}, \text{max}} = 600$ ms.

6.3.3 Experimental Metrics

The metrics used to evaluate the methods build on those introduced in the preceding chapters, counting TPs, FPs, and FNs for each of the three processing stages: radar points, radar clusters, and confirmed radar object hypotheses, hereafter referred to as tracks. These basic metrics are complemented by two additional aspects: The time of first detection t_1 indicates the moment when the VRU is detected for the first time. The object consistency is the percentage of cycles, since the first detection, in which the object is measured, clustered, or maintained as a track. Finally, it is verified whether

the entering phase into the driving tube is successfully detected. For this condition to be satisfied, the set of hypotheses must contain a track for the object at the moment it enters the ego vehicle's driving tube. For all three safety-aware approaches, these metrics are evaluated and compared against the baseline performance of the standard approach, which does not apply any safety-aware methods. Figure 6.7 illustrates the evaluation schematically. It shows the y -position of measurement points over time t . The ego vehicle's driving tube is represented by the blue opaque bar, while the VRU's trajectory is indicated by the red opaque bar. The driving tube is laterally bounded by the vehicle's outer left position y_l and outer right position y_r . Green lines mark measurement cycles where a point can be associated with the VRU, whereas red points represent measurements that cannot be associated with any object in the scenario and are therefore classified as noise. The dashed line indicates measurement cycles in which no points correspond to the VRU. Consequently, these cycles are counted as FNs.




Figure 6.7: Perception flaws on the point level in the experiment. Measurement points evaluated as TPs, FPs, and missing points (FNs) on the VRU.

Table 6.3 summarizes the methods compared in this experiment. For clarity, the methods are itemized below. While method A applies at the radar point level, method B and method C apply at the cluster and track levels. Both use the filtered point cloud complemented by all points evaluated as critical, thereby influencing the generation of clusters and tracks.

The effectiveness of the safety-aware methods depends on the correct identification of critical points. In the radar-point evaluation, each point is annotated with the information indicating whether it corresponds to a reflection from a VRU target. However, this annotation alone is insufficient to assess the performance of the criticality metric. Therefore, the recall is used as a simplified measure. Points on the VRU with a criticality value of $crit > 0.0$ are counted as TPs, while points which are not assessed critically ($crit = 0$) are counted as FNs. To compute the precision metric, all FP points must be counted. This requires deciding, for each radar point in the dataset, whether it represents an existing object. Radar points are generated at parked vehicles, beams,

Table 6.3: List of the perception-related safety-aware modifications evaluated in this experiment.

Item	Description
A	Creation of spatial regions in subsequent point clouds and suspending the filtering in these regions, as outlined in Section 5.2.
B	Modifying the distance metric for the DBSCAN clustering, as outlined in Subsection 6.2.1.
C	Creating object hypotheses directly from single radar points if they are critical enough, as outlined in Subsection 6.2.2

pillars, and fences within the garage. This labeling process requires an enhanced setup, potentially incorporating additional sensor data and substantial manual effort. However, within the AVP function, FP radar points are relevant only when they generate an object hypothesis during signal processing. For these reasons, the precision metric at the radar-point level was not evaluated in this experiment.

6.3.4 Experimental Results

In total, 97 approach sequences toward the VRU were recorded for the dataset. The majority of 49 sequences were recorded with the child target, while both the adult and the toddler were used in 24 sequences each. These 24 sequences result from combinations of two VRU velocities, three ego vehicle velocities, and two distances to the occluding car. Each combination was recorded twice to allow comparison in case one measurement deviated significantly from the other.

First, the criticality evaluation is examined, meaning that the criticality metric is assessed with respect to whether labeled VRU reflections are classified as critical. This serves as a prerequisite for handling these points correctly during signal processing. The mean recall of critical points on the VRU across all sequences was $\mu_{\text{recall}} = 0.332$, with a standard deviation of $\sigma_{\text{recall}} = 0.204$. This metric is most meaningful for comparison with the same experimental setup, since all VRU reflection points are labeled and hence counted as FNs, although the VRU is not yet critical for the function. Considering that the pedestrian's movement lasted approximately 4.5 s, and the last 1.5 s before entering the driving tube should be evaluated as critical, this result suggests that the majority of points are correctly identified as critical.

The precision of the criticality evaluation, however, is comparatively low. The metal structures in the parking garage produce numerous FP points within the driving tube. Across the experiments, the mean precision was $\mu_{\text{precision}} = 0.087$, with a standard deviation of $\sigma_{\text{precision}} = 0.077$. It should be noted that the recordings contain critical

situations in a highly condensed form. In practical applications, the precision could therefore be significantly lower.

Table 6.4 lists the results of method A, which recalls points into the processing pipeline if they occur in regions where critical points were previously measured. The number $\sum \mathcal{P}_{\text{filt}}$ of points removed by the filtering process is only 1.2 % lower in method A compared to the baseline. This highlights that safety-aware filtering works similar to rigid filtering in most regions, but recalls truly critical points into the detection process. As shown, the number of reflection points on the VRU $\sum \mathcal{P}_{TP}$ nearly doubled. Consequently, the number of reflection points erroneously removed from the point cloud was reduced by almost half. The proportion of points filtered by the cascade of rigid filters decreased by 1.2 %, corresponding to approximately 170,000 radar points being recovered by the successive point cloud treatment across all scenarios. Despite this recall, the overall filter rate remained high at approximately 78.3 %, meaning that the majority of points were still removed from the processed point cloud, thereby ensuring limited computational demand for clustering and tracking.

Table 6.4: Detection performance on the point cloud level when successive point cloud treatment is applied.

	$\sum \mathcal{P}_{\text{filt}}$	$\sum \mathcal{P}_{TP}$	$\sum \mathcal{P}_{FN}$
Baseline	10,607,610	9,217	23,333
Method A	10,438,393	20,195	12,355
A to Baseline	−1.2 %	120.3 %	−47.1 %

While the consecutive clustering and tracking level benefits from the improved set of radar points, the behavior of the function depends solely on the object hypotheses (tracked objects) delivered to the automated parking function by the radar perception system. Hence, the next step is to investigate whether these improvements at the point cloud level propagated to the cluster and track levels. Table 6.5 summarizes the TP and FN counts obtained with method A. The increased point cloud density led to a 2.2 % rise in the total number of clusters. The number of clusters associated with the VRU increased by nearly 50 %, while the number of cycles without a cluster from the VRU’s detection points decreased by one third. At the track level, results show that tracking compensates for missing clusters in some measurement cycles, yielding an 18.3 % increase in tracks. However, this improvement is less pronounced than the increase in clusters. Importantly, the number of FN tracks, the number of cycles in which no track was created or maintained for the VRU, decreased by 36.0 %. A drawback of the successive point cloud treatment is that FP tracks accumulate, since more clusters are formed and converted into object hypotheses. As a result, the number of tracks not directly related to the VRU target increased by a factor of approximately 3.5.

Recapturing critical points led to earlier cluster generation in 30.3 % of the sequences, enabling hypotheses to be established earlier. When this occurred, clusters were gen-

Table 6.5: Detection performance on the cluster and track levels when applying method A.

	$\sum \mathcal{C}$	$\sum \mathcal{C}_{TP}$	$\sum \mathcal{C}_{FN}$	$\sum \mathcal{T}$	$\sum \mathcal{T}_{TP}$	$\sum \mathcal{T}_{FP}$	$\sum \mathcal{T}_{FN}$
Baseline	550,457	1,731	2,339	973,467	2,757	793	1,325
Method A	574,356	2,573	1,524	1,023,350	3,255	2,884	842
Difference	2.2 %	49.2 %	−33.3 %	2.1 %	18.3 %	263.7 %	−36.0 %

erated on average 403 ms earlier compared to the baseline method. The success rate, defined as the presence of a VRU hypothesis at the moment the VRU entered the driving tube, increased from 82.6 % to 95.7 %. Consequently, an entering phase without a valid object hypothesis was observed in only one scenario. The reasons for this failure are analyzed in the following Section 6.4.

A qualitative illustration of how recapturing points improves the point density on the VRU target is provided in Figure 6.8. The color of the points represents the assessment provided by the drive-tube-based criticality metric, while red circles indicate the presence of confirmed object hypotheses. These confirmed object hypotheses are essential to guarantee safe behavior of the automated parking function in critical situations. The figure in the upper row shows the camera image at the moment when the VRU target enters the driving tube. The left and right panels in the lower row compare the detection and tracking results of the baseline radar object detection pipeline with the adapted algorithm that recalls points from critical regions. In the baseline, the first track emerges only after the VRU is already in the driving tube. By contrast, with the safety-aware approach this moment is shifted earlier in time, allowing the entering phase to be observed. This capability is essential for ensuring functional safety.

The subsequent analysis evaluates methods B and C. Both approaches operate on the same point cloud, which includes all points that are not removed by the rigid filter and all critical points: $\sum \mathcal{P} = (\mathcal{P} \notin \mathcal{P}_{\text{filt}}) \cap \mathcal{P}_{\text{crit}}$. As a result, the baseline for methods B and C already contains more clusters than the baseline used for method A, as this baseline operates on a point cloud which does not contain critical points. Table 6.6 reports the improvement in the number of clusters.

Table 6.6: Object detection performance of method B on the cluster level.

	$\sum \mathcal{C}$	$\sum \mathcal{C}_{TP}$	$\sum \mathcal{C}_{FN}$
Baseline	592,881	1,762	2,392
Method B	609,781	2,164	2,013
Difference	2.9 %	22.8 %	−15.8 %

Again, the number of TP clusters associated with the VRU target increased dispropor-




Figure 6.8: Example of denser point cloud after applying successive point cloud treatment. Qualitative visualization showing higher point density and earlier track initiation with safety-aware mechanisms (right) compared to the baseline without them (left).

tionately compared to the total number of clusters. Another positive effect is that the number of cycles without a VRU cluster decreased by 15.8 %. Table 6.7 presents the results of tracking based on the extended set of clusters (method B) and compares them with method C, which directly generates object hypotheses from critical points.

Methods B and C produce broadly similar outcomes, although method C yields slightly more track hypotheses. However, in both methods the number of FPs increases substantially, which limits practical applicability due to the proliferation of tracks. Furthermore, FNs could not be eliminated entirely. The success rates achieved by methods B and C were 89.6 % and 88.0 %, respectively. Both methods generated clusters earlier in 15.2 % of all sequences. On average, the TP clusters associated with the VRU target were formed 287 ms earlier than with the baseline, and the resulting object hypothesis appeared 253 ms earlier. While these results show that method B and method C both have the potential to make late signal processing more safe, all of these values remained below the success rate and temporal advantage achieved with method A.

Table 6.7: Object detection performance of methods B and C on the tracking level.

	$\sum \mathcal{T}$	$\sum \mathcal{T}_{TP}$	$\sum \mathcal{T}_{FP}$	$\sum \mathcal{T}_{FN}$
Baseline	1,045,662	2,795	888	1,370
Method B	1,084,259	3,086	7,430	1,098
B to Baseline	3.7 %	10.4 %	736.7 %	−19.9 %
Method C	1,102,428	3,087	7,749	1,100
C to Baseline	5.4 %	10.4 %	772.6 %	−19.7 %

6.4 Discussion

This chapter evaluated safety-aware algorithms using data from actually critical situations within the target ODD of an automated parking function. The proposed modifications to clustering and tracking complete the response to (RQ2) for the later stages of radar data processing. Although information may already be lost at the point level, objects based on critical points are more likely to be generated by the safety-aware methods. This reduces the number of missed objects and hazardous situations, thereby enhancing the safety of the automated function.

The dataset, which was recorded customized for the investigation of perception in safety-critical scenarios, provides an opportunity to investigate both the identification of relevant radar points and the evaluation of improved methods for radar point cloud processing. The evaluation was conducted on one representative scenario in which the VRU crosses perpendicularly in front of the ego vehicle during a transfer drive to a parking lot. More complicated scenarios can lead to incorrect conclusions, for example, when reflection points on the VRU are confused with noise or static points. The mechanisms for criticality evaluation, clustering, and tracking operate similarly in other scenarios. Variations of the scenario, such as a turning ego vehicle or VRU, require the use of different motion models during the tracking stage. These more complex trackers may discard a hypothesis even though points remain available for processing, thereby influencing the overall evaluation results. In this experiment, the objective was to evaluate the integration of criticality-aware approaches while minimizing the influence of baseline tracking performance. For this reason, both the tracker and the evaluated scenario were kept simple.

The drive-tube-based criticality assessment method computes a continuous value for each point's relevance to the planned driving task. The parameterization of the algorithms in this chapter, which depend on this assessment, was chosen as conservatively as possible. This means that specific processing was applied even to radar points with criticality values close to zero. In this way, the maximum potential safety effect of the adapted processing was exploited. However, the continuity of the criticality metric also

suggests modifying the processing pipeline at various steps. For example, an object hypothesis could be created more quickly if the initiating points are assessed as highly critical, and the hypothesis could be maintained longer. Conversely, hypotheses of potential objects could be removed more quickly if the points are assessed as barely critical. These techniques would significantly reduce the number of FP tracks that trigger braking actions of the automated function. Nonetheless, the precise criticality-dependent adjustments in object processing must be adapted in consideration of factors such as the ODD and the acceptable FP rate.

Furthermore, continuous adaptation of processing based on assessed criticality is particularly meaningful if the accuracy of the criticality metric can be quantified. However, providing such prove is challenging. The criticality assessment operates on point attributes such as position and Doppler velocity. To compare the criticality value assigned by the metric with the true criticality value, the latter must be determined independently, without relying on the same factors used by the metric itself. One approach is to have each point evaluated by a functional expert, with evaluation and label rules defined in detail. Another approach, used in this experiment, was to assign increasing criticality values to radar points labeled as reflections of the VRU, based on decreasing TTC. The limitation of this procedure is that the comparison is already closely related to the methodology of the criticality evaluation itself. Additionally, in this experiment, most of the FN points — those evaluated as uncritical despite being reflections of the VRU target in a critical situation — were classified as such due to imprecise angle or Doppler velocity measurements. Consequently, evaluating the criticality assessment would reflect sensor performance as much as algorithmic performance. Under ideal sensor measurements, however, the criticality assessment provides a traceable score of each point's criticality.

The results of the experiment showed that recapturing quality decreases at later processing stages. The best results were achieved by recapturing information at the point cloud level by exploiting promising regions in consecutive point clouds. Similar results from modifying the clustering and tracking algorithms indicate that the maximum information contained in a single point cloud was already utilized effectively. At the same time, this observation suggests that point cloud density limited further performance improvement. Qualitative analyses of individual sequences support this assumption. Figure 6.9 illustrates all available points of a VRU target in one sequence as it crosses in front of the ego vehicle.

In this figure, the measurement time of each point is plotted on the x-axis, while the y-axis shows the position in the *vehicle* coordinate system. Ideally, the movement of the pedestrian should form a line-like trace of points. At least two secondary traces can be identified alongside the main trace, which follows the blue line representing the VRU's true movement. These secondary reflections may be caused by the pedestrian's legs. However, beginning at a y-position of $y > -2.0$ m, i.e. at the moment when the VRU approaches the ego vehicle's driving tube, reflections of the VRU target are missing. This absence can be attributed to limited radar sensor performance in the side




Figure 6.9: Exemplary sequence showing a VRU crossing in front of the ego vehicle. The sensor’s detection performance while the VRU approaches the driving tube is visibly limited.

lobe. This supports the earlier assumption that safe radar perception requires not only advanced signal processing but also sensor performance sufficient to detect VRUs even in high-noise environments.

If the point cloud contains no information about the VRU’s presence, further improvements can only be achieved by incorporating historical sensor measurements or even sensor data from other perception domains. The results of this chapter, as well as Chapter 4, indicate that methods based on radar images would have outperformed the radar-point-based approaches. Since image processing requires more time due to the larger amount of information, a practical compromise must be found for real-world applications.

In this chapter, it is argued that radar point clouds can be filtered according to the SNR value, as this attribute is highly significant for determining the existence of a given point. Points with very low SNR are accordingly removed, while points with high SNR remain in the point set for further processing. However, when the SNR of confirmed FP points is examined, numerous points with high SNR values are observed. All of these points, located in the ego vehicle’s driving tube at close range to the ego vehicle, are evaluated as critical or very critical. The prioritization of these points, combined with an unfavorable chaining of FP point positions across consecutive point clouds, exacerbates the problem. During this process, noise points occasionally appear close to each other and are associated with the same object hypothesis. The Kalman filter may then interpret the noise points as a moving object. Consequently, this effect

leads to a large number of FP tracks. This issue is identified as the primary cause of the significant increase in the FP rate and the resulting weakness of the precision metric. A further difficulty in handling critically evaluated FP points is that their origin cannot be fully identified. If the points arise from multipath propagation, they may be anticipated based on the surrounding scenario. By contrast, FP points resulting from statistical effects inside the radar sensor can hardly, if at all, be distinguished from TP reflections. One way to address this problem is through advanced tracking methods based on deep learning models, which can achieve high precision in distinguishing between TPs and FPs by learning typical object-related point representations within point clouds. By contrast, rule-based algorithms such as the Kalman filter, as applied in this chapter, struggle to recognize such point constellation patterns. However, deep learning models make probabilistic and largely opaque decisions. Their performance on underrepresented point patterns in the training dataset may be limited. The uncertainty about their ability to handle unknown situations prevents the determination of reliable safety guarantees.

In the experiment, each radar point in the dataset belonging to the VRU target was labeled as a TP. The evaluation focused on continuous observation of this target. However, the methodology labeled every point not belonging to the target as a FP point. This approach cannot discriminate between spurious reflections and structural elements of the parking garage, such as pillars or parked vehicles. As a result, these radar points were processed into clusters and object hypotheses. While this is generally correct behavior for the Kalman filter, the resulting hypotheses were counted as FP tracks in the statistics, further inflating the FP rate. Many of these tracks can be eliminated by a velocity filter, which removes tracks with longitudinal velocities similar to that of the ego vehicle. However, such a strategy would also remove VRUs that are stationary within the driving tube, making them unavailable to the planner of the automated function. This approach was therefore classified as hazardous and rejected. Drawing conclusions about the number of FP tracks requires labeling each point across all recorded point clouds. For qualitative analysis, however, it is sufficient to compare the number of FP tracks across different in order to understand the general trend.

A high number of FP object hypotheses correlates with erroneous braking and evasive maneuvers, which implies low stability of the downstream automated parking function. This effect has already been shown in Chapter 4, where Figure 4.1 illustrated the trade-off between high recall and high precision. To quantify this stability, the functional characteristics of the automated parking function must be defined. The function may employ different stages of escalation depending on the urgency of taking action, the situational context, and the confidence in object existence. The experiment demonstrated that maintaining an observation hypothesis for more than one cycle is necessary to accumulate sufficient information about a potential object. The point at which the automated function takes an erroneous action is also the point at which the function user becomes aware of a malfunction. Consequently, assessing robustness requires considering both perception and planning. In addition to the acceptance criteria for

functional safety that can be derived from the corresponding metrics, as described in Subsection 2.4.1, further criteria must be defined regarding the acceptable robustness of the complete function. This includes the rate for faulty activations under normal conditions. For SAE Level 2 and 3 automated functions, excessive FP actions may cause user dissatisfaction, since the passenger still witnesses such behavior. For SAE Level 4 automated functions, the passenger may have already left the vehicle. Nevertheless, the malfunction rate must be minimized to avoid deadlocks requiring human intervention to resume operation of the vehicle. A second consideration is that the vehicle interacts with other traffic participants, such as manually driven vehicles, and must maintain smooth general traffic flow.

7 Application in Radar Odometry Estimation¹³

So far, this work has discussed a safety-aware radar data processing pipeline aimed at reducing object detection errors in an automated parking function. However, the functionality of radar sensors is not limited to object detection alone. In addition to this task, radar sensors can perform other functions relevant to automated parking functions, such as localization and motion estimation. Still, the risk of safety-critical information being removed during radar data processing remains. The objective of this chapter is to demonstrate that the concept of critical radar data is not restricted to object detection tasks, but is equally applicable to other functionalities performed by the radar sensor. As an example of extending the concept to additional functional domains, this chapter employs radar-based motion estimation. The importance of motion estimation for the overall automated parking function motivates a safety-aware adaptation of the baseline motion estimation algorithm, for which the concepts developed in the preceding chapters provide a suitable foundation.

Section 7.1 elaborates on the relevance of motion estimation and the potential advantages of employing radar sensors for this task. It also describes the fundamental mathematical principles underlying radar motion estimation. Since the notion of unsafe radar data is less intuitive in the context of motion estimation than in object detection, this section also establishes the definition of unsafe radar motion estimates. In Section 7.2, the concept of safety-aware signal processing is transferred to this new task. This includes evaluating the criticality of the current driving state within traffic and applying methods to refine the estimation process in order to mitigate critical situations. The introduced approaches are evaluated using the RadarScenes dataset, which was already introduced in Chapter 5. The experiment and its results are presented in Section 7.3. Finally, Section 7.4 concludes the chapter by addressing the final research question and discussing potential limitations of the proposed approach.

7.1 Motion Estimation for Automated Parking Functions

This section introduces radar-based ego vehicle motion estimation, beginning with the relevance of this crucial subfunction in Subsection 7.1.1. Further details about the methodology are provided in Subsection 7.1.2, including the derivation of interrelationships between the measured Doppler velocity and the ego vehicle motion. The section

¹³ Chapter builds on [Brü+24b; Brü+25a].

concludes with Subsection 7.1.3, which highlights the necessity of approaching radar motion estimation from a safety-aware perspective.

7.1.1 Relevance of Motion Estimation

Motion estimation is essential for the operation of an automated function, as the velocity components (v_x, v_y) as well as the yaw rate ω are required for precise control. Furthermore, it complements the ego vehicle's localization function by operating in combination with algorithms that periodically localize the ego vehicle on a map. This map can either be created by a Simultaneous Localization And Mapping (SLAM) algorithm utilizing radar, lidar, or camera data, or through the use of proprietary equipment, such as visual markers [Kal+21] or beacon solutions [HOM12]. As the localization step is computationally expensive, it is not performed at a high frequency. Additionally, localization may return solutions that are slightly fluctuating in their guessed position over time, which can result in jerky behavior of the ego vehicle during control. For this reason, the ego vehicle's position is determined through the fusion of localization estimates and odometry observations in factor graphs [DKP17], propagating the positional shift from the last localization step. This process smooths the resulting localization and reduces the susceptibility for erroneous positional estimates.

Conventionally, motion estimation is performed using data from wheel-speed sensors. These sensors, which count the teeth of a rotating gear rim at each wheel using a Hall sensor, are reliable options for determining the ego vehicle's motion [Win+15]. However, imprecise measurements can result from wheel spin and variations in wheel diameter, caused by uneven tire pressure or deviations from the assumed tire size. Experiments show that the average velocity error during a drive with varying speeds is 0.040 m/s for high-grip asphalt, increasing to 0.077 m/s for low-grip scenarios [TSC06]. The error may be higher on icy surfaces or during braking, where a wheel slip of up to 21 % is desired to achieve an optimal braking distance [CAS13]. As motion estimation using radar sensors is not affected by these factors, it is reasonable to apply radar-based motion estimation in combination with, or even as a replacement for, wheel-speed-sensor-based motion estimation. Research in the area of radar ego-motion estimation indicates that a positional accuracy of 0.75 % after 140 m can be achieved [Kel+14].

Radar motion estimation can be categorized into two approaches that differ significantly in their operational principles. Both require radar perception data in point cloud format as input. The first approach is based on point set registration of two point clouds, recorded consecutively at times $t[k]$ and $t[k + 1]$. These point clouds exhibit high similarity but are offset by a positional shift. Point set registration, also called scan matching, seeks a transformation of the point cloud $\mathcal{P}[k + 1]$ that minimizes a metric describing its congruence with the previous point cloud $\mathcal{P}[k]$. One method for point set registration is the Normal Distributions Transform (NDT), which was first proposed for lidar scan matching [BS03] and later adapted to radar point clouds. This technique

models the two-dimensional radar FoV as a grid, with the occupancy of each cell in this grid represented as a normal distribution. The resulting probability densities of the two point clouds can then be aligned using Newton's method by optimizing a loss function that represents the differences between the first and the second point clouds. Another method is Iterative Closest Point (ICP) [BM92], which iteratively estimates a suitable transformation by repeatedly finding the closest corresponding point in the target set for each point in the source set, computing the transformation, and applying it. These steps are repeated until a mean-square error cost function for the congruence falls below a predefined threshold.

A drawback of radar motion estimation via point set registration is that the ego vehicle's motion must be inferred from the derivative of the positional shift. Moreover, outliers and noise radar points can lead to significant errors if no corresponding point for the FP point is found in the target set, which is a common occurrence in radar point clouds.

In contrast to point set registration methods, motion estimation based on Doppler measurements of radar points provides an alternative [Kel+13]. This approach leverages the fact that the Doppler velocity for stationary targets corresponds directly to the ego vehicle's motion. The algorithmic details of this method are discussed in the following subsection.

7.1.2 Doppler-Based Motion Estimation

The Doppler-based motion estimation applies the mathematical relationship between a stationary target and a sensor that moves relative to this target with the velocity $(v_{x,S}, v_{y,S})$, described by

$$-v_{\text{rel}} = \cos \theta \cdot v_{x,S} + \sin \theta \cdot v_{y,S}, \quad (7.1)$$

with the measured Doppler velocity of this point v_{rel} and its azimuth angle θ in the *sensor* coordinate system. Extended to all N points in the cloud, the problem can be formulated by

$$\underbrace{\begin{bmatrix} -v_{\text{rel},1} \\ \vdots \\ -v_{\text{rel},N} \end{bmatrix}}_{\mathbf{v}_{\text{rel}}} = \underbrace{\begin{bmatrix} \cos \theta_1 & \sin \theta_1 \\ \vdots & \vdots \\ \cos \theta_N & \sin \theta_N \end{bmatrix}}_M \underbrace{\begin{bmatrix} v_{x,S} \\ v_{y,S} \end{bmatrix}}_{\mathbf{v}_S}. \quad (7.2)$$

This equation is solved by a least-squares estimation method. If this concept is expanded and data from multiple radar sensors with varying positions are available, the yaw

rate ω can be estimated. The lever length $(x_{C,S}, y_{C,S})$ between the center point C of the vehicle and the sensor position S is used for estimation. The measurements of two sensors $S1$ and $S2$ with different lever lengths $\|x_{C,S1} + y_{C,S1}\| \neq \|x_{C,S2} + y_{C,S2}\|$ are combined to avoid linear dependency. Then, the relationship between the sensor's velocity $(v_{x,S}, v_{y,S})$ and the ego vehicle's motion state vector $\mathbf{m}$ can be determined with

$$\underbrace{\begin{bmatrix} v_{x,S1} \\ v_{y,S1} \\ v_{x,S2} \\ v_{y,S2} \end{bmatrix}}_{[\mathbf{v}_{S1}, \mathbf{v}_{S2}]^\top} = \underbrace{\begin{bmatrix} -y_{c,S1} & 1 & 0 \\ x_{c,S1} & 0 & 1 \\ -y_{c,S2} & 1 & 0 \\ x_{c,S2} & 0 & 1 \end{bmatrix}}_{[\mathbf{S}_1, \mathbf{S}_2]^\top} \underbrace{\begin{bmatrix} \omega \\ v_x \\ v_y \end{bmatrix}}_{\mathbf{m}}. \quad (7.3)$$

In a final step, the equation to estimate the motion vector $\mathbf{m}$ from the measurement matrix $\mathbf{M}$ containing cosine and sine of the measured azimuth angle and the Doppler measurement $\mathbf{v}_{\text{rel}}$ can be formulated by left-multiplying Equation 7.3 by $\mathbf{M}$, yielding

$$\underbrace{\begin{bmatrix} \mathbf{v}_{\text{rel},S1} \\ \mathbf{v}_{\text{rel},S2} \end{bmatrix}}_{\mathbf{y}} = \underbrace{\begin{bmatrix} \mathbf{M}_1 \mathbf{S}_1 \\ \mathbf{M}_2 \mathbf{S}_2 \end{bmatrix}}_{\mathbf{X}} \underbrace{\begin{bmatrix} \omega \\ v_x \\ v_y \end{bmatrix}}_{\hat{\mathbf{m}}}. \quad (7.4)$$

As this equation system is usually overdetermined and not solvable, the Moore–Penrose inverse can be applied to estimate the least-squares solution for $\hat{\mathbf{m}}$:

$$\hat{\mathbf{m}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}. \quad (7.5)$$

Figure 7.1 illustrates the method to determine all variables of the motion vector $\mathbf{m} = [\omega, v_x, v_y]^\top$.

These steps are discussed in detail in established research [Kel+13; Kel+14]. Radar motion estimation — as described in this work — uses this concept, but introduces two improvements of the described method to achieve a more robust estimate:

- *Outlier Detection*: The standard method [Kel+14] recommends detecting outliers that are not taken into the estimation process. The Random Sample Consensus (RANSAC) algorithm [FB81] is recommended for the outlier detection. It iteratively finds the consensus supported by the points drawn from a sample. However, this RANSAC algorithm weights outliers independently of their actual error. In contrast, the M-Estimator Sample Consensus (MSAC) [TZ00] applied in this work




Figure 7.1: Ego vehicle with center point C and two radar sensors with different lever arms, suitable for determining the motion vector $\mathbf{m}$.

minimizes a cost function which punishes the point's error ϵ to the consensus instead of assuming this error to be zero if the point lies within a predefined threshold.

- *Estimation Process:* The least squares estimation incorporates only errors in the Doppler velocity measurement vector $\mathbf{y}$, ignoring those in the azimuth angle measurement, contained in the matrix $\mathbf{X}$. Research [Kel+14] recommends to apply the Orthogonal Distance Regression (ODR) algorithm that computes the optimal solution when both errors, $\epsilon_{v_{rel}}$ and ϵ_{θ} , are considered. However, this algorithm underperforms compared to the ordinary least-squares estimation method in environments with a limited number of radar points. As a consequence, this work implements automatically switching between both methods depending on the number of points available for the estimation step.

A multiplexer selects two of the input sensors, each with potentially fresh time stamps and distinct lever arms. The MSAC algorithm is executed for 17 iterations, as proposed in the literature [FB81]. The time validity module verifies whether the two measurement time stamps differ by less than 30 ms and groups measurements from a single capture that were transmitted in separate data packets. ODR is applied in addition to the least-squares estimator if the number of raw points is sufficient. An overview of the complete radar-based ego vehicle motion estimation is presented in Figure 7.2.

7.1.3 Erroneous Motion Estimation

Radar-based motion estimation assumes that most detected points are stationary and that FP outliers are sparse relative to TP detections. The least-squares estimator determines the motion vector from multiple radar detections in the point cloud and is




Figure 7.2: Framework for radar-based motion estimation. Evaluation of measurements from at least two radar sensors, including outlier detection and estimation of the motion vector $\hat{\mathbf{m}} = [\omega, v_x, v_y]^T$.

therefore less susceptible to noise. However, the previous chapters have shown that radar sensors may produce erroneous measurements, particularly in noisy environments such as tunnels or parking garages. In another case, if most of the measured radar points are not static but instead are reflections from moving objects such as vehicles, the radar-based motion estimation algorithm tends to estimate the motion of the ego vehicle relative to these objects rather than determining the absolute velocity of the ego vehicle. The number of erroneous measurements may increase to a level that complicates or even prevents accurate ego vehicle's motion estimation. Consequently, an incorrectly estimated motion vector $\hat{\mathbf{m}}$ may lead to unsafe actions by the automated vehicle, potentially endangering other traffic participants.

One objective of this chapter is to investigate the impact of erroneous motion estimation on functional safety. In addition, this chapter also aims to propose methods that prevent motion estimates from leading to critical situations, as such estimates can induce hazardous behavior in the controller that relies on them. These estimates are hereafter referred to as critical estimates.

7.2 Safe Radar Motion Estimation

This section addresses the prevention of critical control actions induced by erroneous radar-based motion estimations, as outlined in Subsection 7.1.3. The core concept is to examine the impact of each estimate on the overall process within the automated parking function. Erroneous motion estimates can lead to incorrect position states if used in short-term ego vehicle position estimation. This chapter focuses on incorrect controller behavior caused by erroneous measurements in the feedback loop. To investigate this effect, a simple control process is established, as introduced in Subsection 7.2.1. Analogous to the preceding chapters, the criticality of the controller output, potentially leading to high-risk situations, must be evaluated. This evaluation strategy is detailed in Subsection 7.2.2. As radar-based motion estimation incorporates filtering, the motion estimation framework is examined for options to reconsider and correct estimates by re-evaluating previously flagged critical measurements. These corrective options are presented in Subsection 7.2.3.

7.2.1 Control Process

In automated functions, trajectory control is typically achieved using Model Predictive Control (MPC). Previous work has presented solutions employing MPC for path-planning control on highways [Ras+17] and in parking scenarios [GGW17]. The MPC framework includes a model block that predicts the system output based on past and current inputs. This enables the prediction of future control errors, which serve as input to the optimizer block. The optimizer block computes future inputs that optimize system behavior with respect to the cost function [CA13]. Since this chapter introduces a core concept of critical motion estimation and its prevention, a simplified control strategy is employed, which aims to regulate the longitudinal velocity v_x using a Proportional–Integral (PI) controller. The concept can be extended to more complex control structures.

A PI controller comprises proportional and integral components. The integral component augments the proportional term to eliminate steady-state errors. For the continuous case, the correcting variable can be determined with the current control error $\epsilon(t)$ by

$$u(t) = K_p \cdot \epsilon(t) + K_i \cdot \int_{\tau} e(\tau) d\tau. \quad (7.6)$$

The controller parameters are the proportional gain K_p and the integral gain K_i . The control output corresponds to the power applied to the drivetrain, resulting in an acceleration according to the vehicle model causing a change in velocity v_x . The discrete control loop is illustrated in Figure 7.3.




Figure 7.3: Closed-loop control of the longitudinal velocity v_x of the ego vehicle using a PI controller.

The controller parameterization depends on the dynamics of the controlled system, as represented by the vehicle model in Figure 7.3. The vehicle model used to evaluate the methodology is described in Section 7.3. As the controller applies a power P to the vehicle, the output is the acceleration at the subsequent time step $a_x[k + 1]$, and the resulting velocity $v_x[k + 1]$ is computed by integrating the acceleration a_x .

7.2.2 Controller Action Evaluation

The controller action evaluation assesses each controller output with respect to the criticality induced by that action. The evaluation aims to prohibit all unsafe control actions. In the simplified scenario, an unsafe control action is defined as uncontrolled acceleration of the ego vehicle. Such control actions may result from various contributing factors:

- *Desired value $r[k]$* : An erroneous planner policy may request a high velocity as the desired value for the next state S_{k+1} of the trajectory, resulting in a large control error $\epsilon[k]$ and, consequently, a high output $y[k]$.
- *Controller*: Internal controller faults may produce large output values $y[k]$. For instance, integrator wind-up can cause a disproportionately large integral component.
- *Measurement of Output Value $v_x + \epsilon$* : The control error $\epsilon[k]$ is computed as $\epsilon[k] = r[k] - (v_x + \epsilon_{\text{obs}})$, where ϵ denotes the observation error. A significant observation error ϵ_{obs} caused by an erroneous measurement leads to a correspondingly high control error. Assuming that the desired value $r[k]$ and the current velocity v_x are close, the approximation $\epsilon[k] \approx \epsilon_{\text{obs}}$ holds.

This work primarily focuses on the third factor, although controller action evaluation is capable of identifying errors from all the sources listed above.

The HARA conducted for many automated functions classifies uncontrolled acceleration as a safety risk, which corresponds to an ASIL D classification. This level implies that uncontrolled acceleration must be prevented with an accepted FIT rate of 10 FIT [Int18]. This suggests that the resulting acceleration $a_x[k]$, induced by the controller action $u[k]$, should be monitored. Conversely, the controller should not be overly restricted, as excessive limitation could prevent the compensation of large control errors.

The trajectory planner should optimize its policy with respect to a cost function that minimizes vehicle acceleration and deceleration, prohibiting too high acceleration values. The reason is that high accelerations are uncomfortable for the passengers [Int97], though this is not relevant if AVP is activated and no passengers are seated in the vehicle. Besides, abrupt actions cause increased wear of the material such as brakes, driveshafts, and tires, and make it more difficult for other traffic participants to react in good time. Various thresholds for moderate acceleration are reported in the literature. According to ISO 2631, accelerations greater than $a_x > 1.25 \text{ m/s}^2$ are classified as "very uncomfortable" and $a_x > 2.00 \text{ m/s}^2$ as "extremely" uncomfortable [Int97]. A study involving sinusoidal longitudinal acceleration to a test group of 23 participants found that participants rated the range $1.23 \text{ m/s}^2 < a_x \leq 1.89 \text{ m/s}^2$ as "bad" using verbal classification [Win+23]. Consequently, in this work, the intended longitudinal acceleration a_x of the ego vehicle is considered critically if it exceeds 1.5 m/s^2 , as a conservative estimate. Additionally, acceleration is considered critical if its

corresponding jerk $j_x = \dot{a}_x$ exceeds 2.5 m/s^3 in accordance with a study that statistically quantifies observed jerks in traffic in the range $-2.60 \text{ m/s}^3 \leq j_{0.01} \leq 0.66 \text{ m/s}^3$, where $j_{0.01}$ represents values between the 1st and the 99th percentiles [Fen+17]. Both large positive and negative (braking) accelerations and jerks are considered critical, although negative acceleration may have fewer consequences and may be necessary to prevent collisions.

7.2.3 Motion Estimate Reprocessing

A critical motion estimate does not necessarily have to be discarded. The correct estimate may still be extracted from the measurements if noise points are identified and excluded from the estimation process. Each motion estimate is derived from a specific estimation set.

Definition 7.1 (Estimation Set $\mathcal{E}$)

The estimation set $\mathcal{E}$ is a subset of the measured radar point cloud $\mathcal{P}$, containing only those points assumed to be static and free of noise.

$$\mathcal{E} \subseteq \mathcal{P}. \quad (7.7)$$

Static points support the correct motion hypothesis and enable successful motion estimation, as long as their angle and Doppler velocity value is measured correctly.

This subsection proposes three methods to scrutinize the measurement and extract an alternative estimate. Each method addresses a distinct root cause of an erroneous estimate. The measurement is re-evaluated by selecting an alternative estimation set $\mathcal{E}_{\text{alt}}$.

Physics-Based Point Cloud Filter

The physics-based point cloud filter selects from the point cloud $\mathcal{P}_k$ of the current state $\mathcal{S}_k$ those points that are physically reachable from the previous state $\mathcal{S}_{k-1}$, and includes them in the estimation set $\mathcal{E}$.

Definition 7.2 (Physically reachable motion vector $\hat{\mathbf{m}}$)

A motion vector $\hat{\mathbf{m}}[k]$ is considered physically reachable if the resulting state is attainable from the previous state $\hat{\mathbf{m}}[k - 1]$, given the ego vehicle's dynamic constraints.

All radar measurement points that support a hypothesis of a non-reachable motion vector $\hat{\mathbf{m}}$ are classified as noise. When estimating the longitudinal velocity $v_x[k]$, the points typically align along a cosine curve similar in shape to that of the previous state S_k , if their relative velocity v_{rel} is plotted over their azimuth angle θ . The tolerance region around the current cosine curve is derived from the vehicle's dynamic constraints. In this case, a maximum longitudinal deceleration of $a_{x,\text{min}} = -15.0 \text{ m/s}^2$ is assumed, corresponding to ideal emergency braking conditions. The maximum longitudinal acceleration is defined as $a_{x,\text{min}} = 10.0 \text{ m/s}^2$. Given that the typical Doppler measurement precision reported by most radar sensor manufacturers is approximately $v_{\text{rel}} = 0.2 \text{ m/s}$ and angular measurement errors must also be considered, a tolerance band of $v_{\text{tol}} = 0.5 \text{ m/s}$ is applied symmetrically above and below the expected cosine curve.

Accordingly, the upper filter bound is calculated with

$$v_{\text{rel, up}}(\theta) = \cos(\theta) \cdot (v_x[k - 1] + a_{x,\text{max}} \cdot \Delta t) + v_{\text{tol}}, \quad (7.8)$$

and the lower filter bound can be determined by

$$v_{\text{rel, low}}(\theta) = \cos(\theta) \cdot (v_x[k - 1] + a_{x,\text{min}} \cdot \Delta t) - v_{\text{tol}}. \quad (7.9)$$

Figure 7.4 illustrates the estimation and the expected regions of TP points, bounded by dashed lines. In this example, the velocity $v_x[k - 1]$ in the previous state is assumed to be 2.0 m/s^2 . Points within this region support the physically reachable motion vector and are therefore included in the estimation set $\mathcal{E}$. These points are marked in blue in Figure 7.4. Points outside this region are physically unreachable and are thus classified as noise, shown in red.

Cluster-Based Point Selection

The core assumption of radar-based motion estimation is that most radar points measured around the ego vehicle are static and therefore suitable for use in the estimation process. Typically, those radar points originating from reflections on moving objects can be removed by the MSAC outlier detection algorithm. However, if the number of




Figure 7.4: The dashed filter channel removes noise points (red) that do not correspond to physically reachable motion states. Points supporting reachable motion states are shown in blue.

moving points is no longer negligible compared to the number of stationary points, the outlier detection process tends to fail consistently, which may result in an erroneous motion estimate. An exemplary situation is illustrated in Figure 7.5.




Figure 7.5: Radar-based motion estimation failure caused by a large dynamic object with multiple reflection points (red), which may not be identified as outliers.

The red points represent a dynamic object with multiple reflection points, moving parallel to the ego vehicle. This dynamic object suggests an alternative velocity hypothesis $v_{x,alt}$ which deviates from the correct estimate v_x , supported by the blue points.

By grouping points according to the motion estimates they support, distinct hypotheses for v_x can be extracted. This observation forms the base of the cluster-based point selection method. In the first step, a per-point velocity estimate $v_{x,pt}$ is computed for

each radar point in the current point cloud $\mathcal{P}_k$ based on its azimuth angle θ and Doppler velocity v_{rel} :

$$v_{x,\text{pt}} = \frac{v_{\text{rel}}}{\cos(\theta)}. \quad (7.10)$$

This leads to a number of velocity estimates, which are then clustered using the DBSCAN algorithm. The DBSCAN parameters are determined empirically, with the ϵ -distance set to $\epsilon = 1.0$ and the minimum number of neighbor points set to $\text{min}_{\text{pts}} = 2$. For each cluster $i = 1, \dots, N_C$, the mean velocity $\bar{v}_{x,i}$ is computed. The two largest clusters, C_1 and C_2 , are selected as potential candidates. The cluster whose mean velocity is closest to the previous motion estimate $k-1$ is identified by $C_{\text{min}} = \arg \min(\bar{v}_{x,C_1}, \bar{v}_{x,C_2})$. Finally, the selected cluster C_{min} is accepted as the estimation set $\mathcal{E}$.

Adapted Least-Squares Estimation

As described in Subsection 7.1.2, the motion vector $\hat{\mathbf{m}}$ is estimated from the set of measurements using the least-squares method. This method identifies the optimal minimum under the assumption that all points in the estimation set support the correct motion hypothesis, that is, they originate from static structures, and errors in the measurements of Doppler velocity v_{rel} and azimuth angle θ follow the expected error distribution. However, especially in noisy environments, this assumption may not hold for all points in the original estimation set. Additionally, the total sum of squared errors in an alternative least-squares solution $\hat{\mathbf{m}}_{\text{alt}}$ may deviate only marginally from the global minimum. Nevertheless, the corresponding motion vectors may differ significantly. As the adapted least squares estimation is applied in two passes, the second-pass solution can be modified using *ridge regression*. This method penalizes large magnitudes in the parameter vector, thereby encouraging solutions within a preferred value range. A regularization factor Λ controls the strength of this penalization. The *ridge pseudoinverse* is based on the Moore–Penrose inverse, with an added penalty term $\Lambda \cdot E$ that forces lower values in the estimation of the motion vector. Given the measurement matrix $\mathbf{X}$ and the Doppler velocity vector $\mathbf{y}$, the alternative motion vector is computed as

$$\hat{\mathbf{m}}_{\text{alt}} = (\mathbf{X}^\top \mathbf{X} + \Lambda \cdot E)^{-1} \mathbf{X}^\top \mathbf{y}. \quad (7.11)$$

To guide the estimate toward the expected direction, the initially estimated motion vector $\hat{\mathbf{m}}_{\text{init}}$ is compared with the most recent state estimate. If the estimated longitudinal velocity v_x is deemed too high, Λ is assigned a negative sign, thereby promoting a lower-valued alternative solution $\hat{\mathbf{m}}_{\text{alt}}$. Ridge regression is performed using two values of Λ : a low value Λ_{low} and a high value Λ_{high} . The two resulting estimates are compared by computing the difference between them:

$$d_{\Lambda} = v_{x,\text{est}}(\Lambda = \Lambda_{\text{low}}) - v_{x,\text{est}}(\Lambda = \Lambda_{\text{high}}). \quad (7.12)$$

The alternative estimate $v_{x,\text{est}}(\Lambda = \Lambda_{\text{low}})$ is accepted if the difference d_{Λ} falls below a predefined threshold $d_{\Lambda,\text{thresh}}$. This comparison mechanism is necessary to validate the alternative solution. For instance, two candidate solutions with similarly low residual errors compete, and the initially chosen one may be incorrect. In such cases, ridge regression may successfully select the correct, alternative estimate. In other cases, multiple low-error solutions may exist, none of which represent the true motion state. In such scenarios, the returned solution depends on a stochastic process influenced by the chosen penalty parameter Λ . As a result, the two estimates may differ significantly, and the alternative solution is not accepted. Accordingly, the adapted least-squares estimation can often identify a correct alternative solution, however, it is also possible for the algorithm to fail due to an unfavorable solution space.


7.3 Evaluation on RadarScenes Dataset

The previous section introduced a safety-aware concept to reconsider radar point clouds, with the potential for identifying and correcting erroneous radar-based motion estimates. The effectiveness of this potential is evaluated in the following section. Subsection 7.3.1 describes the dataset used and the methodology for isolating scenarios that are prone to critical motion estimations. Two experiments are performed: the first investigates the quality of the controller criticality assessment, and the second compares three methods for reprocessing the motion estimation. These experiments, including details about their preconditions and the parameterization of the PI controller, are explained in Subsection 7.3.2. Finally, Subsection 7.3.3 presents the experimental results.


7.3.1 Dataset

To evaluate the proposed methods, the dataset for the experiment must fulfill several preconditions. In general, radar-based motion estimation is a robust algorithm that fails only rarely as long as no interference conditions are present. Radar motion estimation requires point clouds as input. Ideally, a GT motion measurement should be available to investigate the radar motion estimates. Although wheel-speed sensors may not be reliable in few individual cases for the reasons mentioned, they can still serve as GT motion measurement. Ideally, a system combining measurements from Global Navigation Satellite Systems (GNSS) and an Inertial Measurement Unit (IMU) is used to generate GT measures. A sufficient amount of problematic radar measurement frames were required to evaluate the behavior of the safety-aware methods if standard motion estimation fails. The RadarScenes dataset [Sch+21b] was promising for fulfilling

these requirements. Accordingly, the camera images in the dataset were examined for scenarios that are typically challenging for radars due to a high environmental noise figure. These sequences included urban drives with many other moving objects, driving through tunnels, or driving in situations with a high metal density, for example, in construction sites. The ego vehicle motion, given by wheel-speed sensor data and IMU data in the dataset, was computed for these promising sequences and compared with the GT velocity estimation. From this process, the sequences with index [2, 9, 94, 126, 132] were identified as the evaluation set. All of these sequences feature a total distance driven of more than 1,000 m per sequence and contain at least 10 measurement frames in which the standard radar-based motion estimation failed. Figure 7.6 contrasts a measurement frame with accurate motion estimation against one recorded in a tunnel, where noise points led to an erroneous motion estimate. In the first scenario, the radar points align with the shape of a cosine curve, and the parameters can be estimated with small error rates. In the second, erroneous measurement, only a few points are on or near this curve. The few TP points can hardly be distinguished from FP noise points. This situation is relevant for detecting a critical controller action and reprocessing the estimate.




(a) Camera image of a frame in sequence 2, time $t = 22.1$ s [Sch+21b]. Radar motion estimation works robustly in this urban scenario, as expected.



(b) Radar measurement points of the scenario. The points arrange in the shape of a cosine curve.



(c) Camera image of a frame in sequence 2, time $t = 107.5$ s [Sch+21b]. Radar motion estimation fails due to the challenging operating conditions for the radars in the tunnel.



(d) Radar measurement points of this frame. Various points appear to be noise points that are not in the shape of the cosine curve.

Figure 7.6: Comparison between a successful radar motion estimation (first row) and a failed estimation (second row).

7.3.2 Experimental Setup

This subsection formulates three hypotheses, each addressed by a corresponding experiment:

1. Rigid filtering reduces the runtime of the radar-based motion estimation and improves its performance, resulting in lower error rates and smaller standard deviations of the estimated longitudinal velocity v_x .
2. The controller action evaluation can detect controller outputs that may lead to critical situations, thereby triggering reprocessing of the motion estimation when necessary.
3. Motion estimation reprocessing can significantly reduce the number of incorrect estimates.

The controller action evaluation requires a planner output in the form of an intended velocity, which serves as the input to the controller. However, the RadarScenes dataset provides only measured radar points and odometry data. Consequently, the current motion state, as measured by the wheel-speed sensors, is used as input for a potential trajectory. In this experiment, the task of the controller is to reproduce the vehicle motion as it occurred during the manual driving recorded in the dataset.

The controller action is estimated based on an ego vehicle dynamics model. Since this experiment focuses exclusively on the longitudinal velocity v_x , only longitudinal dynamics are modeled. The parameters of the model, representing a contemporary battery-electric sedan in the premium segment, are listed in Table 7.1. This parameter set is exemplary for the proposed setup but can be adapted to represent other vehicles. The air density is assumed to be $\rho = 1.204 \text{ kg/m}^3$.

Table 7.1: Vehicle model parameters used to derive the vehicle's dynamic response to controller actions.

Maximum Drivetrain Power P	$\pm 400 \text{ kW}$
Vehicle mass m	$2,295 \text{ kg}$
Drag coefficient c_w	0.22
Cross sectional area A	2.33 m^2
Rolling resistance coefficient C_{rr}	0.015
Acceleration capability a	$\pm 10.0 \text{ m/s}^2$

Using this vehicle model as the controlled system, the PI controller is parameterized with the objective of eliminating the control error $\epsilon(t)$ rapidly and compensating for persistent deviations via the integral term. The controller is tuned using different step functions. A proportional gain $K_p = 0.8$ and an integral gain $K_i = 0.0004$ are identified

as suitable. The controller's behavior for two desired-velocity stimuli, $r_1(t)$ and $r_2(t)$, is illustrated in Figure 7.7.




Figure 7.7: Response of the PI-controller with parameters $K_p = 0.8$ and $K_i = 0.0004$ on the two demanded velocity profiles $r_1(t)$ and $r_2(t)$.

For motion estimation, radar measurements from the four RadarScenes sensors are aggregated to form the radar point cloud: left, front left, front right, and right. The motion estimation is applied to three different point sets, two of which apply rigid filtering by removing points with a weak RCS from the set:

- $-10 \text{ dB m}^2 / 40$: Removes all radar points with $\text{RCS} < -10 \text{ dB m}^2$ if the total number of remaining points exceeds 40; otherwise, all available points are used for estimation.
- $-10 \text{ dB m}^2 / 60$: Removes all radar points with $\text{RCS} < -10 \text{ dB m}^2$ if the total number of remaining points exceeds 60; otherwise, all available points are used for estimation.

The third point set includes all points measured by the four radar sensors in this frame. The objective of this experiment is to determine whether rigid filtering of the point cloud improves both the performance and the runtime of the motion estimation compared to the unfiltered point cloud.

To evaluate the criticality of the controller action, the acceleration-based approach introduced in Subsection 7.2.2 is compared to the deviation between the estimated velocity $v_{x, \text{est}}$ and the GT velocity $v_{x, \text{gt}}$ given by the dataset. This difference can be interpreted as the GT criticality, which reflects the actual criticality of the controller action. The safety tolerance zone of the AVP function is the allowable distance within which the vehicle can operate in the presence of an erroneous motion estimation. A typical safety tolerance distance is $d = 0.3 \text{ m}$, which corresponds to the range resolution of a radar with chirps of a bandwidth $B = 500 \text{ MHz}$. As outlined in Chapter 3, the

reaction time of a braking system is approximately 300 ms. The velocity required to traverse the safety tolerance zone within this reaction time is approximately 2.0 m/s. Therefore, any deviation between the radar-estimated longitudinal velocity $v_{x, \text{radar}}$ and the GT velocity $v_{x, \text{gt}}$ exceeding 2.0 m/s is considered a critical motion estimate.

7.3.3 Experimental Results

The results of the three performed experiments conducted to verify the hypotheses stated in the previous Subsection 7.3.2 are presented below.

Baseline and Rigid Filter Results

To compare the baseline point set with the point set reduced by rigid filtering, several performance metrics are computed: The Mean Squared Error (MSE) over all motion estimates N in the selected sequences

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (v_{x, \text{gt}}[i] - v_{x, \text{est}}[i])^2, \quad (7.13)$$

the standard deviation of the measurement error

$$\sigma_{\epsilon} = \sqrt{\frac{1}{N} \sum_{i=1}^n (v_{x, \text{est}}[i] - \text{MSE})}, \quad (7.14)$$

the maximum observed motion-estimation error ϵ_{max} , and the mean computation time $\bar{t}_{\text{comp}}$. These metrics, along with the point reduction achieved by the rigid filters, are listed in Table 7.2.

The rigid filter removes approximately one-third of the radar points, resulting in a moderate decrease in runtime when applied. However, the MSE, maximum error, and standard deviation are all significantly higher compared to the unfiltered point cloud.

Table 7.2: Results of the experiment testing the first hypothesis. Results examining whether rigid filters improve the performance of radar-based motion estimation.

	Baseline	-10 dB m ² /40	-10 dB m ² /60
MSE in m ² /s ²	3.564	8.644	4.029
Standard Deviation σ_ϵ in m/s	1.815	2.670	1.870
Max. Error $\epsilon_{\max}$ in m/s	58.711	147.856	71.475
Mean Runtime $\bar{t}_{\text{comp}}$ in ms	1.925	1.776	1.887
Point Reduction	0.00 %	-41.2 %	-31.6 %

Controller Action Evaluation Results

Table 7.3 lists the performance metrics of the controller action evaluation — TPs, FPs, and FNs — for each sequence of the evaluation set. The table also includes sequence-specific statistics, namely the traveled distance and sequence duration. TPs are estimates that are correctly identified as critical by the proposed criticality assessment. FPs are estimates exhibiting only a small difference between the estimated radar velocity $v_{x, \text{radar}}$ and the GT vehicle velocity $v_{x, \text{gt}}$, yet are incorrectly classified as critical. FNs are estimates with a substantial deviation between $v_{x, \text{gt}}$ and $v_{x, \text{radar}}$ that are not identified as critical by the criticality assessment.

Table 7.3: Results of the controller action evaluation. Performance metrics of the controller action evaluated on five challenging RadarScenes sequences.

Sequence Number	Distance in m	Time in s	TPs	FPs	FNs
2	1,046	145.2	51	71	0
9	1,543	92.4	36	7	0
94	1,367	167.2	9	29	4
126	2,508	138.9	24	7	0
132	2,267	116.4	47	7	0

Overall, the sequence length does not correlate with the number of critical actions. With the exception of sequence 94, all critical controller actions are correctly identified. However, several FPs are observed in each sequence. This is attributable to the parameterization of the controller criticality evaluation, which prioritizes safety over minimizing FPs.

Motion Estimation Reprocessing Results

The same metrics used to compare the motion estimation performance of unfiltered and filtered point clouds are applied to evaluate three approaches for measurement data reprocessing: physics-based point cloud filtering (I), cluster-based point selection (II), and adapted least-squares estimation (III). The results are summarized in Table 7.4, with the best-performing method for each sequence highlighted in bold.

Table 7.4: Evaluation of motion estimation enhancement methods. Comparison of three enhancement methods applied to the motion estimation process through iterative revisiting of the estimation.

		MSE in m^2/s^2	σ_ϵ in m/s	$\epsilon_{\max}$ in m/s
Seq. 2	Baseline	3.035	0.053	42.493
	I	0.830	0.911	25.249
	II	1.747	1.322	28.852
	III	1.579	1.255	44.581
Seq. 9	Baseline	4.253	2.062	60.422
	I	0.207	0.453	14.438
	II	0.663	0.814	24.505
	III	1.631	1.277	59.872
Seq. 94	Baseline	1.066	1.033	26.869
	I	0.103	0.321	14.439
	II	0.480	0.692	20.561
	III	0.033	0.003	7.458
Seq. 126	Baseline	2.884	1.696	45.874
	I	0.349	0.065	14.142
	II	0.685	0.825	20.163
	III	1.036	1.017	35.306
Seq. 132	Baseline	4.438	2.105	45.457
	I	0.190	0.423	14.212
	II	1.479	1.215	45.457
	III	0.255	0.492	18.241

All three reconsideration approaches significantly outperform the baseline motion estimation, yielding lower MSE values with reduced variance. In all sequences except sequence 126, method I demonstrates superior performance. The performance of

methods II and III varies depending on the sequence but is generally comparable and inferior to method I. Nevertheless, it should be noted that, in each of the five investigated sequences, at least one estimate with a high error remains after reprocessing. For corrected motion estimates, the error should not exceed $\epsilon_{v_x} \approx 1.0 \text{ m/s}$. This condition is still not met in at least one frame of each of the five sequences, indicating that motion estimation reprocessing does not generally guarantee receiving a correct motion estimate.

7.4 Discussion

This chapter demonstrated that critical radar data can occur independently of the specific function for which the data are used in subsequent processing steps. It further argued that a criticality-aware radar data processing can also be applied to functions beyond object detection. To support this assumption, the concept of erroneously removed radar points was extended to the motion estimation task. Furthermore, the impact of erroneous sensor data processing on the safety of the automated function was evaluated. The results showed that hazardous effects on the control process can indeed arise when incorrect data are used in the motion estimation process. This analysis addressed (RQ4). In the next steps, methods were developed to identify such situations. When safety-aware selection of the radar points is applied, an alternative motion estimation can be derived. Experimental results demonstrated that the use of the safety-aware pipeline significantly reduces the number of safety-critical situations caused by the standard data filtering process and incorrect motion estimates.

The first hypothesis stated that applying rigid filtering to reduce the number of noise radar points (FP points) would improve the performance of motion estimation. As shown in Table 7.2, this assumption was falsified in the experiments. The rigid filters removed a significant portion of the measured points, leading to poorer motion estimation results. The MSE accumulated over the five sequences increased by 143% when 40 points were considered sufficient for a successful motion estimation. Even in the less aggressive filtering configuration, which ensured at least 60 radar points for estimation, the MSE still increased by 13%. Similarly, the standard deviation of the error σ_ϵ and the maximum error $\epsilon_{\max}$ increased compared to the unfiltered baseline. These notable performance losses were accompanied by only minor runtime improvements: the mean runtime decreased by 2% in the less aggressive filtering setup and by 8% in the more aggressive setup. The computational costs of processing more radar points in the estimation is relatively low, and the filtering operation itself incurs a small overhead that must be compensated for. While rigid filtering can be beneficial in increasing the robustness of object detection or localization algorithms, it proved to be non-beneficial for the motion estimation task as defined here. For this reason, all subsequent experiments were conducted using the unfiltered point set as the baseline.

The second experiment defined critical controller actions as those in which the estimated velocity $v_{x, \text{est}}$ deviates significantly from the GT velocity $v_{x, \text{gt}}$. Such critical actions can place the automated function in an unsafe condition, potentially endangering passengers or other traffic participants. The situations identified as critical by this definition were compared with those detected by the controller action evaluation, which assesses the controller's output. In this way, the recall and precision of the method for assessing criticality of the motion estimation from the controller output can be derived. The experiment demonstrated that the controller action evaluation is capable of detecting most of the relevant situations, resulting in a high recall value. The precision across all sequences was approximately 0.580, while the parameterization was tuned for high recall. Consequently, the recall value was found to be 0.977. Most critical estimations were clearly identifiable as outliers, although some contained only minor errors. The FNs in sequence 94 corresponded to estimates that were not pronounced enough to trigger a controller response exceeding the acceleration or jerk thresholds. These thresholds were selected based on the ODD to balance the additional computational effort required for motion estimation reprocessing with the need for thorough identification of critical situations. According to the results, the controller action evaluation appears suitable for identifying those measurement frames with erroneous initial motion estimates, which are warranted for scrutinization.

Critical controller outputs identified in this way can be reprocessed to correct erroneous motion estimates. Three reprocessing methods, introduced in Subsection 7.2.3, were evaluated separately for each sequence. The results show that each method has the potential to reduce the baseline MSE significantly. The physics-based point cloud filter, for example, removes points that are inconsistent with the last observed motion state. This strategy would only lead to incorrect estimates if the cosine shape of the motion estimate were to change significantly between two measurement frames, even when all points fit the new cosine shape. However, such abrupt changes in the cosine shape are not physically plausible. Thus, the method remains safe to apply. This finding supports the observation that TP points exhibit temporal consistency in their spatial distribution, whereas most noise points occur randomly without any discernible pattern. This property can be effectively exploited by the physics-based filtering method.

The second method appeared to be the least effective. It clusters points that support a similar velocity v_x and performs the estimation on the majority of points supporting the same velocity hypothesis. This approach is especially suitable for scenarios involving strong radar moving reflectors. In such cases, the assumption that most measured points are static is no longer valid. However, this algorithm fails if most points are FPs, as this results in clusters forming at random locations. Additionally, the method does not yet prioritize points at perpendicular angles to the sensor. Such points provide more accurate estimates than those with small or no relative velocity components. Figure 7.8 illustrates the mean deviation of the estimated velocity $v_{x, \text{est}}$ per point from the GT velocity $v_{x, \text{gt}}$. If enough points are available, points with perpendicular angles $\theta \approx 0^\circ \cup \theta \approx 180^\circ$ could be assigned a higher weight during clustering.




Figure 7.8: Average motion estimation error over radar point azimuth angle. Average estimation error per radar point plotted against the azimuth angle between the point and the sensor vertical.

The number of points in a cluster indicates how suitable the entire point set is for motion estimation. As the example in Figure 7.6 showed, the ratio of TPs to FPs can reach such a disadvantageous level that safe and correct estimation cannot be guaranteed. Such limitations in measurement quality must be avoided at all costs, for example, by implementing strategies to mitigate intensive interference effects.

In contrast to the first two methods, which select individual, potentially TP points to form a subset for re-estimation, the third method operates on the complete subset. This forces the estimate toward the assumed velocity value, which should be close to the previous velocity value $v_x[k-1]$. If the data quality is sufficiently high, meaning that the least-squares estimation produced an incorrect result only due to an unfavorable configuration of matrices X and y in the first estimate, this method may yield high quality results. On the other hand, it is ineffective in cases of a high FP rate and does not improve the estimation here.

If reprocessing the estimation fails several times consecutively and the controller action remains critical, a replacement value should be chosen. In this case, the wheel-speed sensor value or the last estimate can serve to bridge the time until the next measurement. While analyses showed that false estimations increasingly occurred in challenging environments such as tunnels, acceptable measurement cycles were still available for correct motion estimation even in these situations.

This work on safe radar motion estimation was evaluated only for the longitudinal velocity v_x , although the motion estimation method is capable of providing a complete motion vector $\hat{\mathbf{m}} = [v_x, v_y, \omega]^\top$ using radar sensors. This vector fully describes the ego vehicle's motion in two-dimensional space. The concept of criticality estimation is then transferred from unsafe longitudinal acceleration actions to unsafe steering actions, while the architecture of the safe motion estimation remains as proposed here.

In an advanced implementation, the evaluation of controller actions may be enriched by incorporating objects detected by the perception system as additional inputs. This enhancement enables a more accurate categorization of criticality within the situational context. Extreme control outputs may be reasonable in cases of threatening collisions. Conversely, tolerance for control actions may be increased if sufficient space is available, meaning that an action exceeding the standard threshold may still not be critical.

Similarly to the motion estimation function, safe radar data processing could be applied to other tasks such as localization, as discussed in Subsection 7.1.1, or to free space estimation with minor adaptations. Furthermore, criticality estimation by evaluating controller actions enables sensor fusion with, for example, accelerometers, to validate the estimator's behavior.

8 Conclusion and Outlook

This work presents and discusses a method to introduce a safety-aware component into radar data processing for automated parking functions. Safety-awareness refers to the prevention of incidents, particularly collisions that result in injuries to traffic participants. Compared to incidents of the automated function that cause dissatisfaction, a diminished sense of safety, or even material damage, such as erroneous emergency braking or collisions with static objects, the protection of life has the highest priority in every action. Hence, standards such as ISO 26262 require compliance with defined FIT rates for each incident that puts human life at risk. In the ODD of automated parking, this is potentially the case during collisions between the ego vehicle, which is currently controlled by the automated function, and a VRU.

In addition to high sensor performance, which enables the detection of VRUs, the FIT rates can only be reached if information delivered by the sensor is guaranteed to pass the signal processing pipeline. At the same time, the enormous amount of data delivered by each radar sensor needs to be condensed to a scene representation which can be used by the automated function. The key concept presented in this work is to adapt the processing of information depending on its relevance to the prevention of hazardous situations. Accordingly, this method builds on the current intention of the planner of the parking function, but can also be adapted to manual driving by predicting the most probable vehicle states in the immediate future. The maximum time horizon for this criticality evaluation is the time that is required to set all traffic participants in a safe — i.e. static — motion state while avoiding potential collisions. The criticality estimation which is applied on each sensor information, e.g., a radar point, identifies information whose processing is especially important. Hence, (RQ1) was solved with this method and served as a foundation for the safety-aware radar data processing that was investigated in the following.

The answer for (RQ2) had multiple layers: Action must be taken differently depending on the stages in the processing pipeline. This work contributed methodologies to evaluate information in different radar processing stages and to modify the processing in these stages. In general, the following three stages were considered in detail:

1. *RF images* are one of the earliest stages in radar data processing and contain the signal power for each cell in the measurement range. The planner's intention can already be linked to the relevance of individual RF cells. As this stage's product is a radar point cloud, the main objective is to generate points from cells with a sufficiently high signal power representing a VRU. This work proposed a criticality-aware thresholding method, which facilitates the generation of points

in the neighborhood of cells where high signal values were observed in the past. By scanning the last RF images, a voting system is able to confirm even weak detections.

2. *Radar Point Clouds* are an environment representation that is condensed compared to RF images. As research in point cloud processing recently proceeded, these point clouds are a popular basis for further processing or fusion with other sensor modalities. However, radar point clouds appear noisy and suffer from a low density. A common strategy, filtering the point cloud and accumulating points from multiple radar measurements, can be a risk for safe radar data processing since information may be lost. By region-based recapturing of the filtered points, this work presented an approach to apply point cloud filtering while maintaining all potentially relevant points.
3. *Object Hypotheses* are the final level of radar data processing. An object list is used by the automated function to plan the path, additionally, it offers the opportunity to adapt the trajectory in the short term. Even if the information about the existence of a VRU interfering with the planner's intention reach the radar point stage, it may still be discarded due to a lack of radar points during the clustering and tracking process. Both were points of action for the methods introduced in this work. The generation of both clusters and object hypotheses may be fostered in critical situations by lowering the thresholds for such late-stage object evidence.

The criticality information can be used to adapt each of these steps to foster the measurement data transfer through the processing pipeline if this data is evaluated as sufficiently critical. By this means, the information can contribute to generate an object hypothesis and prevent that critical objects are not considered due to an information loss in the pipeline of radar data processing. The effect of the modified methods can be evaluated methodically using the failure rate estimation proposed in this thesis, that answers (RQ3). Each chapter's methods were evaluated by experiments performed on the respective radar data modality. In Chapter 6, this evaluation required the recording of a custom dataset due to the unavailability of existing ones.

Finally, the concept was transferred from the object detection task to the motion estimation task to show its versatility. Answering (RQ4) confirmed the relevance of this investigation and it can be supposed that each task based on radar sensor data with a relevance for the planner can cause hazardous behavior. In case of motion estimation, the objective was to prevent erroneous estimates that could compromise the safety of the downstream control task. Owing to differences in the underlying processing task, the proposed motion criticality assessment differs from the criticality assessment used for the detection use case, as described in Chapter 3. Nevertheless, the proposed methods successfully identify erroneous measurements and provide the option to correct incorrect estimates to a certain extend.

All experiments with the proposed methods showed that most of the data identified as critical consist of FP information, which is undesirable in the object list as an end

product of the radar data processing pipeline. A crucial finding was that the information provided by individual or few radar points is not sufficient to decide between an existing object (TP) or radar noise (FP). This leaves one major problem, the balance between precision and recall, not completely resolved. However, the methods in this work allow for different balances of both metrics, depending on the specific requirements of each information domain. While the precision at less relevant areas far from the ego vehicle can be increased since a lower recall is acceptable, the processing can be adapted to ensure the required high recall in critical areas. The work proved that integrating the relevance of radar measurements on the planned task demonstrably significantly decreases the likelihood that information is filtered erroneously.

Outlook

Radar data processing is moving toward the application of deep learning algorithms, which increasingly feature black-box characteristics. For the success of highly automated parking and driving functions, including these algorithms is becoming inevitable. It is a future challenge to make these algorithms explainable to provide safety guarantees. The approaches presented in this work, connecting the planner's intention with object detection, need to be combined with these new learning-based algorithms. One alternative that is already discussed in the research community is the application of end-to-end learning, which places higher demands on explainability than combined or hybrid algorithms. This work also proved that the performance of radar sensors is still not sufficient to serve as a singular solution for a SAE Level 4 function. The general measurement concept of radar sensors requires the combination and fusion with other sensor principles such as camera sensors. One approach demonstrated in this work is the use of region proposals, where findings from camera images serve as proposals to identify weak, but critical object evidence in the radar data. The different measurement principles provide additional safety for object detection since the measurement modalities may complement each other.

A Appendices to Chapter 2

A.1 Functional Principle of FMCW Radar Sensors

This section provides a brief overview of the transmission patterns of FMCW radar sensors. These sensors modulate their transmit frequency to obtain the distance d and the relative velocity v_{rel} of reflections by applying two consecutive Discrete Fourier Transforms (DFT)s. The frequency-modulated transmit pattern is referred to as a chirp, analogous to a sound signal. The difference between the start frequency f_0 and the highest transmit frequency is denoted by Δf . Within each chirp, the transmit frequency increases linearly over a bandwidth of $B = 2 \cdot \Delta f$ during the transmission time t_S . The corresponding frequency slope is given by $S = \frac{B}{t_S}$. To distinguish between static and dynamic targets, the radar transmits both up-chirps with a positive frequency slope, and down-chirps, with a negative frequency slope, within a single measurement cycle. This enables the computation of the beat frequency f_b , obtained by mixing the transmitted and received signals. The beat frequency is then extracted via a DFT. The up-chirp beat frequency $f_{b,+}$ and the down-chirp beat frequency $f_{b,-}$ are calculated using

$$f_{b,+} = f_{+, \text{transm}} - f_{+, \text{rec}}, \quad (\text{A.1})$$

$$f_{b,-} = f_{-, \text{rec}} - f_{-, \text{transm}}. \quad (\text{A.2})$$

The range R between the antenna and the reflector, as well as the Doppler shift f_{Dopp} , from which the relative velocity v_{rel} can be derived, can then be computed [Jan18]:

$$R = \frac{c \cdot t_S}{4 \cdot B} \cdot \frac{(f_{b,+} - f_{b,-})}{2} \quad (\text{A.3})$$

$$f_{\text{Dopp}} = -\frac{(f_{b,+} - f_{b,-})}{2}. \quad (\text{A.4})$$

The chirp pattern, including the transmitted and received signals and the resulting beat frequencies, is illustrated in Figure A.1.

In the single-chirp case, the computed range R and Doppler velocity v_{rel} become ambiguous, since a single chirp provides only a linear relationship $\frac{R}{v_{\text{rel}}} = \text{const.}$ between the two parameters. This ambiguity can be resolved by determining the intersection of the relationships obtained from an up-chirp and a down-chirp, a method known




Figure A.1: Chirp Patterns in FMCW and the corresponding beat frequencies for up-chirps and down-chirps, adapted from [Jan18].

as pairing [Gam20]. However, with recent advances in radar technology and faster electronic components, the fast-chirp ramp sequence has become the current standard approach. While fundamental principles for deriving range and Doppler velocity remain unchanged, multiple frequency ramps are transmitted in rapid succession. This enables the use of a two dimensional DFT (2-DFT), which simultaneously yields samples for both range R and Doppler velocity v_{rel} . Finally, the azimuth angle can be estimated when a multi-antenna configuration is employed. The transmitting and receiving antennas are arranged on the sensor surface with a spacing of $\lambda/2$ (half the wavelength). This spatial configuration introduces phase shifts in the received signals depending on the angle of arrival. By processing the signals from multiple antennas using high-resolution direction-of-arrival estimation techniques, such as Multiple Signal Classification (MUSIC), the azimuth angle can be accurately estimated.

A.2 Exemplary STPA for the AVP Function

The goal to this appendix is to identify scenarios that may cause a hazard potentially resulting in a loss during the operation of the AVP function. The tool used for this analysis is the STPA, which is divided into four steps:

1. All potential hazards associated with the use of the AVP function, as specified in [Int23], are identified. Moreover, potential losses corresponding to each hazard are listed. According to STPA, losses can include injury, but also material damage, or loss of public trust. In this work, the investigation of a hazard proceeds only if the suspected loss involves injury or loss of life. The identified hazards and losses are presented in Table A.1.
2. The control structure is modeled. This step is omitted in this work, as the required control input to avoid the identified hazards can be simplified to a braking command.
3. Each behavior within an operational situation in the ODD that can lead to a hazard is referred to as an unsafe control action. The possible unsafe control actions are identified.
4. Causal scenarios are derived from these unsafe control actions. These scenarios should be addressed in the test cases for the AVP function. A causal scenario can consist of one or multiple causal factors leading to the hazardous condition.

In this thesis, the analysis is restricted to unsafe control actions within the perception component of the AVP function. Furthermore, only unsafe control actions leading to hazards classified as relevant are considered. One of these hazards, malfunction in self-diagnosis, covers incorrect observation of sensor impairments, such as sensor blindness or interference. These additional aspects are not covered in this thesis. The remaining two hazards are directly associated with erroneous performance of the perception function, as they involve potential collisions with VRUs due to missed detections. For this analysis, the perception function is considered to comprise three components: peak detection, clustering, and tracking. Four types of hazards may arise at each stage: failure to provide information, provision of incorrect information, provision of information at the wrong time, and premature termination or excessive continuation of information provision [Glö24]. Unsafe control actions that can affect perception-related hazards and their causal factors are summarized in Table A.2. Each of these unsafe control actions can cause hazards such as collisions or insufficient distance from a VRU, thereby posing risks of injury or loss of life. The causal chain defined in the STPA is illustrated in Figure A.2.




Figure A.2: Causal chain in the STPA, illustrating how causal factors lead to hazards.

Table A.1: The first step of STPA: identification of hazards and potential losses resulting from these hazards.

Hazard	Description	Relevant for perception	Loss of life suspected
Exceeding speed	Maximum allowed speed within ODD is exceeded.		X
Min. distance non-VRUs	Minimum distance to non-VRUs is undercut.		X
Min. distance VRUs	Minimum distance to VRUs is undercut.	X	X
Traffic rules	Traffic rules of the ODD are disregarded.		X
Unintended Movement	Vehicle moves unintendedly although in a parked state.		X
VRU detection	VRUs within ODD are not detected.	X	X
Obstacle detection	Appearing obstacles within ODD are not detected.		X
Personnel instructions	Vehicle does not react on instructions sent by the infrastructure backend or teleoperators		
Obstructive parking	Vehicle parks in a way endangering other visitors.		X
Malfunction in self-diagnosis	Vehicle fails in detecting a malfunction which affects the function critically	X	X
Malfunction in emergency situation	Vehicle fails in handling an emergency operation correctly.		X
Leave of ODD	Vehicle leaves ODD.		X

Table A.2: The third and fourth steps of STPA: identification of unsafe control actions and causal factors for the relevant hazards. For each unsafe control action, it is evaluated whether controllability can be achieved by adapting the processing method.

Unsafe control action	Causal Factors	Controllable
Information not provided in radar spectrum	High noise level, low object reflectivity, weak sensor performance	(X)
Information in radar spectrum provided incorrectly	High noise level, interfering signals	(X)
Information in radar spectrum provided at the wrong time	Raw signal processing too slow	
Information in radar spectrum provided too short/long	High noise level, low object reflectivity, weak sensor performance	(X)
Information not provided in radar clusters	Point density too low, clustering parameters incorrect	X
Information in radar clusters provided incorrectly	Point attributes are incorrect (position, velocity), clustering incorporates outliers	X
Information in radar clusters provided at the wrong time	Incorrect spectrum processed	X
Information in radar clusters provided too short/long	Point density too low, clustering parameters incorrect, noise points incorporated in processing	X
Information not provided in object hypothesis	Hypothesis is refused, clustering information is not available, tracking parameters are incorrect	X
Information in object hypothesis provided incorrectly	Incorrect motion model for tracking, wrong information in cluster	X
Information in object hypothesis provided at the wrong time	Incorrect cluster list processed	X
Information in object hypothesis provided too short/long	Hypothesis is refused, clustering information is not available, tracking parameters are incorrect	X

B Appendices to Chapter 4

B.1 Exemplary Projection of Bounding Boxes in the Radar Spectrum

The accuracy of the projection is evaluated for this image. The results are listed in Table B.1. For each object in the image, the Euclidean distance $\epsilon_{img} = \|P_{img} - P_{gt}\|_2$ between the projection coordinate (x_{img}, y_{img}) of the bounding box’s ground point and the GT position (x_{gt}, y_{gt}) is computed. This spatial deviation is mainly caused by the fact that objects must be concentrated to one reference point (x_{gt}, y_{gt}) . This point requires a clear definition on where this point is located at an object (e.g., in the center of the bicycle, at the position of the bicyclist or at the position of the bicycle’s axes). In ROD2021, this reference is not given. Hence, the precision for a pedestrian is higher due to its smaller dimensions compared to a bicycle or a car.

Table B.1: Evaluation of the spatial precision of image-based, three-dimensional object detection. Example from ROD2021, sequence *MLMS003*, image ID 701.

Object	x_{img}	y_{img}	(u, v)	x_{gt}	y_{gt}	ϵ_{img}
Car	12.838 m	9.437 m	(88, 600)	11.876 m	8.316 m	1.477 m
Ped. L	110.767 m	41.632 m	(393, 550)	-	-	-
Bicyclist	14.652 m	3.454 m	(512, 593)	11.166 m	2.171 m	3.715 m
Ped. R	10.909 m	0.516 m	(672, 610)	10.699 m	0.373 m	0.254 m

B.2 Pseudocode for the Threshold Reduction in Critical Cells

This pseudocode illustrates the method of detecting weak peaks in preceding spectra by spatially reducing the detection thresholds using Gaussian dimples and applying a voting process, as outlined in Chapter 4.

Figure B.1: Algorithm illustrating the procedure of the proposed subsequent peak search method.

Algorithm 2 Subsequent peak search algorithm

```

1: Input: Abs. power spectrum  $\mathcal{S}$ , Standard threshold factor  $k_{\text{high}}$ , Reduced threshold
   factor  $k_{\text{low}}$ , Retrospection length  $n_{\text{hist}}$ 
2: Output: List of detections  $\mathcal{D}$ 
3: function INVESTIGATE_PRIOR_FRAME(frame_id)
4:    $\mathcal{S} \leftarrow \text{load\_spectrum}(\text{frame\_id})$ 
5:   Threshold  $\mathcal{T}_0 \leftarrow \text{ca\_cfar}(\mathcal{S})$ 
6:    $\mathcal{D}_{\text{confirm}} \leftarrow \text{peak\_detection}(\mathcal{S}, \mathcal{T}_0, k_{\text{high}})$ 
7:    $\mathcal{D}_{\text{pot}} \leftarrow \text{peak\_detection}(\mathcal{S}, \mathcal{T}_0, k_{\text{low}})$ 
8:   for  $i \leftarrow 0$  to  $-n_{\text{hist}}$  do
9:     Reduction term  $\Delta(i) \leftarrow \text{get\_reduction\_term}(i)$ 
10:    Reduced threshold  $\mathcal{T}_{i, \text{red}} \leftarrow \mathcal{T}_i - \Delta(i)$ 
11:     $\mathcal{D}_{\text{pot}} \leftarrow \text{peak\_detection}(\mathcal{S}_i, \mathcal{T}_{i, \text{red}}, k_{\text{high}})$ 
12:    for each det in  $\mathcal{D}_{\text{pot}}$  do
13:       $d_{\text{eucl}} \leftarrow \text{euclid\_dist}(\text{det}, \text{det}_{\text{confirm}})$ 
14:      if  $d_{\text{eucl}} < 0.5$  and  $\text{is\_critical}(\text{det})$  then
15:        det.voting  $\leftarrow$  det.voting + 1
16:     $\mathcal{D}_{\text{pot}} \leftarrow \mathcal{D}_{\text{pot}}[\text{det.voting} > 2]$ 
17:     $\mathcal{D} \leftarrow [\mathcal{D}_{\text{pot}}, \mathcal{D}_{\text{confirm}}]$  return  $\mathcal{D}$ 

```

Abbreviations and Symbols

Abbreviations

ACC	Adaptive Cruise Control
AD	Automated Driving
ADC	Analog-to-Digital Converter
AEB	Automated Emergency Braking
ASIL	Automotive Safety Integrity Level
AVP	Automated Valet Parking
CA	Cell Averaging
CCL	Connected Component Labeling
CFAR	Constant False Alarm Rate
CNN	Convolutional Neural Network
CTRA	Constant Turn Rate and Acceleration
CTRV	Constant Turn Rate and Velocity
CuT	cell-under-test
CV	Constant Velocity
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DFT	Discrete Fourier Transforms
DGU	distributed generation unit
DSP	Digital Signal Processing
EKF	Extended Kalman Filter
EPS	Electronic Power Steering
ESC	Electronic Stability Control
FFT	Fast-Fourier-Transformation
FIT	Failure-in-Time
FMCW	Frequency-Modulated Continuous-Wave
FN	False Negative
FoV	Field of View
FP	False Positive
GDV	Gesamtverband der Deutschen Versicherungswirtschaft, German Insurance Association
GNSS	Global Navigation Satellite Systems
GT	Ground Truth
HARA	Hazard Analysis and Risk Assessment
HD	High-dimension

HW	Hardware
ICP	Iterative Closest Point
ID	Identifier
IMU	Inertial Measurement Unit
ISO	International Organization for Standardization
KBA	Kraftfahrt-Bundesamt, Federal Motor Transport Authority
KL	Kullback-Leibler
MIMO	Multiple Input Multiple Output
MPC	Model Predictive Control
MSAC	M-Estimator Sample Consensus
MSE	Mean Squared Error
MUSIC	Multiple Signal Classification
NDT	Normal Distributions Transform
NHTSA	National Highway Traffic Safety Administration
$\mathcal{N}$	Noise level, type=latin
ODD	Operational Design Domain
ODR	Orthogonal Distance Regression
OEM	Original Equipment Manufacturer
PI	Proportional-Integral
PRI	Pedestrian Risk Index
RA	Range-Azimuth
RAD	Range-Azimuth-Doppler
RANSAC	Random Sample Consensus
RCS	Radar Cross Section
RD	Range-Doppler
RF	Radio Frequency
RoI	Region(s) of Interest
Rx	Receive
SAE	Society of Automotive Engineers (SAE) International
SLAM	Simultaneous Localization And Mapping
SNR	Signal-to-Noise Ratio
SOTIF	Safety Of The Intended Functionality
STPA	System-Theoretic Process Analysis
TP	True Positive
TTC	Time-to-Collision
Tx	Transmit
VRU	Vulnerable Road User

Latin Letters

Symbol	Description
a	Depending on context: acceleration or attribute
A	Depending on context: point set evaluated critical or system matrix
B	Depending on context: RA blob as a set of cells with detected peaks, set of TP critical points, or control matrix
b	Depending on context: bandwidth or number of cells for CA-CFAR
c	Depending on context: speed of light or camera principle point
C	Depending on context: RA cell in radar spectrum, set of clustered radar points, cross sectional area, or observation matrix
$crit$	Criticality
d	Depending on context: depth, distance, or lateral Frenet coordinate
$\mathcal{D}$	Detection set
e	Ellipsoid parameters
$\mathcal{E}$	Estimation set
F	Set of radar points removed by a rigid filter
f	Depending on context: frequency or camera focal length
g	Gravitational constant
G	Depending on context: antenna gain or process noise matrix
h	Depending on context: height or hypothesis
$\mathcal{H}$	Set of hypotheses
i	Index 1, e.g., range bin index Camera image
j	Depending on context: index 2, e.g., angular bin index, or jerk
k	Depending on context: constant factor or natural number between 1 and n
$\mathcal{K}$	Kernel for CA-CFAR noise floor estimation
l	Length
L	System loss
m	Depending on context: mass or motion state
n	Natural number
N	Total number
P	Depending on context: point 1, power, probability, or error covariance matrix
q	Point 2
r	Depending on context: radial position, radius, desired value, or range cell
s	Longitudinal frenet coordinate
S	Signal magnitude
$\mathcal{S}$	State
T	Threshold for CA-CFAR peak detection
t	Time
$\mathcal{T}$	Trajectory

Symbol	Description
u	Depending on context: controller input or longitudinal pixel position in camera image
v	Depending on context: velocity or lateral pixel position in camera image
w	Width
W	Energy
x	X position/direction
y	Depending on context: y position/direction or controller output
z	Depending on context: z position/direction or observation
Z	Z position direction in world coordinate frame

Greek Letters

Symbol	Description
α	Angle of the circular arc
γ	Rotation angle between <i>world</i> and <i>vehicle</i> coordinate frame
δ	Depending on context: steering wheel angle or front wheel angle
Δ	Gaussian dimple function for threshold reduction
ϵ	Depending on context: distance measure for DBSCAN, error, or permittivity
θ	Angle between a radar point and the moving direction of the ego vehicle
κ	Multiplier for a vector
λ	Depending on context: radar wavelength or multiplier for a vector
Λ	Penalty term
μ	Depending on context: mean value or friction coefficient
ρ	Depending on context: air density or Pearson Correlation Coefficient
σ	Depending on context: RCS or standard deviation
τ	Index of considered measurement frame
φ	Depending on context: azimuth angle or azimuth angle cell
ψ	Steering rack position
ω	Yaw rate, turn rate

Indices and Exponents

Symbol	Description
$\square_A$	Ackermann angle
$\square_{\text{accept}}$	Acceptable
$\square_{\text{adapt}}$	Adapted
$\square_{\square,+}$	Component for addition
$\square_{\text{aero}}$	Aerodynamic
$\square_{\text{alt}}$	Alternative Solution
$\square_{\text{assoc}}$	Associated
$\square_{\text{av}}$	Denominator for averaging cells
$\square_b$	Beat
$\square_{\text{bin}}$	Bins in the RA spectrum
$\square_B$	Blob
$\square_{\text{brake}}$	For braking to standstill
$\square_C$	Cluster
$\square_C$	... in Cartesian coordinates
$\square_{\text{cam}}$	Camera
$\square_{\text{cand}}$	... of the detection candidate
$\square_{\text{class}}$	Semantic class
$\square_{\text{coll}}$	... at collision
$\square_{\text{confirm}}$	Confirmation cycles
$\square_{\square,\times}$	Component for multiplication
$\square_{\text{cur}}$	Current
$\square_{\text{Dopp}}$	Doppler ...
$\square_{\text{ego}}$	... of the ego vehicle
$\square_{FA}$	False alarm
$\square_{\text{filt}}$	Filtered
$\square_{\text{fix}}$	Fixed
$\square_{FN}$	FN
$\square_{FP}$	FP
$\square_{\text{frames}}$	... of the frame rate
$\square_{\text{frenet}}$	... in the Frenet coordinate system
$\square_{\text{ground}}$	Position of pixel which is on the ground level
$\square_{gt}$	GT
$\square_{\text{guard}}$	Denominator for guarding cells
$\square_{\text{hist}}$	Historical
$\square_{\text{img}}$	Image
$\square_{\text{init}}$	Initial
$\square_k$	Arbitrary variable in a sample between 1 and n
$\square_{\mathcal{K}}$	Kernel
$\square_{\text{kin}}$	Kinetic

Symbol	Description
$\square_L$	Left side
$\square_{\text{lat}}$	Lateral component of a vector
$\square_{\text{long}}$	Longitudinal component of a vector
$\square_M$	Mixing
$\square_{\text{margin}}$	Safety margin
$\square_{\text{match}}$	Matching
$\square_{\square, \text{max}}$	(Allowed) maximum of the value
$\square_{\text{meas}}$	Measured
$\square_n$	Number of available samples
$\square_{\mathcal{N}}$	Noise
$\square_{NN}$	Nearest neighbor
$\square_{\text{obj}}$	Object
$\square_{\text{obs}}$	Observed
$\square_{\text{odom}}$	Odometry
$\square_{\text{opt}}$	Optimal ...
$\square_p$	Proportional
$\square_{\text{park}}$	... of a parked vehicle
$\square_{\text{pc}}$	Point cloud
$\square_{\varphi}$	... azimuth angle
$\square_{\text{rec}}$	Required for planning
$\square_{\text{pot}}$	Potential ...
$\square_{\text{pred}}$	Predicted
$\square_{\text{pts}}$	Points
$\square_r$	Range, radius
$\square_{\text{rad}}$	Radial
$\square_{\text{react}}$	required for human or system reaction
$\square_{\text{rec}}$	Received
$\square_{\text{red}}$	Reduced
$\square_{\text{rel}}$	Relative (often the relative velocity that the radar sensor determines from the signal's Doppler shift)
$\square_{\text{rem}}$	Remaining
$\square_{\text{ret}}$	Retarding (due to, e.g., friction)
$\square_R$	Right side
$\square_{\text{RoI}}$	Region of interest
$\square_{\text{safe}}$	Safety boundary
$\square_{\text{stop}}$	... required for stopping the vehicle
$\square_{\text{tang}}$	Tangential
$\square_{\text{thresh}}$	Threshold
$\square_{\text{tol}}$	Tolerance
$\square_t$	Transmitted
$\square_{\text{trust}}$	Trusted
$\square_{\text{veh}}$	Denominator for ego vehicle

Symbol	Description
$\square_{\text{votes}}$	Votes, e.g. for existing evidence about an object in a cycle
$\square_{\text{world}}$	... in the world coordinate system
$\square_x$	... in x-position or x-direction
$\square_y$	... in y-position or y-direction

List of Figures

1.1	Thesis structure	6
2.1	Coordinate Systems in this thesis	11
2.2	Radar data processing pipeline	18
2.3	Environmental effects leading to erroneous detections	19
2.4	Comparison of fixed and noise-adaptive threshold for peak detection	23
2.5	CA-CFAR concept overview	24
2.6	DBSCAN clustering concept showing core and border points	26
2.7	Kalman-filter-based tracking scheme	27
2.8	SOTIF-based approach to reducing both unknown and hazardous ODD areas	31
3.1	Architecture including area criticality estimation	44
3.2	Components of the state vector $\mathcal{S}$	45
3.3	Driving tube construction based on constructing a convex hull	47
3.4	Single-track model with turn radius	48
3.5	Vehicle trajectory in Frenet coordinates	51
3.6	Frenet coordinate calculation from discrete trajectory states $\mathcal{S}$	52
3.7	Longitudinal criticality $crit_{long}$ over the TTC	53
3.8	Lateral criticality $crit_{lat}$ over the lateral distance d_y	54
3.9	Simulated criticality in a VRU crossing scenario	55
3.10	Integration of VRU motion into the criticality metric	57
4.1	Relationship between recall and precision	63
4.2	Error types in CA-CFAR threshold detection.	65
4.3	Signal with two peaks exceeding the noise floor	66
4.4	Comparison of standard and criticality-aware peak detection	67
4.5	Exemplary criticality assessment of range and angle cells	68
4.6	Effect of reduction terms	71
4.7	Gaussian dimple for threshold reduction	73
4.8	Voting mechanism to approve or disprove detection candidates	74
4.9	Scheme for threshold reduction in critical spectrum cells	75
4.10	Camera projection scheme	76
4.11	Camera image of a street scenario showing vehicles and VRUs	78
4.12	Scenario structure of the ROD2021 dataset	79
4.13	Exemplary frame from the ROD2021 dataset	81
4.14	Detection metrics derived from GT annotations	82

4.15	Recall, precision, and F_1 -score over the threshold factor k_{thresh}	86
4.16	Low peak detection using Gaussian dimples	92
5.1	Comparison of conventional filtering and subsequent point cloud treatment .	98
5.2	Moving criticality region modeled by a multivariate normal distribution	99
5.3	Concept of successive point cloud treatment	100
5.4	ODD coverage over time during statistical recording.	102
5.5	Quantification of critical radar-point processing.	107
5.6	Chain of effect from function operation within the ODD to an incident	108
5.7	Method to determine radar point reachability	111
5.8	Exemplary radar point reachability	112
5.9	Framework for safety-aware radar point cloud processing	113
5.10	Detection metrics score over the criticality threshold.	115
5.11	Variation in criticality values among neighboring points	120
6.1	Schematic comparison of levels in late radar data processing	124
6.2	Distance criticality over evaluated point criticality	126
6.3	Mechanisms of track management in radar object tracking	128
6.4	Overview of the scenario used for the real-world experiment	130
6.5	Framework for radar data processing with criticality-aware extensions	133
6.6	SNR distribution for TP and FP detections	134
6.7	Perception flaws on the point level in the experiment	135
6.8	Example of denser point cloud after applying successive point cloud treatment	139
6.9	Measurement sequence of an approaching VRU exhibiting limited sensor performance.	142
7.1	Two radar sensors for vehicle motion estimation	149
7.2	Framework for radar-based motion estimation	150
7.3	Control loop for the ego vehicle's longitudinal velocity	151
7.4	Filter channel for removing physically non-reachable radar points	155
7.5	Radar-based motion estimation failure due to large dynamic objects in FoV	155
7.6	Comparison of successful and failed radar motion estimations	158
7.7	PI-controller behavior for different stimuli	160
7.8	Average motion estimation error over radar point azimuth angle	166
A.1	Chirp Patterns in FMCW	174
A.2	Causal chain in STPA	175
B.1	Subsequent peak search algorithm	180

List of Tables

2.1	Comparison of AVP Type 1 and AVP Type 2	14
2.2	Distribution of environmental information across the six-layer model.	15
2.3	Overview of state-of-the-art radar datasets	38
4.1	Operating figures of the TI AWR 1843 radar sensor.	80
4.2	Acceptance distance for various GT annotations	82
4.3	Variation of the experimental setup	84
4.4	Detection performance for different CA-CFAR parameter combinations.	85
4.5	Comparison of detection performance between baseline and subsequent radar peak search	87
4.6	Comparison of detection performance between baseline and image-based cell proposal	88
5.1	Condensed, relevant scenarios from multiple sources	105
5.2	Comparison of detection metrics for two criticality thresholds	116
5.3	Detection performance without safety-aware filtering.	116
5.4	Detection performance with safety-aware filtering	117
5.5	Comparison of two criticality assessment metrics	118
6.1	Technical data of the corner radar used in the experiment	131
6.2	Scenario parameters modified for the experiment.	132
6.3	Proposed safety-aware modifications for clustering and tracking	136
6.4	Detection performance with successive point-cloud treatment.	137
6.5	Detection performance for method A	138
6.6	Object detection performance of method B on the cluster level.	138
6.7	Object detection performance of methods B and C on the tracking level.	140
7.1	Vehicle model parameters for the controller	159
7.2	Results of the experiment testing the first hypothesis	162
7.3	Results of the controller action evaluation	162
7.4	Evaluation of motion estimation enhancement methods	163
A.1	First step of STPA	176
A.2	Third and fourth steps of STPA	177
B.1	Evaluation of three-dimensional spatial precision in object detection	179

References

Public References

- [Abd17] ABDULKHALEQ, A.: A system-theoretic safety engineering approach for software-intensive systems. en. In: (2017). URL: <http://nbn-resolving.de/urn:nbn:de:bsz:93-opus-ds-90667> (visited on 10/13/2025).
- [Ali+24] ALI, Z.; ELSAYED, M.; TIWARI, G.; AHMAD, M.; LE KERNEC, J.; HEIDARI, H.; GUPTA, S.: Impact of Receiver Thermal Noise and PLL RMS Jitter in Radar Measurements. In: *IEEE Transactions on Instrumentation and Measurement* 73 (2024), pp. 1–10. URL: <https://ieeexplore.ieee.org/abstract/document/10445509> (visited on 09/24/2025).
- [All+19] ALLAND, S.; STARK, W.; ALI, M.; HEGDE, M.: Interference in Automotive Radar Systems: Characteristics, Mitigation Techniques, and Current and Future Research. In: *IEEE Signal Processing Magazine* 36.5 (Sept. 2019), pp. 45–59. URL: <https://ieeexplore.ieee.org/abstract/document/8828037> (visited on 09/24/2025).
- [Ant+23] ANTONANTE, P; VEER, S; LEUNG, K; WENG, X; CARLONE, L; PAVONE, M: Task-aware risk estimation of perception failures for autonomous vehicle. In: *Robotics: Science and Systems (RSS)*. Robotics: Science and Systems (RSS). 2023.
- [ASA22] ASAM E.V.: ASAM OpenSCENARIO® 2.0.0. Tech. rep. 2022. URL: https://www.asam.net/static_downloads/public/asam-openscenario/2.0.0/welcome.html.
- [Bag13] BAGDADI, O.: Estimation of the severity of safety critical events. In: *Accident Analysis & Prevention* 50 (2013), pp. 167–174. URL: <https://www.sciencedirect.com/science/article/pii/S0001457512001406>.
- [Ban+17] BANZHAF, H.; NIENHÜSER, D.; KNOOP, S.; ZÖLLNER, J. M.: The future of parking: A survey on automated valet parking with an outlook on high density parking. In: June 2017, pp. 1827–1834. URL: <https://ieeexplore.ieee.org/abstract/document/7995971> (visited on 06/02/2025).
- [Bar+20] BARNES, D.; GADD, M.; MURCUTT, P.; NEWMAN, P.; POSNER, I.: The Oxford Radar RobotCar Dataset: A Radar Extension to the Oxford RobotCar Dataset. In: ISSN: 2577-087X. May 2020, pp. 6433–6438. URL: <https://ieeexplore.ieee.org/abstract/document/9196884> (visited on 06/10/2025).

- [Bas+15] BASSIL, K.; RILKOFF, H.; BELMONT, M.; BANASZEWSKA, A.; CAMPBELL, M.: *Pedestrian and Cyclist Safety in Toronto*. Tech. rep. 2015.
- [Ber+17] BERTHOLD, P.; MICHAELIS, M.; LUETTEL, T.; MEISSNER, D.; WUENSCH, H.-J.: Radar reflection characteristics of vehicles for contour and feature estimation. In: Oct. 2017, pp. 1–6. URL: <https://ieeexplore.ieee.org/abstract/document/8126352> (visited on 10/15/2025).
- [BF18] BARTON, D. C.; FIELDHOUSE, J. D.: “Steering Systems.” In: *Automotive Chassis Engineering*. Cham: Springer International Publishing, 2018, pp. 45–110. URL: https://doi.org/10.1007/978-3-319-72437-9_2.
- [BFR12] BARTSCH, A.; FITZEK, F.; RASCHOFER, R. H.: Pedestrian recognition using automotive radar sensors. In: *Advances in Radio Science* 10 (2012), pp. 45–55. URL: <https://ars.copernicus.org/articles/10/45/2012/>.
- [Bil23] BILIK, I.: Comparative Analysis of Radar and Lidar Technologies for Automotive Applications. In: *IEEE Intelligent Transportation Systems Magazine* 15.1 (Jan. 2023), pp. 244–269. URL: <https://ieeexplore.ieee.org/abstract/document/9760734> (visited on 12/21/2025).
- [BJT21] BIALER, O.; JONAS, A.; TIRER, T.: Super Resolution Wide Aperture Automotive Radar. In: *IEEE Sensors Journal* 21.16 (Aug. 2021), pp. 17846–17858. URL: <https://ieeexplore.ieee.org/abstract/document/9446133> (visited on 10/15/2025).
- [BM92] BESL, P. J.; MCKAY, N. D.: Method for registration of 3-D shapes. In: *Sensor Fusion IV: Control Paradigms and Data Structures*. Ed. by SCHENKER, P. S. Vol. 1611. International Society for Optics and Photonics. SPIE, 1992, pp. 586–606. URL: <https://doi.org/10.1117/12.57955>.
- [BMM18] BAGSCHIK, G.; MENZEL, T.; MAURER, M.: Ontology based Scene Creation for the Development of Automated Vehicles. In: ISSN: 1931-0587. June 2018, pp. 1813–1820. URL: <https://ieeexplore.ieee.org/abstract/document/8500632> (visited on 10/03/2025).
- [Bor+25] BORBA, T. D.; VACULÍN, O.; MARZBANI, H.; JAZAR, R. N.: Increasing Safety of Vulnerable Road Users in Scenarios With Occlusion: A Collaborative Approach for Smart Infrastructures and Automated Vehicles. In: *IEEE Access* 13 (2025), pp. 8851–8885. URL: <https://ieeexplore.ieee.org/abstract/document/10835073> (visited on 06/05/2025).
- [Bro91] BROOKS, R. A.: New approaches to robotics. In: *Science* 253.5025 (1991), pp. 1227–1232.
- [BS03] BIBER, P.; STRASSER, W.: The normal distributions transform: a new approach to laser scan matching. In: vol. 3. Oct. 2003, 2743–2748 vol.3. URL: <https://ieeexplore.ieee.org/abstract/document/1249285> (visited on 06/25/2025).
- [CA13] CAMACHO, E.; ALBA, C.: *Model Predictive Control*. Advanced Textbooks in Control and Signal Processing. Springer London, 2013.

- [Cae+20] CAESAR, H.; BANKITI, V.; LANG, A. H.; VORA, S.; LIONG, V. E.; XU, Q.; KRISHNAN, A.; PAN, Y.; BALDAN, G.; BEIJBOM, O.: nuScenes: A Multi-modal Dataset for Autonomous Driving. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.
- [Caf+11] CAFISO, S.; GARCIA, A. G.; CAVARRA, R.; ROJAS, M. R.: Crosswalk safety evaluation using a pedestrian risk index as traffic conflict measure. In: *3rd International Conference on Road safety and Simulation*. Vol. 15. 2011.
- [Car+21] CARMONA, J.; GUINDEL, C.; GARCIA, F.; ESCALERA, A. DE LA: eHMI: Review and Guidelines for Deployment on Autonomous Vehicles. en. In: *Sensors* 21.9 (Jan. 2021), p. 2912. URL: <https://www.mdpi.com/1424-8220/21/9/2912> (visited on 09/24/2025).
- [ĆAS13] ĆIROVIĆ, V.; ALEKSENDRIĆ, D.; SMILJANIĆ, D.: Longitudinal wheel slip control using dynamic neural networks. In: *Mechatronics* 23.1 (Feb. 2013), pp. 135–146. URL: <https://www.sciencedirect.com/science/article/pii/S0957415812001778> (visited on 12/21/2025).
- [CC14] CHEN, M.; CHEN, C.-C.: RCS Patterns of Pedestrians at 76-77 GHz. In: *IEEE Antennas and Propagation Magazine* 56.4 (Aug. 2014), pp. 252–263. URL: <https://ieeexplore.ieee.org/abstract/document/6931711> (visited on 04/29/2025).
- [Cha+25] CHAN, P. H.; SHAHBEIGI ROUDPOSHTI, S.; YE, X.; DONZELLA, V.: A Noise Analysis of 4D RADAR: Robust Sensing for Automotive? In: *IEEE Sensors Journal* 25.10 (May 2025), pp. 18291–18301. URL: <https://ieeexplore.ieee.org/abstract/document/10955111> (visited on 05/29/2025).
- [Chu+23] CHU, J.; ZHAO, T.; JIAO, J.; YUAN, Y.; JING, Y.: SOTIF-Oriented Perception Evaluation Method for Forward Obstacle Detection of Autonomous Vehicles. In: *IEEE Systems Journal* 17.2 (June 2023), pp. 2319–2330. URL: <https://ieeexplore.ieee.org/abstract/document/10024520> (visited on 01/03/2024).
- [CM23] CECCARELLI, A.; MONTECCHI, L.: Evaluating Object (Mis)Detection From a Safety and Reliability Perspective: Discussion and Measures. In: *IEEE Access* 11 (2023), pp. 44952–44963. URL: <https://ieeexplore.ieee.org/abstract/document/10115418> (visited on 06/05/2025).
- [CM24] CECCARELLI, A.; MONTECCHI, L.: Better and safer autonomous driving with predicted object relevance. In: ISSN: 2994-810X. Oct. 2024, pp. 329–336. URL: <https://ieeexplore.ieee.org/abstract/document/10771202> (visited on 06/05/2025).
- [CMS13] CAMPELLO, R. J. G. B.; MOULAVI, D.; SANDER, J.: Density-Based Clustering Based on Hierarchical Density Estimates. In: *Advances in Knowledge Discovery and Data Mining*. Ed. by PEI, J.; TSENG, V. S.; CAO, L.; MOTODA, H.; XU, G. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 160–172.

- [Cor+16] CORDTS, M.; OMRAN, M.; RAMOS, S.; REHFELD, T.; ENZWEILER, M.; BENENSON, R.; FRANKE, U.; ROTH, S.; SCHIELE, B.: The Cityscapes Dataset for Semantic Urban Scene Understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- [Dec+24] DECOURT, C.; VANRULLEN, R.; SALLE, D.; OBERLIN, T.: A Recurrent CNN for Online Object Detection on Raw Radar Frames. In: *IEEE Transactions on Intelligent Transportation Systems* 25.10 (Oct. 2024), pp. 13432–13441. URL: <https://ieeexplore.ieee.org/abstract/document/10547638> (visited on 06/13/2025).
- [Dim+20] DIMITRIEVSKI, M.; SHOPOVSKA, I.; HAMME, D. V.; VEELAERT, P.; PHILIPS, W.: Weakly Supervised Deep Learning Method for Vulnerable Road User Detection in FMCW Radar. In: Sept. 2020, pp. 1–8. URL: <https://ieeexplore.ieee.org/document/9294399> (visited on 05/05/2025).
- [Din+23] DING, F.; PALFFY, A.; GAVRILA, D. M.; LU, C. X.: Hidden Gems: 4D Radar Scene Flow Learning Using Cross-Modal Supervision. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023, pp. 9340–9349.
- [DKP17] DELLAERT, F.; KAESSE, M.; PUBLISHERS, N.: *Factor Graphs for Robot Perception*. Foundations and trends in robotics. Now Publishers, 2017.
- [Don82] DONGES, E.: Aspekte der aktiven Sicherheit bei der Führung von Personenkraftwagen. In: *Automob-Ind* 27.2 (1982).
- [Dor+20] DORFF, S. V.; BÖDDEKER, B.; KNEISSL, M.; FRÄNZLE, M.: A Fail-safe Architecture for Automated Driving. In: ISSN: 1558-1101. Mar. 2020, pp. 828–833. URL: <https://ieeexplore.ieee.org/abstract/document/9116283> (visited on 06/02/2025).
- [Dre+23] DREISSIG, M.; SCHEUBLE, D.; PIEWAK, F.; BOEDECKER, J.: Survey on LiDAR Perception in Adverse Weather Conditions. In: ISSN: 2642-7214. June 2023, pp. 1–8. URL: <https://ieeexplore.ieee.org/abstract/document/10186539> (visited on 06/24/2025).
- [Eng+21] ENGELS, F.; HEIDENREICH, P.; WINTERMANTEL, M.; STÄCKER, L.; AL KADI, M.; ZOUBIR, A. M.: Automotive Radar Signal Processing: Research Directions and Practical Challenges. In: *IEEE Journal of Selected Topics in Signal Processing* 15.4 (June 2021), pp. 865–878. URL: <https://ieeexplore.ieee.org/abstract/document/9369027> (visited on 05/29/2025).
- [ESE06] EICHLER, S.; SCHROTH, C.; EBERSPÄCHER, J.: *Car-to-Car Communication*. de. 2006. URL: <https://www.alexandria.unisg.ch/handle/20.500.14171/82385>.
- [Est+96] ESTER, M.; KRIEGEL, H.-P.; SANDER, J.; XU, X., et al.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *kdd*. Vol. 96. 34. 1996, pp. 226–231.

- [ETS20] ETSI: *System Reference document (SRdoc);Transmission characteristics;Technical characteristics for radiodetermination equipmentfor ground based vehicular applicationswithin the frequency range 77 GHz to 81 GHz*. Tech. rep. 2020. URL: https://www.etsi.org/deliver/etsi_tr/103500_103599/103593/01.01.01_60/tr_103593v010101p.pdf.
- [Fan+24a] FAN, L.; WANG, J.; CHANG, Y.; LI, Y.; WANG, Y.; CAO, D.: 4D mmWave Radar for Autonomos Driving Perception: A Comprehensive Survey. In: *IEEE Transactions on Intelligent Vehicles* 9.4 (Apr. 2024), pp. 4606–4620. URL: <https://ieeexplore.ieee.org/abstract/document/10477463> (visited on 06/08/2025).
- [Fan+24b] FANG, Y.; WANG, W.; YANG, C.; ZHANG, Y.; CHEN, Z.: Brake-by-wire architecture design and analysis in accordance with functional safety standard. EN. In: *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering* 238.12 (Oct. 2024), pp. 3794–3808. URL: <https://doi.org/10.1177/09544070231185192> (visited on 06/02/2025).
- [FB81] FISCHLER, M. A.; BOLLES, R. C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In: *Commun. ACM* 24.6 (June 1981), pp. 381–395. URL: <https://dl.acm.org/doi/10.1145/358669.358692> (visited on 06/26/2025).
- [Fen+17] FENG, F.; BAO, S.; SAYER, J. R.; FLANNAGAN, C.; MANSER, M.; WUNDERLICH, R.: Can vehicle longitudinal jerk be used to identify aggressive drivers? An examination using naturalistic driving data. In: *Accident Analysis & Prevention* 104 (July 2017), pp. 125–136. URL: <https://www.sciencedirect.com/science/article/pii/S0001457517301409> (visited on 06/27/2025).
- [FPC24] FENT, F.; PALFFY, A.; CAESAR, H.: DPFT: Dual Perspective Fusion Transformer for Camera-Radar-based Object Detection. In: *IEEE Transactions on Intelligent Vehicles* (2024), pp. 1–11. URL: <https://ieeexplore.ieee.org/abstract/document/10769556> (visited on 06/08/2025).
- [Gam20] GAMBA, J.: *Radar Signal Processing for Autonomous Driving*. Signals and communication technology. Springer, 2020.
- [Gao+21] GAO, X.; XING, G.; ROY, S.; LIU, H.: RAMP-CNN: A Novel Neural Network for Enhanced Automotive Radar Object Recognition. In: *IEEE Sensors Journal* 21.4 (Feb. 2021), pp. 5119–5132. URL: <https://ieeexplore.ieee.org/document/9249018> (visited on 06/09/2025).
- [GCT25] GYLLENHAMMAR, M.; CAMPOS, G. RODRIGUES DE; TÖRNGREN, M.: The Road to Safe Automated Driving Systems: A Review of Methods Providing Safety Evidence. In: *IEEE Transactions on Intelligent Transportation Systems* 26.4 (Apr. 2025), pp. 4315–4345. URL: <https://ieeexplore.ieee.org/document/10858599> (visited on 06/16/2025).

- [Gee21] GEELY AUTO INTERNATIONAL CORPORATION: *Geely Launches Newest High-End Flagship SUV – Xingyue L*. 2021. URL: <https://www.geely.com/en/news/2021/geely-launches-suv-xingyuel> (visited on 12/21/2025).
- [Ger23] GERMAN ASSOCIATION OF THE AUTOMOTIVE INDUSTRY (VDA): *Position — Automated Valet Parking Systems — Requirements for automated valet parking systems*. Tech. rep. 2023. URL: https://www.vda.de/dam/jcr:116de148-db90-4b5c-b723-f9efb0276903/Automated_Valet_Parking.pdf.
- [GGW17] GUTJAHR, B.; GRÖLL, L.; WERLING, M.: Lateral Vehicle Trajectory Optimization Using Constrained Linear Time-Varying MPC. In: *IEEE Transactions on Intelligent Transportation Systems* 18.6 (June 2017), pp. 1586–1595. URL: <https://ieeexplore.ieee.org/abstract/document/7593251> (visited on 06/27/2025).
- [Gha+24] GHARI, B.; TOURANI, A.; SHAHBAHRAMI, A.; GAYDADJIEV, G.: Pedestrian detection in low-light conditions: A comprehensive survey. In: *Image and Vision Computing* 148 (Aug. 2024), p. 105106. URL: <https://www.sciencedirect.com/science/article/pii/S0262885624002105> (visited on 06/24/2025).
- [Gho+22] GHORAI, P.; ESKANDARIAN, A.; KIM, Y.-K.; MEHR, G.: State Estimation and Motion Prediction of Vehicles and Vulnerable Road Users for Cooperative Autonomous Driving: A Survey. In: *IEEE Transactions on Intelligent Transportation Systems* 23.10 (Oct. 2022), pp. 16983–17002. URL: <https://ieeexplore.ieee.org/abstract/document/9756845> (visited on 06/05/2025).
- [Gir+14] GIRSHICK, R.; DONAHUE, J.; DARRELL, T.; MALIK, J.: Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.
- [GK88] GANDHI, P.; KASSAM, S.: Analysis of CFAR processors in nonhomogeneous background. In: *IEEE Transactions on Aerospace and Electronic Systems* 24.4 (July 1988), pp. 427–445. URL: <https://ieeexplore.ieee.org/abstract/document/7185> (visited on 04/28/2025).
- [GO 12] GO BOULDER - CITY OF BOULDER TRANSPORTATION DIVISION: *Safe Streets Boulder - Striving to Make Boulder Streets Even Safer*. Tech. rep. 2012. URL: <https://documents.bouldercolorado.gov/WebLink/edoc/122645/2C.pdf?dbid=0&repo=LF8PROD2>.
- [Gro+08] GROVER, C; KNIGHT, I; OKORO, F; SIMMONS, I; COUPER, G; MASSIE, P; SMITH, B: *Automated Emergency Brake Systems: Technical requirements, costs and benefits*. Tech. rep. TRL Limited, 2008.

- [Gu+25] GU, Z.; MA, J.; HUANG, Y.; WEI, H.; CHEN, Z.; ZHANG, H.; HONG, W.: HGSFusion: Radar-Camera Fusion with Hybrid Generation and Synchronization for 3D Object Detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 39.3 (2025), pp. 3185–3193. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/32328>.
- [Guo+24] GUO, A.; ZHOU, Y.; TIAN, H.; FANG, C.; SUN, Y.; SUN, W.; GAO, X.; LUU, A. T.; LIU, Y.; CHEN, Z.: SoVAR: Build Generalizable Scenarios from Accident Reports for Autonomous Driving Testing. In: *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering. ASE '24*. Sacramento, CA, USA: Association for Computing Machinery, 2024, 268–280. URL: <https://doi.org/10.1145/3691620.3695037>.
- [Hab+08] HABIB, M. A.; BARKAT, M.; AISSA, B.; DENIDNI, T.: Ca-cfar detection performance of radar targets embedded in “non centered chi-2 gamma” clutter. In: *Progress In Electromagnetics Research* 88 (2008), pp. 135–148.
- [HB25] HAITMAN, Y.; BIALER, O.: DoppDrive: Doppler-Driven Temporal Aggregation for Improved Radar Object Detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2025, pp. 26085–26094.
- [HOM12] HEISSMEYER, S.; OVERMEYER, L.; MÜLLER, A.: Indoor positioning of vehicles using an active optical infrastructure. In: Nov. 2012, pp. 1–8. URL: <https://ieeexplore.ieee.org/abstract/document/6418914> (visited on 06/25/2025).
- [HP16] HARRER, M.; PFEFFER, P.: *Steering Handbook*. Springer International Publishing, 2016.
- [HSE22] HOSS, M.; SCHOLTES, M.; ECKSTEIN, L.: A Review of Testing Object-Based Environment Perception for Safe Automated Driving. en. In: *Automotive Innovation* 5.3 (Aug. 2022), pp. 223–250. URL: <https://doi.org/10.1007/s42154-021-00172-y> (visited on 01/03/2024).
- [Hua+17] HUANG, X.; KWIATKOWSKA, M.; WANG, S.; WU, M.: Safety Verification of Deep Neural Networks. en. In: *Computer Aided Verification*. Ed. by MAJUMDAR, R.; KUNČAK, V. Cham: Springer International Publishing, 2017, pp. 3–29.
- [Hus+19] HUSSAIN, Q.; FENG, H.; GRZEBIETA, R.; BRIJS, T.; OLIVIER, J.: The relationship between impact speed and the probability of pedestrian fatality during a vehicle-pedestrian crash: A systematic review and meta-analysis. In: *Accident Analysis & Prevention* 129 (Aug. 2019), pp. 241–249. URL: <https://www.sciencedirect.com/science/article/pii/S0001457519301058> (visited on 06/03/2025).
- [Hyd87] HYDÉN, C.: The development of a method for traffic safety evaluation: The Swedish Traffic Conflicts Technique. In: *Bulletin Lund Institute of Technology, Department 70* (1987).

- [IHFZ21] ISELE, S. T.; HAAS-FICKINGER, F.; ZÖLLNER, J. M.: SERALOC: SLAM on semantically annotated radar point-clouds. In: Sept. 2021, pp. 2917–2924. URL: <https://ieeexplore.ieee.org/abstract/document/9564693> (visited on 06/10/2025).
- [Int11] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *ISO 8855:2011 - Road vehicles - Vehicle dynamics and road-holding ability*. Tech. rep. 2011. URL: <https://www.iso.org/standard/77490.html>.
- [Int18] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *ISO 26262:2018 - Road vehicles - Functional safety*. Tech. rep. 2018. URL: <https://www.iso.org/standard/68383.html>.
- [Int22] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *ISO 21448:2022 - Road vehicles - Safety of the intended functionality*. Tech. rep. 2022. URL: <https://www.iso.org/standard/77490.html>.
- [Int23] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *ISO 23374-1:2023 - Intelligent transport systems — Automated valet parking systems (AVPS)*. Tech. rep. 2023. URL: <https://www.iso.org/standard/78420.html>.
- [Int97] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *ISO 2631-1:1997 - Mechanical vibration and shock — Evaluation of human exposure to whole-body vibration*. Tech. rep. 1997. URL: <https://www.iso.org/standard/7612.html>.
- [Jan05] JANSSON, J: Collision Avoidance Theory with Application to Automotive Collision Mitigation Dissertation No. 950. In: *Linköping Studies in Science and Technology, SE-581 83* (2005).
- [Jan18] JANKIRAMAN, M.: *FMCW Radar Design*. Artech House radar library. Artech House, 2018.
- [Jia+23] JIANG, B.; CHEN, S.; XU, Q.; LIAO, B.; CHEN, J.; ZHOU, H.; ZHANG, Q.; LIU, W.; HUANG, C.; WANG, X.: VAD: Vectorized Scene Representation for Efficient Autonomous Driving. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2023, pp. 8340–8350.
- [Jun+18] JUNIETZ, P.; BONAKDAR, F.; KLAMANN, B.; WINNER, H.: Criticality Metric for the Safety Validation of Automated Driving using Model Predictive Trajectory Optimization. In: ISSN: 2153-0017. Nov. 2018, pp. 60–65. URL: <https://ieeexplore.ieee.org/abstract/document/8569326> (visited on 06/04/2025).
- [Kal+21] KALAITZAKIS, M.; CAIN, B.; CARROLL, S.; AMBROSI, A.; WHITEHEAD, C.; VITZILAIOS, N.: Fiducial Markers for Pose Estimation. en. In: *Journal of Intelligent & Robotic Systems* 101.4 (Mar. 2021), p. 71. URL: <https://doi.org/10.1007/s10846-020-01307-9> (visited on 06/25/2025).
- [Kal60] KALMAN, R. E.: A new approach to linear filtering and prediction problems. In: (1960).

- [KC20] KAO, M.; CHANG, C.: *Understanding Electromagnetic Waves*. Springer International Publishing, 2020.
- [Kel+13] KELLNER, D.; BARJENBRUCH, M.; KLAPPSTEIN, J.; DICKMANN, J.; DIETMAYER, K.: Instantaneous ego-motion estimation using Doppler radar. In: ISSN: 2153-0017. Oct. 2013, pp. 869–874. URL: <https://ieeexplore.ieee.org/abstract/document/6728341> (visited on 06/25/2025).
- [Kel+14] KELLNER, D.; BARJENBRUCH, M.; KLAPPSTEIN, J.; DICKMANN, J.; DIETMAYER, K.: Instantaneous ego-motion estimation using multiple Doppler radars. In: ISSN: 1050-4729. May 2014, pp. 1592–1597. URL: <https://ieeexplore.ieee.org/abstract/document/6907064> (visited on 06/26/2025).
- [Kim+23] KIM, Y.; KIM, S.; CHOI, J. W.; KUM, D.: CRAFT: Camera-Radar 3D Object Detection with Spatio-Contextual Fusion Transformer. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 37.1 (2023), pp. 1160–1168. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/25198>.
- [Kin+00] KINSLER, L. E.; FREY, A. R.; COPPENS, A. B.; SANDERS, J. V.: *Fundamentals of acoustics*. John Wiley & sons, 2000.
- [KLD21] KHAITAN, S.; LIN, Q.; DOLAN, J. M.: Safe Planning and Control Under Uncertainty for Self-Driving. In: *IEEE Transactions on Vehicular Technology* 70.10 (Oct. 2021), pp. 9826–9837. URL: <https://ieeexplore.ieee.org/abstract/document/9525177> (visited on 06/02/2025).
- [Kne+21] KNEISSL, M.; MADHUSUDHANAN, A. K.; MOLIN, A.; ESEN, H.; HIRCHE, S.: A Multi-Vehicle Control Framework With Application to Automated Valet Parking. In: *IEEE Transactions on Intelligent Transportation Systems* 22.9 (Sept. 2021), pp. 5697–5707. URL: <https://ieeexplore.ieee.org/abstract/document/9091894> (visited on 06/16/2025).
- [KR96] KONG, T.; ROSENFELD, A.: *Topological Algorithms for Digital Image Processing*. Machine Intelligence and Pattern Recognition. North Holland, 1996.
- [Kra22] KRAFTFAHRT-BUNDESAMT: *Technical catalogue of requirements for the fully automated driving function „Automated Valet Parking (AVP)“*. Tech. rep. 2022. URL: https://www.kba.de/EN/Themen_en/Typgenehmigung_en/Zum_Herunterladen_en/autonomes_automatisiertes_Fahren_en/dl_anforderungskatalog_AVP_en.pdf.
- [Kus+24] KUSANO, K. D.; JOHN M., S.; YIN-HSIU, C.; TIMOTHY L., M.; RUOSHU, C.; TILIA, G.; VICTOR, T.: Comparison of Waymo rider-only crash data to human benchmarks at 7.1 million miles. In: *Traffic Injury Prevention* 25.sup1 (Nov. 2024), S66–S77. URL: <https://doi.org/10.1080/15389588.2024.2380786> (visited on 06/09/2025).

- [LBR22] LUONG, D.; BALAJI, B.; RAJAN, S.: Performance Prediction for Coherent Noise Radars Using the Correlation Coefficient. In: *IEEE Access* 10 (2022), pp. 8627–8633. URL: <https://ieeexplore.ieee.org/document/9686579> (visited on 06/05/2025).
- [Lev16] LEVESON, N. G.: *Engineering a Safer World: Systems Thinking Applied to Safety*. English. The MIT Press, 2016. URL: <https://library.oapen.org/handle/20.500.12657/26043> (visited on 10/13/2025).
- [Li+20] LI, Y.; LIU, Y.; WANG, Y.; LIN, Y.; SHEN, W.: The Millimeter-Wave Radar SLAM Assisted by the RCS Feature of the Target and IMU. en. In: *Sensors* 20.18 (Jan. 2020), p. 5421. URL: <https://www.mdpi.com/1424-8220/20/18/5421> (visited on 06/05/2025).
- [Lin+14] LIN, T.-Y.; MAIRE, M.; BELONGIE, S.; HAYS, J.; PERONA, P.; RAMANAN, D.; DOLLÁR, P.; ZITNICK, C. L.: Microsoft COCO: Common Objects in Context. In: *Computer Vision – ECCV 2014*. Ed. by FLEET, D.; PAJDLA, T.; SCHIELE, B.; TUYTELAARS, T. Cham: Springer International Publishing, 2014, pp. 740–755.
- [Liu+16] LIU, W.; ANGUELOV, D.; ERHAN, D.; SZEGEDY, C.; REED, S.; FU, C.-Y.; BERG, A. C.: SSD: Single Shot MultiBox Detector. en. In: *Computer Vision – ECCV 2016*. Ed. by LEIBE, B.; MATAS, J.; SEBE, N.; WELLING, M. Cham: Springer International Publishing, 2016, pp. 21–37.
- [LLK18] LIM, S.; LEE, S.; KIM, S.-C.: Clustering of Detected Targets Using DBSCAN in Automotive Radar Systems. In: ISSN: 2155-5753. June 2018, pp. 1–7. URL: <https://ieeexplore.ieee.org/document/8448228> (visited on 07/14/2025).
- [LS22] LI, Y.; SHANG, X.: Multipath Ghost Target Identification for Automotive MIMO Radar. In: ISSN: 2577-2465. Sept. 2022, pp. 1–5. URL: <https://ieeexplore.ieee.org/abstract/document/10012904> (visited on 06/07/2025).
- [LSA23] LIU, S.; SHA, L.; ABDELZAHER, T.: “Criticality-Based Data Segmentation and Resource Allocation in Machine Inference Pipelines.” en. In: ed. by SRIVATSA, M.; ABDELZAHER, T.; HE, T. Cham: Springer International Publishing, 2023, pp. 335–352. URL: https://doi.org/10.1007/978-3-031-40787-1_11 (visited on 09/09/2024).
- [Lud13] LUDLOFF, A.: *Handbuch Radar und Radarsignalverarbeitung*. Vieweg+Teubner Verlag, 2013.
- [Lys+22] LYSENKO, M.; GLADISCH, C.; HEINZEMANN, C.; WOHRLE, M.; TRIEBEL, R.: Towards Safety-Aware Pedestrian Detection in Autonomous Systems. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. ISSN: 2153-0866. Oct. 2022, pp. 293–300. URL: <https://ieeexplore.ieee.org/abstract/document/9981309> (visited on 09/11/2024).

- [Mac+18] MACHARDY, Z.; KHAN, A.; OBANA, K.; IWASHINA, S.: V2X Access Technologies: Regulation, Research, and Remaining Challenges. In: *IEEE Communications Surveys & Tutorials* 20.3 (2018), pp. 1858–1877. URL: <https://ieeexplore.ieee.org/abstract/document/8300313> (visited on 09/24/2025).
- [Maj+19] MAJOR, B.; FONTIJNE, D.; ANSARI, A.; TEJA SUKHAVASI, R.; GOWAIKAR, R.; HAMILTON, M.; LEE, S.; GRZECHNIK, S.; SUBRAMANIAN, S.: Vehicle Detection With Automotive Radar Using Deep Learning on Range-Azimuth-Doppler Tensors. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*. 2019.
- [Mar22] MARKEL, M.: *Radar for Fully Autonomous Driving*. Artech House, 2022.
- [Mei+24] MEISSNER, D.; JORDAN, R.; QUACH, M.; DOMOKOS, C.: Object Relevance for Radar Object Detection on Autonomous Driving Scenarios. In: Oct. 2024, pp. 1–6. URL: <https://ieeexplore.ieee.org/abstract/document/10993665> (visited on 06/14/2025).
- [Mer24] MERCEDES-BENZ GROUP: *PARKTRONIC - Improved automated parking*. 2024. URL: <https://group.mercedes-benz.com/innovations/product-innovation/autonomous-driving/automated-parking.html> (visited on 12/21/2025).
- [Mil+19] MILIOTO, A.; VIZZO, I.; BEHLEY, J.; STACHNISS, C.: RangeNet ++: Fast and Accurate LiDAR Semantic Segmentation. In: ISSN: 2153-0866. Nov. 2019, pp. 4213–4220. URL: <https://ieeexplore.ieee.org/abstract/document/8967762> (visited on 06/09/2025).
- [MKT21] MEYER, M.; KUSCHK, G.; TOMFORDE, S.: Graph Convolutional Networks for 3D Object Detection on Radar Data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*. 2021, pp. 3060–3069.
- [MNK12] MATSUNAMI, I.; NAKAMURA, R.; KAJIWARA, A.: RCS measurements for vehicles and pedestrian at 26 and 79GHz. In: Dec. 2012, pp. 1–4. URL: <https://ieeexplore.ieee.org/abstract/document/6508004> (visited on 04/29/2025).
- [Mus+24] MUSIAT, A.; REICHARDT, L.; SCHULZE, M.; WASENMÜLLER, O.: RadarPillars: Efficient Object Detection from 4D Radar Point Clouds. In: ISSN: 2153-0017. Sept. 2024, pp. 1656–1663. URL: <https://ieeexplore.ieee.org/abstract/document/10919920> (visited on 12/21/2025).
- [Nat23] NATIONAL CENTER FOR STATISTICS AND ANALYSIS: Pedestrians: 2021 data. In: *Traffic Safety Facts. ReportNo. DOT HS 813 458* (2023). URL: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813458.pdf>.
- [Nav25] NAVTECH RADAR: *CTS350 Technical Data Sheet*. Tech. rep. 2025. URL: <https://navtechradar.com/wp-content/uploads/2024/12/CTS350-DATA-SHEET-V2.2.pdf>.

- [Nee19] NEES, M. A.: Safer than the average human driver (who is less safe than me)? Examining a popular safety benchmark for self-driving cars. In: *Journal of Safety Research* 69 (June 2019), pp. 61–68. URL: <https://www.sciencedirect.com/science/article/pii/S0022437518304511> (visited on 06/09/2025).
- [Nor+20] NOROUZIAN, F.; MARCHETTI, E.; GASHINOVA, M.; HOARE, E.; CONSTANTINOU, C.; GARDNER, P.; CHERNIAKOV, M.: Rain Attenuation at Millimeter Wave and Low-THz Frequencies. In: *IEEE Transactions on Antennas and Propagation* 68.1 (Jan. 2020), pp. 421–431. URL: <https://ieeexplore.ieee.org/abstract/document/8826596> (visited on 06/24/2025).
- [NP33] NEYMAN, J.; PEARSON, E. S.: The testing of statistical hypotheses in relation to probabilities a priori. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 29.4 (1933), 492–510.
- [Ort+20] ORT, T.; MURTHY, K.; BANERJEE, R.; GOTTIPATI, S. K.; BHATT, D.; GILITSCHENSKI, I.; PAULL, L.; RUS, D.: MapLite: Autonomous Intersection Navigation Without a Detailed Prior Map. In: *IEEE Robotics and Automation Letters* 5.2 (Apr. 2020), pp. 556–563. URL: <https://ieeexplore.ieee.org/abstract/document/8936918> (visited on 05/27/2025).
- [Oua+21] OUAKNINE, A.; NEWSON, A.; REBUT, J.; TUPIN, F.; PÉREZ, P.: CARRADA Dataset: Camera and Automotive Radar with Range- Angle- Doppler Annotations. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. 2021, pp. 5068–5075.
- [Pal+20] PALFFY, A.; DONG, J.; KOOIJ, J. F. P.; GAVRILA, D. M.: CNN Based Road User Detection Using the 3D Radar Cube. In: *IEEE Robotics and Automation Letters* 5.2 (Apr. 2020), pp. 1263–1270. URL: <https://ieeexplore.ieee.org/abstract/document/8962258> (visited on 06/09/2025).
- [Pal+22] PALFFY, A.; POOL, E.; BARATAM, S.; KOOIJ, J. F. P.; GAVRILA, D. M.: Multi-Class Road User Detection With 3+1D Radar in the View-of-Delft Dataset. In: *IEEE Robotics and Automation Letters* 7.2 (Apr. 2022), pp. 4961–4968. URL: <https://ieeexplore.ieee.org/abstract/document/9699098> (visited on 06/10/2025).
- [Pau+21] PAULY, A.; GRUBER, S.; SENDLER, J.; VOLK, H.; SEETHALER, L.; SATTLER, M.; VIERTLBÖCK, M.; SAGAN, E.; SCHWARZ, M.; UNTERSTRASSER, T.: “Fahrwerk.” In: *Vieweg Handbuch Kraftfahrzeugtechnik*. Ed. by PISCHINGER, S.; SEIFFERT, U. Wiesbaden: Springer Fachmedien Wiesbaden, 2021, pp. 261–459. URL: https://doi.org/10.1007/978-3-658-25557-2_5.
- [Phi+22] PHILIPP, R.; REHBEIN, J.; GRÜN, F.; HARTJEN, L.; ZHU, Z.; SCHULDT, F.; HOWAR, F.: Systematization of Relevant Road Users for the Evaluation of Autonomous Vehicle Perception. In: *2022 IEEE International Systems Conference (SysCon)*. ISSN: 2472-9647. Apr. 2022, pp. 1–8. URL: <https://ieeexplore.ieee.org/abstract/document/9773877> (visited on 09/11/2024).

- [PKF20] PHILION, J.; KAR, A.; FIDLER, S.: Learning to Evaluate Perception Models Using Planner-Centric Metrics. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.
- [PNS20] PADILLA, R.; NETTO, S. L.; SILVA, E. A. B. DA: A Survey on Performance Metrics for Object-Detection Algorithms. In: ISSN: 2157-8702. July 2020, pp. 237–242. URL: <https://ieeexplore.ieee.org/abstract/document/9145130> (visited on 06/12/2025).
- [Ras+17] RASEKHIPOUR, Y.; KHAJEPOUR, A.; CHEN, S.-K.; LITKOUHI, B.: A Potential Field-Based Model Predictive Path-Planning Controller for Autonomous Road Vehicles. In: *IEEE Transactions on Intelligent Transportation Systems* 18.5 (May 2017), pp. 1255–1267. URL: <https://ieeexplore.ieee.org/abstract/document/7576661> (visited on 06/27/2025).
- [Ras+22] RASKIN, J.; CHATTERJEE, K.; DOYEN, L.; MAJUMDAR, R.: *Principles of Systems Design: Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*. Lecture Notes in Computer Science. Springer Nature Switzerland, 2022. URL: <https://books.google.de/books?id=QIekEAAQBAJ>.
- [RC24] RICCARDI, G.; CIFRA, B. A.: *Multidimensional Differential and Integral Calculus*. 1st ed. 2024. Cham : Springer Nature Switzerland, 2024. URL: <https://doi.org/10.1007/978-3-031-70326-3>.
- [Reb+22] REBUT, J.; OUAKNINE, A.; MALIK, W.; PÉREZ, P.: Raw High-Definition Radar for Multi-Task Learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 17021–17030.
- [Ric22] RICHARDS, M. A.: *Fundamentals of Radar Signal Processing, 3rd Edition*. McGraw Hill, 2022. URL: <https://www.accessengineeringlibrary.com/content/book/9781260468717>.
- [Rie+20] RIEDMAIER, S.; PONN, T.; LUDWIG, D.; SCHICK, B.; DIERMEYER, F.: Survey on Scenario-Based Safety Assessment of Automated Vehicles. In: *IEEE Access* 8 (2020), pp. 87456–87477. URL: <https://ieeexplore.ieee.org/abstract/document/9090897> (visited on 10/03/2025).
- [Rob25] ROBERT BOSCH GMBH: *Integrated power brake*. Tech. rep. 2025.
- [SAC19] SALAY, R.; ANGUS, M.; CZARNECKI, K.: A Safety Analysis Method for Perceptual Components in Automated Driving. In: ISSN: 2332-6549. Oct. 2019, pp. 24–34. URL: <https://ieeexplore.ieee.org/abstract/document/8987559> (visited on 06/07/2025).
- [SAE21] SAE INTERNATIONAL.: *J3016 - Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Tech. rep. 2021. URL: https://www.sae.org/standards/content/j3016_202104/s.
- [SAE22] SAE INTERNATIONAL.: *J670 - Vehicle Dynamics Terminology*. Tech. rep. 2022. URL: https://www.sae.org/standards/content/j670_202206/.

- [Sch+20a] SCHEINER, N.; SCHUMANN, O.; KRAUS, F.; APPENRODT, N.; DICKMANN, J.; SICK, B.: Off-the-shelf sensor vs. experimental radar - How much resolution is necessary in automotive radar classification? In: July 2020, pp. 1–8. URL: <https://ieeexplore.ieee.org/abstract/document/9190338> (visited on 06/24/2025).
- [Sch+20b] SCHÜLLER, H.; PLESKER, M.; BÄRWOLFF, M.; SCHREIBER, M.: Unfallrisiko Parken für schwächere Verkehrsteilnehmer. In: *Gesamtverband der Deutschen Versicherungswirtschaft e. V.* (2020). URL: <https://www.udv.de/resource/blob/79194/7d478d2e7ff20210dcd1d746a619a1d3/66-unfallrisiko-parken-fuer-schwaechere-verkehrsteilnehmer-data.pdf>.
- [Sch+21a] SCHOLTES, M.; WESTHOFEN, L.; TURNER, L. R.; LOTTO, K.; SCHULDES, M.; WEBER, H.; WAGENER, N.; NEUROHR, C.; BOLLMANN, M. H.; KÖRTKE, F.; HILLER, J.; HOSS, M.; BOCK, J.; ECKSTEIN, L.: 6-Layer Model for a Structured Description and Categorization of Urban Traffic and Environment. In: *IEEE Access* 9 (2021), pp. 59131–59147. URL: <https://ieeexplore.ieee.org/abstract/document/9400833> (visited on 05/27/2025).
- [Sch+21b] SCHUMANN, O.; HAHN, M.; SCHEINER, N.; WEISHAUP, F.; TILLY, J. F.; DICKMANN, J.; WÖHLER, C.: RadarScenes: A Real-World Radar Point Cloud Data Set for Automotive Applications. In: Nov. 2021, pp. 1–8. URL: <https://ieeexplore.ieee.org/abstract/document/9627037> (visited on 05/20/2025).
- [SDW19] SCHÖNEMANN, V.; DUSCHEK, M.; WINNER, H.: Maneuver-based Adaptive Safety Zone for Infrastructure-Supported Automated Valet Parking. In: 2019.
- [She+21] SHEENY, M.; DE PELLEGRIN, E.; MUKHERJEE, S.; AHRABIAN, A.; WANG, S.; WALLACE, A.: RADIATE: A Radar Dataset for Automotive Perception in Bad Weather. In: ISSN: 2577-087X. May 2021, pp. 1–7. URL: <https://ieeexplore.ieee.org/abstract/document/9562089> (visited on 06/10/2025).
- [Shi17] SHINAR, D.: *Traffic Safety and Human Behavior: Second Edition*. Emerald Publishing Limited, 2017.
- [Sid+19] SIDDIQ, K.; HOBDEN, M. K.; PENNOCK, S. R.; WATSON, R. J.: Phase Noise in FMCW Radar Systems. In: *IEEE Transactions on Aerospace and Electronic Systems* 55.1 (Feb. 2019), pp. 70–81. URL: <https://ieeexplore.ieee.org/abstract/document/8387490> (visited on 05/29/2025).
- [Sko+80] SKOLNIK, M. I. et al.: *Introduction to radar systems*. Vol. 3. McGraw-hill New York, 1980.
- [SM23] SRIVASTAV, A.; MANDAL, S.: Radars for Autonomous Driving: A Review of Deep Learning Methods and Challenges. In: *IEEE Access* 11 (2023), pp. 97147–97168. URL: <https://ieeexplore.ieee.org/abstract/document/10242101> (visited on 09/03/2024).

- [Sun+20] SUN, P.; KRETZSCHMAR, H.; DOTIWALLA, X.; CHOUARD, A.; PATNAIK, V.; TSUI, P.; GUO, J.; ZHOU, Y.; CHAI, Y.; CAINE, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 2446–2454.
- [TLF14] THOMAS, L.; LEVITT, D.; FARLEY, E.: *North Carolina Pedestrian Crash Types 2008 - 2012*. Tech. rep. North Carolina Department of Transportation (NC-DOT), 2014.
- [TSC06] TANELLI, M.; SAVARESI, S. M.; CANTONI, C.: Longitudinal vehicle speed estimation for traction and braking control systems. In: ISSN: 2165-302X. Oct. 2006, pp. 2790–2795. URL: <https://ieeexplore.ieee.org/abstract/document/4777080> (visited on 12/21/2025).
- [Tuz02] TUZLUKOV, V.: *Signal Processing Noise*. Electrical Engineering & Applied Signal Processing Series. CRC Press, 2002.
- [TZ00] TORR, P. H. S.; ZISSERMAN, A.: MLESAC: A New Robust Estimator with Application to Estimating Image Geometry. In: *Computer Vision and Image Understanding* 78.1 (Apr. 2000), pp. 138–156. URL: <https://www.sciencedirect.com/science/article/pii/S1077314299908329> (visited on 06/26/2025).
- [Ulr+22] ULRICH, M.; BRAUN, S.; KÖHLER, D.; NIEDERLÖHNER, D.; FAION, F.; GLÄSER, C.; BLUME, H.: Improved Orientation Estimation and Detection with Hybrid Object Detection Networks for Automotive Radar. In: Oct. 2022, pp. 111–117. URL: <https://ieeexplore.ieee.org/document/9922457> (visited on 06/14/2025).
- [Vas+17] VASWANI, A.; SHAZEER, N.; PARMAR, N.; USZKOREIT, J.; JONES, L.; GOMEZ, A. N.; KAISER, L. U.; POLOSUKHIN, I.: Attention is All you Need. In: *Advances in Neural Information Processing Systems*. Ed. by GUYON, I.; LUXBURG, U. V.; BENGIO, S.; WALLACH, H.; FERGUS, R.; VISHWANATHAN, S.; GARNETT, R. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- [VM+24] VIADERO-MONASTERIO, F.; ALONSO-RENTERÍA, L.; PÉREZ-ORIA, J.; VIADERO-RUEDA, F.: Radar-Based Pedestrian and Vehicle Detection and Identification for Driving Assistance. en. In: *Vehicles* 6.3 (Sept. 2024), pp. 1185–1199. URL: <https://www.mdpi.com/2624-8921/6/3/56> (visited on 10/03/2025).
- [Wan+21a] WANG, Y.; HWANG, J.-N.; WANG, G.; LIU, H.; KIM, K.-J.; HSU, H.-M.; CAI, J.; ZHANG, H.; JIANG, Z.; GU, R.: ROD2021 Challenge: A Summary for Radar Object Detection Challenge for Autonomous Driving Applications. In: *Proceedings of the 2021 International Conference on Multimedia Retrieval*. 2021, pp. 553–559.

- [Wan+21b] WANG, Y.; WANG, G.; HSU, H.-M.; LIU, H.; HWANG, J.-N.: Rethinking of Radar's Role: A Camera-Radar Dataset and Systematic Annotator via Coordinate Alignment. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 2021, pp. 2815–2824.
- [Wan+24a] WANG, A.; CHEN, H.; LIU, L.; CHEN, K.; LIN, Z.; HAN, J., et al.: Yolov10: Real-time end-to-end object detection. In: *Advances in Neural Information Processing Systems* 37 (2024), pp. 107984–108011.
- [Wan+24b] WANG, Y.; HUANG, Y.; LIU, J.; ZHANG, R.; ZHANG, H.; HONG, W.: Interference Mitigation for Automotive FMCW Radar With Tensor Decomposition. In: *IEEE Transactions on Intelligent Transportation Systems* 25.8 (Aug. 2024), pp. 9204–9223. URL: <https://ieeexplore.ieee.org/abstract/document/10493188> (visited on 05/29/2025).
- [Way25] WAYMO LLC: *Waymo Driver*. Website. 2025. URL: <https://waymo.com/waymo-driver/>.
- [WDE21] WOLF, M.; DOUAT, L. R.; ERZ, M.: Safety-Aware Metric for People Detection. In: *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. Sept. 2021, pp. 2759–2765. URL: <https://ieeexplore.ieee.org/abstract/document/9564734> (visited on 09/11/2024).
- [Wei+21] WEIGL, K.; SCHARTMÜLLER, C.; RIENER, A.; STEINHAUSER, M.: Development of the Questionnaire on the Acceptance of Automated Driving (QAAD): Data-driven models for Level 3 and Level 5 automated driving. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 83 (2021), pp. 42–59. URL: <https://www.sciencedirect.com/science/article/pii/S1369847821002151>.
- [Wes+23] WESTHOFEN, L.; NEUROHR, C.; KOOPMANN, T.; BUTZ, M.; SCHÜTT, B.; UTESCH, F.; NEUROHR, B.; GUTENKUNST, C.; BÖDE, E.: Criticality Metrics for Automated Driving: A Review and Suitability Analysis of the State of the Art. en. In: *Archives of Computational Methods in Engineering* 30.1 (Jan. 2023), pp. 1–35. URL: <https://doi.org/10.1007/s11831-022-09788-7> (visited on 06/04/2025).
- [Wil90] WILMSHURST, T.: *Signal Recovery from Noise in Electronic Instrumentation (2nd ed.)* CRC Press, 1990.
- [Win+15] WINNER, H.; HAKULI, S.; LOTZ, F.; SINGER, C.: *Handbook of Driver Assistance Systems: Basic Information, Components and Systems for Active Safety and Comfort*. Handbook of Driver Assistance Systems. Springer International Publishing, 2015.
- [Win+23] WINKEL, K. N. DE; IRMAK, T.; HAPPEE, R.; SHYROKAU, B.: Standards for passenger comfort in automated vehicles: Acceleration and jerk. In: *Applied Ergonomics* 106 (Jan. 2023), p. 103881. URL: <https://www.sciencedirect.com/science/article/pii/S0003687022002046> (visited on 06/27/2025).

- [Win+24] WINNER, H.; DIETMAYER, K.; ECKSTEIN, L.; JIPP, M.; MAURER, M.; STILLER, C.: *Handbuch Assistiertes und Automatisiertes Fahren: Grundlagen, Komponenten und Systeme für assistiertes und automatisiertes Fahren*. ATZ/MTZ-Fachbuch. Springer Fachmedien Wiesbaden, 2024.
- [WPW22] WANG, C.; POPP, C.; WINNER, H.: Acceleration-Based Collision Criticality Metric for Holistic Online Safety Assessment in Automated Driving. In: *IEEE Access* 10 (2022), pp. 70662–70674. URL: <https://ieeexplore.ieee.org/abstract/document/9808147> (visited on 06/04/2025).
- [Wu+23] WU, Z.; CHEN, G.; GAN, Y.; WANG, L.; PU, J.: MVFusion: Multi-View 3D Object Detection with Semantic-aligned Radar and Camera Fusion. In: May 2023, pp. 2766–2773. URL: <https://ieeexplore.ieee.org/abstract/document/10161329> (visited on 06/08/2025).
- [WWN20] WANG, Z.; WU, Y.; NIU, Q.: Multi-Sensor Fusion in Automated Driving: A Survey. In: *IEEE Access* 8 (2020), pp. 2847–2868. URL: <https://ieeexplore.ieee.org/abstract/document/8943388> (visited on 06/24/2025).
- [Xie+22] XIE, S.; CHEN, S.; ZHENG, J.; TOMIZUKA, M.; ZHENG, N.; WANG, J.: From Human Driving to Automated Driving: What Do We Know About Drivers? In: *IEEE Transactions on Intelligent Transportation Systems* 23.7 (July 2022), pp. 6189–6205. URL: <https://ieeexplore.ieee.org/abstract/document/9447799> (visited on 10/03/2025).
- [Xu+21] XU, B.; ZHANG, X.; WANG, L.; HU, X.; LI, Z.; PAN, S.; LI, J.; DENG, Y.: RPFA-Net: a 4D RaDAR Pillar Feature Attention Network for 3D Object Detection. In: Sept. 2021, pp. 3061–3066. URL: <https://ieeexplore.ieee.org/abstract/document/9564754> (visited on 06/08/2025).
- [Yao+24] YAO, S.; GUAN, R.; HUANG, X.; LI, Z.; SHA, X.; YUE, Y.; LIM, E. G.; SEO, H.; MAN, K. L.; ZHU, X.; YUE, Y.: Radar-Camera Fusion for Object Detection and Semantic Segmentation in Autonomous Driving: A Comprehensive Review. In: *IEEE Transactions on Intelligent Vehicles* 9.1 (Jan. 2024), pp. 2094–2128. URL: <https://ieeexplore.ieee.org/abstract/document/10225711> (visited on 06/24/2025).
- [Zel+23] ZELLER, M.; BEHLEY, J.; HEIDINGSFELD, M.; STACHNISS, C.: Gaussian Radar Transformer for Semantic Segmentation in Noisy Radar Data. In: *IEEE Robotics and Automation Letters* 8.1 (Jan. 2023), pp. 344–351. URL: <https://ieeexplore.ieee.org/abstract/document/9968089> (visited on 06/09/2025).
- [Zha+21] ZHAO, H.; JIANG, L.; JIA, J.; TORR, P. H.; KOLTUN, V.: Point Transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021, pp. 16259–16268.

- [Zha+23a] ZHANG, J.; ZHUGE, H.; WU, Z.; PENG, G.; WEN, M.; LIU, Y.; WANG, D.: 4DRadarSLAM: A 4D Imaging Radar SLAM System for Large-scale Environments based on Pose Graph Optimization. In: May 2023, pp. 8333–8340. URL: <https://ieeexplore.ieee.org/abstract/document/10160670> (visited on 12/21/2025).
- [Zha+23b] ZHANG, X.; TAO, J.; TAN, K.; TÖRNGREN, M.; SÁNCHEZ, J. M. G.; RAMLI, M. R.; TAO, X.; GYLLENHAMMAR, M.; WOTAWA, F.; MOHAN, N.; NICA, M.; FELBINGER, H.: Finding Critical Scenarios for Automated Driving Systems: A Systematic Mapping Study. In: *IEEE Transactions on Software Engineering* 49.3 (Mar. 2023), pp. 991–1026. URL: <https://ieeexplore.ieee.org/abstract/document/9763411> (visited on 10/03/2025).
- [Zhe+24] ZHENG, L.; LONG, J.; LOPS, M.; LIU, F.; HU, X.; ZHAO, C.: Detection of Ghost Targets for Automotive Radar in the Presence of Multipath. In: *IEEE Transactions on Signal Processing* 72 (2024), pp. 2204–2220. URL: <https://ieeexplore.ieee.org/abstract/document/10488858> (visited on 05/29/2025).
- [ZNL21] ZHANG, A.; NOWRUZI, F. E.; LAGANIERE, R.: RADDet: Range-Azimuth-Doppler based Radar Object Detection for Dynamic Road Users. In: May 2021, pp. 95–102. URL: <https://ieeexplore.ieee.org/abstract/document/9469418> (visited on 06/10/2025).

Own Publications and Conference Proceedings

- [Brü+24a] BRÜHL, T.; BLAHAK, F.; SCHWAGER, R.; EWECKER, L.; SOHN, T. S.; HOHMANN, S.: A Framework for Localization in a Ground Plan Map based on Radar Perception and Odometry Data. In: *2024 IEEE Intelligent Vehicles Symposium (IV)* (2024). ISSN: 2642-7214, pp. 892–899.
- [Brü+24b] BRÜHL, T.; EBERHARDT, T. D.; SCHWAGER, R.; EWECKER, L.; SOHN, T. S.; HOHMANN, S.: Odometry Estimation by Fusing Multiple Radar Sensors and an Inertial Measurement Unit. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)* (2024), pp. 14109–14115.
- [Brü+24c] BRÜHL, T.; EWECKER, L.; SCHWAGER, R.; SOHN, T. S.; HOHMANN, S.: Making Radar Detections Safe for Autonomous Driving: A Review. In: *10th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS)* (2024), pp. 299–310.
- [Brü+24d] BRÜHL, T.; GLÖNKLER, J.; SCHWAGER, R.; SOHN, T. S.; EBERHARDT, T. D.; HOHMANN, S.: SAFERad: A Framework to Enable Radar Data for Safety-Relevant Perception Tasks. In: *Under Review: IEEE Transactions on Intelligent Vehicles* (2024). under review.

- [Brü+24e] **BRÜHL, T.**; SHANMUGANATHAN, A. K.; EWECKER, L.; SCHWAGER, R.; SOHN, T. S.; VACULIN, O.; HOHMANN, S.: Consideration of Safety Aspects in a Camera-Aided, Radar-Based Free Space Detection. In: *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)* (2024). ISSN: 2153-0017, pp. 3997–4004.
- [Brü+25a] **BRÜHL, T.**; SOHN, T. S.; EBERHARDT, T. D.; SCHWAGER, R.; HOHMANN, S.: Safer Radar Motion by Scrutinizing Critical Velocity Estimates. In: *2025 IEEE Intelligent Vehicles Symposium (IV)* (June 2025), pp. 1868–1874.
- [Brü+25b] **BRÜHL, T.**; VICO, A.; GOLDSCHIEDER, J.; SCHWAGER, R.; SOHN, T. S.; EBERHARDT, T. D.; HOHMANN, S.: Adaptive Radar Clustering and Tracking based on Point Criticality Assessment. In: *2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (Oct. 2025), pp. 490–496.
- [Brü+26] **BRÜHL, T.**; SCHWAGER, R.; SOHN, T. S.; DIETER EBERHARDT, T.; HOHMANN, S.: CFAR++: Region-Aware Noise Thresholding for Safe Radar Detections. In: *2026 IEEE Intelligent Vehicles Symposium (IV)* (June 2026), pp. 1740–1746.

Supervised Student Theses

- [Bla23] BLAHAK, F.: *Lokalisierung in einer Indoor-Umgebungskarte mittels Radar- und Kamerawahrnehmungsdaten für eine automatisierte Parkfunktion*, Master's thesis. Institute of Automotive Engineering (IFS), University of Stuttgart, 2023.
- [Ebe22] EBERHARDT, T. D.: *Implementierung einer Fahrzeug-Eigenbewegungsschätzung unter Nutzung von Sensordaten aus verschiedenen Domänen*, Master's thesis. University of Applied Sciences (HTW) Berlin, 2022.
- [Fig24] FIGLESTAHLER, M.: *A Concept to Guarantee Functional Safety in Perception Functions*, Master's thesis. Institute of Vehicle System Technology (FAST), Karlsruhe Institute of Technology (KIT), 2024.
- [Glö24] GLÖNKLER, J.: *Method to Identify and Process SOTIF-relevant Scenarios of a Radar- and Camera-Based Free Space Detection*, Master's thesis. Faculty of Mobility and Technology, Esslingen University of Applied Sciences, 2024.
- [Sha24] SHANMUGANATHAN, A. K.: *Safe Free Space Detection Based on Radar-Camera Fusion in the Parking Domain*, Master's thesis. Faculty of Electrical Engineering and Information Technology, Technische Hochschule Ingolstadt (THI), 2024.
- [Wei23] WEISSER, J.: *Kameragestützte Freiraumerkennung für Anwendungen des hochautomatisierten Parkens*, Bachelor's thesis. Baden-Württemberg Cooperative State University (DHBW) Friedrichshafen, 2023.

