


MAXIMILIAN NAUMANN

Probabilistic Motion Planning for Automated Vehicles



Maximilian Naumann

**Probabilistic Motion Planning
for Automated Vehicles**

Schriftenreihe
Institut für Mess- und Regelungstechnik,
Karlsruher Institut für Technologie (KIT)
Band 049

Eine Übersicht aller bisher in dieser Schriftenreihe erschienenen
Bände finden Sie am Ende des Buchs.

Probabilistic Motion Planning for Automated Vehicles

by
Maximilian Naumann

Karlsruher Institut für Technologie
Institut für Mess- und Regelungstechnik

Probabilistic Motion Planning for Automated Vehicles

Zur Erlangung des akademischen Grades eines Doktors der
Ingenieurwissenschaften von der KIT-Fakultät für Maschinenbau des
Karlsruher Instituts für Technologie (KIT) genehmigte Dissertation

von Maximilian Naumann

Tag der mündlichen Prüfung: 27. August 2020

Hauptreferent: Prof. Dr.-Ing. Christoph Stiller

Korreferent: Prof. Dr. Mykel J. Kochenderfer

Impressum



Karlsruher Institut für Technologie (KIT)
KIT Scientific Publishing
Straße am Forum 2
D-76131 Karlsruhe

KIT Scientific Publishing is a registered trademark
of Karlsruhe Institute of Technology.
Reprint using the book cover is not allowed.

www.ksp.kit.edu



*This document – excluding the cover, pictures and graphs – is licensed
under a Creative Commons Attribution-Share Alike 4.0 International License
(CC BY-SA 4.0): <https://creativecommons.org/licenses/by-sa/4.0/deed.en>*



*The cover page is licensed under a Creative Commons
Attribution-No Derivatives 4.0 International License (CC BY-ND 4.0):
<https://creativecommons.org/licenses/by-nd/4.0/deed.en>*

Print on Demand 2020 – Gedruckt auf FSC-zertifiziertem Papier

ISSN 1613-4214

ISBN 978-3-7315-1070-3

DOI 10.5445/KSP/1000126389

Preface

The research leading to this work was conducted during my time as scientific researcher at FZI Research Center for Information Technology and at the Institute of Measurement and Control (MRT) at Karlsruhe Institute of Technology (KIT) in Karlsruhe, Germany.

In the following, I would like to appreciate the support that I received on the way towards this thesis and that made this work possible.

Sincere thanks to Prof. Stiller, who gave me the opportunity to conduct research in this very interesting area under his supervision and in his research group. I had both, the freedom to pursue own ideas and new approaches, but also profound scientific support, which I highly appreciate.

Special thanks to Prof. Kochenderfer for serving as co-examiner of this dissertation. That is a great honor to me.

I would also like to thank Prof. Tomizuka for hosting me at MSCL at UC Berkeley. I had a nice and fruitful time there, and we were able to start ongoing cooperations, which I also want to thank Prof. de La Fortelle for. Further, I would like to thank the Karlsruhe House of Young Scientists for their support of this research visit.

Sincere thanks also to my colleagues at FZI and MRT. I enjoy working in and learning from our diverse group, and I appreciate the friendships I have made. Regarding the scientific discussions targeted towards this work, I particularly want to thank Hendrik, Martin, Philipp and Sven. For the great coding and tooling support, I want to thank Fabian, Johannes and Ole.

Beyond the group, I want to thank Constantin, Daniel, Stephan and Thomas for their feedback to drafts of this work.

At this point, I also would like to appreciate the support of my friends, particularly Team Oswald & Co. I am happy to know you.

The biggest thanks goes to my recently grown and hopefully yet growing family. There is too much to thank for. I am endlessly grateful.

Karlsruhe, May 2020

Maximilian Naumann

Zusammenfassung

Diese Dissertation beschäftigt sich mit der Problemstellung der Bewegungsplanung für automatische Fahrzeuge. Als Voraussetzung für den Einsatz im realen Straßenverkehr müssen automatische Fahrzeuge ein angemessenes und zuverlässiges Fahrverhalten im Mischverkehr mit menschgeführten Fahrzeugen aufweisen. Neben den Unsicherheiten, welche aus fehlerbehafteter maschineller Wahrnehmung, Verdeckungen und begrenzter Sensorreichweite resultieren, müssen dabei auch Unsicherheiten im Verhalten anderer Verkehrsteilnehmer berücksichtigt werden.

Verwandte Ansätze zur Bewegungsplanung im Mischverkehr formulieren oft ein deterministisches Optimierungsproblem. Dessen Lösung ist auf die Bestimmung einer einzelnen Trajektorie beschränkt. Fehlerhafte Vorhersagen des Verhaltens anderer Verkehrsteilnehmer werden im Rahmen von fortlaufender Neuplanung korrigiert, während größere Unsicherheiten in konservativem, übervorsichtigem Fahrverhalten resultieren. Auf Grund der Unzulänglichkeiten dieser Problemformulierung in kooperativen und stark unsicherheitsbehafteten Szenarien werden zunehmend probabilistische Ansätze vorgestellt. Die vollumfängliche Unsicherheitsbetrachtung führt im Kontext der Echtzeitanforderung jedoch häufig zu starken Einschränkungen in der Modellierung der Handlungsmöglichkeiten automatischer Fahrzeuge. Darüber hinaus werden Sicherheitsaspekte und Verkehrsregelkonformität oft nicht betrachtet.

In der vorliegenden Arbeit werden daher drei Bewegungsplanungsansätze vorgestellt, welche den unterschiedlichen dominanten Unsicherheiten in verschiedenen Szenarien Rechnung tragen, ohne dabei die Handlungsmöglichkeiten des automatischen Fahrzeugs stark einzuschränken. Darüber hinaus wird ein szenariobasierter Sicherheitsansatz vorgestellt. Dieser baut auf einem bestehenden Ansatz auf, welcher garantiert, dass ein automatisches Fahrzeug nie einen Unfall verursachen wird. Dieser Ansatz wird um die Betrachtung

von Verkehrsregeln für kreuzenden und zusammenlaufenden Verkehr, sowie Verdeckungen, begrenzte Sensorreichweite und Fahrstreifenwechsel erweitert.

Für nicht-interaktive Szenarien mit definiertem Vorfahrtsrecht wird ein probabilistischer Bewegungsplanungsansatz vorgestellt, welcher auf einer vorgelegerten Prädiktion anderer Verkehrsteilnehmer aufbaut. Das Planungsproblem wird als teilweise beobachtbarer Markov-Entscheidungsprozess modelliert. Im Gegensatz zu bisherigen Formulierungen wird jedoch davon ausgegangen, dass die Voraussage möglicher zukünftiger Verläufe der Unsicherheit über die Bewegung anderer Verkehrsteilnehmer unabhängig von der Bewegung des automatischen Fahrzeugs getroffen werden kann. Zusätzlich zu dieser Prädiktion sichtbarer Verkehrsteilnehmer wird der Einfluss von Verdeckungen und begrenzter Sensorreichweite betrachtet. Mit dem vorgeschlagenen Ansatz ist trotz der umfänglichen Berücksichtigung der Unsicherheiten die Planung in einem kontinuierlichen Aktionsraum möglich.

Zwei weitere Ansätze zielen auf die dominanten Unsicherheiten in interaktiven Szenarien ab. Zur Ermöglichung von Fahrstreifenwechseln in dichtem Verkehr wird die Unsicherheit darüber, ob andere Verkehrsteilnehmer bereitwillig Raum für einen Wechsel bieten, in einem regelbasierten Ansatz aktiv reduziert. Die dabei erzeugten Trajektorien sind ebenfalls sicher und verkehrsregelkonform, gemäß des vorgestellten Sicherheitsansatzes. Für die Ermöglichung von Kooperation in Szenarien ohne definiertes Vorfahrtsrecht wird ein Multiagentenansatz vorgestellt. Dabei wird zunächst die global optimale Lösung des vorliegenden Multiagentenproblems hinsichtlich ihrer Eindeutigkeit überprüft. Ist die Lösung eindeutig, so wird diese verfolgt. Trotzdem wird geprüft, ob die anderen Verkehrsteilnehmer sich entsprechend des angenommenen Kooperationsmodells verhalten, um im Falle einer Verletzung der Modellannahmen auf eine konservative Alternativtrajektorie wechseln zu können.

Die Leistungsfähigkeit der vorgestellten Ansätze wird in verschiedenen Szenarien mit kreuzenden und zusammenlaufenden Fahrstreifen, teils mit Verdeckungen und begrenzter Sicht, sowie Fahrstreifenwechseln und einer Engstelle ohne definiertes Vorfahrtsrecht gezeigt.

Abstract

This thesis targets the problem of motion planning for automated vehicles. As a prerequisite for their on-road deployment, automated vehicles must show an appropriate and reliable driving behavior in mixed traffic, i.e. alongside human drivers. Besides the uncertainties resulting from imperfect perception, occlusions and limited sensor range, also the uncertainties in the behavior of other traffic participants have to be considered.

Related approaches for motion planning in mixed traffic often employ a deterministic problem formulation. The solution of such formulations is restricted to a single trajectory. Deviations from the prediction of other traffic participants are accounted for during replanning, while large uncertainties lead to conservative and over-cautious behavior. As a result of the shortcomings of these formulations in cooperative scenarios and scenarios with severe uncertainties, probabilistic approaches are pursued. Due to the need for real-time capability, however, a holistic uncertainty treatment often induces a strong limitation of the action space of automated vehicles. Moreover, safety and traffic rule compliance are often not considered.

Thus, in this work, three motion planning approaches and a scenario-based safety approach are presented. The safety approach is based on an existing concept, which targets the guarantee that automated vehicles will never cause accidents. This concept is enhanced by the consideration of traffic rules for crossing and merging traffic, occlusions, limited sensor range and lane changes. The three presented motion planning approaches are targeted towards the different predominant uncertainties in different scenarios, while operating in a continuous action space.

For non-interactive scenarios with clear precedence, a probabilistic approach is presented. The problem is modeled as a partially observable Markov decision process (POMDP). In contrast to existing approaches, the underlying assumption is that the prediction of the future progression of the uncertainty

in the behavior of other traffic participants can be performed independently of the automated vehicle's motion plan. In addition to this prediction of currently visible traffic participants, the influence of occlusions and limited sensor range is considered. Despite its thorough uncertainty consideration, the presented approach facilitates planning in a continuous action space.

Two further approaches are targeted towards the predominant uncertainties in interactive scenarios. In order to facilitate lane changes in dense traffic, a rule-based approach is proposed. The latter seeks to actively reduce the uncertainty in whether other vehicles willingly make room for a lane change. The generated trajectories are safe and traffic rule compliant with respect to the presented safety approach. To facilitate cooperation in scenarios without clear precedence, a multi-agent approach is presented. The globally optimal solution to the multi-agent problem is first analyzed regarding its ambiguity. If an unambiguous, cooperative solution is found, it is pursued. Still, the compliance of other vehicles with the presumed cooperation model is checked, and a conservative fallback trajectory is pursued in case of non-compliance.

The performance of the presented approaches is shown in various scenarios with intersecting lanes, partly with limited visibility, as well as lane changes and a narrowing without predefined right of way.

Contents

Preface	i
Zusammenfassung	iii
Abstract	v
Notation and Symbols	xi
Notation	xi
Variables, Functions and Operators	xii
Common Subscripts and Superscripts	xiii
1 Introduction	1
1.1 Problem Statement	1
1.2 Goal Description and Limitations	2
1.3 Contributions	3
1.4 Outline	4
2 Fundamentals and Related Work	5
2.1 Fundamentals of Robot Motion Planning and Decision Making	5
2.1.1 Basic Components of Planning	6
2.1.2 Properties of Planning Algorithms	8
2.1.3 Planning in Discrete State Spaces	9
2.1.4 Planning in Continuous State Spaces	13
2.1.5 Planning under Uncertainty	16
2.1.6 Planning under Differential Constraints	25
2.1.7 Multi-Agent Planning	27
2.1.8 Replanning	28

- 2.2 Related Work on Motion Planning for Automated Driving . . . 29
 - 2.2.1 System Architectures of Automated Vehicles 30
 - 2.2.2 Deterministic and Reactive Motion Planning 32
 - 2.2.3 Probabilistic and Interactive Motion Planning 36
 - 2.2.4 Safety in Decision Making and Motion Planning 42

- 3 Formulating the Decision Problem of Motion Planning in Mixed Traffic 45**
 - 3.1 Underlying Model Assumptions 46
 - 3.2 Perfect Knowledge 48
 - 3.3 Uncertainty in Future Trajectories of Others 49
 - 3.4 Uncertainty in State of Others 54
 - 3.5 Uncertainty in Existence of Others 55
 - 3.6 Occlusions and Limited Sensor Range 56
 - 3.7 Interaction 58

- 4 Solution to the Decision Process 59**
 - 4.1 Prioritizing Safety over Comfort and Convenience 60
 - 4.2 Ensuring Safety: Incorporating Traffic Rules 61
 - 4.2.1 Single Lane 63
 - 4.2.2 Intersecting Lanes 65
 - 4.2.3 Parallel Lanes 76
 - 4.2.4 Further Scenarios 79
 - 4.3 Probabilistic Motion Planning based on Independent Predictions 79
 - 4.3.1 General Approach 79
 - 4.3.2 Single Lane 86
 - 4.3.3 Intersecting Lanes 87
 - 4.4 Courtesy-Aware Lane Changes 93
 - 4.5 Cooperative Global Optimum Approach with Model Compliance Consideration 96
 - 4.6 Transitions between Scenarios 99

5	Evaluation	101
5.1	Implementation	101
5.1.1	On-Road Testing	102
5.1.2	Simulation	103
5.2	Scenario-Based Evaluation	103
5.2.1	Clear Precedence	104
5.2.2	Lane Change	129
5.2.3	Mutual Cooperation	132
6	Conclusions and Future Directions	137
6.1	Conclusions	137
6.2	Future Directions	140
A	Appendix	141
A.1	Safe Distance for Merging	141
A.2	Evaluation Parameters	143
A.2.1	Safety Parameters	143
A.2.2	Comfort and Convenience Parameters	143
	Bibliography	147
	Publications by the Author	159
	Supervised Theses	161

Notation and Symbols

Notation

ACC	Adaptive Cruise Control
CZ	Conflict Zone
GPS	Global Positioning System
IDM	Intelligent Driver Model
LIDAR	Light Detection And Ranging (Sensor)
LQR	Linear–Quadratic Regulator
MCMC	Markov Chain Monte Carlo
MCTS	Monte Carlo Tree Search
MDP	Markov Decision Process
MEU	Principle of Maximum Expected Utility
MIQP	Mixed-Integer Quadratic Programming
MMDP	Multi-Agent Markov Decision Process
MPC	Model Predictive Control
PDF	Probability Density Function(al)
PET	Post-Encroachment Time
POMDP	Partially Observable Markov Decision Process
PRM	Probabilistic Roadmap
PVD	Path-Velocity Decomposition
QMDP	Approximation technique for solving POMDPs
QP	Quadratic Programming
RL	Reinforcement Learning
ROS	Robot Operating System
RRT	Rapidly Exploring Random Tree
RSS	Responsibility Sensitive Safety
SQP	Sequential Quadratic Programming

TZC	Time of Zone Clearance
V2X	Vehicle to Vehicle or Infrastructure Communication
VRU	Vulnerable Road User

Variables, Functions and Operators

Some letters are used with different meanings, such as a for acceleration and action, or s for arc length and state. While in Sec. 2.1, the notation is introduced as part of the fundamentals, the following list denotes the usage throughout the rest of this work, including figures.

a	Acceleration
b	Belief
\mathcal{B}	Belief space
d	Distance (arc length)
\mathcal{D}	Vehicle dimensions
$\mathbf{E}(\cdot)$	Expected value
J	Cost functional (at one point in time)
\mathbf{J}	Cost functional (integrated over time)
l	Length
p	Probability (as variable or function with custom arguments)
$\text{Pr}(\cdot)$	Probability (of an event)
ρ	Reaction time
s	Position (arc length)
t	Time
\mathbf{u}	Input (control theory)
v	Velocity
\mathbf{x}	Trajectory
$\underline{\mathbf{x}}$	Trajectory ensemble (multiple trajectories)
\mathbf{X}	Stochastic process describing a trajectory
\mathcal{X}	Trajectory set
$\underline{\mathcal{X}}$	Trajectory ensemble set
\mathbb{X}	Belief progression over trajectory ensemble sets

Common Subscripts and Superscripts

*	Optimal
'	Next time step
accel	Acceleration
avg	Average
brake	Deceleration
cfb	Conservative fallback
cons	Considered
crit	Critical
exp	Expected
go	Globally optimal
init	Initial
LC	Lane change
mc	Model correct
prio	Prioritized
react	Reaction/reactive
row	Right of way
suff	Sufficient
vis	Visible
e	Ego vehicle
o	Other vehicles
oi	Particular other vehicle <i>i</i>

1 Introduction

The large potential of automated vehicles to change the future of mobility in terms of safety, energy efficiency, accessibility and comfort has been known for a long time. In recent years, the feasibility of automated driving along with regular human-driven vehicles, so-called *mixed traffic*, has been investigated and shown in various experimental vehicles.

The problem statement of motion planning for automated vehicles in mixed traffic, the goal of this work and its outline are introduced in the following.

1.1 Problem Statement

Motion planning is key to every automated physical system, and a well investigated subject in the field of robotics. Analyses of human drivers show that for humans, the driving task is commonly split into three subtasks: Navigation, guidance and stabilization [Don16, p.22]. Assistance in navigation, e.g. through in-car or mobile navigation systems, and assistance in stabilization, e.g. through anti-lock braking systems or electronic stability control, have already been available in series cars for decades. Guidance describes the task of determining an appropriate path and velocity based on the road layout and the motion of other traffic participants. In the field of automated vehicles, a similar task distribution is common. The general goal of transportation, which is to reach a certain position in the world, is commonly processed by a navigational layer and turned into a route towards the destination. Along this route, the task is to determine an appropriate *trajectory*, i.e. the state of the vehicle as a function of time. The latter task is referred to as *motion planning* for automated vehicles. The decision is based on previously acquired knowledge about the environment, such as the drivable area and detected objects, but also traffic rules. The resulting trajectory is then passed to a controller, which determines

the input to the actuators in order to track the trajectory closely and compensate for disturbance.

The key challenges in motion planning for automated driving arise from the collision risk with other traffic participants. Since human lives are at risk, safety takes the top priority. On the other hand, over-cautious behavior, sometimes referred to as “driving like a learner”¹, is not only inconvenient but can also cause dangerous situations, as this behavior might cause misunderstanding by humans or entice them to risky overtaking maneuvers, for example.

1.2 Goal Description and Limitations

The goal of this work is to design motion planning approaches for automated vehicles in mixed traffic. As the safety of passengers and other traffic participants must never be put at risk, the approaches must allow for a **safety** verification w.r.t. reasonable assumptions. Still, they should yield **convenient** motion plans, i.e. plans that are comfortable for the passenger but also useful w.r.t. the desired destination. Thus, the motion plan should not be over-cautious, also in order to not obstruct the traffic.

Throughout this work, the presence of a high definition map of the road layout is assumed, including the traffic rules. Further, all objects within the sensor range that are not occluded are assumed to be detected by the perception module, and assigned a non-zero existence probability. Also, an independent probabilistic prediction of the future trajectory of other traffic participants is required, assuming they are not influenced by the ego vehicle. Since the first decades of automated driving will be taking place alongside human traffic, the approach is targeted towards this mixed traffic and does not rely on vehicle to vehicle communication. Also, the approach is targeted towards road traffic that is structured into lanes, excluding for example parking lots. Scenarios with special traffic rules, such as zipper merges, and overtaking with oncoming traffic are not considered.

¹ Cf. Ziegler et al. [ZBS⁺14]: “Some passengers thus compared the robot to a human learner taking driving lessons.”

1.3 Contributions

The contributions of this thesis are as follows:

- The concept of responsibility-sensitive safety (RSS) [SSS18], which is targeted towards the guarantee that automated vehicles will never cause a collision, is extended by considering of right-of-way rules for scenarios with crossing and merging traffic, occlusions, limited sensor range and lane changes.
- For driving scenarios with clear precedence, a new problem formulation is introduced, exploiting the independence of the future trajectories of prioritized traffic participants on the behavior of the ego vehicle. The resulting formulation corresponds to a partially observable Markov decision process (POMDP) with an ego-action-independent belief over the trajectories of other traffic participants.
- In order to solve the introduced action-independent POMDP in real time, a heuristic is proposed, employing driver models to generate reference trajectories and a reaction analysis w.r.t. unfavorable behavior of prioritized traffic participants. The approach determines safe trajectories without the limitation to a discrete action space while considering uncertainty in the current state of other traffic participants, their future route, their behavior along the route and their existence. The consideration particularly includes limited sensor range and occlusions.
- For lane changes, an approach that incorporates uncertainty regarding the courtesy of other traffic participants is proposed. The approach facilitates safe lane changes also in dense traffic.
- Lastly, an approach for scenarios that require mutual cooperation is presented. As a cooperative motion plan, the solution to a multi-agent formulation is determined. However, since the cooperative solution is not always accomplished, the success probability and the convenience of the fallback plan is incorporated into the decision of whether to attempt the cooperative solution or to switch to the conservative fallback plan.

1.4 Outline

The remainder of this work is structured as follows: In Chapter 2, the fundamentals of decision making and robot motion planning are explained, followed by a review of related work in the field. Subsequently, in Chapter 3, the problem formulation of the underlying decision problem is deduced, before approaches to solve the decision problem are presented in Chapter 4. These approaches are evaluated in several scenarios in Chapter 5. Finally, the work is concluded and future research directions are given in Chapter 6.

2 Fundamentals and Related Work

Since the term *motion planning* describes various problems in the field of robotics and beyond, Sec. 2.1 starts with an overview of existing formulations of motion planning problems and algorithms to solve these. This includes the definition of the terms and variables used in the remainder of this work. Further, these problems are put in context to related areas, such as decision making and control theory. Subsequently, related work from the area of automated driving is presented in Sec. 2.2.

2.1 Fundamentals of Robot Motion Planning and Decision Making

In robotics, the term *motion planning* is often used to describe the task of finding the *optimal* sequence of *configurations* or *states* that move a robot from a given source configuration/state to a destination. Similarly, the term *decision making* is commonly used to describe the task of selecting the action out of a (potentially infinite) set of actions, or deriving the policy to select this action, depending on the current state, which ultimately leads to the largest reward. In the course of a decision process, the previous task commonly results in a sequence of actions. *Control theory*, as a third term, targets the task of finding a control law which yields optimal inputs to a (dynamical) system, such that the cost induced by both input to and state of the system is minimized. The previous definitions presage that the field of motion planning overlaps, or is at least related to, decision making and control theory. In the following, these three fields are subsumed under the term *planning*.

Starting with basic components of planning and properties of planning algorithms, the remainder of this section introduces planning in different state

spaces, under uncertainty and differential constraints, as well as planning with multiple agents and replanning.

2.1.1 Basic Components of Planning

In order to further distinguish between the problem formulations that are common in the respective fields, we start by defining their basic components. The following definitions are based on the book of LaValle [LaV06].

Agent The entity that takes decisions, pursues motion plans or controls is commonly referred to as *agent*, *robot* or *controller*. Also, problem formulations can include multiple agents, robots or controllers. In the latter case, the decision, plan or control can be made centralized or decentralized. Decentralized approaches can include explicit communication between entities.

State Problem formulations in the above fields are commonly defined including a *state*, which is a distinct configuration of the *world*, meaning the part of the world that is relevant to the problem. The set of all possible states is called *state space*. It can be discrete, i.e. finite or countably infinite, or continuous, i.e. uncountably infinite. The definition of the state space is crucial to the problem formulation, as it largely affects the choice or design of algorithms to solve it. In motion planning and control, this state should be chosen to cover at least the configuration of the agent or robot. Thus, it is often called configuration space. In decision making, it commonly also covers the state of the world around the agent.

State Transitions, Limitations and Constraints Arising from the different state definitions, also the definitions of *state transitions* and their limitations differ between the fields. In decision making, the transition from one state to another is realized using actions. The action space might be limited, based on the current state. Furthermore, as the state includes the description of the world, an action does in general not lead to a deterministic state transition. A common simplification is the Markov assumption, stating that the conditional probability distribution of future states of the decision process only depends

on the present state and not on previous states. In motion planning, both the state space and the state transitions are in general subject to constraints. An example for constraints that limit the state spaces itself are obstacles, while the transitions might be limited by non-holonomic¹ constraints. In control theory, the state is manipulated by inputs and the state transition is described by differential equations. Constraints can limit both the state and the input.

Sequences and Time While the solutions to the above problem formulations might be abstract policies or control laws, their realization mostly involves a sequence of states, actions or inputs. This sequence might be continuous or discrete. Within this sequence, time might be explicitly modeled, especially if the state of the world changes over time, but also if time is part of the objective. Or time can be implicit, i.e. only the order of actions/inputs is important. In case of continuous sequences, the parametrization of the solution is then time independent. A sequence of states is referred to as *path* in case of time independence. In case the sequence is parametrized using the time, it is referred to as *trajectory*. Sequences that do not violate any constraint are called *feasible*.

Initial and Goal States Usually, problem formulations start with the robot or agent being in an *initial state*. In decision making, the definition of one or a set of *goal states* is common. In case a set of goal states is defined, reaching it is an additional condition for feasibility, i.e. solutions that lead from the initial state to a goal state without violating any constraint are called feasible.

Optimality Criterion In order to distinguish the quality of different feasible solutions, the problem formulations commonly include an optimality criterion. This criterion can include multiple objectives, which might even be contradictory, such as “be fast and consume little energy”. Mathematically, the criterion can be formulated as cost, which is to be minimized, or as reward, which is to be maximized.

¹ Constraints that are only expressible via differential equations are called non-holonomic [CHL⁺05, p.48].

Depending on the constraints implied by the robot, the world around it and the definition of *optimality*, the challenges differ largely. In robot motion planning, high dimensional configuration spaces are mostly the challenging part, leading to the necessity of discretization and approximative solutions. In decision making, the unknown evolution of the future, commonly modeled via transition probabilities, is often key to the problem formulation. In optimal control, complex dynamics of the system often make analytical solutions intractable, resulting in helping frameworks, such as model predictive control. Some common problem formulations of these fields are discussed in the following sections.

2.1.2 Properties of Planning Algorithms

Having defined the basic components of the planning problem formulation, this subsection introduces properties of planning algorithms. In the following, the term *planning algorithm* is used to describe a sequence of instructions with the goal of solving a planning problem. The quality and even applicability of planning algorithms strongly depend on the underlying planning problem. In order to distinguish planning algorithms later on, we start by introducing common characteristics.

Completeness An algorithm is said to be *complete*, if it finds a solution in case one exists, and is able to tell that no solution exists otherwise. For sampling-based approaches, the relaxed property *probabilistic completeness* is used for algorithms with which the probability of finding a solution (if it exists) converges to one for infinite runtime.

Feasibility vs. Optimality Algorithms for feasible planning solely focus on finding a feasible solution while not further distinguishing between multiple feasible solutions. In optimal planning, additionally, the goal is to find the feasible solution which is optimal w.r.t. a defined optimality criterion. While solutions can be optimal within a local neighborhood of solutions (so-called *local optimality*), the term optimality commonly refers to *global optimality*, if not further specified. Planning algorithms focusing on finding local optima are sometimes referred to as *local planning algorithm* or *local planner*.

Online/Offline In general, an *online* algorithm is able to process parts of the input, i.e. to start finding a solution without the need for the entire input. On the other hand, *offline* algorithms compute the entire solution at once. In motion planning, online algorithms allow the robot to execute motion before the entire trajectory has been computed, i.e. the trajectory is updated during motion. This is particularly important if the input is received from on-board sensors of a robot, because this inherently only allows to gather necessary input after the robot has moved. Online algorithms that focus on the solution for only a certain horizon are said to operate with a *receding horizon*.

Open-Loop/Closed-Loop Algorithms that operate *open-loop* do not use feedback for planning. Thus, all offline algorithms are by definition open-loop. *Closed-loop* algorithms, on the other hand, do use feedback.

Replanning The term *replanning* is closely related to the online/offline and open-loop/closed-loop properties: Replanning means that a found solution is updated (*replanned*) over and over again. This is a common way to include feedback into planning (closed-loop). It can be used to account for changes in the environment that are not or hardly predictable, or to account for disturbance or control errors. Receding horizon algorithms must replan in order to update or at least extend the solution for the new horizon.

Anytime An *anytime* or interruptible algorithm can yield a valid solution prior to its termination. Anytime algorithms commonly start by finding an initial solution and improving it over time. The anytime property only makes sense for optimal planning, as in feasible planning, there is no need to improve an initial solution.

2.1.3 Planning in Discrete State Spaces

The most basic planning problems are those in discrete state spaces, i.e. with a finite or countable infinite number of states. The action space is finite and state transitions are deterministic. No uncertainty is incorporated in this section.

Since the states along with their transitions can be represented in a graph, the presented algorithms in this section are graph search algorithms.

Notation

For problem and algorithm formulation, we use the following notation: The state space, i.e. the set of possible states s is denoted \mathcal{S} with every distinct state s being part of this set $s \in \mathcal{S}$. The action space, i.e. the set of possible actions a is denoted \mathcal{A} , and might be limited depending on the current state $\mathcal{A}_s \subseteq \mathcal{A}$, again with $a \in \mathcal{A}$ and furthermore $\mathcal{A} = \bigcup_{s \in \mathcal{S}} \mathcal{A}_s$. The state transition, which is deterministic for now, can be expressed by a state transition function $f : \mathcal{S} \rightarrow \mathcal{S}$. We write $s' = f(s, a)$ meaning that action a applied in state s leads to the new state s' . The set of goal states that are to be reached is denoted \mathcal{S}_G with $\mathcal{S}_G \subset \mathcal{S}$. Sequences of states and actions are denoted with \mathbf{s} and \mathbf{a} .

Discrete Feasible Planning

For discrete feasible planning, the task is to find a sequence of actions \mathbf{a} to reach a goal state $s_G \in \mathcal{S}_G$ from the initial state $s_I \in \mathcal{S}$. In general, there are three ways to search for a solution in the graph: Forward search from the initial state s_I , backward search from goal states s_G or the combination of both in a bidirectional search. In the following, we focus on the forward search methods.

Breadth First The breadth first search algorithm grows the search tree in a way that the depth difference in-between different branches stays minimal. Only when all branches have the same depth k , meaning that all plans with the same number of actions k have been investigated, the depth is increased, meaning that plans with $k + 1$ actions are considered.

Depth First The opposite to breadth first is depth first, where the first branch is fully investigated up to the point from where no new state can be reached anymore, and only then another search branch is opened.

Discrete Optimal Planning

In this section, feasible plans are additionally distinguished regarding their quality. In other words, we are not looking for any action sequence that leads to the goal region, but want the *optimal* one, w.r.t. some criterion. The term $l(s, s')$ denotes the cost that is assigned to every state traversal $s \rightarrow s'$. Due to the deterministic state transition, it can also be expressed as $l(s, a)$. This cost is equivalent to edge cost in the graph representation. Given this traversal cost, we can define *optimal* cost² to reach state s from the initial state s_I as $C^*(s)$. This yields $C^*(s_I) = 0$. Cost that is not (yet) known to be optimal is denoted $C(s)$. The cost computation can be done via $C(s') = C^*(s) + l(s, a)$. Determining $C^*(s')$ is generally not trivial. As further notation, we enumerate the sequences of states, starting with $s_1 = s_I$, such that an arbitrary action sequence (a_1, a_2, \dots, a_K) leads to a state sequence $(s_1, s_2, \dots, s_{K+1})$, where s_{k+1} can be iteratively computed from s_k using f : $s_{k+1} = f(s_k, a_k)$. In discrete optimal planning, the following principle is key to most planning algorithms:

Definition 1 (Bellman’s principle of optimality). According to Bellman, “an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision” [Bel57, p.83].

This facilitates breaking down optimization problems into simpler subproblems. Methods which exploit this principle are called *dynamic programming* methods. In the following, we introduce some basic algorithms. The state space is assumed to be finite.

Fixed-Length Forward Value Iteration A straightforward algorithm to compute the cost to reach states s starting at s_I is to iteratively compute the optimal cost $C_{k+1}^*(s)$ to reach s in $k + 1$ steps starting at the optimal cost $C_k^*(s)$ for k steps. This step being computed for all s in step k before moving on to $k + 1$ induces the limitation to finite state spaces. This dynamic programming method thus computes $C_1^* \rightarrow C_2^* \rightarrow \dots \rightarrow C_K^*$.

² Cost to reach a state s from an initial state s_I is sometimes referred to as *cost-to-come* in order to distinguish it from cost to go to a goal state s_G from the state s , which is then denoted as *cost-to-go*.

General Forward Value Iteration The generalization of fixed-length plans to plans of arbitrary length requires the definition of a terminal action a_T as a special treatment for plans that have reached a goal state already. This terminal action a_T must lead to zero transition cost in goal states $l(s_G, a_T) = 0$. With this, instead of requiring plans with fixed length K , we can consider any plan of length K or less. As an example, the plan (a_1, a_2) is equal to (a_1, a_2, a_T, a_T) . The optimal plan can now be found as soon as the costs become stationary, i.e. as they no longer change with ongoing value iteration stages. For this to occur, the cost must be defined in a way that negative-cost cycles do not exist.³

Dijkstra's Algorithm This algorithm also employs dynamic programming, but its general approach is different: Instead of repeatedly revisiting states, it systematically searches for optimal paths in a graph. In Dijkstra's algorithm, as with breadth first and depth first search, visited states are tracked. More precisely, states whose cost is optimal, i.e. stationary already, are no longer visited. The basic idea behind the algorithm is to always explore the edge which leads to the vertex with minimal cost/distance (starting at the initial vertex). Once the distance to a vertex is computed, it is never updated again. This is why, in contrast to value iteration, where only negative cycles are prohibited, Dijkstra's algorithm requires all costs to be non-negative.

A-Star As an extension of Dijkstra's algorithm, A* uses a heuristic to search more efficiently, i.e. with visiting less states. The heuristic h used in A* is an under-approximation of the cost from an arbitrary state to the goal state, $h(s \rightarrow s_G) \leq C^*(s \rightarrow s_G)$. With this, instead of visiting the state (vertex) with the least cost $C(s_I \rightarrow s)$, the state with the smallest expected total cost $C(s_I \rightarrow s) + h(s \rightarrow s_G)$ is visited. In other words, Dijkstra's algorithm corresponds to A* without heuristic: $h(s \rightarrow s_G) = 0 \forall s$.

³ Otherwise, they would be exploited repeatedly, leading to cost of $-\infty$.

2.1.4 Planning in Continuous State Spaces

With the generalization from planning in discrete state spaces to planning in continuous state spaces, several challenges arise: The number of available states is uncountably infinite, i.e. value iteration or other graph-based algorithms are not straightforward applicable. Moreover, the state space, in this context more often called configuration space, is at the same time multi-dimensional and constrained by (moving) obstacles. In this section, we focus on problems without differential constraints. Here, sampling-based methods for continuous spaces are introduced, which discretize the state space efficiently but are generally not complete. Further, an introduction to non-stationary planning with time discretization is given. Problems with differential constraints and continuous time will be covered in Sec. 2.1.6.

Sampling-Based Planning

In the context of planning, the term *sampling* can be used with two meanings: Sampling as in signal processing means the discretization of a continuous signal. Sampling as in statistics means selecting individual samples from a potentially infinite set. For planning in a continuous but bounded state space, such as a room or a parking lot, the first type of sampling makes the graph-search methods from Sec. 2.1.3 applicable. As this involves building a graph which can then be reused or queried multiple times, LaValle refers to such methods as *multiple-query* methods [LaV06, p.217]. Two relevant alternatives will be explained in the first two paragraphs. Contrary to this two-step approach, *single-query* methods do not precompute a graph. With this, they additionally cover open spaces. These methods are also called *incremental* sampling-based methods, rather referring to the second meaning of sampling. The well-known method RRT along with its variant RRT* will be shortly introduced later on.

Grid-Based Planning Grid-based planning is the most straightforward way to apply existing graph search methods to bounded continuous state spaces. The state space is subdivided into (multi-dimensional) grid cells and subsequently, graph search is applied within the grid. While such grid-based methods cannot

be complete, they can be *resolution complete*, i.e. guarantee to find a solution if it exists and if the grid cells are sufficiently small.

Roadmaps In the concept of so-called roadmaps, which gained attention through the work about probabilistic roadmaps (PRMs) [KSLO96], the graph is created by random sampling (with subsequent collision checking), and attempts to connect the vertices using a local planner, again checking for collisions. For roadmap-based methods, *probabilistic completeness*, which was defined in Sec. 2.1.2, can be reached.

Rapidly Exploring Random Tree (RRT) In contrast to grid-based planning and roadmaps, RRT does not precompute the graph. Furthermore, as the name suggests, the graph is a tree. It is grown by drawing random samples from the configuration space. RRTs are *probabilistically complete*. This is reached as the sample is always connected with its nearest neighbor in the tree, i.e. the probability of expanding a state is proportional to its Voronoi region. In order to prevent many far away but unreachable samples, the new state is commonly chosen to be at a maximum distance from the nearest state. In case it is further away than this distance, the random sample is replaced by the state that is at this maximum distance in the direction of the random sample. This new state is then connected to the tree. RRTs have been introduced for feasible planning and are generally not optimal. This motivated RRT*, explained in the next paragraph.

RRT-Star The suboptimality of RRT is overcome by the extension RRT*: Instead of connecting the new state s_{new} to its nearest neighbor in the tree, all states within a certain distance (respectively, all states within a set S_{near}) around the new state are checked, and it is connected to “the vertex that incurs the minimum accumulated cost up until” this new state s_{new} [KF11]. Further, the new state is extended to those vertices/states in S_{near} that can be reached with smaller cost through s_{new} , which is called *rewiring*. With these additional steps, RRT* is *asymptotically optimal*.

Time Discrete Non-Stationary Planning

For non-stationary planning, the time has to be considered additionally. Thus, the resulting plans of this section will be trajectories instead of paths. Neglecting differential constraints, which will be covered in Sec. 2.1.6, non-stationary planning is particularly interesting if goal or obstacle regions change over time. Generally, these changes are assumed to be predictable and thus known in the planning problem. Approaches to overcome these challenges can be split into methods decoupling path planning and velocity planning, and direct methods, which solve the problem without decomposing it.

Direct Methods Many incremental sampling-based methods can easily be adapted to non-stationary problems. Here it is crucial that, during sampling or during verification of the sample (collision checking), it must be ensured that no trajectories travel backwards in time. Further, traveling at infinite velocities is not meaningful, leading to the requirement of strictly monotonically increasing time for every valid trajectory.

Decoupled Trajectory Planning As the time adds a further dimension to the problem, it increases complexity. This can be reduced again by decoupling path planning and velocity planning. Kant and Zucker proposed a well-known approach to solve such non-stationary problems with a heuristic, called *Path-Velocity Decomposition (PVD)* [KZ86]: PVD is the decomposition into first path planning to avoid stationary obstacles and then velocity planning to avoid moving obstacles. While this heuristic is not generally optimal, it allows for the use of time-efficient algorithms within the subproblems. As the previously presented incremental sampling-based planners are only asymptotically optimal, PVD might even lead to better results given limited runtime. Note that it is also possible to improve the computed path and velocity profile iteratively, for example. The iterative computation is referred to only as *decoupled trajectory planning*, while PVD refers to the previously introduced heuristic.

2.1.5 Planning under Uncertainty

While the previous sections modeled problems in a deterministic way, for both the knowledge about the state and the state transition, this assumption is not always reasonable. Planning under uncertainty, which is also referred to as *decision theoretic planning* or *decision making*, focuses on finding optimal decisions based on models that account for uncertainty. In order to allow for the use of the Bellman equation (in other words: to allow for dynamic programming), the state is commonly chosen such that the Markov assumption holds⁴, i.e. that the probabilistic state transition function only depends on the current state and the current action. The decision process is then called *Markov decision process (MDP)*, it will be further explained in the next section. Its extension to also include uncertain knowledge about the state itself, called *partially observable Markov decision process (POMDP)*, is introduced subsequently. In this section, the state space is again assumed to be discrete. Note that uncertainty introduced by multiple agents and their unknown behavior is only dealt with in Sec. 2.1.7.

Markov Decision Processes (MDP)

As explained above, the formulation of MDPs differs from those of Sec. 2.1.3 in the state transition. Instead of modeling a deterministic transition, probabilistic transitions are considered. The transition probability is described as $\Pr(s'|s, a) = \Pr(s_{i+1} = s' | s_i = s, a_i = a)$, i.e. $\Pr(s'|s, a)$ is the probability that action a applied in state s will lead to a transition to state s' in the next step, where integer i enumerates steps. Also, the cost has to be redefined to account for this probabilistic transition: The term $l_a(s, s')$ denotes the immediate cost for transitioning from s to s' taking action a . In the decision making literature, the cost is commonly turned into a reward $r_a(s, s')$, which is to be maximized for optimality.⁵ Yet, it can be reformulated to $r(s)$, only depending on the current state, without fundamentally changing the problem [RN16, pp.647]. With this probabilistic formulation, an optimal sequence of actions, also called

⁴ At least such that this assumption appears reasonable.

⁵ As suggested by LaValle, this does not change the core problem to any extend, and the transfer can for example be done choosing $r_a(s, s') = -l_a(s, s')$ [LaV06, p.439].

optimal plan, can no longer be determined, as it depends on the actual transition after each action in the sequence. Thus, the problem is reformulated from finding a sequence of actions to finding a sequence of decisions, described by a *policy* $\pi : \mathcal{S} \rightarrow \mathcal{A}$, that determines which action a to choose when in a specific state s . Note that the same policy starting at the same initial state might lead to a different state and action sequence, due to the stochastic nature of the problem. Consequently, policies cannot be compared by their future reward, as the latter is not unique. As Bellman wrote, “it is now on the whole meaningless to speak of maximizing *the* return⁶” [Bel57, p.85]. Rather, some *expected* value $\mathbf{E}(\cdot)$ has to be considered:

Definition 2 (Utility of a policy, expected utility). The utility U^π of a policy π in a state is defined as the expected utility of executing the policy:

$$U^\pi(s) = \mathbf{E}(U(\tilde{\mathbf{s}})) \quad (2.1)$$

i.e. of the random sequence $\tilde{\mathbf{s}} = (s, s', \dots)$ that arises from repeatedly applying actions according to policy π starting in the initial state s .

This utility can then be maximized. However, it is noteworthy that “with a finite horizon, the optimal action in a given state could change over time”, i.e. the optimal policy is not stationary in this case [RN16, pp.648-649]. As an example, consider you have to reach a train departing at a specific time. At a point 200 m from the train station, you would probably walk if the departing time was in 5 min, but rather run if it was in 2 min. For further details, we refer to the book of Russell and Norvig [RN16]. Thus, we focus on problems with an infinite horizon, where the optimal policy is not obviously non-stationary. Under the reasonable demand of stationarity for optimal policies⁷, the utility definition $U(\tilde{\mathbf{s}})$ for infinite sequences of states $\tilde{\mathbf{s}} = (s_0, s_1, s_2, \dots)$ is limited to two options, according to Russell and Norvig [RN16, p.649]:


1. Additive rewards: $R_{\text{add}}(\tilde{\mathbf{s}}) = \lim_{N \rightarrow \infty} \sum_{i=0}^N r(s_i)$.
2. Discounted rewards: $R_{\text{disc}}(\tilde{\mathbf{s}}) = \lim_{N \rightarrow \infty} \sum_{i=0}^N \gamma^i r(s_i)$ with discount factor γ , $0 \leq \gamma < 1$.

⁶ Return is a different word for reward.

⁷ That is, the action selection does not depend on the current step/time i .

While the discounted reward is bounded for $\gamma < 1$ if the single rewards are bounded, the additive reward is generally not bounded⁸. If the additive reward is not bounded, we would need to compare infinite rewards, which is not meaningful (other than saying that $+\infty$ is better than $-\infty$). One solution to circumvent this, is to consider the average reward $R_{\text{avg}}(\xi) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=0}^{N-1} r(s_i)$ as utility, which is also bounded given that the single rewards are bounded. Another solution is to employ a limited but receding horizon. For every decision, only a finite number N_H of future steps are considered as utility: $R_{\text{RH}}(\xi) = \sum_{i=0}^{N_H} r(s_i)$, which is bounded by $(N_H + 1) \cdot r_{\text{max}}$. This finite receding horizon is not to be confused with the previously discussed finite horizon. The optimal policy for this problem will still be stationary, as the horizon is always the same and therefore does not depend on i . However, similar to the discounted reward, the optimal policy for this problem might differ from the actual optimal policy, as illustrated in the example below.

Example 2.1.1. Consider a problem with 3 states: A with reward $r_A = r(A) = 1$, B with reward $r_B = r(B) = 0$ and C with reward $r_C = r(C) = 2$. The available actions are to stay in the state or to change from A to B or vice versa and from B to C and vice versa, where the state transitions themselves do not affect the reward. The transitions are deterministic.



The optimal policy is: “When in A move to B, when in B move to C and when in C stay there” (denoted π_2), leading to an average reward of 2 for

⁸ The additive reward is bounded if a terminal state exists and the policy is guaranteed to find one (then called *proper policy*).

$N \rightarrow \infty$. The optimal policy for a discounted reward with $\gamma < \frac{1}{2}$ is “Stay in A when in A” (denoted π_1), as

$$R^{\pi_1}(s_0 = A) = \lim_{N \rightarrow \infty} \sum_{i=0}^N \gamma^i r_A = r_A \frac{1}{1 - \gamma} = \frac{r_A}{1 - \gamma}$$

and

$$\begin{aligned} R^{\pi_2}(s_0 = A) &= r_A + 0 + \left(\lim_{N \rightarrow \infty} \sum_{i=2}^N \gamma^i r_C \right) \\ &= r_A + 0 + \left(\frac{r_C}{1 - \gamma} - r_C - \gamma r_C \right) \\ &= \frac{r_A - \gamma r_A + \gamma^2 r_C}{1 - \gamma} \end{aligned}$$

lead to

$$\frac{R^{\pi_1}(s_0 = A)}{R^{\pi_2}(s_0 = A)} = \frac{r_A}{r_A - \gamma r_A + \gamma^2 r_C}.$$

Thus, policy π_1 is preferred if

$$\begin{aligned} \frac{R^{\pi_1}(s_0 = A)}{R^{\pi_2}(s_0 = A)} &= \frac{r_A}{r_A - \gamma r_A + \gamma^2 r_C} > 1 \\ &\Leftrightarrow \gamma^2 r_C - \gamma r_A < 0 \\ &\Leftrightarrow \frac{r_A}{\gamma r_C} > 1 \end{aligned}$$

which is true for all $0 < \gamma < \frac{1}{2}$ with $r_A = 1$ and $r_C = 2$.

The policy π_1 , however, leads to a lower overall average reward, as the average reward is 1 for $N \rightarrow \infty$ and $s_0 = A$. The same suboptimal policy is generated with a finite receding horizon of $N_H = 1$.

While for many applications, neither $N_H = 1$ nor $\gamma < \frac{1}{2}$ are meaningful choices, the example shows that with these alternative reward formulations, “optimal” policies might lead to suboptimal results w.r.t. the original cumulative reward.

The definition of the expected utility leads to the following principles and formulations:

Definition 3 (Principle of maximum expected utility (MEU)). The principle of maximum expected utility expresses that a rational agent should maximize its expected utility when choosing actions:⁹

$$a_{\text{select}} = \arg \max_{a \in \mathcal{A}} \mathbf{E}(U(a)). \quad (2.2)$$

Definition 4 (The optimal policy). According to our previous definitions, determining the optimal policy π^* can now be formulated as

$$\pi^* = \arg \max_{\pi} U^{\pi}(s). \quad (2.3)$$

With the above definitions, we can use Bellman's principle of optimality to state that

$$\pi^*(s) = \arg \max_{a \in \mathcal{A}(s)} \sum_{s'} \Pr(s'|s, a) U^{\pi^*}(s') \quad (2.4)$$

which is key to solving MDPs.

In the following, we give a short introduction to the application of value iteration in the context of MDPs, to policy iteration, and to reinforcement learning.

Value Iteration The value iteration algorithm, as introduced in Sec. 2.1.3, is also applicable to MDPs. As the probabilistic transition prohibits the calculation of some optimal reward that will be reached with certainty, the algorithm is adapted to iteratively compute an approximation for the utility $\tilde{U}(s)$ of states, starting at their reward value $\tilde{U}_0(s) = r(s)$. In step k , this update works as follows:

$$\tilde{U}_{k+1}(s) \leftarrow r(s) + \max_{a \in \mathcal{A}(s)} \sum_{s'} \Pr(s'|s, a) \tilde{U}_k(s'). \quad (2.5)$$

Russell and Norvig refer to this step as *Bellman Update* [RN16, p.652]. As with the deterministic equivalent, this step being computed for all s in step k

⁹ Note that humans do not always act accordingly, for example when they gamble [RN16, pp.619-620].

before moving on to $k + 1$ induces the limitation to finite state spaces. This step is repeatedly performed until \tilde{U} converges. Given the optimal utility function U^* , the policy can be computed using eq. 2.4, by substituting $U^{\pi^*} = U^*$.

Policy Iteration Another algorithm is to iteratively compute the best policy directly. Starting point is any policy π_0 . This policy is then repeatedly updated in a two-step process: First, evaluate the policy π_k by computing $U_k = U^{\pi_k}$. This can be done by either computing the exact solution of $U_k(s) = r(s) + \sum_{s'} \Pr(s'|s, \pi_k(s))U_k(s')$, or by using value iteration, but now with the fixed policy π_k . Second, improve the policy to π_{k+1} by evaluating eq. 2.4 with the new utility U_k . The repetition is stopped once the policy evaluation does not lead to a change in the utility and thus the improvement does not change the policy anymore.

Reinforcement Learning For the previous algorithms, we assumed $\Pr(s'|s, a)$ to be known. In reinforcement learning (RL), the idea is to combine *learning* $\Pr(s'|s, a)$ and computing the optimal policy. LaValle also refers to this concept as *simulation-based dynamic programming* or *simulation-based planning* [LaV06, p.528], since the required samples are commonly generated using Monte Carlo simulation. Further, even the state space and the reward of unvisited states do not have to be known. Only the set of available actions in a visited state $\mathcal{A}(s)$ is required. While there are plenty of RL algorithms explained in literature, we focus on the basic algorithms that are similar to value iteration and policy iteration and use the model-free *Q-learning*. The *Q-function* $Q(s, a)$ is a state-action utility function, denoting the utility of taking action a in state s . Related to the state utility function U , $U(s) = \max_a Q(s, a)$ must hold. However, Q is not an inefficient way to store U , but it is a means to model-free learning of the right policy, as no transition model $\Pr(s'|s, a)$ is needed.¹⁰ The Bellman equation for state-action utilities is

$$Q^*(s, a) = r(s) + \sum_{s'} \Pr(s'|s, a) \max_{a' \in \mathcal{A}(s')} Q(s', a'). \quad (2.6)$$

¹⁰ It is not beneficial if $\Pr(s'|s, a)$ and $r(s)$ are known.

While this approach again includes $\Pr(s'|s, a)$, the latter can be circumvented using the concept of temporal differences (TD), leading to

$$Q(s, a) \leftarrow (1 - \alpha)Q(s, a) + \alpha(r(s) + \max_{a' \in \mathcal{A}(s')} Q(s', a')) \quad (2.7)$$

with learning or convergence rate $\alpha \leq 1$. This is, $Q(s, a)$ is updated after a was taken in s which leads to s' , without the need to know or estimate the transition model $\Pr(s'|s, a)$. The determination of which action a to choose is a trade-off between exploration and exploitation. A well-known method to tackle this trade-off is the ϵ -greedy algorithm: With a defined probability ϵ , the action is chosen derived from the Q -values, otherwise a random action is chosen. For rather limited state and action spaces, Q can be stored in a table. For large problems, a Q -function can be approximated using regular function approximation or artificial neural networks.

Partially Observable Markov Decision Processes (POMDP)

Partially observable Markov decision processes are the generalization of MDPs to uncertainty in the state. The name stems from the fact that the state is not fully observable. Instead, the agent receives *observations* $o \in \Omega$, where Ω is the space of possible observations, and states are related with observations via the conditional observation probability $\Pr(o|s')$. To solve POMDPs, a *belief space* is commonly introduced. The *belief state* b is a probability distribution over all possible states where $b_i(s) = \Pr(s_i = s)$ is the probability of being in s at step i . The belief state update depends on the current belief $b(s)$, the action a , and the next observation o

$$b'(s') = \alpha \Pr(o|s') \sum_s (\Pr(s'|s, a)b(s)) \quad (2.8)$$

with an additional normalizing factor α that ensures $\sum_s b_i(s) = 1$.

The key difference of a POMDP to an MDP is that the *optimal* action depends only on the current belief state b and not on the actual state s . Thus, the optimal policy maps from belief states to actions $\pi^* : \mathcal{B} \rightarrow \mathcal{A}$. However, POMDPs can be reformulated as MDPs in the continuous belief state space \mathcal{B} . This is also, why the task of optimally solving POMDPs is sometimes referred to as *belief state planning* or *planning in belief space*. For this reformulation, we

need a belief space transition model $\Pr(b'|a, b)$ and a reward function for belief states $\rho(b)$.

The transition model is

$$\Pr(b'|a, b) = \sum_o \left(\Pr(b'|o, a, b) \Pr(o|a, b) \right). \quad (2.9)$$

With the probability that we observe o when taking action a in belief state b

$$\Pr(o|a, b) = \sum_{s'} \left(\Pr(o|s') \Pr(s'|a, b) \right) \quad (2.10)$$

$$= \sum_{s'} \left(\Pr(o|s') \sum_s \Pr(s'|s, a) b(s) \right) \quad (2.11)$$

we obtain

$$\Pr(b'|a, b) = \sum_o \left(\Pr(b'|o, a, b) \sum_{s'} \left(\Pr(o|s') \sum_s \Pr(s'|s, a) b(s) \right) \right) \quad (2.12)$$

where $\Pr(b'|o, a, b)$ is 1 if eq. 2.8 is fulfilled and 0 otherwise.

The reward function $\rho(b)$ is

$$\rho(b) = \sum_s b(s) r(s). \quad (2.13)$$

The Bellman equation for belief states is then

$$U^*(b) = \max_{a \in \mathcal{A}(b)} \left[\rho(b) + \gamma \sum_{b' \in \mathcal{B}} \Pr(b'|b, a) U^*(b') \right] \quad (2.14)$$

regarding the discounted cumulative reward model for the utility.

An optimal policy for this observable MDP in the belief space is also optimal for the underlying POMDP [RN16, p.660]. However, while the original POMDP was on a finite physical state space, the state space of the belief MDP is continuous. Thus, the previously introduced value iteration for finite spaces is no longer applicable here. Apart from toy examples, solving POMDPs was




Figure 2.1: Exemplary belief tree with two possible actions a_1, a_2 and two possible observations o_1, o_2 . The white circles denote belief states, which are updated on new observations.

considered intractable for a long time, amongst others due to the fact that the belief space is continuous and $(|S| - 1)$ -dimensional.

One algorithm to solve large POMDPs was presented by Silver and Veness [SV10]. It is based on Monte Carlo Tree Search (MCTS): Possible future evolutions of the belief are found using a particle filter, where the particles are updated based on sampled observations. The tree structure that is built during this search is called *belief tree* (cf. Fig. 2.1). This concept is used in related work and will be revisited in Chapter 3. One strength of employing a POMDP formulation is the ability to choose *information gathering actions* to actively reduce the uncertainty in the belief. A method to solve POMDPs approximately is called *QMDP*, it is based on the assumption that the state will become fully observable in the next time step [Koc15, p.144]. With this assumption, the state-action function (known from the concept of Q-learning) of the underlying MDP can be used to determine the approximately optimal action. For more algorithms to solve POMDPs, the interested reader is referred to [Koc15], [KLC98] and [RN16].

2.1.6 Planning under Differential Constraints

In this section, we go back to the deterministic world. In Sec. 2.1.3, the state space was discrete, and instead of considering the time, we rather focused on the order of actions in a decision sequence. In Sec. 2.1.4, the state space was continuous, but it was assumed that neighboring configurations in obstacle-free regions could easily be connected, i.e. a path between them was easy to determine. In this section, we focus on problem formulations with differential constraints, which can be seen as an extension of Sec. 2.1.4 to the continuous time domain. This allows for example to consider kinematics, such as that cars cannot move sideways, or dynamics, such as that velocities and accelerations are generally limited.¹¹ In this context, the introduction of a *phase space* is common, which, in addition to the *configuration space* of a robot, also contains derivatives, such as velocities and accelerations. The state transition equation $s_{i+1} = f(s_i, a_i)$ is replaced by the differential equation $\dot{s} = f(s, a)$, which can also include higher order derivatives of s . The theory behind finding optimal plans for problems with differential constraints is also known as *dynamic optimization* or *optimal control*.

As in Sec. 2.1.4, solution methods to optimal control problems can be split into methods operating directly on the trajectory and methods that decouple path planning from velocity planning. In the following, we focus on the non-decoupled methods, as the idea of decoupled planning has already been introduced in Sec. 2.1.4. Yet, some direct methods from this section can also be applied to either path or velocity planning.

The methods for trajectory planning can be split into analytic methods for simple cases, numerical methods from optimal control and enhanced sampling-based methods from Sec. 2.1.4.

Analytical Solutions

The solution of trajectory planning in both continuous space and time is a function $s : \mathbb{R} \rightarrow \mathcal{S}$. In order to define optimality criteria for these trajectories,

¹¹ It is also possible to model this with limited actions per state, to circumvent formulations using differential constraints.

we need functions of functions, called *functionals*. Extrema of functionals can be found using the *calculus of variations*, more specifically using the *Euler-Lagrange equation*. Using this equation, it can for example be shown that jerk-optimal trajectories in free space are fifth-order polynomials [THNS89]. The extension of this calculus to include a system equation that relates input to state is based on the method of *Lagrange multipliers*, that is known from static optimization, i.e. optimization of functions, not functionals. The functional of which the extremum is to be found is then called the *Hamiltonian*¹². The optimal control of a linear system with a quadratic cost functional, i.e. the optimum of the Hamiltonian in this case, can be found solving the *Riccati differential equation*. The respective controller is called Riccati controller or *linear-quadratic regulator (LQR)*. Conditions for the optimal solution to the Hamiltonian under input inequality constraints are known from *Pontryagin's maximum principle* [Pon87, p.17]. While this theory allows for analytic solutions in some cases, it is in practice often used in combination with numerical methods which facilitate at least the approximation of the optimal solution.

Numerical Methods

Numerical methods solve the previously described problem of optimal continuous trajectory planning by approximation. They can be distinguished depending on where this approximation is used. Most significant is the distinction between methods that numerically solve the optimality criteria derived in the previous section (called *indirect* methods), and those that reduce the problem to a static optimization problem (called *direct* methods). A known indirect method is the numerical solution of the Riccati differential equation. For a linear system with a quadratic cost function and linear constraints, the direct method leads to a quadratic program which can be solved very efficiently. In presence of non-linear constraints, the problem can be solved by *sequential quadratic programming (SQP)*, i.e. iteratively linearizing the constraints around the current solution and then refining the solution by quadratic programming. The extension of quadratic programming to combinatorial problems with multiple optima can be realized using *mixed-integer quadratic programming (MIQP)*.

¹² For static optimization it was the *Lagrangian*.

Sampling-Based Methods

The extension of the sampling-based methods of Sec. 2.1.4 to facilitate the incorporation of differential constraints is possible and common in robotics. One option is to discretize the state space into a lattice for subsequent graph search in this *state lattice* [PK05]. Such approaches, however, can reach resolution completeness at most. The integration of dynamically changing constraints, such as dynamic obstacles, is also possible [ZS09]. Another option is the integration of non-holonomic motion primitives into the incremental RRT-based planners, which corresponds to discretized actions. To be able to perform rewiring (cf. Sec. 2.1.4), the challenge is to have a fast local planner to determine feasibility and eventually connect the closest samples in the phase space. For path planning for wheeled robots, this can be realized using *Dubin's curves* [Dub57], for example. Generally, results from sampling-based planning approaches can be smoothed by a local numerical planning algorithm, for example a gradient-based algorithm.

2.1.7 Multi-Agent Planning

The previous sections all had in common, that one agent or robot was considered, while the environment was either considered known, also in the future, or at least the probability over its evolution was known. Sometimes, however, this assumption is not valid or at least questionable, such as in the presence of other agents, robots or also in the presence of humans. In such situations, the agent that is to be controlled is referred to as the *ego* agent. The study of interaction among decision making agents is known as *game theory*. Russell and Norvig state that the multi-agent perspective should be taken if the other agents also try to maximize a reward, and the latter is affected by the ego agent's behavior [RN16, p.43]. Depending on the interdependence of the rewards, such a multi-agent setting can be *competitive*, i.e. the agents have contradicting goals, *cooperative*, i.e. the agents have a common goal, or a combination of both.

In case every agent has its own strategy, John Nash proved that there is at least one *equilibrium* in every game, named *Nash equilibrium* after him, i.e. a combination of policies in which no agent can benefit from switching strategies, given that the others stick with their policy [RN16, p.669]. Often, cooperation

based on mutual trust could lead to superior results. However, the Nash equilibrium should be the choice of a rational agent.

The task can also be to control several agents or even all agents. Here, the planning can be centralized, i.e. one plan is computed that all (controlled) agents have to follow, or decentralized, i.e. agents plan on their own but might be able to communicate with each other and coordinate their actions. Centralized planning for multiple agents is also referred to as *multi-body planning*, which can be considered a single agent problem [RN16, p.425]. For the case of multiple agents that choose their actions individually but share the same cost function and each have perfect state knowledge, Boutilier uses the term *multi-agent Markov decision process (MMDP)* [Bou96]. In case of uncertainty in the state, the problem is called *decentralized partially observable Markov decision process (Dec-POMDP)* [Koc15, p.159].

2.1.8 Replanning

The previous sections dealt with finding single plans or policies. As policies per definition yield an action that is to be taken for every possible state, an unexpected transition to another state does not affect the optimal policy, given that it does not violate the transition model and thus the preliminary for the policy to be found optimal. The open-loop control of a plan, however, easily leads to deviating from it, for example due to disturbance. In addition to that, the environment might be hardly predictable, and wrongly predicted changes can also lead to plans becoming invalid or at least suboptimal.

The problem of disturbance and control errors can be solved with an underlying closed-loop controller that repeatedly tunes the control signal such that the previously computed plan is pursued [LaV06, p.23]. The problem of unpredictability is commonly tackled with replanning. Being aware of the fact that replanning is necessary, this can be further exploited during planning, for example with the following approach.

Model Predictive Control (MPC)

The idea behind model predictive control is to iteratively consider a rather short finite but receding horizon for problems with a long horizon and non-linear dynamics. Therefore, it is also called *receding horizon control*. The plan, however, is only pursued for a small part of the considered horizon. This concept certainly does not guarantee optimality. For non-linear systems, however, it can lead to results that are superior to optimally solving the linearized problem, given that the chosen horizon is large enough. In robot motion planning, a similar concept is applied to deal with unpredictable future constraints, even though the system dynamics themselves are often linearized. The approach is also used in multi-agent problems. If the game theoretic view is intractable, due to complexity or unknown desires of others, they are treated as obstacles rather than agents. Wrong assumptions about their future movement are corrected from new sensor input during replanning.

2.2 Related Work on Motion Planning for Automated Driving

Having reviewed the fundamentals of motion planning and the related areas decision making and control, we now focus on motion planning for automated driving. In general, scenarios that have to be mastered by fully automated vehicles can be split into driving in structured and unstructured environments [Wer11, p.23]. Driving in unstructured or little structured environments is commonly at low velocities, for example reaching a parking lot in a car park. Here, graph search-based approaches are dominating and feasibility is already a challenging goal, due to narrow gaps and the non-holonomic constraints. Motion planning for this type of scenarios is not the focus of this thesis and thus not further considered. Motion planning in regular on-road traffic largely differs from the latter. It takes place in highly structured environments and at higher velocities. Here, comfort and efficiency come into focus, and safety becomes more challenging.

In order to depict the context and integration of motion planning for automated driving, we start by giving a brief insight into system architectures of automated vehicles in Sec. 2.2.1. Subsequently, we review existing approaches to motion

planning for automated driving in Sec. 2.2.2 and Sec. 2.2.3. In Sec. 2.2.4, safety approaches are introduced.

2.2.1 System Architectures of Automated Vehicles

While the design of system architectures for automated driving with an increased need for redundancy and safety is a distinct field of study, we only briefly introduce it in the following, to put the presented motion planning approaches in context. Afterwards, we shortly explain those modules that are closest related to the motion planning module.

Overview Solving the motion planning problem discussed in this work is of course no end in itself. For the resulting motion plan to be pursued, a closed-loop trajectory controller determines the input to the actuators. In order to be able to compute a plan at all, the environment must be perceived. From this data, information about other objects and their state but also about occlusions must be extracted. To gain a better understanding of the scene, this information and the ego vehicle's position are commonly set in relation to an existing high definition map of the road network. Given this information and a desired destination, the route can be planned using off-the-shelf navigation/routing approaches. While there are highly interactive scenarios even in structured road traffic, there are also many scenarios with clear precedence, such that a separation into first prediction and then ego motion planning is applicable. Thus, as already outlined in Sec. 1.2, we assume a prediction that is independent of the ego motion as input to the motion planning module. In this case, the prediction is said to be performed *upstream* of motion planning. An overview of the modules is depicted in Fig. 2.2. The interfaces are intentionally excluded, due to the large variety of possible implementations. Exemplary implementations are presented in [TKZS16]. In the following, we shortly introduce those modules that are most relevant for the motion planning module.

Navigation In the context of this work, the task of a navigation module is to specify the route towards the goal on a lane level. In general, this problem can be considered solved, with many mobile applications and dedicated GPS

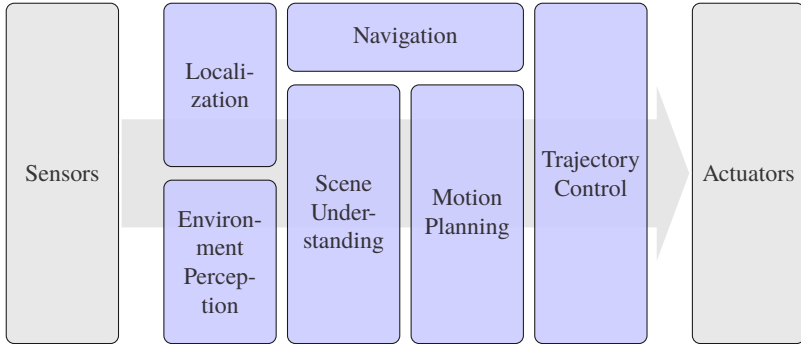


Figure 2.2: Exemplary system architecture of an automated vehicle. Based on sensor data, actions are planned and communicated to the actuators. Commonly used modules are depicted in blue.

navigation devices available. Algorithms to solve such problems are commonly based on graph-search with sophisticated heuristics.

Prediction The prediction module is often considered to be part of the *scene understanding* module which combines sensor information with the map and some *world knowledge* to infer more information about the current situation and also about the possible future evolution. While an independent prediction and subsequent ego motion planning hinder interaction, this approach is well suited for scenarios with clear precedence.

Control In automated driving, it is common to separate the trajectory generation, which is the topic of this work and commonly called *motion planning*, from its tracking, called *low-level control* or *trajectory control*. This allows for a more simple motion model in the motion planning problem formulation. Further, perturbation due to wind or rough road surface can be corrected immediately by a high frequency low-level controller. Note, however, that consistency during replanning has to be ensured in order to facilitate smooth control, i.e. the part of the trajectory that is currently controlled should not be altered.

2.2.2 Deterministic and Reactive Motion Planning

Events that attracted much attention in the field of motion planning for automated vehicles were the DARPA¹³ Grand Challenges. While the first challenge in 2004 could not be completed by any team, five teams could finish the 2005 Grand Challenge, with the winner being “Stanley” from Stanford [TMD⁺06]. The motion planner of Stanley makes use of the fact that dynamic objects did not have to be considered: The concept is based on the path-velocity decomposition approach [KZ86], as explained in Sec. 2.1.4. As a rough global path was predefined by the organizers, only a local planner is needed. To avoid obstacles, several lateral offsets to the offline preprocessed base path are considered. The path is then chosen by evaluating a cost function, accounting, amongst others, for dynamic constraints, obstacles and the defined road border. Replanning is performed with 10 Hz, deviations of the vehicle from the planned path are entirely left to the controller. The velocity is controlled subsequently, based on the map, the terrain and the vehicle dynamics.

In the DARPA Urban Challenge in 2007, “for the first time, 11 full-size autonomous vehicles interacted with each other and human-driven vehicles on a closed course” [FTO⁺08]. Yet, all but the winning team still used modified path planning approaches [Wer11, p.4]. In the following, we introduce motion planning approaches that were employed during this challenge. For BOSS [UAB⁺08], an approach similar to the one of Stanley is used for on-road driving, with the distinction that the velocity is incorporated into planning and not only added subsequently. As explained in more detail in [FHL08], they employ the approach of Howard and Kelly [HK07]: Initialized from a table of motion primitives, the motion is predicted using a non-linear model (MPC), and then optimized using a linearization of the latter prediction. Since the goal is to be reached by following a road network consisting of lanes, the target states are chosen in vicinity of the lane center of the target lane. In order to account for planning and control delays, the common start state of the generated trajectories is chosen by predicting the vehicle state at the time when the new trajectory will be executed, based on the trajectories from previous planning steps. The velocity is chosen from a set of profiles, with the goal to maximize speed while not violating the motion model (too high speed combined with

¹³ US Defense Advanced Research Projects Agency

high curvature) and braking smoothly for stop lines. From the generated set, the best trajectory is chosen, based on metrics such as distance to obstacles and smoothness.

Team AnnieWAY uses a hierarchical state machine [KZP⁺08]. While strongly relying on GPS and the given road network definition files, smoothed with splines, if possible, a tentacle-based approach is used for low speed collision avoidance or in case of poor GPS information [vHHH⁺08]. Intersection scenarios, i.e. scenarios in which AnnieWAY potentially comes into conflict with others, striving for the same space, are treated in a special way with the assumption of fixed and known paths and known velocity profiles of others.

The team from MIT also uses a hierarchical planning approach for their vehicle Talos, with maneuvers and precedence at intersections being decided at a higher level, called navigator, which then passes a desired short term goal to the RRT-based motion planner [LHT⁺08]. The authors vividly describe this concept saying “the goal acts as a ‘carrot’ to motivate the motion planner” [LHT⁺08, Sec. 5.1]. Modifications to the RRT include closed-loop simulation of the inputs using the controller, inherently ensuring dynamically feasible trajectories, and biased sampling, for example based on the lane.

The Stanford team [MBB⁺08] apparently uses the approach of lateral offsets to precomputed paths from the 2005 Grand Challenge. Since there was not a single predefined global path, but a mission goal given by the organizers, a global path planner based on dynamic programming is added to the approach. This global planner operates on a discrete version of the map. Depending on the current position in the road network, the options for the motion planner include several *principle paths*, i.e. distinct route choices, and also several paths with different lateral offset from the base path per principle path. The best path is then chosen based on the cost of the local path with a distinct lateral offset, determined by the local planner, and the cost from the end of this local path to the goal, determined by the global planner. Similar to the approaches of Team MIT and Team AnnieWAY, intersections and merges are treated in a special way: The respective zones are derived from the map in advance and particular states of the finite state machine are used to handle precedence for those zones.

A combined optimization of steering and longitudinal motion based on a prediction of dynamic objects, however, was not employed by any team of

the DARPA Urban Challenge [Wer11, p.4]. In [ZS09], Ziegler and Stiller present such an approach. It is an enhancement of the state lattice concept to include the time domain. In order to limit the effect of the increased dimensionality, they reparametrize the Cartesian space through a lane-based coordinate system, being able to restrict the search to the interesting region effectively. Subsequently, the state lattice created from jerk optimal quintic polynomials is searched exhaustively, because of the lack of good heuristics for this problem formulation.

Werling et al. [WZKT10] propose an optimal control approach to the problem. Reparametrizing the Cartesian space, they propose to decouple lateral and longitudinal motion¹⁴ and to create a set of jerk-optimal trajectories in each dimension. Subsequently, infeasible trajectories are removed per dimension, before they are combined to full trajectories. In the longitudinal direction, the desired value for the computation of the jerk-optimal course is computed multiple times for different assumptions, such as free driving or following a vehicle ahead. As a heuristic to choose the optimal trajectory, they choose the best solution of the created set that does not violate any constraint. Thus, given that the best solution is not invalidated by constraints, they claim to be temporally consistent. Details of the approach are also explained in [Wer11, Chap. 4].

McNoughton et al. [MUDL11] present an approach similar to the one of Ziegler and Stiller [ZS09], also based on a spatiotemporal state lattice. They agree with Ziegler and Stiller that an exhaustive search is preferable, since the worst case scenario is the crucial one for runtime considerations. For high velocities such as on a highway, stopping and further evaluating possible trajectories, which is fine for parking maneuvers, is not an option. In contrast to Ziegler and Stiller, they do not restrict the velocities to discrete values. Instead, they allow for a range of velocities (and thus also a range of time) at which a certain vertex in the lattice can be reached, approximately. By doing so, they aim to facilitate driving at a particular velocity in a gap between two vehicles, which travel at a velocity that might differ significantly from the closest velocity in the precomputed lattice.

¹⁴ Not to be confused with path-velocity decomposition. Here, both the lateral and the longitudinal motion are parametrized in time, such that a change in either motion potentially changes the path.

In contrast to the previously described trajectory planning approaches that use sampling, Ziegler et al. [ZBDS14] exploit the strong lane structure of urban traffic to employ a local, continuous numerical optimization. Thus, errors from discretization into lattices can be avoided. As explained in Sec. 2.1.6, the problem of minimizing a cost functional is reformulated to minimizing a cost function by employing the method of finite differences. Since the cost function is chosen to be quadratic, this non-linear optimization problem with non-linear inequality constraints can be solved using sequential quadratic programming (SQP). The decision of whether to pass an obstacle on the left or on the right is taken in a preprocessing step, yielding only one left and one right bound as constraints to the local planner. In case of a blocked lane, a stop line is added additionally. Consistency during replanning is ensured by *binding* or *pinning*, i.e. not optimizing, points of the trajectory, such that the optimized part of the trajectory is C_2 -continuous to the trajectory that is currently controlled. Further details of this approach and its application during the Bertha Benz Memorial Drive can also be found in [Zie17], [ZBS⁺14]. Note, however, that for some maneuvers or situations, such as stopping or follow driving, the series ACC was used for longitudinal control. Further, the combinatorial aspects which have not been in focus of the previous approach are targeted in [BTZS15].

Schlechtriemen et al. [SWK16] propose a combination of the combinatorial consideration of [BTZS15] and the jerk-optimal trajectory generation in lane-based coordinates [WZKT10]. Assuming a given prediction of all dynamic objects of interest, the areas of interest for certain maneuvers are identified and a set of states within those areas is sampled. Consider for example a lane change, where areas in both lanes must be free for the lane change to be performed. The traversal of those areas as well as the trajectory towards and from there are computed using the concatenation of jerk-optimal quintic polynomials. Finally, the best trajectory is chosen as the one with lowest jerk that does not violate any constraint.

Zhan et al. [ZCC⁺17] argue that the longitudinal motion is more decisive than the lateral motion in structured road traffic. Thus, they propose to first determine a rough longitudinal reference velocity using A^* . They build a topological map from the reference path and the motion prediction of others. With this information, they determine whether boundary constraints are active (such as objects in the ego lane), potentially active (such as, in case of a lane change, objects in either lane), or inactive (such as objects in lanes that the

ego vehicle certainly will not enter). On a shorter horizon, the previously created velocity profile is then smoothed using quadratic programming. Next, the lateral motion is investigated using A* again in a grid of several lateral displacements from the reference path, similar to [TMD⁺06]. The result of the smoothed longitudinal and the A* lateral motion is again smoothed using quadratic programming.

Graf et al. [GSZD18] modify the approach of Ziegler et al. [ZBDS14] to overcome the necessity to use ACC for car following or stopping. Instead of minimizing the deviation from the lane center and a desired speed, they propose to precompute a trajectory from a driver model such as the intelligent driver model (IDM) [THH00], and subsequently minimize the deviation from this trajectory along with minimizing jerk and acceleration.

Further approaches can be found in the comprehensive surveys of motion planning approaches for automated driving [KQCD15], [GPMN16] and [PČY⁺16].

The various approaches presented in this section presage the interest to the field of motion planning for automated driving. Most approaches are based on sampling, sometimes combined with optimal control and simple cost functions. More recent approaches try to minimize more complex, but commonly quadratic cost functions, using (sequential) quadratic programming. Interaction with other traffic participants and the probabilistic nature of predictions, however, is at most addressed by frequent replanning. Thus, maneuvers that require cooperation, such as merging into dense traffic, could at best succeed by coincidence. Also, safety is only considered w.r.t some assumed predictions, but not w.r.t. unforeseen events such as emergency brakings of vehicles ahead.

In the following, approaches that consider uncertainties and interaction are presented.

2.2.3 Probabilistic and Interactive Motion Planning

Deterministic motion planning for automated vehicles is already challenging, especially due to the need of real-time capability. More recent approaches additionally account for the fact that the prediction of other traffic participants and even their current state is subject to uncertainties. Furthermore, this uncertainty is not necessarily independent of the ego vehicle's behavior. Thus,

the approaches can be split into either *deterministic*, considering the state as a function of time $s(t)$ per object, or *probabilistic*, considering multiple functions or a probability distribution of states over time per object. In the probabilistic case, also occlusions and limited sensor range can be modeled. Regarding the consideration of interaction, the approaches can be split into *reactive*, i.e. the future states of other vehicles are independent of the ego actions, or *interactive*, i.e. they depend on the ego action. In the interactive case, courteous and cooperative behavior can be considered.

We start by reviewing probabilistic but reactive approaches, before introducing deterministic interactive and eventually probabilistic and interactive approaches. Lastly, a brief insight into learning-based methods is given. As the interest in this field has drastically grown in recent years, the following summary only briefly overviews different ideas and approaches to modeling and solving the problem of probabilistic and interactive motion planning for automated driving.

Probabilistic Reactive Approaches

Xu et al. [XPWD14] model the sensing and control errors of the ego vehicle. Trajectories of other traffic participants are predicted using a local planner with constant velocity for all but the vehicle to be predicted. Uncertainties for these trajectories are computed using Gaussian error propagation. As a result, larger safety margins are planned for encounters with far away vehicles, since the prediction uncertainty grows over time. On the other hand, one could argue that, at the time of the actual encounter, the uncertainty stemming from control errors is comparably small. Multi-modal uncertainties arising from several maneuver options are not considered.

Zhan et al. [ZLCT16] focus on intersections. They state that “taking an immediate action does not mean [...] to immediately make a final decision among possible future actions”. Similar to Damerow and Eggert [DE15], they consider multiple future trajectories which are equal at least until the subsequently planned trajectory becomes effective. Focusing on binary yield/pass decisions, they employ the SQP approach of Ziegler et al. [ZBDS14]: They plan two deterministic trajectories that are identical until the subsequent replanning step, such that both options remain available. In the subsequent step, certainty about

the decision is assumed, as in the QMDP approach (cf. Sec. 2.1.5). The cost of the diverging trajectory segment is weighted according to the yield/pass probabilities at the time of planning. A similar approach is pursued by Tas et al. [THS18].

Bouton et al. [BCK17] also propose an online POMDP algorithm with focus on the uncertainty in the longitudinal behavior of other vehicles at intersections. An interacting multiple model filter is used to estimate whether a vehicle follows a constant velocity or a constant acceleration assumption. Interaction is not explicitly modeled. The approach is compared to a time to collision (TTC)-based heuristic regarding the time to cross and the collision rate and yields similar results. In [BNFK18], Bouton et al. enhance their work by considering static occlusions. The policy is computed offline, yet the approach scales well with the number of potentially occluded traffic participants.

Banzhaf et al. [BDSZ18] tackle uncertainties in localization and control for maneuvering in narrow scenarios. They extend the steering functions that are commonly used in RRT* to include uncertainty propagation and propose a method for fast probabilistic collision checking. With this approach, they claim to reduce the collision risk by an order of magnitude in narrow scenarios.

Deterministic Interactive Approaches

Cunningham et al. [CGEO15] claim that they are the first to consider “extensively coupled interactions between agents” in their multi-policy decision making approach. They choose the very general decentralized partially observable Markov decision process (Dec-POMDP) formulation to model their problem. For the sampling-based solving of this Dec-POMDP, however, they largely reduce the state space and only consider a set of known policies for the ego and one particular policy per other vehicle. The available policies are then evaluated using deterministic forward simulation. Further, as uncertainty in the state of the ego or other vehicles is not considered when solving the decision problem, the employed formulation is similar to a deterministic multi-agent MDP (MMDP). From the forward simulation results, the policy that leads to the trajectory with minimum cost is chosen, where the cost function is a weighted sum of multiple criteria. Deviations from the presumed policies are handled by replanning (cf. Sec. 2.1.8).

Lenz et al. [LKK16] model the problem also as a deterministic MMDP. In order to solve this MMDP, they apply MCTS where all agents decide rationally, based on a cooperative cost function. The latter consists of cost for the respective vehicle i plus discounted cost for all other relevant vehicles j , i.e. $C_{i,\text{coop}} = C_i + \lambda \sum_{j \neq i} C_j$ with $0 \leq \lambda \leq 1$. The forward simulation is performed with the IDM as default policy. They show the potential of their approach in simulated highway on-ramp scenarios.

Damerow and Eggert [DE15] also target the uncertainty in the future behavior of others. Depending on the observed scene, they predict different situations, such as different future routes for others or possible violations of traffic rules. For each situation that is of interest for the ego vehicle, they perform an interactive trajectory prediction for each entity using the Foresighted Driver Model (FDM) [EDK15]. Based on these predictions, they create risk maps which are then used for the RRT*-based ego trajectory planning. In order to avoid neglecting situations with low probability but high collision risks, the planned trajectory is checked on each risk map, i.e. for each situation. If necessary, an *escape trajectory* is planned, also described as a *plan B*, for the that case this unlikely situation occurs.

Probabilistic Interactive Approaches

Brechtel et al. [BGD14] target the problem of uncertain information about both the current state and the future evolution of the scene by formulating and solving a continuous POMDP. The probabilistic transition model consists of a (presumably almost deterministic) physics model $p(x'_i | x_i, a_i)$ per non-ego-vehicle and a decision making process per vehicle to determine the probability of action a_i . In order to keep the problem tractable, the road layout is considered static knowledge and not covered in the state, and the policy is generated offline. Thus, the problem has to be precomputed for various different road layouts (including occlusions) and numbers of traffic participants in advance. The transfer of policies to similar scenarios was not focused on in this work. Still, they show promising results for a merging scenario under severe occlusions.

Bahram et al. [BLF⁺16] account for the interaction between the ego and other vehicles by planning under the assumption of a fixed prediction and subsequently predicting under the assumption of the ego plan. Planning and

prediction are performed on a maneuver basis. The actual planning is done via evaluation of the planning-prediction-cycle in a forward simulation and determining the probability of such an evolution via collision risk assessment. The most likely prediction for the ego vehicle is then chosen as the ego plan.

Sunberg et al. [SHK17] propose to extend the POMCP algorithm [SV10] by progressive widening in order to solve POMDPs with continuous observation spaces. They show the performance of this approach in a highway lane change scenario. In their simulation, the behavior of other traffic participants is modeled via the IDM and the lane change model MOBIL [KTH07], while the parametrization varies from aggressive to timid driving. Considering these parameters as hidden variables in their POMDP formulation, they show promising results. For a strong correlation of the behavior parameters, they are even close to an omniscient planner. Also, they account for safety w.r.t. the vehicle ahead by guaranteeing to be able to come to a full stop even if the vehicle ahead performs an emergency deceleration at the physical limits.

Hubmann et al. [HBA⁺17], [HSB⁺18] propose an online POMDP algorithm for automated driving. They consider uncertainty in the state of other vehicles and model their intended route as a hidden variable. For the interaction with other vehicles along their potential paths, an interaction model considering the acceleration to follow a reference velocity, an interaction based acceleration and Gaussian noise is used. Even though the set of possible actions for solving the POMDP using MCTS is limited to ensure the online capability, the results in the presented intersection scenarios with up to four vehicles look promising.

Chen et al. [CTX⁺18] focus on longitudinal decision making in the presence of one other vehicle. Similar to Zhan et al. [ZLCT16], the problem is reduced to a binary decision of whether to go first or second through conflict zones. As an extension to the previous work, the trajectory prediction for the other vehicle is not fixed. Deviations from the non-interactive prediction are possible but induce cost for the ego vehicle, which is called *courtesy awareness* and should prevent reckless behavior. The safety guarantee that is addressed in this work is limited to a horizon of 1 s which would not allow coming to a full stop in the simulated examples.

In [HSX⁺18], Hubmann et al. apply their POMDP solver from [HSB⁺18] to merging in congested traffic. Here, the information gathering actions that are central benefit of the POMDP formulation come into play: The uncertainty of

whether a vehicle in the target lane is cooperative or not is actively reduced by approaching the gap in front of it. The path-velocity decomposition is omitted to allow for a lateral deviation from the predefined path. The results in simulation look human-like and highlight the strength of the approach.

In [HQS⁺19], Hubmann et al. apply their POMDP solver from [HSB⁺18] to occlusions in urban intersections. Both occlusions from static but also from dynamic objects are considered. During the forward simulation through Monte Carlo sampling, the evolution of the occlusion situation is anticipated, including the potential discovery of currently occluded vehicles. Depending on the scenario configuration, the approach sometimes performs similar to a planner with access to ground truth state information.

Learning-Based Motion Planning for Automated Driving

Since learning-based approaches achieved impressive results in robotics and multi-agent games, they were also applied to the automated driving domain. Two very common approaches to solve MDPs are reinforcement learning, as discussed in Sec. 2.1.5, and imitation learning, i.e. trying to mimic demonstrated behavior.

An imitation learning approach that gained attention is NVIDIA's end to end learning approach using a convolutional neural network (CNN) [BDTD⁺16]. The drawback of such approaches is that they cannot outperform the demonstrators and that the generalization to unseen scenarios is at least questionable. Details on the generalization issues, for example due to different input distributions between the training set and the real world, can be found in [Sto09]. A known reinforcement learning approach is the one by Shalev-Shwartz et al. from Mobileye [SSS16]. They propose to learn desires that are then pursued by a conventional trajectory planner. Further, they introduce hard constraints outside the learning framework such that unsafe decisions cannot be taken. While this approach is able to achieve super-human performance, the learned policy strongly depends on the environment in which the agent was trained. Being trained in simulation, it only learns to interact with simulation agents. Training on the road is time-consuming, as it cannot be accelerated faster than real-time, and still a potential safety risk, as strange behavior that is still within the safety limits might lead to human drivers acting scared and thus unsafely.

While such approaches have large theoretic potential, a demonstration in public road traffic is – to the best of the author’s knowledge – still to come, presumably due to the above described difficulties.

2.2.4 Safety in Decision Making and Motion Planning

With the obvious motivation of human lives being at risk, safety is crucial to the deployment and success of automated driving in public road traffic. Due to the kinematic and dynamic constraints of a vehicle, decisions that finally lead to a collision might have been taken seconds before. Fraichard and Asama [FA03] define *inevitable collision states* as those states that inevitably lead to a collision, in other words, from which a collision can no longer be avoided. Such states must not be entered by automated vehicles.

In presence of other agents, safety in the meaning of *collision-free operation* also depends on their behavior. Thus, the definition of inevitable collision states is not straightforward applicable. Further, the guarantee of being collision free is not possible for many scenarios, as depicted in Fig. 2.3.




Figure 2.3: Exemplary situation where the ego vehicle (blue) cannot prevent a collision: If the front vehicle (right) decelerates but the rear vehicle (left) does not, the ego vehicle eventually collides with one of the vehicles in case leaving the road is not possible, for example due to jersey barriers.

Bouraine et al. [BFAS14] circumvent this problem by giving a *passive safety guarantee*, i.e. they guarantee that, if a collision occurs, the ego robot is at rest. This approach is applicable to unstructured environments such as parking lots, where the velocity is commonly low. For structured road traffic at higher velocities, however, some roads are prioritized over others and this *right of way* must be accounted for in safety considerations.

The introduction of the notion of *blame* by Shalev-Shwartz et al. [SSS18] is one approach to address this. Instead of guaranteeing that the ego vehicle will not be involved in a collision, they propose to guarantee never to *cause* a

collision, and call this concept *responsibility sensitive safety (RSS)*. Still, if a collision is inevitable for another vehicle, it is everyone's responsibility to try to avoid or at least mitigate the collision, also if they did not cause this situation. The rule behind this is that "right of way is given, not taken" [SSS18]. In case a situation becomes unsafe, they define a *proper response* to regain safety. In RSS, the safety consideration is divided into *longitudinal safety*, along the direction of the lane, and *lateral safety*, perpendicular to the lane boundaries. The proper response in case of a violation of the longitudinal safety distance to a vehicle in front is a deceleration. The proper response in case of a violation of the lateral safety distance is so-called *lateral braking*, which can be steering or deceleration.

The same notion is followed by Althoff and Dolan [AD14]. They compute *reachable sets* for other traffic participants, but discard samples that break traffic rules. Once a rule has been broken, or certainly will be broken soon¹⁵, however, this rule is no longer taken into consideration. This method is close to the idea of demanding to avoid or mitigate collisions even when not being at fault in RSS. Orzechowski et al. [OML18] expand the concept of set-based safety verification by the consideration of occlusions and limited sensor range. Pek et al. [PZA17] expand the concept to lane changes and thus are, to the best of the author's knowledge, the first who proved the safety of lane change maneuvers. In their approach, the responsibility for safety during lane changes lies with the lane changing vehicle until it has fully entered the target lane. Since the reachable set methods are worst case considerations, full acceleration for the following vehicle in the target lane is assumed. This leads to lane changes in dense traffic being intractable with this approach.¹⁶

¹⁵ Observing a vehicle approaching a red light with high velocity which, even with full possible deceleration, will run over the red light shortly, for example.

¹⁶ In their experiments, a fixed prediction for the following vehicle in the target lane is assumed and uncompliant behavior of this vehicle is not considered. Thus, the approach is not suitable for mixed traffic where compliance cannot be ensured.

3 Formulating the Decision Problem of Motion Planning in Mixed Traffic

In this chapter, the problem of motion planning in mixed traffic is framed into an appropriate model. The main contribution of this chapter is a new problem formulation for scenarios with clear precedence, exploiting the independence of the future trajectories of prioritized traffic participants on the behavior of the ego vehicle. To recapitulate from the introduction, the goal of motion planning for automated vehicles is to yield **convenient** motion plans, i.e. plans that are comfortable for the passengers but also useful w.r.t. the desire of reaching a destination within decent time. If the area that is dedicated to public road traffic would be restricted to robots, central planning approaches could yield *the* globally optimal plan which everyone had to follow. As central planning is not an option for humans though, road traffic is commonly structured into roads and lanes, and is controlled by traffic rules such as speed limits, right of way rules and traffic lights, amongst others. Thus, those traffic rules must be incorporated into the planning problem. Finally, the most important constraint is that the safety of passengers and other traffic participants must never be put at risk. With these goals and constraints in mind, the problem formulation for motion planning in mixed traffic is derived, which is dominated by different uncertainties.

After introducing the underlying model assumptions in Sec. 3.1, the formulation for perfect knowledge is presented in Sec. 3.2. From Sec. 3.3 through Sec. 3.7, the assumption of perfect knowledge is successively relaxed to better approximate the actual problem.

3.1 Underlying Model Assumptions

The problem of choosing actions based on states that are derived from sensor measurements is inherently a partially observable decision process, as the state of an object (such as a vehicle or a pedestrian) cannot be measured directly, but is inferred from observations of different sensors with different physical models. In a single agent POMDP formulation, the behavior of other agents is commonly modeled by a probabilistic state transition model. The underlying states can be hidden, such as the desired route of an agent. In this context, a probability distribution over the current state of the world, including all agents, is denoted *belief*. During planning, the state transition model in conjunction with a probabilistic observation model is repeatedly queried to estimate possible future evolutions of the belief, and thus, to choose optimal actions. Here, the uncertainties from perception can be modeled by including sensor models into the observation model, while the underlying state transitions of other agents are modeled in the probabilistic state transition model.

For road traffic, however, we argue that the trajectories of agents are their conscious decisions. Thus, we are faced with a multi-agent setting, as introduced in Sec. 2.1.7. Assuming perfect trajectory control for the ego vehicle, its state transition model can be considered deterministic, i.e. it can decide upon its future trajectory \mathbf{x}^e by influencing its state change $\dot{\mathbf{x}}^e$ through actions or control inputs \mathbf{u}^e :

$$\dot{\mathbf{x}}^e(t) = f(\mathbf{x}^e(t), \mathbf{u}^e(t)). \quad (3.1)$$

This model assumes collision free motion, subject to the constraints that arise from the dynamics and kinematics of a vehicle. The latter constraints cannot actually be violated, but their adherence must be ensured by the planner to ensure the feasibility of the trajectory. Infeasible trajectories lead to inevitable control errors and thus to the realization of a trajectory that was never considered regarding safety and optimality. Effects such as a bad road surface or an altered friction coefficient (e.g. due to rain or snow) can be incorporated here, but are outside the scope of this work. The same transition model is assumed for other traffic participants i , denoted oi : $\dot{\mathbf{x}}^{oi}(t) = f(\mathbf{x}^{oi}(t), \mathbf{u}^{oi}(t))$, again subject to the dynamics and kinematics. This corresponds to a deterministic MMDP formulation.

From the perspective of the ego vehicle, however, the future actions of other traffic participants are unknown and the current state is subject to uncertainties. Yet, for scenarios with clear precedence, the decisions of prioritized vehicles are assumed to be taken independently of the ego trajectory. Only the observations of other traffic participants are partially influenced by the ego trajectory, due to limited sensor range, for example.¹ Thus, in such scenarios, estimates for the probability of future trajectories are determined by an *upstream* prediction module, i.e. a prediction module that operates independently of and provides input to the motion planning module.

In order to distinguish different feasible ego trajectories \mathbf{x}^e regarding their quality, a cost functional J is commonly used, as introduced in Chapter 2. This cost functional is a mapping from ego trajectory \mathbf{x}^e including its time derivatives and the ego dimensions² \mathcal{D}^e , trajectories of other traffic participants $\underline{\mathbf{x}}^o = (\mathbf{x}^{o1}, \mathbf{x}^{o2}, \dots)$ including their time derivatives and dimensions $\underline{\mathcal{D}}^o = (\mathcal{D}^{o1}, \mathcal{D}^{o2}, \dots)$ and the static and dynamic environment (including the road geometry) to a scalar quality indicator. The soft constraints from the traffic rules and the destination are also incorporated in this cost functional.

Throughout this work, it is assumed that existing objects within the sensor range that are not occluded are detected and assigned a non-zero existence probability. Further, information about the traffic rules and the road layout is assumed to be certain and complete.

¹ Apart from sensor limitations, the influence of the ego trajectory on observations that indicate the future behavior of others, such as a change in velocity or the use of the indicator, is assumed to be negligible.

² The dimensions can for example be described through a hull, defined via a list of points. Also, a simplified definition via a bounding box is possible, besides other representations.

3.2 Perfect Knowledge

In the most simple model, we assume to know the future trajectories of all traffic participants and all other obstacles, as depicted in Fig. 3.1. With this knowledge, the safety constraints are reduced to the avoidance of collisions.




Figure 3.1: The most simple problem formulation: The ego vehicle (blue) shall plan its motion along the predefined route (blue), knowing the exact trajectory of the other traffic participants (white). Exemplary scenario inspired by [HSB⁺18, Fig.1].

With the cost functional J , we choose the optimal trajectory \mathbf{x}^{e*} as the one that minimizes the total cost while not violating any constraint:

$$\mathbf{x}^{e*} = \arg \min_{\mathbf{x}^e} \int_{t_0}^{\infty} J(\mathbf{x}^e, \underline{\mathbf{x}}^o, t) dt. \quad (3.2)$$

The dimensions \mathcal{D}^e , $\underline{\mathcal{D}}^o$, the traffic rules, the environment and the destination are considered as parameters of J . This model of knowing the future trajectories, i.e. having a single deterministic prediction per traffic participant, corresponds to the model taken by the approaches of Sec. 2.2.2. Also, this model assumes perfect control. A simple model to incorporate bounded position control errors $e \leq e_{\max}$ in the collision avoidance constraint is to enlarge the distance for collision avoidance from 0 to e_{\max} , which can be done velocity dependent.

3.3 Uncertainty in Future Trajectories of Others

The first assumption to relax is the exact knowledge of the trajectory of other traffic participants. In structured road traffic, the future trajectory can be subdivided into the future route of the traffic participant, commonly determined in a navigational layer, the future path along that route and the velocity along that path. Large parts of everyday driving consist of following a vehicle within a lane, where the route and the path of that vehicle are assumed to be known, but its future velocity is uncertain. This uncertainty is even addressed in the traffic rules, where a safe distance (often a certain time headway) is demanded. In Germany, for example, the rule of thumb is to keep a distance of at least half the current velocity, i.e. $d_{\min} = \frac{v}{2} \frac{\text{m}}{\text{km h}^{-1}}$, which corresponds to a time headway of 1.8 s. The traffic rule behind that rule of thumb is that one should in general be able to stop behind a vehicle, even if it suddenly brakes [StVO (D), §4 I]. To account for this uncertainty, we consider a probabilistic prediction of all traffic participants, as depicted in Fig. 3.2, which is independent of the ego trajectory.




Figure 3.2: Motion planning problem formulation: The ego vehicle (blue) shall plan its motion along the predefined route (blue). While the current pose of other vehicles (white) is known, their future trajectories are subject to uncertainties (depicted with black ellipses: larger corresponds to higher uncertainty).

The course of the state of object i over time is then a stochastic process \mathbf{X}^{oi} . The state for every point in time t is a random variable $\mathbf{X}^{oi}(t)$. Its PDF is denoted $p_{\mathbf{X}^{oi}(t)}$. For each object i , we still assume to know the current state \mathbf{x}_0^{oi} such that

$$p_{\mathbf{X}^{oi}(t_0)}(\mathbf{x}^{oi}(t_0)) = \delta(\mathbf{x}^{oi}(t_0) - \mathbf{x}_0^{oi}), \quad (3.3)$$

with the Dirac delta δ .

Meaningful stochastic processes account for the kinematic and dynamic constraints of the object. Thus, the object states always lie within a bounded set $\mathcal{X}^{oi}(t)$, such that

$$\Pr(\mathbf{x}^{oi}(t) \in \mathcal{X}^{oi}(t)) = \int_{\mathcal{X}^{oi}(t)} p_{\mathbf{X}^{oi}(t)}(\mathbf{x}^{oi}(t)) d\mathbf{x}^{oi}(t) = 1 \quad (3.4)$$

for all t . The stochastic processes of several objects should not be considered independent, for example to exclude mutual collisions. We again employ the notation $\underline{\mathbf{X}}^o = (\mathbf{X}^{o1}, \mathbf{X}^{o2}, \dots)$ for all objects. The PDF of the process for all objects is denoted $p_{\underline{\mathbf{X}}^o}$.

With this probabilistic model, collisions can no longer be analyzed deterministically. Instead, probabilistic collision checking has to be performed. Here, the motion constraints can be used (cf. eq. 3.4). Further, minimizing *the* cost J is then no longer possible, but the *expected* cost

$$J^{\text{exp}}(\mathbf{x}^e, \underline{\mathbf{X}}^o, t) = \mathbf{E}(J(\mathbf{x}^e, \underline{\mathbf{X}}^o, t)) = \int p_{\underline{\mathbf{X}}^o}(\underline{\mathbf{x}}^o) J(\mathbf{x}^e, \underline{\mathbf{x}}^o, t) d\underline{\mathbf{x}}^o \quad (3.5)$$

can be minimized. In order to compare different costs over a time interval \mathcal{T} , we introduce the cost integral $\mathbf{J}(\mathbf{x}^e, \underline{\mathbf{x}}^o, \mathcal{T}) = \int_{\mathcal{T}} J(\mathbf{x}^e, \underline{\mathbf{x}}^o, t) dt$, and the respective expected cost

$$\mathbf{J}^{\text{exp}}(\mathbf{x}^e, \underline{\mathbf{X}}^o, \mathcal{T}) = \int p_{\underline{\mathbf{X}}^o}(\underline{\mathbf{x}}^o) \int_{\mathcal{T}} J(\mathbf{x}^e, \underline{\mathbf{x}}^o, t) dt d\underline{\mathbf{x}}^o. \quad (3.6)$$

Minimizing \mathbf{J}^{exp} with a single deterministic ego trajectory is possible. Yet, consecutive replanning with replanning interval Δt_{plan} , as introduced in Sec. 2.1.8, facilitates better results. It exploits the fact that the uncertainty in the future trajectory of others is likely to be reduced in the future. At the same time, the




Figure 3.3: Replanning: Based on the previously planned trajectory (gray), at t_0 (●) a new trajectory for the interval $[t_0 + \Delta t_{\text{plan}}, t_0 + 2\Delta t_{\text{plan}}]$ is planned in the hatched interval. This plan is based on perception information from $t_0 - \Delta t_{\text{perc}}$ (○). The planned trajectory is depicted in blue. From $t_0 + 2\Delta t_{\text{plan}}$ onwards, different possible trajectories can be considered (blue dotted).

latency and processing time of perception³ Δt_{perc} and planning Δt_{plan} , called delay, must not be neglected. A plan that has been computed in the interval $[t_0, t_0 + \Delta t_{\text{plan}}]$ based on the perception information from $t_0 - \Delta t_{\text{perc}}$ is executed from $[t_0 + \Delta t_{\text{plan}}, t_0 + 2\Delta t_{\text{plan}}]$. For the interval $[t_0, t_0 + \Delta t_{\text{plan}}]$, the previously planned trajectory is assumed to be controlled. Consequently, we only have to decide for a trajectory in $[t_0 + \Delta t_{\text{plan}}, t_0 + 2\Delta t_{\text{plan}}]$ at t_0 and can consider different trajectories from $t_0 + 2\Delta t_{\text{plan}}$ onwards (cf. Fig. 3.3). The decision for one particular trajectory from $[t_0 + 2\Delta t_{\text{plan}}, t_0 + 3\Delta t_{\text{plan}}]$ is taken in the subsequent replanning step.

In this subsequent replanning step, we will again be faced with uncertainty. Simply assuming that the exact trajectory ensemble \underline{x}^o is known by then is a strong simplification.⁴ While the previously introduced stochastic process represents the current belief over the future trajectories of others, this belief is updated in subsequent planning steps due to new observations of the other traffic participants. In addition to the belief over states in classical POMDP formulations, this belief over trajectories also includes knowledge about the state transition and observation probabilities. The motivation behind this

³ Including localization, scene understanding, amongst others.

⁴ This simplification corresponds to the QMDP assumption (cf. Sec. 2.1.5) combined with the assumption of deterministic state transitions.

formulation is that the trajectories of others are conscious decisions of other agents which do not depend on the ego action. Thus, a scene understanding or prediction module should be able to generate appropriate joint estimates of the belief over states and the respective transition probabilities, independent of the planning problem. Different possibilities for new beliefs can then already be integrated into the planning problem formulation, as shown in the example below.

Example 3.3.1. Imagine, there were only two discrete trajectories possible for the white vehicle in Fig. 3.2: One for turning right and one for going straight. In case the velocity profiles of the two trajectories differ significantly at $t_0 + \Delta t_{\text{plan}}$, certainty would be reached in the next planning step. The decision could be postponed by Δt_{plan} , since the probability distribution is believed to become singular in either turning right or going straight. In case the two discrete trajectories are identical until $t_{\text{identical}} \gg t_0 + \Delta t_{\text{plan}}$, certainty about the route of this vehicle will not be reached soon, and the probability distribution does not change. Instead of postponing the decision to the next planning step, both routes have to be taken into consideration also beyond the next planning step.

The belief over the future trajectories of other objects at a particular time corresponds to a probability distribution of $\underline{\mathbf{X}}^o$. The stochastic process with the probability distribution according to the current belief $b_{\underline{\mathbf{x}}^o}$ is denoted $\underline{\mathbf{X}}_b^o$. The set of possible new beliefs $b'_{\underline{\mathbf{x}}^o}$ in the subsequent replanning step (Δt_{plan} later) is denoted \mathcal{B} , the respective PDF is denoted $p_{b'|b}$.

The expected total cost over the entire planning horizon $\mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e, b_{\underline{\mathbf{x}}^o})$ of \mathbf{x}_1^e , which is pursued in $\mathcal{T}_{\mathbf{x}_1^e} = [t_0 + \Delta t_{\text{plan}}, t_0 + 2\Delta t_{\text{plan}})$, is:

$$\begin{aligned} \mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e, b_{\underline{\mathbf{x}}^o}) &= \mathbf{J}^{\text{exp}}\left(\mathbf{x}_1^e, \underline{\mathbf{X}}_b^o, \mathcal{T}_{\mathbf{x}_1^e}\right) \\ &+ \int_{\mathcal{B}} p_{b'|b}\left(b'_{\underline{\mathbf{x}}^o} | b_{\underline{\mathbf{x}}^o}\right) \mathbf{J}_{\text{total}}^{\text{exp}}\left(\mathbf{x}_2^{e*}, b'_{\underline{\mathbf{x}}^o}\right) db'_{\underline{\mathbf{x}}^o}, \end{aligned} \quad (3.7)$$

where \mathbf{x}_2^{e*} is the optimal trajectory continuation of \mathbf{x}_1^e . In words, the expected total cost consists of the expected cost of the pursued trajectory \mathbf{x}_1^e in $\mathcal{T}_{\mathbf{x}_1^e}$, and the expected cost of its optimal continuation for the time beyond $t_0 + 2\Delta t_{\text{plan}}$. To determine \mathbf{x}_2^{e*} and its expected total cost, recursive computation with the




Figure 3.4: Recursion during planning: Based on the previously planned trajectory (gray), at t_0 (•) a new trajectory for the interval $[t_0 + \Delta t_{\text{plan}}, t_0 + 2\Delta t_{\text{plan}})$ is planned (cf. Fig. 3.3). To determine the optimal trajectory, possible future belief updates for every replanning have to be considered. However, new information can only be incorporated at the times of replanning (dashed).

respective $b'_{\underline{x}^o}$ is necessary (cf. Fig. 3.4). This accounts for the fact that we are able to replan repeatedly in intervals of Δt_{plan} , and that the belief will have been updated with new measurements by then. Consequently, the term $\mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e, b_{\underline{x}^o})$ is based on the full possible future progress of the belief, similar to the recursion in the Bellman equation for belief states (cf. eq. 2.14).

Based on the principle of maximum expected utility (cf. eq. 2.2), the optimal trajectory part is to be chosen as

$$\mathbf{x}_1^{e*} = \arg \min_{\mathbf{x}_1^e} \mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e, b_{\underline{x}^o}). \quad (3.8)$$

The pursued trajectory \mathbf{x}^e eventually consists of \mathbf{x}_1^e , determined at t_0 , continued by \mathbf{x}_2^e , determined at $t_0 + \Delta t_{\text{plan}}$, continued by \mathbf{x}_3^e and so forth.

3.4 Uncertainty in State of Others

In the previous section, the current state of other traffic participants $\mathbf{x}^{oi}(t_0)$ was still assumed to be known. At least for far away objects, this assumption is not reasonable. Consequently, we relax the assumption from eq. 3.3. The additional uncertainty is depicted in Fig. 3.5. Also, the dimensions of an object \mathcal{D}^{oi} are subject to uncertainties, such that we need to use a probability distribution there. Note that, while this relaxation appears minor compared to the uncertainty in the future trajectory, it affects the bounds of the possible trajectories $\mathcal{X}^{oi}(t)$, which are critical to the safety constraints.

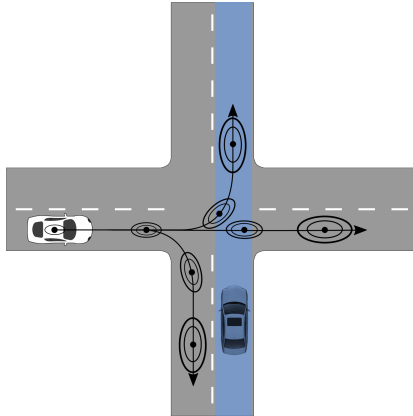


Figure 3.5: Motion planning problem formulation: The ego vehicle shall plan its motion along the predefined route. In addition to Fig. 3.2, the current pose of other vehicles is subject to uncertainties, commonly leading to an increased uncertainty in the future trajectory (depicted with multiple black ellipses).

3.5 Uncertainty in Existence of Others

Another uncertainty regarding the current state is the existence of an object, depicted in Fig. 3.6. Again, especially for far away objects, the existence of an object cannot be guaranteed by the perception of the vehicle. In fact, so-called *ghost objects*⁵ are a common phenomenon and thus, the existence uncertainty should not be neglected. This uncertainty has to be included in the beliefs and the belief updates. If one object exists, another object cannot be at the same position at the same time, at least for meaningful predictions. If the other object was a ghost and does not actually exist, however, it could be there.




Figure 3.6: Motion planning problem formulation: The ego vehicle shall plan its motion along the predefined route. In addition to Fig. 3.5, even the existence of other objects is uncertain (depicted with a question mark).

⁵ One source of ghost objects for RADAR sensors are reflections of radio waves at planes, such as walls, for example.

3.6 Occlusions and Limited Sensor Range

Besides the existence probability of objects due to ambiguous observations, the range of sensors is generally limited. Further, occlusions shadow certain areas from the sensors. While previously, the uncertainties in existence and position could be considered on an object-level, there is an uncountably infinite number of possible vehicle configurations in unobserved areas, each with an infinitesimal existence probability. Exemplary unobservable areas are depicted in Fig. 3.7. Instead of an object-level consideration of hypothetical objects in unobservable areas, we propose to explicitly include the unobservable areas into the problem formulation. Further, while the pursued trajectory of others is still assumed to be independent of the ego trajectory, the belief depends on whether or not an object is observable, and thus on the pose of the ego vehicle. We denote this dependency $p_{b'|b, \mathbf{x}_1^e} \left(b'_{\mathbf{x}^o} | b_{\mathbf{x}^o}, \mathbf{x}_1^e \right)$, similar to $\Pr(b'|b, a)$ in eq. 2.14, which has to be substituted into eq. 3.7.




Figure 3.7: Motion planning problem formulation: The ego vehicle (blue) shall plan its motion along the predefined route (blue). In addition to the uncertainties depicted in Fig. 3.6, occlusions (due to the building depicted in black) and limited sensor range impose additional uncertainties (unobservable areas are shaded).

While for perfect knowledge, it was theoretically possible to precompute the trajectory until the final destination, the consideration of uncertainty in the future trajectory of others already strongly limited the possibilities for safe

open-loop plans. With the consideration of limited observability, however, safety can impossibly be guaranteed for destinations outside the observable area. To overcome this issue, replanning must be used, as formulated in eq. 3.7.

Due to the limited visibility, the computation can further be limited to an appropriate planning horizon⁶ Δt_H with $t_H = t_0 + \Delta t_H$. The cost beyond t_H is approximated as so-called end cost $\mathbf{J}^{\text{end}}(\mathbf{x}^e) \approx \int_{t_H}^{\infty} J(\mathbf{x}^e, t) dt$. With this approximation, the recursive computation of \mathbf{x}_2^{e*} in eq. 3.7 ends at $t = t_H$. Further, J and thus also \mathbf{J}^{end} now are additionally subject to the visible area.

⁶ The planning horizon Δt_H must be larger than the replanning interval Δt_{plan} . Usually, it is chosen to be much larger $\Delta t_H \gg \Delta t_{\text{plan}}$.

3.7 Interaction

Lastly, interaction between the ego vehicle and other traffic participants must be considered, as depicted in Fig. 3.8. In this case, not only the observability of other vehicles, but also their trajectories potentially depend on the ego trajectory.




Figure 3.8: Motion planning problem formulation: The ego vehicle (blue) shall plan its motion along the predefined route (blue). In addition to the uncertainties depicted in Fig. 3.7, the plan might depend on interactions (depicted with blue arrows) with vehicles of uncertain existence (white).

On one hand, interaction cannot be neglected in some situations. On the other hand, deriving general models for the complex interdependence between the ego's and the objects' trajectories, with unknown goals and desires, unknown sensing limitations of other objects and possible future interaction with not yet visible objects, or objects that are visible to others but not to the ego, is intractable. Consequently, further simplifications or heuristics are necessary to solve this motion planning problem, which will be the focus of Chapter 4.

4 Solution to the Decision Process

In this chapter, the problem formulation from Chapter 3 is simplified using further assumptions and subsequently, solutions to the simplified problem are presented. The contributions of this chapter are the extension of the RSS-framework [SSS18] to right-of-way rules and the proposal of three motion planning approaches, targeted towards scenarios with clear precedence, scenarios that require courtesy of other traffic participants and scenarios that require mutual cooperation with other traffic participants. The need for further assumptions or simplifications can be derived from eq. 3.7, where the integration over the set of possible beliefs is required. In addition to that, the dependence of the cost function J on the observable space needs to be modeled and safety considerations have to be included.

We first review the key challenges of automated driving in mixed traffic and propose to prioritize safety considerations in Sec. 4.1, before we introduce a safety concept employing a decomposition into scenarios in Sec. 4.2. Subsequently, approaches to convenient motion planning are presented in Sec. 4.3 with explicit courtesy and cooperation consideration in Sec. 4.4 and Sec. 4.5. Insights into the transition in-between those approaches are given in Sec. 4.6.

This chapter is based on and was partially previously published in [3, 6, 8, 9]: Prioritizing safety over comfort and convenience (Sec. 4.1) has been proposed in [6]. The safety consideration for intersections with crossing traffic from [8] is extended by merging traffic in Sec. 4.2. The safety consideration for parallel lanes (Sec. 4.2) has been presented in [9]. The courtesy-aware approach to lane changes (Sec. 4.4) has also been introduced in [9]. The approach to mutual cooperation (Sec. 4.5) has been proposed in [3, 6].

4.1 Prioritizing Safety over Comfort and Convenience

The key challenges originate in the collision risk with obstacles and other traffic participants, including vulnerable road users (VRUs). Due to the high velocities in public road traffic, such collisions put human lives at risk. This collision risk can be eliminated in enclosed environments, such as subways with automatic sliding doors, where access to the track and the behavior of the vehicles can be controlled centrally. In public road traffic, however, the space is shared with many other traffic participants who plan their actions individually. Regarding their detection, we face challenges due to imperfect perception regarding other objects' state and even their existence, limited sensor range and occlusions, as outlined in the previous chapter. Additionally, we have to cope with uncertainties regarding their route, their behavior along that route and their willingness to be courteous or cooperative. Even their traffic rules compliance is not ensured, non-compliance can for example be caused by inattention. One attempt to reduce the collision risk or at least the collision severity in open environments would be to lower the velocity. While this might help in certain scenarios, such as on parking lots, it is not an option for on-road driving. There, it is inconvenient, but also might even increase the collision risk, as it might entice people into risky overtaking maneuvers, for example. Also, the behavior of the ego vehicle should be comprehensible for other traffic participants, in the meaning of not obviously confusing them. Driving at very low velocities would violate this objective. Furthermore, this behavior might even be unlawful, as vehicles must not hinder traffic flow by driving slowly without compelling reason [StVO (D), §3 II].

Convenience and safety appear to be mutually exclusive, because safety is easier to ensure at very low velocities, but the latter is not convenient, at least not as a default behavior. The goals of safety and convenience are inherently very different. Safety needs guarantees, because we certainly do not want to put human lives at risk. Thus, we are limited to approaches that are complete and real-time capable. Optimality, however, is not necessary, one feasible safe plan is sufficient. Convenience, on the other hand, should consider probabilities: Occasional uncomfortable behavior is acceptable, guarantees are not necessary, but on average, the behavior should yield convenient results.

These severe differences can be accounted for with different approaches as follows. Safety is targeted by one approach that only intervenes in case safety is at risk. The approach to ensure convenience can operate within the safe bounds that are determined by the safety approach. This facilitates running the safety planner with a higher frequency, leading to a smaller reaction time and consequently a larger scope for comfort planners. Most importantly, safety can be guaranteed with a complete and verifiable approach, as requested in Sec. 1.2, while incomplete methods without guarantees remain applicable for the independent comfort planners.

4.2 Ensuring Safety: Incorporating Traffic Rules

At first glance, there appear to be multiple approaches to validate the safety of a motion planning and decision making approach for automated vehicles. The safety of road traffic is commonly measured in fatalities per distance or time traveled. However, the comparability of such measures is only given for similar traffic volume and road layout, as driving through intersections with unclear or even confusing layout is more challenging than driving along straight roads, for example. Besides driving many miles, either on-road or in simulation, and measuring accidents or disengagements therein, formal safety validation is an option. Shalev-Shwartz et al. [SSS18] elaborate that driving enough miles to statistically prove reasonable safety claims is intractable, especially because any code change would result in the need for reevaluation. Regarding simulation, they argue that validating the simulator is as challenging as validating the approach itself: Even if the real world accident probability can be reproduced for some driving policies, a different policy might introduce behaviors that lead to different results in simulation and in real world traffic. Thus, to guarantee safety, formal methods should be used.

While the introduction and formal proof of a comprehensive safety concept that complies with all aspects of traffic law is beyond the scope of this thesis, we propose enhancements to the RSS concept of Shalev-Shwartz et al. [SSS18], as introduced in Sec. 2.2.4. The enhancements include the consideration of traffic rules at intersections with crossing and merging traffic, occlusions and parallel lanes. To shortly recapitulate [SSS18]: Instead of guaranteeing that the ego vehicle will not be involved in a collision, which is intractable for real road

traffic, the guarantee is never to *cause* a collision. This guarantee is given w.r.t. any other vehicle, but the relation to every vehicle is considered separately to allow for scalability. Yet, the ego vehicle should behave cautious enough to be able to “compensate for reasonable mistakes of other drivers” [SSS18]. Shalev-Shwartz et al. claim that their concept is *sound*, i.e. complies with human understanding of traffic rules, *useful*, i.e. not overly defensive, and *efficiently verifiable*. We enhance their concept and partially deviate from it, based on our interpretation of the German traffic law. Note that we can of course not guarantee conformity with the traffic rules, but only propose one possibility to translate the traffic rules into machine readable rules.

Shalev-Shwartz et al. list five so-called *common sense rules* on which RSS is based (direct quote from [SSS18]):

1. Do not hit someone from behind.
2. Do not cut-in recklessly.
3. Right-of-way is given, not taken.
4. Be careful of areas with limited visibility.
5. If you can avoid an accident without causing another one, you must do it.

Rule 1 is easy to formulate mathematically, at least for structured roads, as we will see in the next subsection. For rule 4, Shalev-Shwartz et al. define two so-called *unreasonable situations*. The first one is to exceed certain speeds in areas with limited visibility. They assume that occlusions are symmetric, and request all vehicles to drive within *reasonable limits*. These limits are not further specified, but should depend on the priority rules, amongst others. While the definition of such limits is not part of this work, we follow a similar concept, as outlined in the next section. The second unreasonable situation is when vehicles that are visible to us violate safe distances to other vehicles. Rule 2, the non-reckless cut-in, is not further specified, but it is mentioned that, for a *safe cut-in*, the vehicle behind is expected to adjust its velocity. In case the safe cut-in refers to a cut-in where the vehicle behind has just more than the safe distance to the ego vehicle, we deviate. While this would indeed still guarantee collision free traffic if all traffic participants would follow this rule, it does not reflect the traffic rules. Essentially, vehicles that have the right of way

“must neither be endangered nor be significantly impeded” [StVO (D), §8 II 2]. A vehicle entering a priority road with 10 km/h just before a vehicle that drives 50 km/h, however, clearly impedes this vehicle by forcing it to sharply brake in order to maintain a safe longitudinal distance, and thus violates its right of way. Rule 3, to not take right of way, is essentially motivated by rule 5: Do avoid accidents if possible, also if you had right of way. Waiving the right of way is unusual and requires a prior arrangement with the prioritized traffic participant [StVO (D), §11 III]. This is challenging for automated vehicles, especially for perception, and not focused on in this work. Finally, rule 5 should be extended to: If you can mitigate an accident without causing another one, you must do it.

As motivated in the previous chapter, the road layout and traffic laws are man-made rules that have evolved over time to make traffic safer and more efficient. Since the responsibility for accidents depends on compliance with the traffic rules, the proposed safety concept must account for the traffic rules. Notationwise, a *lane* describes the part of a road or carriageway that is to be used by a single line of vehicles in a particular direction. Thus, roads with traffic in both directions commonly have at least two lanes. In the following, we distinguish different scenarios based on the lane structure of roads. We start with the single lane consideration in Sec. 4.2.1, before we consider intersecting lanes in Sec. 4.2.2 and parallel lanes in Sec. 4.2.3. In Sec. 4.2.4, we briefly discuss further scenarios.

4.2.1 Single Lane

The simplest case, which is also the basic example in [SSS18], is driving within a lane without the necessity to consider traffic in other lanes or outside of lanes. In this case, path-velocity decomposition, as proposed by Kant and Zucker [KZ86] and presented in Sec. 2.1.4, can be applied. Lateral safety can be ensured by staying within the lane boundaries, longitudinal safety can be ensured by guaranteeing RSS rule 1, not to hit someone from behind. To reformulate Definition 1 of [SSS18], this can be ensured for the ego vehicle by guaranteeing to not collide with its predecessor in the following worst case: The predecessor, previously traveling with velocity v_p , decelerates with a defined maximum deceleration $a_{\max, \text{brake}}$ until it reaches a full stop. During a reaction time ρ_e , the ego vehicle, previously traveling at v_e , accelerates

with a maximum acceleration $a_{\max, \text{accel}}$, and subsequently it decelerates with a minimal assured deceleration $a_{\min, \text{brake}}$. If, in this worst case, the ego car does not collide with its predecessor, the distance to the predecessor is considered safe. This minimum safe distance according to this scenario can be computed as

$$d'_{\min, \text{safe}} = v_e \rho_e + \frac{1}{2} a_{\max, \text{accel}} \rho_e^2 + \frac{(v_e + \rho_e a_{\max, \text{accel}})^2}{-2a_{\min, \text{brake}}} - \frac{v_p^2}{-2a_{\max, \text{brake}}}, \quad (4.1)$$

where the first three terms denote the distance until the ego vehicle is at a full stop, and the last term is the distance in which the predecessor is at a full stop. If, due to a very large velocity of the predecessor, $d'_{\min, \text{safe}}$ is negative, the safe distance must still be non-negative, as a negative distance would mean a collision. Thus, minimum safe distance is generally $d_{\min, \text{safe}} = \max(d'_{\min, \text{safe}}, 0)$. For the proof by induction, we refer to [SSS18]. In contrast to [SSS18], we always use $a < 0$ for decelerations, and thus, the stopping distance is $s_{\text{stop}} = \frac{v_0^2}{-2a_{\text{decel}}}$.

Regarding occlusions and limited sensor range, we must always be able to come to a full stop within our visible range [StVO (D), §3 I 4]. This corresponds to expecting a static obstacle just behind the visible range d_{vis} . With $a_{\max, \text{accel}} = 0$, because we may no longer accelerate when at the maximum allowed velocity $v_{e, \max}$, eq. 4.1 leads to

$$d_{\text{vis}} = d_{\min, \text{safe}} = v_{e, \max} \rho_e + \frac{v_{e, \max}^2}{-2a_{\min, \text{brake}}}, \quad (4.2)$$

with solution

$$v_{e, \max} = a_{\min, \text{brake}} \rho_e \pm \sqrt{(a_{\min, \text{brake}} \rho_e)^2 - 2a_{\min, \text{brake}} d_{\text{vis}}}. \quad (4.3)$$

With $d_{\text{vis}} > 0$ we know $-2a_{\min, \text{brake}} d_{\text{vis}} > 0$ and thus, subtracting the root leads to a negative maximum safe velocity, which is not reasonable. Consequently, the maximum safe velocity is

$$v_{e, \max} = a_{\min, \text{brake}} \rho_e + \sqrt{(a_{\min, \text{brake}} \rho_e)^2 - 2a_{\min, \text{brake}} d_{\text{vis}}}. \quad (4.4)$$

When driving within a lane with a predecessor that drives in the same direction, there is typically no interaction and the state uncertainty is generally low. Even though the uncertainty in the future behavior of others is high, assuming that the predecessor will not decelerate uncomfortably without compelling reason, we can simply follow this vehicle.¹ In case of occlusions and limited visibility, however, driving at the maximum safe velocity necessitates sharp braking maneuvers whenever a vehicle appears at the edge of the visible range. This issue will be targeted in Sec. 4.3. The special case of narrow roads without right of way, where cooperation is necessary, is covered in Sec. 4.5.

Even if no other lanes intersect our lane, it is still possible that other traffic participants enter our lane, such as cars parked at the side of the road, or pedestrians crossing the street. Thus, also for driving within a lane, we need to consider lateral evasive maneuvers, as described in [SSS18]. Moreover, for areas where pedestrians have priority, we have to consider that they enter the lane from occluded areas. Shalev-Shwartz et al. proposed in version 5 of the RSS whitepaper [SSS18, v5], that a vehicle is not responsible for an accident arising from such a situation in the following case: It did not accelerate from the time the pedestrian became visible, braked at least after the reaction time ρ_e and was on average slower than the pedestrian in the time between exposure and collision. This proposal, has been removed in version 6 [SSS18, v6]. In this work, we do agree that the proper behavior in such cases must incorporate a limitation of the ego velocity. This limitation should depend on the danger arising from an occlusion, which again depends on the lateral distance of the ego vehicle to the relevant occlusion. The trade-off between the law to not hinder traffic flow by driving slowly without compelling reason [StVO (D), §3 II] and lowering velocity to preempt endangering children [StVO (D), §3 IIa], is not a pure technical question and therefore not further investigated throughout this work.

4.2.2 Intersecting Lanes

At points where lanes intersect each other, we distinguish between *crossing traffic*, where vehicles originating from different lanes continue in different

¹ Advanced cruise control, which provides this functionality, is already available in series vehicles.

lanes, and *merging traffic*, where vehicles originating from different lanes continue in the same lane. Note that at most junctions, both crossing and merging traffic is possible. For *diverging traffic*, where vehicles originating from the same lane continue in different lanes, the previously introduced single-lane consideration is sufficient. Special road layouts, such as roundabouts, can be broken down into these three basic categories. As the overlapping area between lanes is an area of potential conflict between traffic in these lanes, we call this area *conflict zone*. For crossing and merging traffic, the right of way has to be defined. The concept of path-velocity decomposition is generally well applicable to scenarios with intersecting lanes. Thus, at intersections, we focus on longitudinal safety, i.e. on the distance to the conflict zone along the path.

From the perspective of a vehicle with right of way, vehicles that enter its path could be crossing or merging. This ambiguity has to be accounted for in the safety consideration. As previously stated, vehicles that have the right of way “must neither be endangered nor be significantly impeded” [StVO (D), §8 II 2]. Thus, disregarding occlusions, we propose that prioritized vehicles have to regularly maintain a safe longitudinal distance to objects entering their lane only in the following case: If, at the time when the intersecting traffic enters the prioritized lane, the prioritized vehicle can come to a safe stop in front of the conflict zone when constantly decelerating with a *reduced required deceleration* $a_{\min, \text{brake, row}} > a_{\min, \text{brake}}$ after a reaction time ρ , similar to eq. 4.1 with $v_p = 0$. Additionally, once an intersecting vehicle is observed to not be able to ensure safety, the prioritized vehicle has to maintain safety according to RSS rule 5. This is assumed if the intersecting vehicle either merges into the lane of the prioritized vehicle or, in case of crossing, under both the assumption of constant velocity and constant acceleration, it will not have left the conflict zone by the time the prioritized vehicle arrives. This rule is illustrated in the following examples.

Example 4.2.1. A vehicle enters the conflict zone such that the required deceleration to stop in front of the conflict zone for the prioritized vehicle is only $a_{\text{req}} \geq a_{\min, \text{brake, row}}$. The prioritized vehicle has to maintain a safe longitudinal distance to this vehicle, independent of whether this vehicle is crossing or merging.

Example 4.2.2. A vehicle starts crossing closely such that the required deceleration to stop in front of the conflict zone for the prioritized vehicle is $a_{\text{req}} < a_{\text{min,brake,row}}$. The crossing vehicle ensures that it can leave the conflict zone before the prioritized vehicle enters. The prioritized vehicle does not have to react to this vehicle.

Example 4.2.3. A vehicle starts crossing closely such that the required deceleration to stop in front of the conflict zone for the prioritized vehicle is $a_{\text{req}} < a_{\text{min,brake,row}}$, but then, the crossing vehicle decelerates and stops within the conflict zone. The prioritized vehicle only has to react if under both, the assumption of constant velocity for the crossing vehicle and the assumption of constant acceleration for the crossing vehicle, the latter will not have left the conflict zone by the time the prioritized vehicle arrives. If a collision can only be mitigated but no longer be avoided, only the crossing, but not the prioritized vehicle is to blame for the accident.

Example 4.2.4. A vehicle enters the conflict zone closely such that the required deceleration to stop in front of the conflict zone for the prioritized vehicle is $a_{\text{req}} < a_{\text{min,brake,row}}$, and then merges in front of the prioritized vehicle. The prioritized vehicle only has to react once the non-prioritized vehicle leaves the conflict zone in direction of the prioritized vehicle's lane. If a collision can only be mitigated but no longer be avoided, only the merging, but not the prioritized vehicle is to blame for the accident.

In order to prevent marginal cases, we further define an *expectable deceleration* $a_{\text{brake,row,exp}}$ that is at most assumed by non-prioritized vehicles, with $a_{\text{min,brake,row}} < a_{\text{brake,row,exp}}$. For safety distances that would lead to a deceleration a such that $a_{\text{min,brake,row}} < a < a_{\text{brake,row,exp}}$, the non-prioritized vehicle has to ensure safety, but also the prioritized vehicle has to maintain a safe distance.

In the following, we further distinguish between crossing and merging traffic.

Crossing

We start with the consideration from the perspective of a non-prioritized vehicle that intends to cross. When intending to cross, safety can be ensured by either of two conditions:

C1 Ensure to come to a safe state before the conflict zone.

C2 Ensure to safely pass the conflict zone.

The traversal from condition C1 to condition C2 leads through their intersection, in which both conditions are fulfilled. This intersection has to be non-empty: $C1 \cap C2 \neq \{\}$.²

To satisfy condition C2, one option is to check the required deceleration of prioritized vehicles at the time the non-prioritized vehicle enters the conflict zone, leading to:

C2(a) Entering the conflict zone such that the prioritized vehicle is at sufficient distance, i.e. its required deceleration to stop in front of the conflict zone is acceptable $a_{\text{req}} > a_{\text{brake, row, exp}}$, is considered safe for the non-prioritized vehicle.

For large velocities, however, condition C2(a) leads to large safety distances, as shown in the example below.

Example 4.2.5. Assuming a reaction time for the prioritized vehicle $\rho_{\text{row}} = 1 \text{ s}$ and an expectable deceleration $a_{\text{brake, row, exp}} = -1 \text{ m/s}^2$. If the vehicle on the prioritized road is traveling at the speed limit $v_{\text{row}} = 28 \text{ m/s} \approx 100 \text{ km/h}$, safely entering according to condition C2(a) is only possible if the distance of the prioritized vehicle to the conflict zone is larger than $d_{\text{suff, row}} \approx 28 \text{ m} + 392 \text{ m} = 420 \text{ m}$.

If the vehicle on the prioritized road is traveling at $v_{\text{row}} = 14 \text{ m/s} \approx 50 \text{ km/h}$, the distance is $d_{\text{suff, row}} \approx 14 \text{ m} + 98 \text{ m} = 112 \text{ m}$, and with $v_{\text{row}} = 9 \text{ m/s} \approx 32 \text{ km/h}$ it is only $d_{\text{suff, row}} \approx 9 \text{ m} + 41 \text{ m} = 50 \text{ m}$.

² Empty intersections might arise in presence of severe occlusions. Here, a third safety condition is necessary to allow for carefully advancing into a conflict zone, but the latter is out of the scope of this work.




Figure 4.1: The desired paths (green) of two vehicles overlap (a). The conflict zone (CZ) is depicted in gray. The time of zone clearance (TZC) or post-encroachment time (PET) is the time that elapses between one vehicle leaves the conflict zone (b) and the subsequent vehicle enters it (c).

Under condition C2(a), entering the conflict zone is safe, since the prioritized vehicle has to react in case the non-prioritized vehicle is unable to leave the conflict zone. In some cases, however, the intersection is well visible and clear of other traffic participants, such that the non-prioritized vehicle can guarantee to clear the conflict zone even prior to entering it. Thus, we propose a further option for safely crossing a conflict zone:

- C2(b) Entering the conflict zone is considered safe if the non-prioritized vehicle can guarantee to have left the conflict zone by the time the prioritized vehicle can enter it the earliest, w.r.t. $a_{\max, \text{accel}}$.

While condition C2(b) is sufficient for safety, drivers of prioritized vehicles might still feel endangered. Thus, we propose to additionally consider criticality measurements as constraint for reasonable behavior of automated vehicles. As such, the time between the first vehicle leaves the conflict zone and the second vehicle arrives at it, called *post-encroachment time (PET)* or *time of zone clearance (TZC)* and visualized in Fig. 4.1, is used. It indicates how closely the vehicles approached each other [ASC78].

The TZC accounts for the velocity of the second vehicle through a spatial distance that is proportional to its velocity. Given a possible ego trajectory

and a reasonable prediction of other vehicles, such as constant velocity on the prioritized lane, we propose to only cross the intersection if the expected TZC is larger than a predefined $t_{TZC,\min}$. In other words, we propose that if we plan considering this minimum TZC, we did not endanger the prioritized vehicle and are not responsible for overreactions to our crossing, even if violating C2(a).

Example 4.2.6. The vehicle on the prioritized road is traveling at $v_{\text{row}}(t_0) = 28 \text{ m/s} \approx 100 \text{ km/h}$. The non-prioritized ego vehicle is at a full stop $v_{\text{ego}}(t_0) = 0$.

Further assume a width of the conflict zone of $s_{cz} = 4 \text{ m}$ from the perspective of the ego vehicle, and a length of the ego vehicle of $l_{\text{ego}} = 5 \text{ m}$. If the ego vehicle intends to accelerate with $a_{\text{ego}} = 1.8 \text{ m/s}^2$, it takes $\Delta t = \sqrt{2 \frac{s_{cz} + l_{\text{ego}}}{a_{\text{ego}}}} \approx 3.2 \text{ s}$ to pass the conflict zone.

With a minimum required time of zone clearance of $t_{TZC,\min} = 2.5 \text{ s}$, the ego vehicle can start to cross if the prioritized vehicle is further than $d_{\text{suff},TZC} \approx 28 \text{ m/s} \cdot 5.7 \text{ s} \approx 160 \text{ m}$ away, if it can ensure to leave the conflict zone in 3.2 s. The prioritized vehicle may ignore the longitudinal safe distance to the ego vehicle in this case (cf. ex. 4.2.2).

If the vehicle on the prioritized road is traveling at $v_{\text{row}} = 14 \text{ m/s} \approx 50 \text{ km/h}$, the distance is $d_{\text{suff},TZC} \approx 80 \text{ m}$, and with $v_{\text{row}} = 9 \text{ m/s} \approx 32 \text{ km/h}$ it is $d_{\text{suff},TZC} \approx 52 \text{ m}$. For large velocities of the prioritized vehicles, option C2(b) facilitates crossing with less distance. For small velocities, however, this option is not beneficial.

Merging

For merging traffic, the paths overlap from a certain point on. Consequently, in case the non-prioritized vehicle has a lower desired velocity than the prioritized vehicle, a deceleration of the latter is inevitable, no matter how large the time gap was at the time of merging. In such cases, instead of evaluating the overall impact on the prioritized vehicle for its entire journey, we can only evaluate the impact at or around the merging point. Generally, for accelerating from zero up to the maximum velocity on roads without on-ramps, which is up to

100 km/h \approx 28 m/s in Germany for example, times of 15 s are not unusual. During this time, however, a vehicle driving at 100 km/h travels more than 400 m. Thus, observing those vehicles prior to merging is unrealistic.

Condition C1, being able to ensure to come to a safe state in front of the conflict zone, remains unchanged for merging. To safely enter the conflict zone, first of all, a merge maneuver must not be so close that the longitudinal safety distance of the prioritized vehicle is violated.

Example 4.2.7. The vehicle on the prioritized road is traveling at $v_{\text{prio}}(t_0) = 28 \text{ m/s} \approx 100 \text{ km/h}$. The ego vehicle that intends to merge is at a full stop $v_{\text{merge}}(t_0) = 0$ just at the merge point. Assuming a reaction time for the prioritized vehicle $\rho_{\text{prio}} = 1 \text{ s}$, and a minimum required deceleration $a_{\text{min,brake}} = -7 \text{ m/s}^2$, and further assuming that the latter vehicle may no longer accelerate as it is already at the speed limit, the required safe distance at the point where the ego vehicle merges (with approximately velocity zero) can be computed as

$$d_{\text{min,safe}} = v_{\text{prio}}(t_0)\rho_{\text{prio}} + \frac{(v_{\text{prio}}(t_0))^2}{-2a_{\text{min,brake}}} = 28 \text{ m} + 56 \text{ m} = 84 \text{ m}.$$

Further, as the requirement to obey the right of way is to not significantly impede prioritized vehicles, the imposed deceleration for those vehicles must not exceed the previously defined expectable deceleration $a_{\text{brake,row,exp}}$. The computation of this deceleration for the case of the almost stopped merger is of theoretic nature, as the desire of the merging vehicle is to accelerate to a certain velocity after merging.³ Thus, we propose the following concept to ensure smooth merges.

The merging vehicle intends to accelerate with some minimum acceleration $a_{\text{min,accel,merge}}$, which would only be violated due to the behavior of other vehicles.⁴ The prioritized vehicle facilitates the merging by decelerating with

³ The consideration of the merging vehicle (constantly) driving approx. 0 km/h leads to large required distances, such as 420 m when the prioritized vehicle is driving at 100 km/h (cf. ex. 4.2.5).

⁴ Such vehicles, for example another vehicle merging in front performing an emergency braking maneuver, would also affect the prioritized vehicle. Thus, we disregard reactive decelerations of the merging vehicle here.

$a_{\text{brake,row,exp}}$ after the merging vehicle has merged plus the usual reaction time ρ_{prio} . The time of merging can be defined in multiple ways. If merging is the only available maneuver for the merging vehicle, the maneuver is obvious from the time when the merging vehicle enters the prioritized lane (cf. Fig. 4.2b). If, from the perspective of the prioritized vehicle, the merging vehicle could also intend to cross, the maneuver is obvious from the time when the merging vehicle enters the common path (cf. Fig. 4.2c).⁵

The safety reserve for the prioritized rear vehicle along its lane to the front vehicle, assuming the front vehicle is already in this lane, is

$$d_{\text{reserve}}(t) = s_{\text{merge}}^-(t) - s_{\text{prio}}^+(t) - d_{\text{safe,prio}}(t), \quad (4.5)$$

where $d_{\text{safe,prio}}$ is the longitudinal safe distance required by the prioritized vehicle, and s_{merge}^- and s_{prio}^+ are visualized in Fig. 4.2d.

The critical time t_{crit} of minimum safety reserve over the merging maneuver can be found via differentiating $d_{\text{reserve}}(t)$ w.r.t. t and setting this term equal to zero:

$$\frac{\partial d_{\text{reserve}}(t)}{\partial t} \stackrel{!}{=} 0. \quad (4.6)$$


The minimum required initial distance of the prioritized vehicle $d_{\text{init,prio,min}}^+$ can be found by requiring that the reserve at the critical point is zero

$$d_{\text{reserve}}(t_{\text{crit}}) \stackrel{!}{=} 0 \quad (4.7)$$


and rearranging the equation for $d_{\text{init,prio}}^+$. Thus, the condition for safe merging is

- C2(c) If the actual distance $d_{\text{init,prio}}^+$ at the time of merging is larger than $d_{\text{init,prio,min}}^+$, merging in front of the prioritized vehicle is considered safe and not significantly impeding when accelerating with at least $a_{\text{min,accel,merge}}$.

⁵ Beyond that, definitions including the point of no return for the merging vehicle, or the merging vehicle being fully in the common area are possible, but not further considered throughout this work. Adopting the following formulas to those considerations is straightforward.




(a) Vehicle configuration before the merging.



(b) One possibility to define the time of merging: The time when the non-prioritized vehicle enters the prioritized lane. Applicable if merging is obvious, for example when the merging vehicle has no option to cross the prioritized lane.

(c) Another possibility to define the time of merging: The time when the non-prioritized vehicle enters the area from where prioritized lane and merging lane are equal. Applicable if merging is not already obvious by the time the non-prioritized vehicle enters the prioritized lane, for example when the merging vehicle could also cross the prioritized lane (crossing option not visualized here).



(d) Vehicle configuration after the merging: Both vehicles are in the prioritized lane. The origin for the s -coordinate is the point from where prioritized lane and merging lane are equal.


Figure 4.2: Positions and distances during merging: The straight lane is prioritized, vehicles in the lane arriving at an angle have to yield.

Thus, in case this vehicle hits us from behind, it is to blame even though it had right of way. By implication, the prioritized vehicle can assume that we accelerate at least such that comfortably decelerating with $a_{\min, \text{brake}, \text{row}}$ suffices for not violating the longitudinal safe distance at any time. For the detailed calculation we refer to Appendix A.1. An example is given below.

Example 4.2.8. The vehicle on the prioritized road is traveling at $v_{\text{prio}}(t_0) = 28 \text{ m/s} \approx 100 \text{ km/h}$. The ego vehicle is at a full stop $v_{\text{ego}}(t_0) = 0$ just before entering the prioritized lane. The safety parameters are $\rho_{\text{prio}} = 1 \text{ s}$, $a_{\min, \text{brake}} = -7 \text{ m/s}^2$, $a_{\max, \text{brake}} = -8 \text{ m/s}^2$, $a_{\min, \text{accel}, \text{merge}} = 1.8 \text{ m/s}^2$, $a_{\text{brake}, \text{row}, \text{exp}} = -1 \text{ m/s}^2$. Both merging and crossing are possible from the perspective of the prioritized vehicle, the latter only has to react ρ_{prio} after the merging vehicle entered the common path (cf. Fig. 4.2c). The distance along the path of the ego vehicle to this merge point is $d_{\text{init}, \text{merge}}^+ = 7 \text{ m}$ and the length of the ego vehicle is $l_{\text{ego}} = 5 \text{ m}$. When starting to accelerate, the ego vehicle reaches the merge point in approx. 2.8 s at approx. 5 m/s. During this time, the prioritized vehicle travels approx. 78 m. The critical time is $t_{\text{crit}} \approx 5.8 \text{ s}$ after the merge point, independent of the position of the prioritized vehicle. The safe distance for the prioritized vehicle at this time is $d_{\text{safe}, \text{rear}}(t_{\text{crit}}) \approx 47 \text{ m}$, with the prioritized vehicle traveling at approx. 23 m/s and the merging vehicle at approx. 15 m/s. Thus, demanding that $d_{\text{reserve}}(t_{\text{crit}}) > 0$, we obtain $d_{\text{init}, \text{prio}}^+ > 143 \text{ m}$. With the distance traveled until the merging vehicle reaches the merge point, this maneuver can be safely started if the initial distance of the prioritized vehicle to the point from where the paths overlap is larger than 222 m (cf. Fig. 4.3a).

If crossing is not an option, and thus, the merging is obvious from the time the ego vehicle enters the prioritized lane (cf. Fig. 4.2b), this distance is reduced to 194 m (cf. Fig. 4.3b).

With the prioritized vehicle driving at $v_{\text{prio}} = 20 \text{ m/s} \approx 70 \text{ km/h}$, these distances reduce to 131 m and 112 m, and with $v_{\text{prio}} = 14 \text{ m/s} \approx 50 \text{ km/h}$ to 79 m and 66 m, respectively.



- (a) Merging is only obvious once the merging vehicle enters the common path (cf. Fig. 4.2c). The first dashed vertical line is at the time of merging, the second is at ρ_{prio} after, the third line is the time of minimal safety reserve t_{crit} .
- (b) Merging is obvious once the merging vehicle enters the prioritized lane (cf. Fig. 4.2b) at $t = 0$. The first dashed vertical line is at ρ_{prio} after, the second line is the time of minimal safety reserve t_{crit} .

Figure 4.3: Positions including the safe distance during a merging maneuver, as visualized in Fig. 4.2.

Approaching an Intersection

In order to surpass the maximum allowed velocity according to condition C1⁶, i.e. to enter the conflict zone without stopping beforehand, the safe passage according to the above requirements for condition C2(a/b/c) must be guaranteed for some time in the future. Here again, reasonable worst case assumptions for the behavior of other vehicles have to be considered, such as full acceleration up to the speed limit or even some margin above it. Given the current position and velocity of a prioritized vehicle, the result is a minimum average ego velocity, for which a safe passage can be guaranteed w.r.t. this vehicle. Further given a guaranteed minimum ego acceleration $a_{\text{min, accel}}$ up to a certain velocity v_{max} , a minimum ego velocity for which a safe passage can be guaranteed can be computed. Note that, to be able to guarantee the minimum ego acceleration $a_{\text{min, accel}}$, we must be able to guarantee that there will be no need for a deceleration as a reaction to the behavior of other traffic participants.

⁶ Violating condition C1 is also called *passing the point of no return*.

Several overlapping conflict zones can be treated as follows: Condition C1 is taken from the earliest conflict zone, such that stopping in front of one conflict zone, but within another, is prevented. To pass all the overlapping conflict zones, all conditions C2 must be satisfied simultaneously.

Occlusions and Limited Sensor Range

Traffic participants that are currently occluded or outside the sensor range have to be incorporated in the safety consideration. Condition C1 remains a limitation on the maximum safe velocity, as previously introduced. For condition C2, the worst case assumption is a non-observable vehicle at the maximum considered velocity⁷ just at the edge of the occlusion, again resulting in a minimum velocity for the non-prioritized that allows safe passing. Traveling along the edge of condition C1 consists of traveling at constant velocity at the speed limit when far enough away from the conflict zone and full deceleration just before the conflict zone. This behavior is obviously uncomfortable. The probability-based comfort consideration will be covered in Sec. 4.3.

4.2.3 Parallel Lanes

The remaining scenario besides single-lane driving and intersecting lanes is driving in parallel lanes, where lane changes are possible and sometimes necessary. RSS rule 2 requires to “not cut-in recklessly”. Neglecting lane changes of other vehicles, the safe distances to four vehicles have to be considered, as depicted in Fig. 4.4. While path-velocity decomposition (PVD) is not straightforward applicable for lane changes, it can be facilitated by a velocity-dependent transition path between the source and the target lane. PVD within the lanes before and after the lane change is straightforward again.

As previously stated, we propose to require more than just the longitudinal safe distance as cut-in margin for performing a lane change. According to the German law, a lane change may only be performed if the risk of endangering other traffic participants can be preempted [StVO (D), §7 V]. Further, a

⁷ Speed limit or some margin above.

lane change has to be indicated, including the use of the *direction-indicator lamps* [StVO (D), §7 V]. Pek et al. propose for example that the follower in the target lane is responsible from the time where the lane changing vehicle is entirely within the target lane [PZA17]. This, however, would motivate fast lane changes, as they require a smaller distance to the follower in the target lane to be able to start the lane change. Instead, we propose to grant an additional response time ρ_{LC} for the following vehicle in the target lane, such that a safe, non-reckless lane change can be performed via:

- Indicating the lane change desire using the direction-indicator lamps.
- Monitoring that a safe longitudinal distance to the ego vehicle was maintained by the following vehicle in the target lane for at least a lane change response time ρ_{LC} , before the lane change is performed.
- Maintaining a safe longitudinal distance to the leading vehicle in the source lane until the lane change is performed.
- Maintaining a safe longitudinal distance to the leading vehicle in the target lane from the time when the lane change is performed.

As the lane change might need to be abandoned during the lane change response time, for example due to tailgating of the following vehicle in the target lane, the longitudinal safe distance to the following vehicle in the source lane is still necessary. Thus, in order to prevent the following vehicle in the source lane from the belief that the lane change of the ego vehicle is certain, and that a lateral safe distance is sufficient, we propose to restrict omitting longitudinal safety because of lateral safety to the case where the lane changing vehicle has left the lane with a certain proportion of its geometry (cf. Fig. 4.5). By implication, the lane changing vehicle must only leave the lane with that proportion, if the lane change will be completed certainly, in order to be able to return safely. In case the gap is large enough, such that the safe distance for the following vehicle in the target lane is also satisfied for the extended reaction time $\rho_{\text{extended}} = \rho + \rho_{LC}$, the lane change is safe from the start and can be commenced straightaway. The influence of limited sensor range might be that an immediate lane change is not possible, as vehicles with a safe distance w.r.t. ρ_{extended} are outside the sensor range. If the sensor range is so low, that even vehicles for which the regular safe distance would be satisfied are not visible, a lane change is not possible with this sensor setup. The same holds for limited visibility due to occlusions. In this case, however, the visibility




Figure 4.4: Safe longitudinal distances throughout the lane change process: The red safety distances are the responsibility of the ego vehicle (blue) throughout the lane change. The following vehicle in the source lane is responsible to maintain the green safety distance. The responsibility for the yellow safety distance changes throughout the lane change: For a certain time, it is to be maintained by the ego vehicle, subsequently, the responsibility is with the following vehicle in the target lane.




Figure 4.5: Continuation of Fig. 4.4: From the perspective of the following vehicle in the source lane, the situation is both longitudinally (green) and laterally safe (transparent extension). However, omitting longitudinal safety through lateral safety is only allowed once the ego vehicle (blue) has left the source lane with a certain proportion of its geometry, because returning back might still be necessary at the start of a lane change.

might be increased by moving to the very edge and maybe even slightly across the lane boundary.

In reverse, the ego vehicle must react to lane changing vehicles. Thus, a safe longitudinal distance must be maintained to vehicles that indicate their lane change, at the latest when the lane change has been indicated for the time ρ_{LC} , incorporating the own processing delay.

The planning approach considering interaction and potential courtesy of other vehicles is presented in Sec. 4.4. The uncertainty about whether a particular traffic participant is courteous, which is linked to the uncertainty in its future trajectory, is the dominant uncertainty in challenging dense traffic.

4.2.4 Further Scenarios

The previously introduced scenarios cover large parts of the road network. Whether or not additional scenarios are necessary for specific areas can be checked offline in the map. Further scenarios include the merging of parallel lanes, called *zipper merge*, which can be symmetric or asymmetric, and where all vehicles are required to facilitate the zipper merge for vehicles from the other lane, such that the progress alternates between the two lanes. The execution of a lane change on a multi-lane road where additionally yielding to intersecting traffic is required is not intended for safety reasons, thus, the reaction to other lane changers is enough in such combined scenarios.

4.3 Probabilistic Motion Planning based on Independent Predictions

In this section, we describe the approach to deal with uncertainties that do not depend on interaction between the ego vehicle and other traffic participants. Thus, the prediction of other traffic participants can be performed in an upstream module, independent of the planning problem formulation. While the trajectory planning approach is generally backed up by a safety planner (cf. Sec. 4.1), an intervention of the latter mostly involves high decelerations and thus is very uncomfortable. The goal of the probability consideration is to plan comfortable and convenient trajectories, without causing safety planner interventions. In the following, after introducing the basic idea to deal with these uncertainties, we focus on the different sources of uncertainty and their treatment.

4.3.1 General Approach

The underlying assumption for the uncertainty considerations in this section is that the pursued trajectories of prioritized traffic participants do not depend on the ego trajectory. Due to the uncertainties from the unknown future behavior of others, but also due to the lack of knowledge from occlusions and limited sensor range, replanning (cf. Sec. 2.1.8) is crucial to motion planning for

automated vehicles. This consideration leads to the problem formulation of eq. 3.8, as introduced in Chapter 3.

Solving eq. 3.8 for \mathbf{x}_1^{e*} in the general case involves considering infinitely many belief updates during the recursive computation of alternative trajectory continuations \mathbf{x}_2^{e*} . Further, the evaluation of \mathbf{x}_1^e involves the integration over all possible trajectory ensembles for others, weighted with a PDF (cf. eq. 3.6).

To solve the POMDP efficiently, we propose to focus on several likely *trajectory ensemble sets* \underline{X} for other traffic participants. Instead of describing the trajectory of an object as stochastic process \mathbf{X}^{oi} , i.e. as a mapping from time t to a random variable $\mathbf{X}^{oi}(t)$ for the state, a *trajectory set* \mathcal{X}^{oi} is used. That is, a trajectory is described with a mapping from time t to a bounded set of states $\mathcal{X}^{oi}(t)$ per object i , without any information regarding the probability distribution within this set. Thus, a trajectory ensemble set \underline{X}^o contains information about the bounds of the states of all objects over time. Initial beliefs are described as trajectory ensemble sets \underline{X}_k^o . As in Sec. 3.3, in addition to the belief over states in classical POMDP formulations, these beliefs over trajectories also include knowledge about the state transition probabilities. A trajectory ensemble set \underline{X}_l^o , along with the timestamp t_l that specifies the start of its consideration and a probability p_l that it has to be considered, is called *belief update* $(t_l, p_l, \underline{X}_l^o)$ w.r.t. a prior trajectory ensemble \underline{X}_k^o . That is, with probability p_l , we have to consider \underline{X}_l^o instead of \underline{X}_k^o , and the information leading to this update is available at t_l . The set of all considered belief updates w.r.t. a base set \underline{X}_k^o is denoted \mathcal{L}_k . A base trajectory ensemble set \underline{X}_k^o along with possible belief updates is denoted *belief progression* $\mathbb{X}_k^o = \{\underline{X}_k^o, (t_l, p_l, \underline{X}_l^o), l \in \mathcal{L}_k\}$. The set of all considered belief progressions k is denoted \mathcal{K} .

This simplification, which can be performed independently of the planning problem by the upstream scene understanding module, corresponds to selecting possible paths through a belief tree, while not allowing recursive branching for now. Unfavorable ensembles can often be subsumed because they lead to the same result, such as stopping in front of an intersection. Further, the cost estimation is only performed up to a certain planning horizon t_H .

For a base trajectory ensemble set \underline{X}_k^o along with belief updates to \underline{X}_l^o at time $t_l > t_0$ with probability p_l , $l \in \mathcal{L}_k$, the cost estimation is

$$\begin{aligned} \mathbf{J}_{\text{total},k}^{\text{exp}}(\mathbf{x}_1^e) &= \int_{t_0 + \Delta t_{\text{plan}}}^{t^H} J(\mathbf{x}_1^e, \underline{X}_k^o, t) dt \\ &+ \sum_{l \in \mathcal{L}_k} p_l \int_{t_l + \Delta t_{\text{plan}}}^{t^H} J(\mathbf{x}_l^{e*}, \underline{X}_l^o, t) - J(\mathbf{x}_1^e, \underline{X}_k^o, t) dt, \end{aligned} \quad (4.8)$$

where $\Pr(\mathbf{x}_1^e) + \sum_{l \in \mathcal{L}_k} p_l = 1$ must hold, such that all possible future trajectory ensemble evolutions are considered.

In words, the cost of \mathbf{x}_1^e over the full planning horizon is only considered weighted with probability $1 - \sum_{l \in \mathcal{L}_k} p_l$. For each possible reaction l , the cost of \mathbf{x}_1^e is only considered as long as it is expected to be pursued, i.e. up to $t_l + \Delta t_{\text{plan}}$. For the remaining time within the planning horizon, the cost of the alternative continuation \mathbf{x}_l^{e*} is considered. The cost for each reaction l is weighted with probability p_l .

Recursive branching can be incorporated as follows. Instead of updating to trajectory ensemble sets \underline{X}_l^o in the belief progressions \mathbb{X}_k^o , the updates include another belief progression, such that $\mathbb{X}_k^o = \{\underline{X}_k^o, (t_l, p_l, \mathbb{X}_l^o), l \in \mathcal{L}_k\}$. For the computation of the term $\int_{t_l + \Delta t_{\text{plan}}}^{t^H} J(\mathbf{x}_l^{e*}, \underline{X}_l^o, t) dt$ in eq. 4.8, we recursively evaluate eq. 4.8, now starting at t_l instead of t_0 . The recursion ends when the belief progression consists only of a trajectory ensemble set, i.e. once $\mathcal{L}_k = \{\}$.

The different considered belief progressions \mathbb{X}_k^o , $k \in \mathcal{K}$ are independent simplifications of the belief tree. Since they each underestimate the available information from the belief tree, the expected cost can be approximated as the minimum of the cost estimations over all considered belief progressions, i.e.

$$\mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e) = \min_{k \in \mathcal{K}} \mathbf{J}_{\text{total},k}^{\text{exp}}(\mathbf{x}_1^e). \quad (4.9)$$

Finally, the optimal trajectory \mathbf{x}_1^{e*} is the one that minimizes this cost, i.e.

$$\mathbf{x}_1^{e*} = \arg \min_{\mathbf{x}_1^e} \mathbf{J}_{\text{total}}^{\text{exp}}(\mathbf{x}_1^e). \quad (4.10)$$

For the probability of a belief update p_t , the history of previous belief updates has to be accounted for, as illustrated in the example below.

Example 4.3.1. Recall ex. 3.3.1: Imagine, there were only two discrete trajectories possible for the white vehicle in Fig. 3.2: One for turning right \underline{X}_r^o and one for going straight \underline{X}_s^o . They diverge in 4 s, such that after 4 s, we have certainty about the path. Let the probability that the vehicle turns right be 50%.


When attempting a trajectory based on the right turn assumption, there is one possible belief update that the vehicle does not turn right, but goes straight instead, with $p = 50\%$ at $t = 4$ s:

$$\left\{ \underline{X}_r^o, \left(4 \text{ s}, 50\%, \{ \underline{X}_s^o, () \} \right) \right\}.$$

Further knowing that 90% of the vehicles that turn right indicate within 2 s from now, and all vehicles that indicate turn certainly, another possible belief update is that we do not have certainty that the vehicle turns right after 2 s, with $p = 100\% - 90\% \cdot 50\% = 55\%$. In the latter case, the probability of the belief update after 4 s changes to $\frac{50}{55} \approx 91\%$ for \underline{X}_s^o :

$$\left\{ \underline{X}_r^o, \left(2 \text{ s}, 55\%, \left\{ \underline{X}_r^o, \left(4 \text{ s}, 91\%, \{ \underline{X}_s^o, () \} \right) \right\} \right) \right\}.$$

The belief tree for the latter scenario is depicted in the following:



The planning horizon t_H should be chosen such that trade-offs of whether to stop at or continue through an intersection, or obligations such as to stop at a red light, are realized early enough to allow for comfortable decelerations.⁸ On the other hand, Ziegler proposed to treat stopping outside classical optimization, with a control approach [Zie17, p.93]. Following this proposal, stopping trajectories along with their duration can be computed by forward integration of such a control approach. Further, in the relevant scenarios for this probability consideration, path-velocity decomposition is applicable, and we propose its usage to reduce the dimension of the optimization problem.

Comparing the cost of a trajectory that likely allows convenient passing of an intersection but might result in an uncomfortable deceleration with the cost of a trajectory that certainly stops comfortably in front of the intersection is not trivial. Stopping at an intersection is convenient if prioritized vehicles are currently passing, but it is not convenient if continuing would be safely possible, for example. In the analysis of human driving behavior in [11], we argued that cost functions that do not yield zero cost for some trajectory parts can hardly be recovered using inverse reinforcement learning. Similarly, it is questionable whether the stopping/deceleration behavior of a human is different when having to wait for 10 s compared to having to wait 15 s at a stop sign, which would be the case if waiting would induce cost. With the previously stated insights from [Zie17], we propose the following heuristic instead of performing this trade-off between comfort and convenience based on cost: Based on *reference trajectories* \mathbf{x}_1^e , that are computed using forward integration of a driver model, and that are safe and traffic rule compliant w.r.t. trajectory ensemble set \mathcal{X}_k^o , we compute the necessary decelerations as reactions to belief updates $(t_l, p_l, \mathbb{X}_l^o)$ and verify the safety of those reactions. In presence of intersections, the *point of no return* is the focus of the safety consideration, since this point may only be passed if the intersection can be entered safely. The pursued trajectory is chosen as the fastest feasible reference trajectory where the maximum necessary decelerations do not violate *comfort margins*. The comfort margins are defined via the following rules:

⁸ The availability of emergency maneuvers is always monitored by the safety planner, including the execution of those maneuvers, if necessary. Thus, a lower bound for Δt_H is not mandatory to ensure safety. Yet, the horizon to allow for comfortable decelerations is commonly larger than this fictive lower bound.

Rule 1. Given a function $p_{\max} : \mathbb{R}^- \rightarrow [0, 1]$ that defines the maximum allowed probability over decelerations. A reactive deceleration a_{react} with probability p_{react} is only accepted if

$$p_{\text{react}} \leq p_{\max}(a_{\text{react}}). \quad (4.11)$$

Since the reference trajectories can already include decelerations themselves, only reactive decelerations with larger absolute value than the decelerations of the reference trajectories are considered. Also, the additionally required decelerations ($a_{\text{react,add}} = a_{\text{react,max}} - a_{\text{reference,max}}$) are checked as follows.

Rule 2. Given a function $p_{\max,\text{add}} : \mathbb{R}^- \rightarrow [0, 1]$ that defines the maximum allowed probability over decelerations. An additional deceleration $a_{\text{react,add}}$ with probability p_{react} is only accepted if

$$p_{\text{react}} \leq p_{\max,\text{add}}(a_{\text{react,add}}). \quad (4.12)$$

Further, the comfort margins reflect elements of perceived safety:

Rule 3. Comfortable trajectories must not violate the minimum time of zone clearance.

Other comfort aspects such as a maximum lateral and longitudinal acceleration, a maximum jerk and a minimum time and space headway to other vehicles are already considered in the creation of the reference trajectories, such that no reference trajectory violates them. The result is basically a trade-off between smaller, but potentially unnecessary early decelerations in the reference trajectory, and potentially larger reactive decelerations in the future.

Using this heuristic, the belief updates $(t_l, p_l, \mathbb{X}_l^o)$ can be limited to those that are unfavorable compared to \mathbb{X}_k^o , i.e. that require a reactive deceleration. In case a belief update allows for a reactive acceleration, this does not violate the defined comfort margins. Further, very rare cases are always covered by the safety planner and thus do not necessarily need to be considered here.⁹

⁹ If, for example, an emergency braking maneuver as reaction to the emergency braking of the vehicle ahead would be considered uncomfortable even if it only occurs with probability 10^{-10} , driving with reasonable time or space headway would not be possible at all.

Note that the knowledge about the future progression of uncertainties is particularly important to probabilistic motion planning, as illustrated in the example below.

Example 4.3.2. Assume a pure V2X-based perception. We know that there is a persistent adversarial attack with one false object per existing object in the object list. Even though we know that this lowers the existence probability of every object to 50%, we have to act as if the existence probability of every object was 100%. This is due to the fact that, besides knowing the current existence probability of every object is 50%, we also know that this probability will not change within the planning horizon, due to the adversarial attack. In other words, since there is no belief update available, we can only consider the belief progression $\{\underline{X}_{\text{all}}^o\}$ with a single trajectory ensemble set that contains all objects.

Regarding the future path of an object, reactions to the other objects' trajectories along all possible paths have to be included via belief updates $(t_l, p_l, \underline{X}_l^o)$ up to points in time where it is expected that some paths can be excluded with certainty. From these points on, reactions to particular behavior of others can be excluded: If the initial trajectory \mathbf{x}_1^e is still pursued by then, these behaviors did not occur. Otherwise, the ego vehicle would already have reacted with a new trajectory, and \mathbf{x}_1^e would no longer be valid. The decision to exclude paths can for example be based on or derived from the (missing) usage of the direction indicator lamps. As, however, this decision is safety relevant, these conclusions have to be drawn with care and the responsibility in case of false or missing indicator usage has to be agreed upon as part of the general safety concept.

In contrast to the proposed approach, some approaches that perform *decision postponing* only consider the current probability estimate and implicitly expect reaching certainty at the time of the subsequent replanning, corresponding to a QMDP approach. For specific combinations of cost weighting and probability progression, this yields promising results, as in [ZLCT16]. However, cases where certainty is never reached or reached too late are not considered. In edge cases, for example where the probability of a particular path for an object is constantly very low but non-zero, such approaches only react once the safety constraints due to an imminent collision come into play. Under these hard constraints, the cost is neglected and the reaction is most likely very harsh

and uncomfortable. If this constant uncertainty over time was known from the beginning, as in the previous example, the resulting behavior is suboptimal.

4.3.2 Single Lane

The dominating uncertainty in single-lane scenarios is the existence uncertainty. Frequent decelerations due to so-called *ghost objects*, i.e. objects that do not actually exist, but are detected by the perception module, have to be avoided to ensure a comfortable ride. A similar uncertainty arises from the missing information outside the visible range. As previously said, the uncertainty in the current state or the future trajectory of other objects is commonly not hampering a comfortable ride, as presented by series ACC systems. The latter, however, deal with ghost objects simply by tuning for a low false-positive rate, such that almost no ghost detections occur, while false-negatives, i.e. missing detections, are to be dealt with by the driver.

With the presented approach, the existence uncertainty can be treated by considering several reference trajectories between ignoring and incorporating the potential ghost object in planning. Given the sufficient comfort of lane following models such as ACC or IDM, reference trajectories can be planned by forward integration of such models for the cases of non-existence and existence, and linearly interpolating in-between both cases. As explained before, the heuristic is then to choose the fastest reference trajectory that does not violate comfort margins (cf. Sec. 4.3.1, rules 1-3). Note that the important information are belief updates that lead to certainty about the non-existence of the potential ghost. If the existence probability remains non-zero, since safety is never put at risk, the safety planner would eventually react and cause a clearly uncomfortable emergency deceleration.

Limited Visibility

With this approach, uncertain information about far away objects, for example from camera only, can be used for comfortable deceleration also to objects outside the so-called *visible range*. To recapitulate, in the visible range we require the detection of actually existing objects with a non-zero probability for safety reasons. Still, information beyond the visible range might be available.

Note, however, that this planning approach in general assumes valid existence probabilities without systematic errors, while the determination of those probabilities is to be done externally, independent of the planning problem. Due to the fact that safety distances must be maintained to all objects with non-zero existence probability, volatile close objects inevitably lead to a jumpy and jerky behavior of an automated vehicle.

Regarding limited visibility beyond the sensor range or due to occlusions, such that no sensor information is available, general traffic data can be used. Such data can include the knowledge about frequent traffic jams behind certain bends, but also live traffic data, allowing for an early and more comfortable reaction to reported traffic jams. In addition to that, information about frequently occurring slow traffic, such as trucks leaving a construction site, can be made use of. Including this information in the approach is realized by adding possible belief updates $(t_l, p_l, \mathbb{X}_l^o)$. The reaction to such possible belief updates is an appropriate reduction of the velocity.

4.3.3 Intersecting Lanes

At intersecting lanes, multiple route options of others enlarge the uncertainty in their future behavior. Further, the influence of occlusions and limited sensor range along prioritized roads must be considered.

For a vehicle that has unrestricted right of way along its desired route¹⁰, the single-lane consideration is applicable. Additional probabilistic reasoning is necessary for traffic turning away from this lane, traffic merging onto this lane and potential violation of the right of way. For all of those cases, eq. 4.8 is applicable.

Further, analytic solutions to simplified problem formulations, for example assuming constant acceleration for a vehicle merging in front and constant deceleration for the ego vehicle, as requested by the safety concept, can be computed, as introduced in the previous section. Similarly, for vehicles that

¹⁰ Vehicles turning from a priority lane commonly have right of way over vehicles entering an intersection from other lanes, except if they turn through oncoming prioritized traffic. In the latter case, they have to yield to oncoming vehicles, so their priority is not unrestricted.

can be ignored according to the safety concept¹¹ unless an accident is imminent, dangerous misbehavior of the crossing vehicle can be considered as a possible belief update with low probability. For diverging lanes, far-away objects that decelerate to turn away from the prioritized lane should not cause an over-cautious reaction in case it is likely that they have left the lane anyway by the time the ego vehicle would arrive there. This consideration is already included in eq. 4.8.

For vehicles that have to yield at intersecting lanes, meaningful reference trajectories only enter the conflict zone at times when it is not occupied by prioritized vehicles. Apart from this additional requirement for the reference trajectory generation, the general approach as in eq. 4.8 is applicable. Yet, the safety verification w.r.t. all vehicles along their potential routes is more complex. For intersections with many vehicles involved, a meaningful restriction to likely but also distinct trajectory ensemble sets along with possible belief updates is crucial to avoid a multitude of similar computations.

Limited Visibility

Considering limited visibility at intersections, an object-based approach is not expedient. An appropriate reaction to occlusions, even if no vehicle actually appears from behind them, is commonly a deceleration below the desired velocity for the non-occluded case. This ensures a comfortable stop in case an object appears. The deceleration is followed by an acceleration back to the desired velocity, which is only pursued in case no object appeared. The reaction to objects appearing at the sensing edge is then a further deceleration to let this object pass, in case going first is not safely possible. The maximum velocity with which an occluded intersection can be passed is defined by the safety concept. Traveling along the edge of the safety limit, however, consists of constant velocity at the speed limit, when far enough away from the conflict zone, and full deceleration just before the conflict zone, which is clearly uncomfortable. Once the visibility is sufficient and no object appeared, the vehicle can accelerate again. In case an object occurs which the ego vehicle needs to yield to, the full deceleration has to be continued to a full stop.

¹¹ Vehicles that cross the path of the ego vehicle closely have to ensure that they leave the path before the ego vehicle arrives (cf. ex. 4.2.2).




Figure 4.6: Exemplary occupancy of a conflict zone for crossing traffic (from [8], © 2019 IEEE): Black denotes the times in which the zone is occupied/blocked by prioritized traffic. The time of zone clearance before and after this occupancy is marked in orange. The expected duration of the passage of the ego vehicle is marked in blue. The scope for the ego vehicle’s passage is the remaining time (blank).

A comfortable trajectory should decelerate earlier and to a lower velocity. For the choice of a comfortable trajectory, a trade-off is to be made. In case no object appears, any deceleration to a velocity below the safety constraint was unnecessary. While a gentle deceleration might not cause discomfort, it is certainly inconvenient to frequently decelerate unnecessarily. On the other hand, any deceleration above the one leading to a comfortable stop leads to an uncomfortable deceleration in case we have to yield.

While this consideration is similar to eq. 4.8, the variety of possible object configurations within the occlusion and the dependency of the visibility on the ego motion call for a different treatment. We recall the idea of subsuming configurations of others that lead to the same \mathbf{x}_l^{e*} from eq. 4.8 and use a *simplified reaction model*. Further, in addition to the approach presented in [8], the knowledge of the replanning moments in time is incorporated. These are the only moments when newly detected objects can be considered. This reaction model corresponds to a belief update containing a virtual object that blocks the intersection.

In order to estimate the probability that we have to react in a certain replanning step, the occupancy density of the conflict zone can be employed, i.e. the fraction of time the conflict zone is occupied by prioritized vehicles. This information can be inferred from monitoring particular intersections or from fleet data, for example. From every gap in the occupancy of the conflict zone, we subtract the desired time of zone clearance for the prioritized vehicles and the ego vehicle. Subsequently, we subtract the time that the ego vehicle needs to pass the conflict zone. If the gap time is still positive, it now represents the potential scope $\Delta t_{\text{scope gap } i}$ for our traversal of gap i , as visualized in Fig. 4.6.

With this information, the probability that an arbitrary trajectory \mathbf{x}_1^e can be pursued unaffectedly is

$$\Pr(\mathbf{x}_1^e \text{ is unaffected}) = \frac{\sum_i \Delta t_{\text{scope gap } i}}{\Delta t_{\text{total}}}, \quad (4.13)$$

where the numerator denotes the potential scope for the traversal from all available gaps i in Δt_{total} . The probability that we have to deviate from our desired trajectory,

$$p_{\text{react}} = \Pr(\text{ego must react}) = 1 - \Pr(\mathbf{x}_1^e \text{ is unaffected}), \quad (4.14)$$

can be subdivided into probabilities that we have to react per replanning step: With the probability $p_{\text{vis},k}$ that a *relevant vehicle*¹² is visible at replanning step k , and the general reaction probability p_{react} , the probability of a reaction in a certain replanning step can be computed as

$$p_{\text{react},k} = p_{\text{react}} \cdot (p_{\text{vis},k} - p_{\text{vis},k-1}), \quad (4.15)$$

where $(p_{\text{vis},k} - p_{\text{vis},k-1})$ is the probability that a relevant vehicle *becomes* visible at replanning step k . In the following, an exemplary approach to compute $p_{\text{vis},k}$ is presented.

The computation is realized using information about the visible area at a certain replanning moment, and information about where relevant vehicles might be located. The latter consideration is limited to vehicles in a certain velocity range $[v_{\text{min}}, v_{\text{max}}]$. While v_{max} can be taken from the safety concept, for example some margin above the speed limit, v_{min} should lie below the speed limit. Choosing v_{min} too high, however, only leads to the estimation of later reactions and thus to an over-approximation of the probabilities of higher decelerations. From this velocity range, we can compute the bounds of relevant areas, which are visualized in Fig. 4.7. Close to the conflict zone, there is an area where vehicles do not affect the ego vehicle, because even slow vehicles in this area will already have left the conflict zone early enough (light gray in Fig. 4.7). Its bound depends on v_{min} , the maximum considered vehicle length l_{max} , and the time of zone clearance for the ego vehicle. Similarly, there is a distance from where even vehicles at v_{max} will only arrive at the conflict zone

¹² A vehicle that causes a reaction.

at the required time of zone clearance¹³ after the ego vehicles has left it (dark gray in Fig. 4.7). The relevant area is in-between those bounds. Far away from the conflict zone, we most likely have no visibility of the relevant area. At the latest just before the conflict zone, we should have full visibility of the relevant area¹⁴, such that a reaction at a subsequent replanning step can be excluded.

Neglecting the length of objects, which again leads to an over-approximation¹⁵, and assuming a uniform distribution over the relevant area, the probability that the relevant object is visible in step k is

$$p_{\text{vis},k} = \frac{d_{\text{visible},k}}{d_{\text{relevant},k}}, \quad (4.16)$$

with the area $d_{\text{visible},k}$ that is relevant and visible at step k (green in Fig. 4.7) and the area $d_{\text{relevant},k}$ that is relevant at step k (orange and green in Fig. 4.7).¹⁶

The computed reactive decelerations and their estimated probability are then added to the reactive decelerations due to belief updates. The presented heuristic to determine the optimal trajectory remains unchanged (cf. Sec. 4.3.1, rules 1-3).

¹³ In case of merging instead of crossing, the distance is computed using the required initial distance for merging with v_{max} as explained in Sec. 4.2.2.

¹⁴ Otherwise, approaches to slowly advance into the conflict zone to gain more visibility have to be applied.

¹⁵ The rear end of an object, and thus its length, is only relevant for those objects that leave the conflict zone before the ego vehicle enters. If this length is neglected, we assume that we could only react once the rear end of the object becomes visible. Consequently, a later reaction is assumed, leading to a stronger deceleration.

¹⁶ Better estimates of the distribution of relevant vehicles, if available, can be incorporated here.




Figure 4.7: Relevant visibility over the course of the trajectory of the blue ego vehicle through an intersection occluded by a building (black): The current visibility of the ego vehicle is depicted with a dashed line, originating in the vehicle center. Only vehicles in specific areas affect the planned ego trajectory. The part of the relevant area that is visible is depicted in green, the occluded part is depicted in orange. Vehicles that are further away (dark gray zone) would not reach the conflict zone before the ego vehicle has left, even at the maximum considered velocity. Vehicles that are closer (light gray zone) will have left the conflict zone by the time the ego vehicle arrives, even at the minimum considered velocity.

4.4 Courtesy-Aware Lane Changes

In scenarios where courtesy is necessary, such as lane changes in dense traffic, an independent prediction of other traffic participants is no longer expedient. On the other hand, since waiving the right of way is rather unusual and requires a prior arrangement with the prioritized traffic participant [StVO (D), §11 III], the courtesy consideration is limited to lane changes.

The decision for a lane change is assumed to be taken by an upstream module and not part of this work. Examples for approaches to determine whether or not optional lane changes are beneficial are given in [9]. Mandatory lane changes are determined in the navigational layer. Assuming the need for a lane change was determined, the lane change process itself is commonly divided into gap selection, gap approach, gap evaluation and the subsequent merge into the gap [HSX⁺18]. The focus of this section is the gap evaluation and the subsequent merge into the gap, while a simple heuristic for gap selection and gap approach is implemented to show the potential of the approach.¹⁷

The basic idea behind the proposed solution is as follows. The key uncertainty in lane change scenarios is whether or not vehicles in the target lane are courteous. While the courtesy of others certainly also depends on the behavior of the lane changing vehicle, we assume that this influence is negligible if the lane change approach is well indicated and would not induce a large deceleration on the suspectedly courteous vehicle. Further, we want to ask for courtesy, but not force it. Thus, cutting in just in front of other vehicles, even if we could anticipate that it would be safe since other vehicles certainly want to avoid collisions and would react appropriately, is not regarded. As a result, the lane change desire is clearly communicated. Subsequently, unless the target lane is free, the reaction is investigated in order to find out whether the respective traffic participant is courteous or not, i.e. in order to reduce the key uncertainty. Moreover, granting an additional reaction time is part of the safety concept for lane changes, as proposed in Sec. 4.2.3.

¹⁷ As the gap selection is not safety critical, but the probability that a certain traffic participant is courteous potentially depends on various features, machine learning might be well-suited for the upstream module.

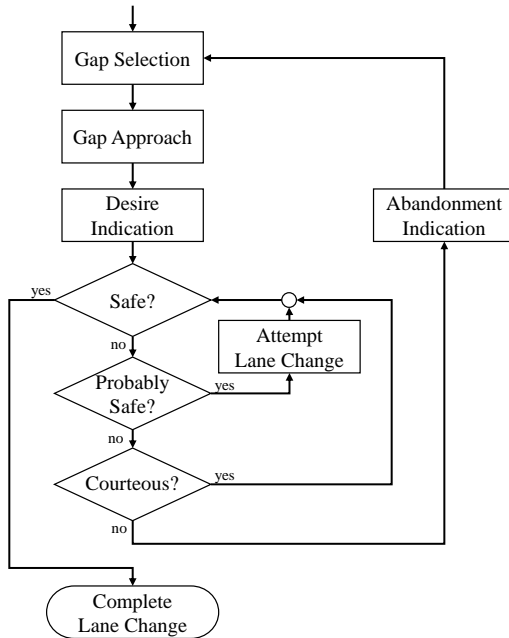


Figure 4.8: Lane change process overview (from [9], © 2019 IEEE).

An overview over this rule-based approach is given in Fig. 4.8. After a gap has been approached longitudinally, the lane change desire is indicated. While the use of the direction-indicator lamps is mandatory according to the traffic rules, additional communication, such as moving towards the target lane within the source lane, can be used.¹⁸ After the desire has been indicated, the safety of the lane change has to be investigated. Here, three situations can be distinguished:

1. The situation is safe, even w.r.t. the enhanced lane change reaction time $\rho + \rho_{LC}$ (cf. Sec. 4.2.3). We can perform the lane change immediately.

¹⁸ A thorough investigation of how the desire can be communicated effectively from a psychological perspective is outside the scope of this work. Yet, additional hints from this field can be integrated into the approach.

2. The situation is currently safe w.r.t. the target lane, and assuming the most likely prediction (independent of the ego vehicle) it will still be safe in ρ_{LC} . This situation is called *probably safe*. We still have to wait but presumably do not need courteous behavior of the following vehicle in the target lane. We can already attempt the lane change, for example by moving further towards the target lane.¹⁹
3. The situation is currently not safe w.r.t. the target lane. We rely on the following vehicle in the target lane being courteous. In this case, a cooperative prediction approach must be used, with a repeated check whether the respective traffic participant complies with that model. In case of compliance, the situation will eventually become *probably safe*. Otherwise, the abandonment of the lane change into this gap has to be communicated, and the approach has to be repeated for another gap.

While we hope for courtesy of other traffic participants to facilitate lane change maneuvers, the same courteous behavior is expected by the ego vehicle. Using lane following models, an exemplary solution is to check the necessary deceleration, regarding the vehicle with lane change desire as the leading vehicle. If this deceleration is within a comfortable range, we can decelerate accordingly. Note, however, that the desire of others, along with an estimation of whether they actually intend to merge just in front of us, is a prerequisite to facilitate courtesy of the ego vehicle.

To summarize, the approach is based on trial and error regarding the courtesy estimation of others. Regarding the lane change itself, it is rule-based, such that the behavior can be designed in a way that is comprehensible and pleasant for humans. Exploiting the sense of security of other traffic participants in order to force “courtesy”, which could be the result of egocentric reinforcement learning or the solution to egocentric POMDP formulations, is precluded.

¹⁹ Here again, a thorough investigation of how the attempt can be communicated effectively from a psychological perspective, in contrast to simply communicating the desire, is outside the scope of this work. Yet, additional hints from this field can be integrated into the approach.

4.5 Cooperative Global Optimum Approach with Model Compliance Consideration

In some scenarios, where the right of way is not clearly defined, mutual cooperation between traffic participants is required to find convenient solutions, similar to the previous section, where courtesy was required. Again, as waiving the right of way is rather unusual and requires a prior arrangement with the prioritized traffic participant [StVO (D), §11 III], we do not focus on cooperation that violates the prescribed right of way.

Since the interaction with other traffic participants is key to cooperative scenarios, the prediction has to be integrated into the planning problem. As known from game theory, globally optimal solutions can be unveiled if the agents trust each other. Thus, one of the goals of this approach is, similar to the previous sections, to behave comprehensibly for others. In order to model the behavior of other traffic participants, we assume that they are rational agents that also seek for a convenient, comfortable and safe trajectory. With this assumption, we strive to find the globally optimal solution to the multi-agent planning problem. Yet, we have to keep in mind that, if this solution is ambiguous, or if we misestimate the behavior or intend of other vehicles, others might deviate from our notion of the globally optimal solution. The same might happen if our behavior is misinterpreted by others. Thus, a fallback needs to be considered, along with the probability that it becomes effective.

In the following, we first derive the solution to the ideal case of cooperation based on a known multi-agent cost function, before we generalize the approach. As generalizations, we consider uncertainties regarding the current state and the ambiguity of the global optimum. To deal with these uncertainties, we include a model compliance check with fallback plan consideration.

Global Optimum

The multi-agent cost function²⁰ for the globally optimal case, mapping from a trajectory ensemble $\underline{\mathbf{x}}$ to a scalar, is defined as

$$\mathbf{J}_{\text{total}}^{\text{go}}(\underline{\mathbf{x}}) = \sum_i \left(\mathbf{J}_{\text{total}}^{i,0}(\mathbf{x}^i) + \sum_{j \neq i} \mathbf{J}_{\text{total}}^{i,j}(\mathbf{x}^i, \mathbf{x}^j) \right) \quad (4.17)$$

with cost term $\mathbf{J}_{\text{total}}^{i,0}$ for the singleton trajectory cost of vehicle i and cost term $\mathbf{J}_{\text{total}}^{i,j}$ for pairwise cost for vehicle i due to vehicle j . Here, the cost $\mathbf{J}_{\text{total}}^{\text{go},e} = \mathbf{J}_{\text{total}}^{e,0} + \sum_{j \neq e} \mathbf{J}_{\text{total}}^{e,j}$ corresponds to the cost $\mathbf{J}_{\text{total}}$ for the ego vehicle $i = e$, that was used in previous sections.

The globally optimal trajectory ensemble is the one that minimizes $\mathbf{J}_{\text{total}}^{\text{go}}(\underline{\mathbf{x}})$. For complex cost functions $\mathbf{J}_{\text{total}}^{i,0}$ and especially $\mathbf{J}_{\text{total}}^{i,j}$, analytic solutions to the minimization might not exist. Further, since the decision to be taken is most likely of combinatorial nature (“Who goes first?”), local optimization approaches are not applicable. On the other hand, strict optimality in the solution of this ideal case is not necessary, since we can only trust in this solution in case it is unambiguous, as explained in the next section. Thus, approximate solutions, such as Monte Carlo sampling approaches, suffice: The approximate solution k^* out of K samples $\{\underline{\mathbf{x}}^1, \dots, \underline{\mathbf{x}}^K\}$ must fulfill

$$\mathbf{J}_{\text{total}}^{\text{go}}(\underline{\mathbf{x}}^{k^*}) \leq \mathbf{J}_{\text{total}}^{\text{go}}(\underline{\mathbf{x}}^k) \forall k \in [1, K]. \quad (4.18)$$

Impact of Uncertainties

Previously, we assumed that the global optimum is unique and will be pursued, such that the cost is definitely $\mathbf{J}_{\text{total}}^{\text{go}}$, respectively the proportion of the ego vehicle $\mathbf{J}_{\text{total}}^{\text{go},e}$. Uncertainties were not considered. In mixed traffic, however, we face several uncertainties that might jeopardize the optimal solution, as outlined in Chapter 3. In ambiguous situations, even slight differences in the sensed state might lead to the global optimum lying within a different homotopy

²⁰ Reminder: $\mathbf{J}_{\text{total}}$ denotes the cost over the whole planning horizon.

class. Further reasons for deviations from the previously approximated globally optimal solution include:

- A traffic participant does not act according to a stationary policy.
- Wrong estimation of the quality measure or destination of a traffic participant by the ego vehicle.
- Wrong estimation of the quality measure or destination of the ego vehicle by a traffic participant.

As a consequence of this non-compliance with our model, the optimal cost is not reached with certainty.

Reaction to Uncertainties

Knowing that we might have to deviate from the globally optimal plan, we have to consider how likely this deviation is, and what cost we face in case of a deviation. Denoting the probability that the globally optimal plan will be pursued with p_{go} , the expected cost can be defined as

$$\mathbf{E}(\mathbf{J}_{\text{total}}^e) = p_{\text{go}}\mathbf{J}_{\text{total}}^{\text{go},e} + (1 - p_{\text{go}})\mathbf{E}(\mathbf{J}_{\text{total}}^{-\text{go},e}). \quad (4.19)$$

However, p_{go} as well as $\mathbf{E}(\mathbf{J}_{\text{total}}^{-\text{go},e})$ are unknown. The probability of exactly a certain trajectory ensemble $\underline{\mathbf{x}}^k$ is infinitesimal. Thus, instead of distinguishing whether or not $\underline{\mathbf{x}}^{\text{go}}$ is exactly followed, we propose to consider the probability that our model is correct in the sense that it yields a result in the correct homotopy class. This probability is denoted by p_{mc} . For the expected cost in case our model is incorrect, we consider a *conservative fallback plan* (cfb). The latter is scenario-dependent. The cost estimation with this model compliance consideration is then

$$\mathbf{E}(\mathbf{J}_{\text{total}}^e) = p_{\text{mc}}\mathbf{J}_{\text{total}}^{\text{go},e} + p_{-\text{mc}}\mathbf{E}(\mathbf{J}_{\text{total}}^{\text{cfb},e}). \quad (4.20)$$

We argue that trust and comprehensibility are crucial to cooperation. Thus, in contrast to Sec. 4.3, where we could modify our ego trajectory to minimize the expected cost, we only choose between pursuing the globally optimal solution $\underline{\mathbf{x}}^{\text{go}}$ or switching to the fallback. The idea behind this distinct choice is that

the latter is believed to be obvious to the other traffic participants such that, for example, in case we switch to the fallback, they can use this information to avoid deadlocks. When following the global optimum $\underline{x}^{\text{go}}$, at a certain *point of no return* the original conservative fallback plan becomes infeasible, for example because braking in front of the conflict zone is no longer physically possible. Still, because we never enter states that are unsafe, we can decelerate and come to a full stop within the conflict zone. This behavior involves the same uncomfortable full braking, but is even less convenient because one of the vehicles would have to back up out of the conflict zone and this takes significantly more time than just waiting for the other to pass. Hence, the trade-off is whether (a) to hope for the cooperative solution but risk an inconvenient stop (potentially even within the conflict zone) or (b) to certainly pursue the also inconvenient stop in front of the conflict zone.

For the determination of the probability p_{mc} whether our model is correct and the expected cost of the reaction $\mathbf{E}(\mathbf{J}_{\text{total}}^{\text{cfb,e}})$, many approaches are conceivable. In order to show the potential of the proposed method, we present one heuristic approach, in addition to the one presented in [6], in Sec. 5.2.3.

4.6 Transitions between Scenarios

In the previous sections, we presented approaches to enhance the concept of responsibility sensitive safety and to plan comfortable and convenient trajectories in different scenarios. The transition between scenarios can be implemented using state machines or arbitrators, but is outside the scope of this work. Yet, we want to briefly explain why their interoperation should not impose any problems.

Firstly, the influence of upcoming scenarios is limited by the maximum allowed velocity and the comfortable deceleration. In other words, knowing that there is a stop sign in 1 km does not affect the motion planning decisions in the very moment. Further, the single-lane scenario is the basis of both multi-lane and intersecting lane scenarios, such that single-lane safety is considered in both other scenarios. The combination of multi-lane and intersection scenarios does occur, but often such that the multi-lane road has right of way or such that the intersection is controlled by traffic lights. In scenarios where prioritized traffic is intersecting with the ego vehicle's multi-lane road, ego lane changes can be

avoided while the reaction to other lane changers is still possible. Lastly, as already mentioned for the safety concept in Sec. 4.2.4, all information about the road layout and the traffic rules is available in the map, such that the scenario transitions for any desired route can be checked up front.

5 Evaluation

As described in Sec. 1.2, the goal of this work is to design a motion planning framework that yields convenient motion plans for automated vehicles in mixed traffic, while accounting for the need of safety guarantees. In order to show the potential of our approaches, we evaluate them in several scenarios. A key point throughout our reasoning is that striving for strict optimality w.r.t. certain performance indicators is not expedient, as the latter often involves significant simplifications. Rather, an appropriate treatment of the inherent uncertainties is needed, as their influence often jeopardizes seemingly optimal solutions. In conjunction with an appropriate uncertainty consideration, the heuristic comparison of different possible trajectories is sufficient. While uncertainties can easily be reduced to zero in simulation, the performance in deterministic or quasi-deterministic scenarios does not allow for conclusions about the performance in mixed traffic. Thus, in contrast to other control or decision making approaches, a pure numerical evaluation is not possible. Instead, we focus on showing the performance of our approaches in hand-crafted exemplary scenarios. Furthermore, a common prejudice against safety approaches in the area of automated vehicles is that they inevitably lead to conservative and over-cautious behavior. Thus, a secondary objective of this scenario-based evaluation is to show that the presented safety approach allows for convenient behavior.

In the following, we first outline the implementation that was used to evaluate the approaches in Sec. 5.1. Next, the evaluation of different scenarios, separated by their dominating uncertainties, is performed in Sec. 5.2.

5.1 Implementation

In order to show the potential and to evaluate the presented approaches, they were implemented in C++. As stated in the previous chapter, all approaches rely

on high definition semantic map information to extract the route, the driveable area and the traffic rules, such as the speed limit or the right of way. Regarding this map information, the implementation employs the lanelet2 library [7]. For matching the ego vehicle to the lanelet map, we utilize the probabilistic approach presented by Petrich et al. in [PDK⁺13]. For path generation, we employ the method proposed by Thrun et al. [TMD⁺06] as implemented by Poggenhans [Pog19]: The curvature of the path is minimized, while deviations from the (non-smooth) analytically computed centerline are punished. With this approach, the path for the full route can be precomputed and sanity checks can be performed. In presence of static obstacles at the border of the ego lane, circumventing paths can be generated online.

For the lateral and longitudinal trajectory control, we employ the concept presented by Ziegler et al. [ZBS⁺14]. Instead of only providing the trajectory planned within the horizon, we additionally provide the planned future path for lateral control. This way, the steering angle can also be controlled when waiting at a red traffic light, even though the planned trajectory over the planning horizon spatially consists only of a single support point, for example.

The integration into our experimental vehicle (Sec. 5.1.1) and our simulation framework (Sec. 5.1.2) is realized using the middleware Robot Operating System (ROS) [QCG⁺09, HWT⁺16]. Implementation details of the different approaches are explained in Sec. 5.2.

5.1.1 On-Road Testing

On-road tests were performed in our experimental vehicle *BerthaOne* (cf. Fig. 5.1). Besides changing to the middleware ROS, running under Ubuntu 18.04, the vast amount of changes since the comprehensive report in [ZBS⁺14] include using the map library lanelet2 [7] as the successor of the formerly used liblanelet [BZS14], enhancing the visual localization to a multi-drive and multi-camera approach [SS18], and including multiple LIDARs into the perception stack [RWKS19].



Figure 5.1: Experimental vehicle *BerthaOne*.

5.1.2 Simulation

For the development and evaluation of motion planning approaches, simulation frameworks are an essential tool. In contrast to the research area of perception, datasets are not suitable, because the feedback loop is important: For control, but also for cooperative motion planning that focuses on interaction between traffic participants. The ROS-based simulation framework CoInCar-Sim [5] has been developed and implemented particularly for cooperative motion planning. Details about the simulation framework can be found in [5], and the framework is accessible open-source at <https://github.com/coincarsim>.

5.2 Scenario-Based Evaluation

In this section, we start by evaluating the approach to probabilistic motion planning from Sec. 4.3 in Sec. 5.2.1. This approach, that is based on an independent prediction of other traffic participants, covers mostly urban scenarios. Subsequently, the approach to lane changes from Sec. 4.4, applicable to urban multi-lane scenarios and highways, is evaluated in Sec. 5.2.2. The section con-

cludes with the evaluation of the approach to mutual cooperation from Sec. 4.5 in Sec. 5.2.3.

5.2.1 Clear Precedence

The approach presented in Sec. 4.3 is targeted towards scenarios with clear precedence. It is evaluated in single lane scenarios and scenarios with intersecting lanes, such as intersections or roundabouts. The evaluated planning approach relies on the prediction and uncertainty estimation of upstream modules being meaningful, and its performance depends on the latter. Note that, since the times of replanning are explicitly incorporated in the approach, changing those times or the replanning frequency would potentially alter the solutions, which is deliberate.

Implementation Details

Along the defined route, the conflict zones are extracted using the lanelet2 library. As driver model for the trajectory generation via forward integration, the enhanced IDM [KTH10] is used. The desired velocity v_{des} for this driver model is computed as the minimum of the speed limit and a limitation from the lateral acceleration due to the curvature $\kappa(s)$ along the path s as in [LKB⁺13], i.e.

$$v_{\text{des}}(s) = \min \left(v_{\text{speedlim}}(s), \sqrt{\frac{a_{\text{max,lat}}}{\kappa(s)}} \right). \quad (5.1)$$

Further, the course of the desired velocity is smoothed by limiting positive and negative acceleration and jerk. While testing on our experimental vehicle *BerthaOne*, we decided to further remove short phases of acceleration and subsequent deceleration for comfort reasons: If an acceleration started at $v_{\text{accel,start}}$, and the velocity dropped below $v_{\text{accel,start}}$ again in less than $\Delta t_{\text{accel-decel}}$, this acceleration phase is removed. Such behavior occurred for example in roundabouts.

For the generation of reference trajectories, the driver model is forward integrated based on the vehicles in the ego lane of trajectory ensemble set \underline{X}_k^o . In presence of a conflict zone, its occupancy in \underline{X}_k^o is analyzed. Starting at the time when the conflict zone would be reached without considering its occupancy,

possible times of arrival are sampled equidistant in time up to the planning horizon. The trajectories that reach the conflict zone at these desired times are computed by minimizing the squared acceleration while the velocity is limited by the desired velocity from eq. 5.1. The optimization is implemented using Google's Ceres Solver¹. From this arrival at the conflict zone, the trajectory is continued using forward integration of the driver model, incorporating those vehicles in \underline{X}_k^o that leave the conflict zone along the ego vehicle's route. The safety of each trajectory is checked w.r.t. the point of no return regarding the conflict zone. Based on these reference trajectories, the reactions to the belief updates \underline{X}_l are recursively computed from the earliest possible replanning step after the belief update is expected to be available. Similarly, the reactions to potentially occluded vehicles are computed. As heuristic, the trajectory that arrives earliest at the conflict zone without violating the comfort margin is pursued. The comfort margins are defined by a lookup-table, relating acceleration bounds to a maximum probability with which they would still be considered comfortable (cf. Table A.3).

In case no vehicles are present at an intersection, the reference trajectories for the occlusion consideration are computed in a simplified way: A fixed number of trajectories is interpolated between the *best case*, i.e. the approach with no occlusions, as above, and the *worst case*, i.e. stopping at the conflict zone, again computed via forward integration of the driver model. For these trajectories, starting with the best case, the reactive decelerations due to occurring objects and their probabilities are computed. Once a trajectory is found for which the comfort margins are not violated, it is pursued. In presence of vehicles with uncertain existence in the ego lane, the reference trajectories are also generated as interpolation between the best case, i.e. the vehicle does not exist, and the worst case, i.e. the vehicle exists. Based on these reference trajectories, again, the reactive decelerations and their probabilities are computed. The fastest trajectory that does not violate the comfort margins is pursued. The planning parameters can be found in Appendix A.2.



(a) The intersection scenario. Imagery © 2020 Google, map data © 2020 GeoBasis-DE/BKG (©2009). (b) The roundabout scenario. Imagery © 2020 Google, map data © 2020 GeoBasis-DE/BKG (©2009).

Figure 5.2: Aerial images of the evaluation scenarios.

Evaluation Scenarios

Starting with unlimited sensor range and perfect perception, more and more uncertainties are added throughout this section. First, the only uncertainty lies in the unknown destination, i.e. the unknown future path of other vehicles. Next, uncertainty in the existence of particular objects is added, before existence uncertainty due to occlusions and limited sensor range is considered, which is no longer based on particular object configurations. Lastly, the combination of those uncertainties is evaluated in one scenario.

In this section, we consider a four way intersection and a roundabout as basic scenarios. The intersection scenarios are based on the map of the CommonROAD [AKM17] scenario *DEU_Ffb-1*, which is depicted in Fig. 5.2a. The roundabout scenarios are based on the map of the INTERACTION Dataset [10] scenario *DR_DEU_Roundabout_OF*, which is depicted in Fig. 5.2b.

¹ <http://ceres-solver.org> by Agarwal, Mierle et al., last retrieved 2020-05-28.


Trajectory Visualization

The trajectory plots in the following sections visualize the position of the vehicle center of all relevant objects. Further, the velocity and the acceleration of the ego vehicle are visualized. An exemplary plot with a crossing vehicle is Fig. 5.4. The units are $[t] = \text{s}$, $[s(t)] = \text{m}$, $[v(t)] = \frac{\text{m}}{\text{s}}$ and $[a(t)] = \frac{\text{m}}{\text{s}^2}$. The conflict zone, if present, is depicted in gray. Different possible options for the ego trajectory are depicted in blue, solid or dashed. For particular plans or attempts, the attempted ego trajectory is visualized in solid blue. Possible reactions in future replanning steps are visualized in dashed or dotted blue. Alternative trajectories at the current planning step are visualized in dashed or dotted light blue (cf. Fig. 5.5). For other vehicles, the visualization depends on their path. For vehicles in the ego path, the position is visualized without modification. For vehicles in merging paths (cf. obj2 in Fig. 5.6), the position is visualized without modification when in the common path, i.e. beyond the end of the conflict zone. Prior to this point, their fictive position along the ego path is also visualized, but an overlap with the ego trajectory does not indicate a collision. For vehicles in crossing paths (cf. obj1 in Fig. 5.6), the position is visualized opposite to the ego vehicle's driving direction, such that the center of the conflict zone along the conflicting path and along the ego path are at the same position. Since the conflict zones are usually of similar size along both paths, overlapping trajectories in the conflict zone would denote collisions.

Perfect Perception with Unlimited Range

In this first part of the evaluation, perfect perception and perfect prediction of objects along each possible path are assumed, but there is uncertainty in the path that they actually drive. The possible paths for the intersection scenarios are adopted from [HSB⁺18, Fig. 7 and Fig. 10], they are depicted in Fig. 5.3.

The velocity profile for the objects is planned using the enhanced IDM [KTH10], extended by incorporating a maximum lateral acceleration as in [LKB⁺13]. The vehicle configurations are chosen such that the ego vehicle's intended trajectory in case of a free road is impeded by the existing objects.



(a) Crossing scenario with one vehicle: Obj1 can either go straight or turn right. (b) Merging scenario with two vehicles: Obj1 can either go straight or turn right. Obj2 can only go straight. The ego vehicle potentially crosses the path of obj1, and merges into the path of obj2.

Figure 5.3: Possible paths for the intersection scenario: The ego vehicle (blue) has to yield to vehicles from either side.

In the **crossing scenario** (Fig. 5.3a), the setup is chosen such that, if obj1 drives straight, the ego vehicle would have to decelerate to let it pass. If the vehicle turns right, the ego vehicle can unimpededly pass (cf. Fig. 5.4).

As a consequence, the behavior of the ego vehicle depends on both, the probabilities assigned to the different possible paths as well as the time when certainty about the path is expected. In case the right turn is very likely (e.g. 95%), and certainty about the path decision is expected in 1 s, for example due to the use of the indicator², the ego vehicle does not decelerate (cf. Fig. 5.5a). For the same path probability, but with certainty expected only in 3 s, the ego vehicle decides to decelerate as if obj1 would go straight (c.f. Fig. 5.5b). This behavior results from the insight that, if the ego vehicle would not decelerate, but obj1 would still go straight (5% of the cases), the ego vehicle’s reaction would have to be very harsh. Consequently, the deceleration is started straight-away, as certainty will not be reached soon enough to be able to avoid this potentially unnecessary deceleration without risking a very sharp braking. If certainty is only expected to be reached once the paths split up, the behavior is the same. To summarize, even with identical path probabilities, the motion

² As explained in Sec. 4.3, the possibility of the usage of the indicator for safety relevant decisions has to be decided upon in the safety concept.





Figure 5.4: Possible trajectories in the crossing scenario: In case obj1 turns right, the ego vehicle can unimpededly pass (solid line). In case obj1 goes straight, the ego vehicle has to yield (dashed line).


plan strongly depends on the point where certainty about the path is expected to be reached.

If the probability of a right turn is very high (e.g. 99.99%), the deceleration is avoided if certainty is expected sooner than in 4 s (c.f. Fig. 5.5c), because more deceleration is tolerated for very unlikely cases. If certainty is expected at 4 s or later, the ego vehicle will decelerate because then, regularly planning a full stop for the crossing vehicle is no longer possible and the safety planner would prompt an emergency deceleration. In case the right turn is rather unlikely (e.g. 15%), the ego vehicle plans to decelerate, regardless of when certainty will be reached, because already in the next replanning step, the additionally needed deceleration for the reaction to straight driving exceeds the limits of what is tolerated with 85% probability (c.f. Fig. 5.5d).

Compared to the work of Zhan et al. [ZLCT16], the presented approach explicitly considers the current estimate for the uncertainty progression in the future. This enables an early and comfortable reaction to persistent uncer-



(a) The right turn of obj1 is likely and certainty is expected soon, such that an unimpeded passage of the conflict zone is planned. (b) The right turn of obj1 is likely but certainty is not expected early enough, such that yielding to obj1 is planned.



(c) The right turn of obj1 is very likely and certainty is expected late, but still soon enough, such that an unimpeded passage of the conflict zone is planned. (d) The right turn of obj1 is unlikely. Even though certainty is expected soon, yielding to obj1 is planned.

Figure 5.5: Planned ego trajectories for different expected future uncertainty progression along with the possible reactions for the crossing scenario. The planned trajectory is depicted in solid blue. The conflict zone is depicted in gray. The possible reaction in case obj1 does not turn right is depicted in dashed blue, if applicable. The alternative option at the current planning step is depicted in dashed light blue.

tainties, which is not considered in [ZLCT16]. Safety was not part of the optimization problem in their approach. Only the given prediction, but no deviations from it were considered. Consequently, the point of no return is passed regardless of whether safety can be guaranteed for malicious behavior of others.

The POMDP-based approach of Hubmann et al. [HSB⁺18] is inherently considering the uncertainty progression. Due to the lack of so-called information gathering actions for this path uncertainty, however, jointly solving the uncertainty analysis and motion planning does not lead to benefits over the presented approach. Rather, the discrete action space is very limited to ensure real-time capability of the employed sampling-based solver. A trade-off for very unlikely but large ego decelerations is not possible with the action space that was used. Also, the uncertainty estimation is solely based on the vehicle's position and velocity. Information such as the usage of the indicator or traffic fleet data is not incorporated. With a decoupled prediction and uncertainty analysis, as proposed in our approach, the incorporation of such information does not affect the state space of the motion planner. Lastly, the approach of Hubmann et al. does not focus on safety w.r.t. deviations from the expected prediction of others, as opposed to the presented approach.

In the **merging scenario** (Fig. 5.3b), the setup is chosen such that the ego vehicle is able to merge in front of obj2, in case obj1 turns right. Otherwise, it has to yield to obj1 and can no longer merge in front of obj2 (cf. Fig. 5.6). In case the right turn is likely (e.g. 75%), and certainty about the path decision is expected in 1 s, the merge in front is attempted (cf. Fig. 5.7a). If certainty is only expected in 3 s, the merge in front is discarded. Instead, merging behind is planned (cf. Fig. 5.7b). For a very high probability of the right turn (e.g. 99%), the merge in front is still attempted for certainty in 3 s (cf. Fig. 5.7c). If certainty is only expected later, however, the merge is not attempted regardless of the probability (cf. Fig. 5.7d).

Merging scenarios were not considered by Zhan et al. [ZLCT16]. In the approach of Hubmann et al. [HSB⁺18], the reaction of vehicles to others merging in front is only based on the estimated arrival of both vehicles at the conflict point, assuming constant velocity and assuming that the path of the merging vehicle is known. The right of way is not incorporated there. In contrast to this, the presented approach checks for violations of the right of way by considering the imposed deceleration on prioritized vehicles, i.e. the



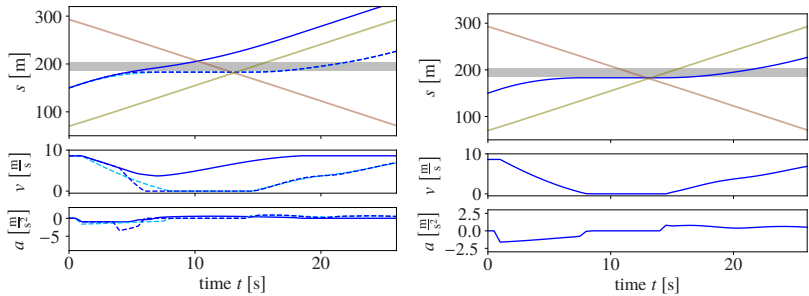


Figure 5.6: Possible trajectories in the merging scenario. In case obj1 turns right, the ego vehicle can merge in front of obj2 (solid line). In case obj1 goes straight, the ego vehicle has to yield (dashed line).

deceleration that is necessary for them to maintain a longitudinal safe distance to the ego vehicle. Also, this deceleration is only expected once the non-prioritized vehicle has merged, i.e. once it has passed the *conflict point* that is used in [HSB⁺18].



- (a) The right turn of obj1 is likely and certainty is expected soon, such that merging in front of obj2 is planned. The possible reactive yielding resembles the current alternative.
- (b) The right turn of obj1 is likely but certainty is not expected early enough, such that yielding is planned. Alternative and reaction are merging in front of obj2.



- (c) The right turn of obj1 is very likely and certainty is expected late, but still soon enough, such that merging in front of obj2 is planned. The possible reaction in case obj1 does not turn right is harsh.
- (d) The right turn of obj1 is very likely but certainty is only expected once the paths diverge, such that yielding is planned. There is no alternative to this plan.

Figure 5.7: Planned ego trajectories for different expected future uncertainty progression along with the possible reactions for the merging scenario. The planned trajectory is depicted in solid blue. The conflict zone is depicted in gray. The possible reaction in case of a belief update is depicted in dashed blue, if applicable. The alternative option at the current planning step is depicted in dashed light blue, if applicable.

The **roundabout scenario** is depicted in Fig. 5.8. In case obj1 leaves the roundabout before our entrance, we can enter unimpededly. In case it continues through the roundabout, we have to yield and merge in behind. Here again, the ego motion plan largely depends on the time and probability of certainty about whether obj1 will leave the roundabout before our entrance. Let the probability that we know within 2 s that obj1 leaves be p_{2s} , and the probability that we have certainty in 4 s in case we did not know it after 2 s be p_{4s} . The probability that we still do not know whether obj1 will continue or leave after 4 s is $1 - p_{2s} - (1 - p_{2s})p_{4s} = (1 - p_{2s})(1 - p_{4s})$, which at the same time is the maximum probability that obj1 will continue in the roundabout. In case $p_{2s} = 60\%$ and $p_{4s} = 50\%$, the ego vehicle attempts to enter, and decelerates again in case certainty is not reached after 2 s (cf. Fig. 5.9a). In case $p_{2s} = 20\%$ and $p_{4s} = 90\%$, entering before obj1 is still attempted, but a short deceleration is needed to still be able to stop in the unlikely event that obj1 does not leave (cf. Fig. 5.9b). In case $p_{2s} = 10\%$ and $p_{4s} = 10\%$, entering before obj1 is not attempted, because the probability of decelerating is too high for the necessary deceleration (cf. Fig. 5.9c).






Figure 5.8: Possible paths for the roundabout scenario: The ego vehicle has to yield in case obj1 stays in the roundabout.



(a) The ego vehicle plans in expectation of certainty after 2 s (*A*).
 (b) With the probability of certainty after 2 s (*A*) being only 20%, the ego vehicle plans for certainty after 4 s (*B*).



(c) With the probability of *A* and *B* being low, the ego vehicle plans to merge behind obj1, as certainty is not expected after 4 s (*C*).

Figure 5.9: Planned ego trajectories for different expected future uncertainty progressions along with the possible reactions for the roundabout scenario. The planned trajectory is depicted in solid blue. The conflict zone is depicted in gray. The possible reaction in case of a belief update is depicted in dashed blue, if applicable. The alternative option at the current planning step is depicted in dashed light blue. The considered events are *A*: certainty that obj1 leaves after 2 s; *B*: certainty that obj1 leaves after 4 s; *C*: no certainty that obj1 leaves after 4 s.

Uncertain Existence

Uncertainty in the existence of objects at intersections is treated similar to the uncertainty in their future path: Both certainly non-existing and certainly non-conflicting objects do not have to be considered. In case of uncertainty, they have to be considered at least in the safety approach. For convenient planning, the probability that they impede the ego motion plan along with the severeness of the reaction is considered. More specifically, in the scenarios depicted in Fig. 5.3 and 5.8, the probability that the non-existence of an object is certain at time t leads to the same motion plan as the probability that we have certainty at time t that the vehicle turns and does not conflict with our path.

The phenomenon of *ghost objects*, however, also occurs in regular single-lane driving, as described in Sec. 4.3.2. In order to show the potential of our approach for such situations, we consider the following scenario: Driving at 70 km/h \approx 20 m/s on a road with speed limit 100 km/h \approx 28 m/s, there is an uncertain detection of a fully stopped vehicle 100 m ahead (cf. Fig. 5.10). The considered options in this case are to continue as if there was no object, to immediately react to the object, or to perform some trade-off in-between, as shown in Fig. 5.11. The chosen trajectory now depends on the probability that we have to react to this object, either because the object actually exists, or because we are unable to reach certainty regarding its non-existence. This information is provided by the perception or the scene understanding module. The reactions to different possible belief updates are provided in Fig. 5.12. To summarize, as for the previous intersection consideration, more deceleration is tolerated as reactive behavior the less likely this reactive behavior is.



Figure 5.10: Ghost object scenario: There is an uncertain vehicle detection ahead in the ego lane.

In [TS20], an approach to dealing with uncertain object existence is presented. However, by allowing for different reactions for either case (object exists or not) in the next replanning step, the authors implicitly assume that certainty


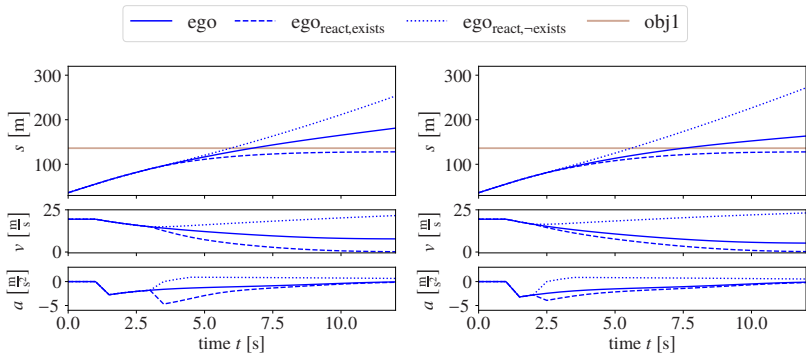



Figure 5.11: Possible trajectories as reaction to obj1 with uncertain existence: Linearly interpolated in-between continuing as if there was no object (uppermost blue trajectory) and reacting to this object according to the driver model (lowermost blue trajectory).

about the object existence is reached in the next replanning step.³ Contrasting to that, the approach proposed in this work accounts for the fact that reaching certainty that an object detection was a false-positive might take several seconds, depending on the sensor configuration and the measurements that lead to the object assumption. Besides [TS20], to the best of the author’s knowledge and also stated in [TS20], there is no comparable approach that incorporates potential future updates of the existence uncertainty of objects, which could be used as a baseline for the evaluation.

³ This simplification corresponds to the QMDP assumption (cf. Sec. 2.1.5), leading to suboptimal results (cf. ex. 4.3.2).



(a) With a reaction probability of 1% in 2 s, (b) With a reaction probability of 20% in 1 s, the ego vehicle decelerates with more than -2 m/s^2 , limiting the reactive deceleration in case of the reaction to -5 m/s^2 . the ego vehicle decelerates with more than -3 m/s^2 , limiting the reactive deceleration in case of the reaction to -4 m/s^2 .



(c) With a reaction probability of 20% in 1 s, a further reaction probability of 1% in 2 s does not alter the velocity profile. The reaction to the latter from the attempted velocity profile is less severe than the reaction when pursuing 5.12a.


Figure 5.12: Attempted velocity profiles in case of an uncertain detection of a fully stopped vehicle 100 m ahead, along with the possible reactions. The planned trajectory is depicted in solid blue. The reaction in case of a certainty that the object was a false detection is depicted in dotted blue. The reaction in case the object exists or certainty regarding its non-existence could not be reached is depicted in dashed blue.

Occlusions and Limited Sensor Range in Single-Lane Scenarios


With the presented approach, general information about potential objects even beyond the sensing horizon can be included: We again consider driving at $70 \text{ km/h} \approx 20 \text{ m/s}$, with limited visibility due to the sensor range. With the visible range being 100 m , we consider a belief update including a stopped vehicle at 100 m from the ego vehicle, with reaction probability 0.01% in 1 s , i.e. in the next replanning step. This corresponds to the notion that a stopped object on a road with speedlimit 70 km/h is unlikely (0.01%), and if there was one, it would be detected until the next replanning step (1 s). This scenario is



- (a) A hypothetical stopped object just behind the edge of the sensing horizon. Certainty about whether this object exists is expected in the next replanning phase.




- (b) A hypothetical stopped object, waiting to turn left at an upcoming intersection. Certainty about whether this object exists is only expected once the ego vehicle is close enough.



- (c) An upcoming traffic light with an unknown state. Certainty about the traffic light state is only expected once the ego vehicle is close enough.

Figure 5.13: Reactions to hypothetical upcoming objects or other events that require a reaction of the ego vehicle. The sensing horizon is depicted in blue.



(a) With a reaction probability of only 0.01% in 1 s, the ego vehicle continues at its desired velocity. The reaction in case obj1 actually exists would be very harsh.

(b) With a reaction probability of 5% in 100 m, the ego vehicle decides to slightly decelerate already. The reaction in case obj1 actually exists would be less strong. In case obj1 does not exist, the ego vehicle can accelerate again.

Figure 5.14: Attempted velocity profiles in case of hypothetical upcoming events. The planned trajectory is depicted in solid blue. The reaction in case there is a stopped vehicle is depicted in dashed blue. The reaction in case there is no stopped vehicle is depicted in dotted blue, if applicable.

depicted in Fig. 5.13a. As visible in Fig. 5.14a, the velocity is not reduced in this case. This mimics human driving behavior in the sense that we do not expect the worst case, i.e. a stopped object, behind every bend.⁴

Similarly, future information that is only available after a certain distance is traveled can be incorporated. Assume we know that there is an unprotected left turn outside the visible range, where vehicles are frequently waiting to turn, due to much oncoming traffic for example (cf. Fig. 5.13b). The knowledge about the future unprotected left turn is incorporated by considering a hypothetical stopped vehicle in 200 m with 5% reaction probability. Further, we know that we can only react in 100 m, because only then it would become visible to us. With this information, the vehicle decides to slightly decelerate already, to minimize the impact of an actual reaction in 100 m, as visible in Fig. 5.14b.

⁴ The visible range for humans is generally larger than 100 m, but sometimes also limited due to occlusions, such as on curvy roads.

Similarly, uncertain information about a traffic light state can be incorporated, where certainty is only available at a closer distance to the traffic light (cf. Fig. 5.13c). Also, with people walking towards a pedestrian crossing, the probability of whether they cross or walk past the crossing can be incorporated. Furthermore, live traffic data can be used for an early and more comfortable reaction to reported traffic jams, where the information about the position is uncertain and can be subdivided into different positions along with a probability.

To the best of the author's knowledge, there is no comparable approach that incorporates information about traffic light states, frequent traffic jams or similar, which could be used as baseline for the evaluation.

Occlusions at Intersections

The approach to occlusions at intersections, explained in Sec. 4.3.3, is evaluated in the scenario proposed by Orzechowski et al. [OML18]: The scenario *DEU_Ffb-1* is enhanced by a building in the south east, which occludes the view on the east branch of the junction (cf. Fig. 5.15). In the work of Orzechowski et al. [OML18], the reaction to occlusions is a deceleration with a priorly defined maximum deceleration. Comfort is only considered by choosing this deceleration as a trade-off between a comfortable and the maximum possible deceleration. Yet, their approach is, to the best of the author's knowledge, the only approach that guarantees safety w.r.t. meaningful assumptions, i.e. accounting for the traffic rules. Thus, it is considered as baseline for the presented approach. The trajectory for the baseline approach is visualized in Fig. 5.16: Independent of whether we approach a priority road with much traffic, or a rural road where only one vehicle exits per day, the vehicle continues at constant velocity and only starts to decelerate with -4 m/s^2 once necessary to maintain safety.

With the proposed approach, the attempted trajectory depends on the probability that the attempt will be impeded by vehicles that are currently occluded. The considered velocity profiles, visible in Fig. 5.16, vary from ignoring the occlusion to expecting that a full stop is necessary. For all these options, starting with the fastest, the probability and severity of the reaction to potentially occluded objects are analyzed. In case the attempted velocity profile violates




Figure 5.15: Scenario *DEU_F/b-1*, enhanced by occlusions, as proposed by Orzechowski et al. [OML18]. Note that the actual occlusions at this intersection are even more severe. Potentially as a result, the real intersection is controlled by traffic lights.

the safety constraints, the emergency planner will take over with certainty, which corresponds to an emergency deceleration with a probability of 100%.

With a reaction probability of 1%, the vehicle already slows down to approximately half the initial velocity, yet risking decelerations of approx. -4 m/s^2 (cf. Fig. 5.17a). With a reaction probability of 10%, the vehicle slows down even further, reducing the risked decelerations to approx. -2 m/s^2 (cf. Fig. 5.17b). With a reaction probability of 80%, the vehicle actually attempts a full stop. As, however, passing is safe from a certain point on, the vehicle attempts to accelerate again from this point, in case no vehicle appears behind the occlusion (cf. Fig. 5.17c). If a vehicle appears, the deceleration is continued, and the driven velocity profile is equal to the one of approaching a stop sign, for example. Contrasting to that, if the probability is very low, such as $\frac{10\text{s}}{\text{day}} \approx 0.01\%$, i.e. only one or two vehicles drive this road per day, the velocity is not slowed down at all, risking a potential emergency brake (cf. Fig. 5.17d). Note that, in contrast to a violation of the safety constraints, which leads to an emergency brake with a probability of 100%, in the presented approach, the probability of an emergency brake is $< 0.01\%$, i.e. only if a vehicle appears unfavorably. Again, the presented behavior shows the key idea of the probability treatment within


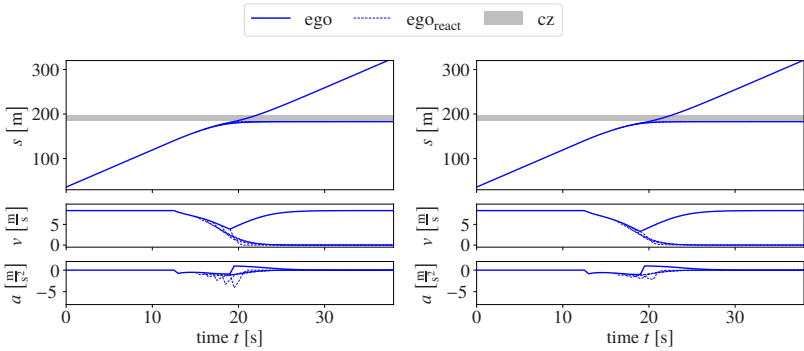



Figure 5.16: Possible trajectories for the occlusion scenario: Linearly interpolated in-between continuing as if there was no occlusion (uppermost blue trajectory) and reacting to this occlusion by attempting to stop in front of the intersection according to the driver model (lowermost blue trajectory). Once the occlusion does no longer compromise the safe passage of the conflict zone, the ego vehicle accelerates again. As visible in the acceleration diagram, this decision is only possible at the predefined replanning times in 1 s distances. In the baseline approach by Orzechowski et al. [OML18], a predefined deceleration of -4 m/s^2 is used to maintain a safe velocity, independent of the probability that the ego vehicle actually has to react to currently occluded vehicles.

this work: While safety is never put at risk, occasional high decelerations are tolerated for the benefit of mostly convenient and not overly conservative motion plans.



(a) With the reaction probability being only 1%, the ego vehicle only decelerates slightly, keeping the reactive deceleration at all replanning times when a reaction might still be necessary below -5 m/s^2 . (b) With the reaction probability being 10%, the ego vehicle decelerates, keeping the reactive deceleration at all replanning times when a reaction might still be necessary below -2.5 m/s^2 .



(c) With the reaction probability being 80%, the ego vehicle plans to stop in front of the conflict zone and only accelerates again once a safe passage is possible. In the expected case that a currently occluded vehicle hinders its passage, it stops smoothly. (d) With the reaction probability being $< 0.01\%$, the ego vehicle continues at its desired velocity. In the very unlikely case that the ego vehicle has to stop, the necessary deceleration might exceed -5 m/s^2 .

Figure 5.17: Attempted velocity profiles in presence of occlusions. The attempted velocity profile is visualized in solid blue, the reactions at the different replanning steps in dashed blue. The conflict zone is depicted in gray.

All Uncertainties Combined

Finally, we present a scenario with all the different types of uncertainties. In the previously evaluated occluded intersection scenario, we now consider five certainly existing vehicles and one potential ghost object (cf. Fig. 5.18). Obj1 and obj6 use the indicator, such that their paths are assumed to be certain. Obj2 is a potential ghost object. If it exists, it is assumed to be at full stop. Obj4 is already in the intersection, at a velocity such that only going straight is feasible. For obj3 and obj5, multiple paths are considered. Beyond that, there might be vehicles occluded behind the building, with the same route choices as obj5. As explained in Sec. 4.3.3, those potentially occluded objects are not considered on an object basis but with a simplified reaction model. The probability of being impeded by occluded vehicles is set to 1%. Since obj6 has to yield to the ego vehicle, obj6 is not considered in any trajectory ensemble set \underline{X}^o . Unfavorable trajectory ensemble sets, that lead to the intersection being blocked, are subsumed by introducing a virtual object objV along the ego path just before the intersection.




Figure 5.18: Possible paths of objects in a crowded intersection scenario. The existence of obj2 is uncertain. Objects 3-5 are on the prioritized road. Objects 1 and 6 are using the indicator. There might be further vehicles, occluded by the building (dark gray).


The considered possible belief progressions $\{\underline{X}_k^o, (t_l, p_l, \underline{X}_l^o), l \in \mathcal{L}_k\}$, that are to be provided by an upstream prediction module, are hand-crafted for this scenario, based on the following assumptions: Obj3 goes straight with 80%, left and right with 10% each. Obj2 exists with 10%. Obj5 turns left with 30%.

With these assumptions, the three considered possible belief progressions are as follows:

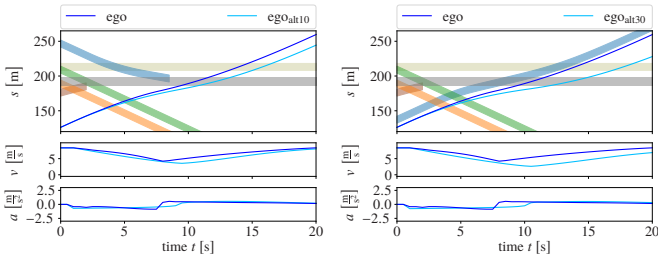
1. \underline{X}_{10}^o : Obj3 turns right, obj4 and obj5 go straight, obj1 turns unimpededly, obj2 does not exist. $p_{10} = 6\%$. Considered possible updates:
 - \underline{X}_{11}^o : Obj3 does not turn right: unfavorable, approximated with objV, leading to a full stop. $p_{11} = 90\%$, $t_{11} = 2$ s.
 - \underline{X}_{12}^o : Obj5 wants to turn left and has to wait: unfavorable, approximated with objV, leading to a full stop. $p_{12} = 3\%$, $t_{12} = 1$ s.
 - \underline{X}_{13}^o : Obj2 exists: obj2 is added to the objects in \underline{X}_{10}^o . $p_{13} = 1\%$, $t_{13} = 4$ s.
2. \underline{X}_{20}^o : Obj3 goes straight, obj4 and obj5 go straight, obj1 turns unimpededly, obj2 does not exist. $p_{20} = 36\%$. Considered possible updates:
 - \underline{X}_{21}^o : Obj3 does not go straight: unfavorable, approximated with objV, leading to a full stop. $p_{21} = 20\%$, $t_{21} = 2$ s.
 - \underline{X}_{22}^o : Obj5 wants to turn left and has to wait: unfavorable, approximated with objV, leading to a full stop. $p_{22} = 24\%$, $t_{22} = 1$ s.
 - \underline{X}_{23}^o : Obj1 yields to obj3: unfavorable, approximated with a full stop of obj1. $p_{23} = 16\%$, $t_{23} = 1$ s.
 - \underline{X}_{24}^o : Obj2 exists: obj2 is added to the objects in \underline{X}_{20}^o . $p_{24} = 4\%$, $t_{24} = 4$ s.
3. \underline{X}_{30}^o : Obj3 turns left, obj4 and obj5 go straight, obj1 turns unimpededly, obj2 does not exist. $p_{30} = 19\%$. (Obj3 turning right is subsumed here.) Considered possible updates:
 - \underline{X}_{31}^o : Obj3 does not turn left: unfavorable, approximated with objV, leading to a full stop. $p_{31} = 80\%$, $t_{31} = 2$ s.

- \underline{X}_{32}^o : Obj3 has to wait: unfavorable, approximated with objV, leading to a full stop. $p_{32} = 1\%$, $t_{32} = 4$ s.

As visible in Fig. 5.19, the set \underline{X}_{20}^o leads to the fastest approach. Since the reactive decelerations along with their probabilities are considered comfortable, this approach is pursued. The considered decelerations include those to belief updates w.r.t. certainly existing objects and also those to potentially occurring objects behind the occlusion. The runtime of the planner for this scenario with the presented belief progressions on a single core of an Intel® Core™ i7-8565U was below 100 ms. Parallelized computation for the different belief progressions could further lower the computation time. Yet, runtime analysis or optimization is not the focus of this work. To summarize, the approach works well in presence of many objects, given that the prediction module provides suitable belief progressions.



(a) Attempted velocity profile, assuming obj3 goes straight and obj2 does not exist (\mathcal{X}_{20}^o). The reactions to the possible belief updates are depicted dashed. The two reactions that lead to a full stop in front of the CZ ($\mathcal{X}_{21}^o, \mathcal{X}_{22}^o$) are almost alike, as they start only 1 s apart from each other (slight deviations in a from 2 s to 6 s).



(b) Alternative velocity profile, assuming obj3 turns right ($\mathcal{X}_{10}^o, \text{ego}_{\text{alt}10}$). Since certainty about this right turn is only reached once it turns, a change to going straight is not neglected in the safety checks, leading to a slow down compared to the attempted profile.

(c) Alternative velocity profile, assuming obj3 turns left ($\mathcal{X}_{30}^o, \text{ego}_{\text{alt}30}$). The ego vehicle has to slow down further to merge in behind obj3.

Figure 5.19: Attempted velocity profile and alternatives for the scenario from Fig. 5.18. Since sets are considered, the objects' positions are depicted with an area instead of a line.

5.2.2 Lane Change

In this section, the approach to courtesy-aware lane changes from Sec. 4.4 is evaluated in an urban lane change scenario.

Implementation Details

The implementation presented in [9] was integrated into a lanelet2-based routing state machine and the previously explained modified control (cf. Sec. 5.1) has been employed. Further, cases with only one or no vehicle in the target lane have been added to the lane change state machine. As heuristic for the gap selection, the gap that can be reached earliest in space along the target lane with a trajectory within the comfortable acceleration limits is selected. The approach to the gap is computed analytically as a sequence of constant acceleration, constant velocity and constant acceleration. If the gap is currently ahead, regarding position and velocity difference, the vehicle accelerates up to a certain velocity v_{peak}^+ to catch up, before decelerating to the gap velocity. If the gap is behind, the vehicle decelerates to a velocity v_{peak}^- before accelerating to the gap velocity. The path for the lane change is computed as transition from the path in the source lane to the path in the target lane using the sigmoid function \tanh . The planning parameters can be found in Appendix A.2.

Evaluation Scenario

The general performance of the approach to facilitate lane changes is shown in a mandatory lane change at Haid-und-Neu-Straße in Karlsruhe, Germany (cf. Fig. 5.20), as presented in [9]. The success of approaches relying on cooperation or courtesy of other traffic participants inherently depends on whether the respective agents are willing to cooperate or not. Instead of providing velocity or acceleration profiles for this multi-lane scenario, we perform a qualitative evaluation of the lane change behavior: If the vehicle behind us in the target lane is courteous, the lane change is assumed to be safe after the additional lane change response time has been granted and thus can be performed. Courtesy of the vehicle behind was modeled using the enhanced IDM [KTH10] for the motion planning of this vehicle and considering the ego vehicle as front vehicle as soon as it indicated its lane change desire. In case

the respective vehicle does not open a gap of sufficient size for RSS safety, even though the desire to change the lane has been communicated using the indicator and moving left within the current lane, the lane change is aborted after waiting for Δt_{wait} . The non-courteous behavior was modeled by refraining from considering vehicles in other lanes as front vehicle.

This experiment shows the general applicability to scenarios where approaches with an independent prediction would lead to overly conservative behavior, i.e. waiting for large gaps without “asking for them” via clear communication of the lane change desire. Besides this advantage in comfort and convenience, safety is ensured w.r.t. the defined assumptions of the presented RSS concept for *non-reckless* lane changes. Since the reaction time for the ego vehicle can be significantly lower than the commonly granted one second for human drivers, the safety distance to the front vehicle can be significantly smaller than the one that has to be granted to the vehicle behind (cf. Fig. 5.21). This leads to less additional space needed for lane changes. Both lane change attempts are also shown in the video⁵ provided along with [9].

The runtime of the planner for this scenario on a single core of an Intel® Core™ i7-8565U was below 2 ms for all plans throughout the simulation. While runtime analysis or optimization is not the focus of this work, the computation time underlines the general suitability of the approach for real-time planning in automated vehicles. In order to further show the real-time capability and applicability of the approach, it has also been implemented on our experimental vehicle *BerthaOne*.

⁵ Also available via http://www.mrt.kit.edu/z/publ/download/2019_LaneChange_Naumann.mp4, last retrieved 2020-05-28.



Figure 5.20: Aerial image of the lane change scenario. Imagery © 2020 Google, map data © 2020 Stadt Karlsruhe VLW, GeoBasis-DE/BKG (©2009).




Figure 5.21: Screenshot of the lane change in the simulation framework CoInCar-Sim [5] (from [9], © 2019 IEEE). The other vehicles are depicted in blue. The ego vehicle is depicted in black. Its trajectory is visualized with the colored circles at equidistant times.

5.2.3 Mutual Cooperation

In addition to the sampling-based implementation presented in [3, 6], we show the performance of the approach to mutual cooperation from Sec. 4.5 with a C++ implementation in a narrowing scenario.

Implementation Details

In contrast to the implementation presented in [3, 6], the global optimum is no longer found via sampling, but via employing the enhanced IDM [KTH10]. For the globally optimal trajectory ensemble, vehicles are no longer expected to accelerate beyond the velocity that they would have chosen if no other vehicle was present. The actual human behavior in cooperative situations is to be investigated further by experts in this field and is not part of this work. Yet, this second implementation shows the flexibility of the approach regarding different models of cooperation.

In order to investigate whether cooperation is possible, i.e. whether the globally optimal solution is not ambiguous, the homotopy classes of the solutions for different driver models are analyzed. In case they unambiguously yield that the ego vehicle should go first, a predefined model compliance probability p_{mc} is assumed. As model compliance check for other drivers, the deviation from the expected velocity Δv_{mc} is considered. A model compliance violation is expected to be measurable by the time when the predictions for the different homotopy classes of the cooperative vehicle differ by Δv_{mc} . For this point in time, we compute the necessary deceleration to switch to going second. With this information, the heuristic presented in Sec. 4.3 is employed to check whether attempting the globally optimal solution violates the comfort margin, as the expected deceleration for the fallback plan is too high. In case the comfort margin is not violated, the globally optimal solution is attempted. The expected velocity profile for the other vehicle is stored.

If the ego vehicle is still far away from the conflict zone, however, the situation might be currently ambiguous, or the velocity difference in-between the homotopy classes might be not detectable, but the situation might become clear later. Thus, it is checked whether continuing in the hope of the cooperative solution is convenient. Since the time and the probability of a positive decision

for the cooperative solution are hard to estimate, the decision is only postponed if (a) the necessary deceleration for going second in the next step is below a certain value $a_{\text{decel,low}}$ or (b) the maximum deceleration remains unchanged, for example as a deceleration due to significant curvature is upcoming anyway.

If neither attempting to go first, nor decision postponing is considered comfortable, the ego vehicle attempts to go second. In case the model compliance is violated, i.e. the velocity of the cooperative vehicle is more than Δv_{mc} higher than expected, the ego vehicle switches to go second as the globally optimal solution is no longer believed to be realized.

The planning parameters can be found in Appendix A.2.

Evaluation Scenario

The approach is evaluated at an artificial narrowing in Karlsruhe (cf. Fig. 5.22, Fig. 5.23). The vehicle configuration is chosen such that the desired trajectories for both vehicles, each assuming the other vehicle was not there, would result in a collision inside the conflict zone (cf. Fig. 5.24). For such scenarios, mutually predicting the other vehicle independently, for example with constant velocity or with the IDM, commonly results in a deadlock, at least until one vehicle comes to a full stop. Fig. 5.25 shows the planning result of both vehicles with independent prediction after 8 s. A similar approach was for example used by Ziegler et al. [ZBS⁺14]. With the presented approach, cooperation is possible in this scenario, in case the other driver behaves cooperatively. Meaningful malicious behavior, i.e. egoistic behavior, can also be detected and dealt with comfortably (cf. Fig. 5.26). Adversarial behavior of others, such as provoking deadlocks, is still possible, but not expected and not the focus of this work. The runtime of the planner for this scenario on a single core of an Intel® Core™ i7-8565U was below 20 ms for all plans throughout the simulation. While runtime analysis or optimization is not the focus of this work, the computation time underlines the general suitability of the approach for real-time planning in automated vehicles.



Figure 5.22: Aerial image of the narrowing scenario. Imagery © 2020 Google, map data © 2020 Stadt Karlsruhe VLW,GeoBasis-DE/BKG (©2009).




Figure 5.23: Paths in the narrowing scenario.




Figure 5.24: Possible solutions from the ego perspective: Expected is the solution depicted in solid, with the ego vehicle going first. The trajectory for the other vehicle acting egoistic is depicted in brown dashed. The reaction to this behavior, by the time a model violation is expected to be detectable, is depicted in blue dashed. The alternative of assuming straightaway that the other will act egoistic, and thus attempting to go second, is depicted in light blue dashed.




Figure 5.25: Trajectories for mutually predicting each other with constant velocity: The already driven part is visualized with a solid line, the prediction for the other vehicle is dotted, the own plan is dashed.




Figure 5.26: Driven trajectories with the proposed approach: In case obj1 behaves cooperatively as expected, the ego vehicle drives first (solid). In case obj1 acts egoistically, the ego vehicle detects this and reacts early and comfortably (dashed).

6 Conclusions and Future Directions

6.1 Conclusions

In this work, several contributions to motion planning for automated vehicles in mixed traffic have been presented. At first, the different sources of uncertainties and common problem formulations in the context of motion planning for automated vehicles are outlined. This consideration exposed that, while optimal control approaches with one deterministic trajectory as result are applicable, they can be outperformed by approaches that account for the repeated replanning. Particularly in cases where the possible future evolutions of a scene largely differ, for example due to multiple route options for other traffic participants, considering multiple future trajectories is beneficial. Yet, probabilistic approaches that attempt a comprehensive uncertainty treatment with limited runtime are often operating on a largely restricted action space. Also, many approaches do not consider safety and compliance with the traffic rules.

Regarding the safety verification, recent frameworks focus on the notion of responsibility. That is, instead of guaranteeing to never be involved in an accident, automated vehicles should guarantee that they will never cause an accident. The traffic rules and the right of way defined therein have to be accounted for when defining responsibility. Thus, the safety of a motion plan does not only depend on the road layout and the current vehicle configurations, but also on the priorities of lanes and roads over each other. This motivates the decomposition of the safety concept into scenarios: As a contribution to safety in motion planning, in Sec. 4.2, the existing responsibility sensitive safety (RSS) framework [SSS18] is enhanced by the consideration of traffic rules at intersections with crossing and merging traffic, including occlusions and limited sensor range, as well as lane changes. In particular, the consideration

of right-of-way rules at intersections leads to an upper bound for the velocity to ensure safety in presence of occlusions and limited sensor range.

Given these hard constraints from safety and the traffic rules, the goal is to generate convenient motion plans, i.e. plans that are comfortable but not over-cautious. Towards this goal, this work relies on the assumption that, due to the complexity of road traffic, different scenarios can comprise inherently different types of uncertainties that have to be taken into account. These differences largely affect the suitability of different problem formulations. While some uncertainties have to be actively reduced, others can hardly be influenced, for example. Thus, the main contribution of this work is the introduction of three motion planning approaches, targeted towards the different predominant uncertainties in different scenarios.

For scenarios with unimpeded sight and unambiguous routes for all traffic participants, where mutual cooperation is required, a global optimum approach with model compliance consideration is presented in Sec. 4.4. The approach is based on a multi-agent formulation, where other vehicles are expected to cooperate and to have a similar understanding of a globally optimal solution. At the same time, it is checked whether the global optimum is unique, also for different measures of optimality, being aware that the cooperation model might be violated. In case the global optimum is unambiguous and the other traffic participant behaves within our model bounds, cooperation is facilitated, resulting in globally optimal trajectories. Otherwise, the ego vehicle reacts early and pursues a comfortable fallback solution.

For lane change scenarios, in which courtesy of others is required in case of dense traffic, a second approach is presented in Sec. 4.5. The approach is rule-based, and also explicitly modeling other drivers' behavior as part of the decision problem. Even for gaps that are currently too narrow for a safe lane change, the lane change desire is indicated and the response of the potentially courteous vehicle is awaited. In case this vehicle opens a gap for us that is large enough for a safe lane change according to the enhanced RSS framework, the lane change is pursued. Otherwise, the next available gap is to be approached. The approach facilitates the design of gap approach and desire indication according to results from human behavior studies. Furthermore, to the best of the author's knowledge, the presented approach is the first lane change approach that examines safety also w.r.t. unexpected but lawful human behavior, such as emergency decelerations or closing a gap that would have

allowed a lane change. It facilitates lane changes on urban multi-lane roads as well as on highways.

Large parts of the urban and rural road network consist of a single lane per driving direction and intersections with predefined right of way. The most significant contributions of this thesis are the formulation of this problem as a partially observable Markov decision process (POMDP) with an ego-action-independent belief over the trajectories of other traffic participants in Chapter 3 and the presentation of an approach to simplify and solve this problem formulation in Sec. 4.3. The approach generates trajectories based on uncertainty estimates for the trajectories of other traffic participants. The uncertainties include their current state and their desired route, for example. Not only the current uncertainty estimate, but also possible future progressions are incorporated. Furthermore, occlusions and limited sensor range, along with an estimate of their future evolution, and even potential ghost objects or estimates of objects behind the sensing horizon are considered. Also, appropriate reactions to uncertain traffic light state classifications or pedestrians that walk towards a crosswalk can be determined by the approach. Safety is ensured based on the enhanced RSS framework. Related approaches that solve this POMDP in a belief space with algorithms based on Monte Carlo tree search (MCTS) build an action-dependent belief tree during planning. In this process, knowledge about the behavior of others is commonly modeled via a probabilistic transition model, hidden states, and a probabilistic observation model. The assumption behind the presented approach is that the behavior of prioritized vehicles is independent of the ego behavior. The observations that facilitate a behavior prediction are only influenced by occlusions and the sensor range, which again depend on the position of the ego vehicle along its route. Yet, the belief tree can be built independently of the ego motion plan in an upstream prediction module. Here, more sophisticated behavior and observation models can be employed, since they are not repeatedly queried during planning. Instead, only the result of those sophisticated models, provided through selected belief progressions, is used during planning. To the best of the author's knowledge, the presented approach is the first to comprehensively consider uncertainties regarding the current state and future trajectory of other traffic participants, including multiple routes, and regarding their existence, based on object detections within the field of view, but also beyond the sensor range and in occluded areas, while operating in a continuous action space and while guaranteeing safety w.r.t. to the enhanced RSS framework. Lastly, the

approach is suitable for scenarios with several traffic participants, since the combinatorial burden can be largely reduced through the limitation to selected belief progressions.

6.2 Future Directions

On the way towards safe and convenient motion planning for automated vehicles in mixed traffic, several research directions are promising. The presented approaches can be extended to scenarios that are subject to special traffic laws, such as the zipper merge, or overtaking with oncoming traffic.

Further, to allow for more efficient cooperation, the models of human behavior can be improved using naturalistic driving datasets, such as the recently presented INTERACTION dataset [10]. Fundamental work on how to reveal cost functions from parts of human-driven trajectories has been presented in [11].

Another research direction is a planning-focused scene understanding and prediction. Computing a full trajectory for all potential trajectories of others, along with the fallback trajectories in case of deviations, is infeasible. Thus, the behavior of others was grouped into belief progressions with a set-based motion prediction. Here, the right granularity has to be found as a trade-off between keeping the problem formulation feasible and solvable in real-time, while not discarding promising maneuver ensembles. With this input, the presented belief-dependent planning approach can be evaluated in experimental vehicles.

Further work in the area of scene understanding and prediction includes heuristics for cooperation probabilities with other traffic participants, in order to improve the gap choice for desired lane changes, but also to get a better confidence estimate for mutually cooperative scenarios.

A Appendix

A.1 Safe Distance for Merging

As stated in eq. 4.5, for the time $t > 0$ after the front vehicle has merged, the safety reserve for the prioritized rear vehicle to the front vehicle is

$$d_{\text{reserve}}(t) = s_{\text{merge}}^-(t) - s_{\text{prio}}^+(t) - d_{\text{safe,prio}}(t). \quad (\text{A.1})$$

It is assumed that the rear vehicle no longer accelerates from the time the front vehicle merges, visualized in Fig. 4.2, as this maneuver is foreseeable. Still, the usual reaction time ρ_{prio} is granted before the vehicle is expected to start decelerating.

For $t > \rho_{\text{prio}}$, the position of the front bumper of the prioritized vehicle is

$$s_{\text{prio}}^+(t) = -d_{\text{init,prio}}^+ + v_{0,\text{prio}}t + \frac{1}{2}a_{0,\text{prio}}(t - \rho_{\text{prio}})^2. \quad (\text{A.2})$$

The position of the rear bumper of the merging vehicle is

$$s_{\text{merge}}^-(t) = -d_{\text{init,merge}}^- + v_{0,\text{merge}}t + \frac{1}{2}a_{0,\text{merge}}t^2. \quad (\text{A.3})$$

The safe distance for the prioritized vehicle is

$$\begin{aligned} d_{\text{safe,prio}}(t) &= \rho(v_{0,\text{prio}} + a_{0,\text{prio}}(t - \rho_{\text{prio}})) & (\text{A.4}) \\ &+ \frac{(v_{0,\text{prio}} + a_{0,\text{prio}}(t - \rho_{\text{prio}}))^2}{-2a_{\text{min,brake}}} - \frac{(v_{0,\text{merge}} + a_{0,\text{merge}}t)^2}{-2a_{\text{max,brake}}}. \end{aligned}$$

The time derivatives of those terms are

$$\dot{s}_{\text{prio}}^+(t) = v_{0,\text{prio}} + a_{0,\text{prio}}(t - \rho_{\text{prio}}), \quad (\text{A.5})$$

$$\dot{s}_{\text{merge}}^-(t) = v_{0,\text{merge}} + a_{0,\text{merge}}t \quad (\text{A.6})$$

and

$$\begin{aligned} d_{\text{safe,prio}}(t) = & a_{0,\text{prio}}\rho_{\text{prio}} + \frac{2(v_{0,\text{prio}} + a_{0,\text{prio}}(t - \rho_{\text{prio}}))a_{0,\text{prio}}}{-2a_{\text{min,brake}}} \quad (\text{A.7}) \\ & - \frac{2(v_{0,\text{merge}} + a_{0,\text{merge}}t)a_{0,\text{merge}}}{-2a_{\text{max,brake}}}. \end{aligned}$$

Setting the time derivative of the safety reserve to zero yields

$$t_{\text{crit}} = \frac{-v_{0,\text{merge}} + v_{0,\text{prio}} - \frac{v_{0,\text{prio}}a_{0,\text{prio}}}{a_{\text{min,brake}}} + \frac{a_{0,\text{prio}}^2\rho_{\text{prio}}}{a_{\text{min,brake}}} + \frac{v_{0,\text{merge}}a_{0,\text{merge}}}{a_{\text{max,brake}}}}{a_{0,\text{merge}} - a_{0,\text{prio}} + \frac{a_{0,\text{prio}}^2}{a_{\text{min,brake}}} - \frac{a_{0,\text{merge}}^2}{a_{\text{max,brake}}}}. \quad (\text{A.8})$$

From demanding the critical, i.e. smallest, safety reserve to be at least zero

$$d_{\text{reserve}}(t_{\text{crit}}) \stackrel{!}{>} 0, \quad (\text{A.9})$$

such that no emergency maneuver is to be performed, we can compute the required initial distance of the rear vehicle $d_{\text{init,prio}}^+$, such that the merging can be performed safely when accelerating with at least $a_{0,\text{merge}}$ from the time of merging $t = 0$.

A.2 Evaluation Parameters

A.2.1 Safety Parameters

Table A.1: RSS Parameters.

Ego response time ρ_{ego}	0.3 s
Others' response time ρ_{other}	1 s
Additional lane change response time ρ_{LC}	2 s
Max. accel. during response time $a_{\text{max,accel,ego}}$	2 m/s^2
Max. accel. during response time $a_{\text{max,accel,other}}$	3 m/s^2
Max. emergency decel. $a_{\text{max,brake}}$	-8 m/s^2
Min. emergency decel. $a_{\text{min,brake}}$	-7 m/s^2
Min. req. decel. for prior. vehicles $a_{\text{min,brake,row}}$	-1.5 m/s^2
Exp. decel. for prior. vehicles $a_{\text{brake,row,exp}}$	-1 m/s^2
Time of zone clearance granted to prior. vehicles TZC_{row}	3 s
Time of zone clearance to drive behind others TZC_{ego}	2 s

A.2.2 Comfort and Convenience Parameters

Table A.2: General Planning Parameters.

Time between trajectory support points Δt	0.5 s
Replanning interval Δt_{plan}	1 s
Planning horizon Δt_{H}	20 s-40 s

Table A.3: Accepted deceleration probabilities.

More than a_{react}	with max. p_{react}	More than $a_{\text{react,add}}$	with max. p_{react}
-2 m/s^2	10%	-0.5 m/s^2	50%
-3 m/s^2	2%	-1 m/s^2	10%
-4 m/s^2	1%	-1.5 m/s^2	5%
-5 m/s^2	0.2%	-2 m/s^2	1%
-6 m/s^2	0.1%		
-7 m/s^2	0.01%		
-8 m/s^2	0%		

Table A.4: Parameters for the computation of v_{des} .

Desired velocity lookahead $\Delta t_{\text{lookahead, vDes}}$	1 s
Max. lateral accel. $a_{\text{max,lat}}$	1.5 m/s^2
Max. free road accel. $a_{\text{max,free}}$	1.5 m/s^2
Max. free road decel. $a_{\text{brake,free}}$	-1.5 m/s^2
Max. pos. jerk $j_{\text{max,pos}}$	5 m/s^3
Max. neg. jerk $j_{\text{max,neg}}$	-5 m/s^3
Min. duration for accel.-decel.-phases $\Delta t_{\text{accel-decel}}$	4 s

Table A.5: IDM Parameters.

Desired velocity v_{des}	see Tab. A.4
Max. accel. a	1 m/s^2
Desired time gap T	2 s
Desired decel. b	$(-)2 \text{ m/s}^2$
Jam distance s_0	2 m
Coolness factor c	0.99
Accel. exponent δ	4

Table A.6: Updated parameters for the lane change experiment, since lane changes in dense traffic require driving more dynamically.

Max. accel. a	1.4 m/s ²
Desired time gap T	0.7 s
Ego response time ρ_{ego}	0.1 s
Waiting time for courtesy Δt_{wait}	8 s

Table A.7: Additional parameters for the narrowing experiment.

Assumed probability that the model is correct p_{mc}	98%
Margin for deviation from expected velocity Δv_{mc}	3.6 km/h
Tolerated decel. $a_{\text{decel,low}}$	-0.5 m/s ²

Bibliography

- [AD14] M. Althoff and J. M. Dolan, “Online Verification of Automated Road Vehicles Using Reachability Analysis,” in *IEEE Transactions on Robotics*, vol. 30, no. 4, pp. 903–918, Aug. 2014, DOI: <http://dx.doi.org/10.1109/TRO.2014.2312453>.
- [AKM17] M. Althoff, M. Koschi, and S. Manzingler, “CommonRoad: Composable Benchmarks for Motion Planning on Roads,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA, Jun. 2017, pp. 719–726, DOI: <http://dx.doi.org/10.1109/IVS.2017.7995802>.
- [ASC78] B. L. Allen, B. T. Shin, and P. J. Cooper, “Analysis of Traffic Conflicts and Collisions,” in *Transportation Research Record*, no. 667, pp. 67–74, 1978, ISSN: 0361-1981.
- [BCK17] M. Bouton, A. Cosgun, and M. J. Kochenderfer, “Belief State Planning for Autonomously Navigating Urban Intersections,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA, Jun. 2017, pp. 825–830, DOI: <http://dx.doi.org/10.1109/IVS.2017.7995818>.
- [BDSZ18] H. Banzhaf, M. Dolgov, J. Stellet, and J. M. Zöllner, “From footprints to beliefprints: Motion planning under uncertainty for maneuvering automated vehicles in dense scenarios,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 1680–1687, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569897>.
- [BDTD⁺16] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, *End to End Learning for Self-Driving Cars*, Apr. 2016, *arXiv:1604.07316v1*, <https://arxiv.org/abs/1604.07316v1>, last retrieved 2020-05-28.
- [Bel57] R. Bellman, *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press, 1957, ISBN: 978-06910-79516.

- [BFAS14] S. Bouraine, T. Fraichard, O. Azouaoui, and H. Salhi, “Passively Safe Partial Motion Planning for Mobile Robots with Limited Field-of-Views in Unknown Dynamic Environments,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Hong Kong, China, May 2014, pp. 3576–3582, DOI: <http://dx.doi.org/10.1109/ICRA.2014.6907375>.
- [BGD14] S. Brechtel, T. Gindele, and R. Dillmann, “Probabilistic Decision-Making under Uncertainty for Autonomous Driving Using Continuous POMDPs,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Qingdao, China, Oct. 2014, pp. 392–399, DOI: <http://dx.doi.org/10.1109/ITSC.2014.6957722>.
- [BLF⁺16] M. Bahram, A. Lawitzky, J. Friedrichs, M. Aeberhard, and D. Wollherr, “A Game-Theoretic Approach to Replanning-Aware Interactive Scene Prediction and Planning,” in *IEEE Transactions on Vehicular Technology*, vol. 65, no. 6, pp. 3981–3992, Jun. 2016, DOI: <http://dx.doi.org/10.1109/TVT.2015.2508009>.
- [BNFK18] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, “Scalable Decision Making with Sensor Occlusions for Autonomous Driving,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, Australia, May 2018, pp. 2076–2081, DOI: <http://dx.doi.org/10.1109/ICRA.2018.8460914>.
- [Bou96] C. Boutilier, “Planning, Learning and Coordination in Multiagent Decision Processes,” in *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge*, ser. TARK ’96. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1996, pp. 195–210, also available via: <https://www.cs.toronto.edu/~cebly/Papers/tark96.pdf>, last retrieved 2020-05-28.
- [BTZS15] P. Bender, Ö. Ş. Taş, J. Ziegler, and C. Stiller, “The Combinatorial Aspect of Motion Planning: Maneuver Variants in Structured Environments,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Seoul, South Korea, Jun. 2015, pp. 1386–1392, DOI: <http://dx.doi.org/10.1109/IVS.2015.7225909>.
- [BZS14] P. Bender, J. Ziegler, and C. Stiller, “Lanelets: Efficient map representation for autonomous driving,” in *Proceedings of the IEEE*

- Intelligent Vehicles Symposium (IV)*, Dearborn, MI, USA, Jun. 2014, pp. 420–425, DOI: <http://dx.doi.org/10.1109/IVS.2014.6856487>.
- [CGEO15] A. G. Cunningham, E. Galceran, R. M. Eustice, and E. Olson, “MPDM: Multipolicy Decision-Making in Dynamic, Uncertain Environments for Autonomous Driving,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015, pp. 1670–1677, DOI: <http://dx.doi.org/10.1109/ICRA.2015.7139412>.
- [CHL⁺05] H. M. Choset, S. Hutchinson, K. M. Lynch, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementation*. Cambridge, MA, USA: The MIT press, 2005, ISBN: 978-02620-33275.
- [CTX⁺18] J. Chen, C. Tang, L. Xin, S. E. Li, and M. Tomizuka, “Continuous Decision Making for On-Road Autonomous Driving under Uncertain and Interactive Environments,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Changshu, China, Jun. 2018, pp. 1651–1658, DOI: <http://dx.doi.org/10.1109/IVS.2018.8500605>.
- [DE15] F. Damerow and J. Eggert, “Risk-Aversive Behavior Planning under Multiple Situations with Uncertainty,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Las Palmas, Spain, Sep. 2015, pp. 656–663, DOI: <http://dx.doi.org/10.1109/ITSC.2015.113>.
- [Don16] E. Donges, “Driver Behavior Models,” in *Handbook of Driver Assistance Systems: Basic Information, Components and Systems for Active Safety and Comfort*, H. Winner, S. Hakuli, F. Lotz, and C. Singer, Eds. Cham, Switzerland: Springer, 2016, pp. 19–33, DOI: <http://dx.doi.org/10.1007/978-3-319-12352-3>.
- [Dub57] L. E. Dubins, “On Curves of Minimal Length with a Constraint on Average Curvature, and with Prescribed Initial and Terminal Positions and Tangents,” in *American Journal of Mathematics*, vol. 79, no. 3, pp. 497–516, Jul. 1957, DOI: <http://dx.doi.org/10.2307/2372560>.
- [EDK15] J. Eggert, F. Damerow, and S. Klingelschmitt, “The Foresighted Driver Model,” in *Proceedings of the IEEE Intelligent Vehicles*

- Symposium (IV)*, Seoul, South Korea, Jun. 2015, pp. 322–329, DOI: <http://dx.doi.org/10.1109/IVS.2015.7225706>.
- [FA03] T. Fraichard and H. Asama, “Inevitable Collision States. A Step towards Safer Robots?” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA, Oct. 2003, pp. 388–393, DOI: <http://dx.doi.org/10.1109/IROS.2003.1250659>.
- [FHL08] D. Ferguson, T. M. Howard, and M. Likhachev, “Motion Planning in Urban Environments,” in *Journal of Field Robotics*, vol. 25, no. 11-12, pp. 939–960, Nov. 2008, DOI: <http://dx.doi.org/10.1002/rob.20265>.
- [FTO⁺08] L. Fletcher, S. Teller, E. Olson, D. Moore, Y. Kuwata, J. How, J. Leonard, I. Miller, M. Campbell, D. Huttenlocher *et al.*, “The MIT–Cornell Collision and Why It Happened,” in *Journal of Field Robotics*, vol. 25, no. 10, pp. 775–807, Oct. 2008, DOI: <http://dx.doi.org/10.1002/rob.20266>.
- [GPMN16] D. González, J. Pérez, V. Milanés, and F. Nashashibi, “A Review of Motion Planning Techniques for Automated Vehicles,” in *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 4, pp. 1135–1145, Apr. 2016, DOI: <http://dx.doi.org/10.1109/TITS.2015.2498841>.
- [GSZD18] M. Graf, O. Speidel, J. Ziegler, and K. Dietmayer, “Trajectory Planning for Automated Vehicles Using Driver Models,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 1455–1460, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569373>.
- [HBA⁺17] C. Hubmann, M. Becker, D. Althoff, D. Lenz, and C. Stiller, “Decision Making for Autonomous Driving Considering Interaction and Uncertain Prediction of Surrounding Vehicles,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA, Jun. 2017, pp. 1671–1678, DOI: <http://dx.doi.org/10.1109/IVS.2017.7995949>.
- [HK07] T. M. Howard and A. Kelly, “Optimal Rough Terrain Trajectory Generation for Wheeled Mobile Robots,” in *The International*

- Journal of Robotics Research*, vol. 26, no. 2, pp. 141–166, Feb. 2007, DOI: <http://dx.doi.org/10.1177/0278364906075328>.
- [HQS⁺19] C. Hubmann, N. Quetschlich, J. Schulz, J. Bernhard, D. Althoff, and C. Stiller, “A POMDP Maneuver Planner For Occlusions in Urban Scenarios,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, Jun. 2019, pp. 2172–2179, DOI: <http://dx.doi.org/10.1109/IVS.2019.8814179>.
- [HSB⁺18] C. Hubmann, J. Schulz, M. Becker, D. Althoff, and C. Stiller, “Automated Driving in Uncertain Environments: Planning With Interaction and Uncertain Maneuver Prediction,” in *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 1, pp. 5–17, Mar. 2018, DOI: <http://dx.doi.org/10.1109/TIV.2017.2788208>.
- [HSX⁺18] C. Hubmann, J. Schulz, G. Xu, D. Althoff, and C. Stiller, “A Belief State Planner for Interactive Merge Maneuvers in Congested Traffic,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 1617–1624, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569729>.
- [HWT⁺16] A. M. Hellmund, S. Wirges, Ö. Ş. Taş, C. Bandera, and N. O. Salscheider, “Robot Operating System: A Modular Software Framework for Automated Driving,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, Nov. 2016, pp. 1564–1570, DOI: <http://dx.doi.org/10.1109/ITSC.2016.7795766>.
- [KF11] S. Karaman and E. Frazzoli, “Sampling-Based Algorithms for Optimal Motion Planning,” in *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, Jun. 2011, DOI: <http://dx.doi.org/10.1177/0278364911406761>.
- [KLC98] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and Acting in Partially Observable Stochastic Domains,” in *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, May 1998, DOI: [http://dx.doi.org/10.1016/S0004-3702\(98\)00023-X](http://dx.doi.org/10.1016/S0004-3702(98)00023-X).
- [Koc15] M. J. Kochenderfer, *Decision Making under Uncertainty: Theory and Application*. Cambridge, MA, USA: The MIT Press, Jul. 2015, ISBN: 978-02620-29254.

- [KQCD15] C. Katrakazas, M. Quddus, W.-H. Chen, and L. Deka, “Real-Time Motion Planning Methods for Autonomous on-Road Driving: State-of-the-Art and Future Research Directions,” in *Transportation Research Part C: Emerging Technologies*, vol. 60, pp. 416–442, Nov. 2015, DOI: <http://dx.doi.org/10.1016/j.trc.2015.09.011>.
- [KSLO96] L. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars, “Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces,” in *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, Aug. 1996, DOI: <http://dx.doi.org/10.1109/70.508439>.
- [KTH07] A. Kesting, M. Treiber, and D. Helbing, “General Lane-Changing Model MOBIL for Car-Following Models,” in *Transportation Research Record*, vol. 1999, no. 1, pp. 86–94, 2007, DOI: <http://dx.doi.org/10.3141/1999-10>.
- [KTH10] A. Kesting, M. Treiber, and D. Helbing, “Enhanced Intelligent Driver Model to Access the Impact of Driving Strategies on Traffic Capacity,” in *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, Oct. 2010, DOI: <http://dx.doi.org/10.1098/rsta.2010.0084>.
- [KZ86] K. Kant and S. W. Zucker, “Toward Efficient Trajectory Planning: The Path-Velocity Decomposition,” in *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 72–89, Sep. 1986, DOI: <http://dx.doi.org/10.1177/027836498600500304>.
- [KZP⁺08] S. Kammel, J. Ziegler, B. Pitzer, M. Werling, T. Gindele, D. Jagzent, J. Schröder, M. Thuy, M. Goebel, F. von Hundelshausen *et al.*, “Team AnnieWAY’s Autonomous System for the 2007 DARPA Urban Challenge,” in *Journal of Field Robotics*, vol. 25, no. 9, pp. 615–639, Sep. 2008, DOI: <http://dx.doi.org/10.1002/rob.20252>.
- [LaV06] S. M. LaValle, *Planning Algorithms*. New York, NY, USA: Cambridge University Press, May 2006, ISBN: 978-0-521-86205-9.
- [LHT⁺08] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman *et al.*, “A

- Perception-Driven Autonomous Urban Vehicle,” in *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, Oct. 2008, DOI: <http://dx.doi.org/10.1002/rob.20262>.
- [LKB⁺13] M. Liebner, F. Klanner, M. Baumann, C. Ruhhammer, and C. Stiller, “Velocity-Based Driver Intent Inference at Urban Intersections in the Presence of Preceding Vehicles,” in *IEEE Intelligent Transportation Systems Magazine*, vol. 5, no. 2, pp. 10–21, 2013, DOI: <http://dx.doi.org/10.1109/MITS.2013.2246291>.
- [LKK16] D. Lenz, T. Kessler, and A. Knoll, “Tactical Cooperative Planning for Autonomous Highway Driving Using Monte-Carlo Tree Search,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Gothenburg, Sweden, Jun. 2016, pp. 447–453, DOI: <http://dx.doi.org/10.1109/IVS.2016.7535424>.
- [MBB⁺08] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke *et al.*, “Junior: The Stanford Entry in the Urban Challenge,” in *Journal of Field Robotics*, vol. 25, no. 9, pp. 569–597, Sep. 2008, DOI: <http://dx.doi.org/10.1002/rob.20258>.
- [MUDL11] M. McNaughton, C. Urmson, J. M. Dolan, and J.-W. Lee, “Motion Planning for Autonomous Driving with a Conformal Spatiotemporal Lattice,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 2011, pp. 4889–4895, DOI: <http://dx.doi.org/10.1109/ICRA.2011.5980223>.
- [OML18] P. F. Orzechowski, A. Meyer, and M. Lauer, “Tackling Occlusions Limited Sensor Range with Set-Based Safety Verification,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 1729–1736, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569332>.
- [PČY⁺16] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, “A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles,” in *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, Mar. 2016, DOI: <http://dx.doi.org/10.1109/TIV.2016.2578706>.

- [PDK⁺13] D. Petrich, T. Dang, D. Kasper, G. Breuel, and C. Stiller, “Map-Based Long Term Motion Prediction for Vehicles in Traffic Environments,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, The Hague, Netherlands, Oct. 2013, pp. 2166–2172, DOI: <http://dx.doi.org/10.1109/ITSC.2013.6728549>.
- [PK05] M. Pivtoraiko and A. Kelly, “Efficient Constrained Path Planning via Search in State Lattices,” in *Proceedings of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Munich, Germany, Sep. 2005, p. n/a, also available via: https://www.ri.cmu.edu/pub_files/2005/9/ISAIRAS05.pdf, last retrieved 2020-05-28.
- [Pog19] F. Poggenhans, “Generierung hochdetaillierter Karten für das automatisierte Fahren,” doctoral dissertation, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany, 2019, DOI: <http://dx.doi.org/10.5445/IR/1000100719>.
- [Pon87] L. S. Pontryagin, *Mathematical Theory of Optimal Processes*. Boca Raton, FL, USA: CRC Press, Mar. 1987, ISBN: 978-2-88124-077-5.
- [PZA17] C. Pek, P. Zahn, and M. Althoff, “Verifying the Safety of Lane Change Maneuvers of Self-Driving Vehicles Based on Formalized Traffic Rules,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA, Jun. 2017, pp. 1477–1483, DOI: <http://dx.doi.org/10.1109/IVS.2017.7995918>.
- [QCG⁺09] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “ROS: an open-source Robot Operating System,” in *Open-source software workshop of the IEEE International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, May 2009, p. n/a, also available via: <http://robotics.stanford.edu/~ang/papers/icraoss09-ROS.pdf>, last retrieved 2020-05-28.
- [RN16] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. Harlow, UK: Pearson Education Limited, May 2016, ISBN: 978-12921-53964.
- [RWKS19] S. Richter, S. Wirges, H. Königshof, and C. Stiller, “Fusion of Range Measurements and Semantic Estimates in an Evidential

- Framework / Fusion von Distanzmessungen und semantischen Größen im Rahmen der Evidenztheorie,” in *tm - Technisches Messen*, vol. 86, no. s1, pp. 102–106, Sep. 2019, DOI: <http://dx.doi.org/10.1515/teme-2019-0052>.
- [SHK17] Z. N. Sunberg, C. J. Ho, and M. J. Kochenderfer, “The Value of Inferring the Internal State of Traffic Participants for Autonomous Freeway Driving,” in *Proceedings of the American Control Conference (ACC)*, Seattle, WA, USA, May 2017, pp. 3004–3010, DOI: <http://dx.doi.org/10.23919/ACC.2017.7963408>.
- [SS18] M. Sons and C. Stiller, “Efficient Multi-Drive Map Optimization towards Life-Long Localization Using Surround View,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 2671–2677, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8570011>.
- [SSS16] S. Shalev-Shwartz, S. Shammah, and A. Shashua, *Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving*, Oct. 2016, *arXiv:1610.03295v1*, <https://arxiv.org/abs/1610.03295v1>, last retrieved 2020-05-28.
- [SSS18] S. Shalev-Shwartz, S. Shammah, and A. Shashua, *On a Formal Model of Safe and Scalable Self-Driving Cars*, Oct. 2018, *arXiv:1708.06374v6*, <https://arxiv.org/abs/1708.06374v6>, last retrieved 2020-05-28.
- [Sto09] A. Storkey, “When Training and Test Sets are Different: Characterising Learning Transfer,” in *Dataset Shift in Machine Learning*, J. Quionero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, Eds. Cambridge, MA, USA: The MIT Press, 2009, ch. 1, pp. 3–28, ISBN: 978-0-262-17005-5.
- [StVO (D)] *German law: Straßenverkehrs-Ordnung vom 6. März 2013 (BGBl. I S. 367)*, die zuletzt durch Artikel 4a der Verordnung vom 6. Juni 2019 (BGBl. I S. 756) geändert worden ist.
- [SV10] D. Silver and J. Veness, “Monte-Carlo planning in large POMDPs,” in *Advances in Neural Information Processing Systems, Proceedings of the Neural Information Processing Systems Conference*, Vancouver, Canada, 2010, pp. 2164–2172, also

- available via: <https://papers.nips.cc/paper/4031-monte-carlo-planning-in-large-pomdps.pdf>, last retrieved 2020-05-28.
- [SWK16] J. Schlechtriemen, K. P. Wabersich, and K.-D. Kuhnert, “Wiggling through Complex Traffic: Planning Trajectories Constrained by Predictions,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Gothenburg, Sweden, Jun. 2016, pp. 1293–1300, DOI: <http://dx.doi.org/10.1109/IVS.2016.7535557>.
- [THH00] M. Treiber, A. Hennecke, and D. Helbing, “Congested Traffic States in Empirical Observations and Microscopic Simulations,” in *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000, DOI: <http://dx.doi.org/10.1103/PhysRevE.62.1805>.
- [THNS89] A. Takahashi, T. Hongo, Y. Ninomiya, and G. Sugimoto, “Local Path Planning And Motion Control For Agv In Positioning,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshops*, Tsukuba, Japan, Sep. 1989, pp. 392–397, DOI: <http://dx.doi.org/10.1109/IROS.1989.637936>.
- [THS18] Ö. Ş. Taş, F. Hauser, and C. Stiller, “Decision-time postponing motion planning for combinatorial uncertain maneuvering,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 2419–2425, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569580>.
- [TKZS16] Ö. Ş. Taş, F. Kuhnt, J. M. Zöllner, and C. Stiller, “Functional System Architectures towards Fully Automated Driving,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Gothenburg, Sweden, Jun. 2016, pp. 304–309, DOI: <http://dx.doi.org/10.1109/IVS.2016.7535402>.
- [TMD⁺06] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann *et al.*, “Stanley: The Robot That Won the DARPA Grand Challenge,” in *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, Sep. 2006, DOI: <http://dx.doi.org/10.1002/rob.20147>.
- [TS20] Ö. Ş. Taş and C. Stiller, *Tackling Existence Probabilities of Objects with Motion Planning for Automated Urban Driving*, Feb.

- 2020, *arXiv:2002.01254v1*, <https://arxiv.org/abs/2002.01254v1>, last retrieved 2020-05-28.
- [UAB⁺08] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer *et al.*, “Autonomous Driving in Urban Environments: Boss and the Urban Challenge,” in *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, Aug. 2008, DOI: <http://dx.doi.org/10.1002/rob.20255>.
- [vHHH⁺08] F. von Hundelshausen, M. Himmelsbach, F. Hecker, A. Mueller, and H.-J. Wuensche, “Driving with Tentacles: Integral Structures for Sensing and Motion,” in *Journal of Field Robotics*, vol. 25, no. 9, pp. 640–673, Sep. 2008, DOI: <http://dx.doi.org/10.1002/rob.20256>.
- [Wer11] M. Werling, *Ein Neues Konzept für die Trajektoriengenerierung und -Stabilisierung in Zeitkritischen Verkehrsszenarien*, ser. Schriftenreihe des Instituts für Angewandte Informatik - Automatisierungstechnik, Karlsruher Institut für Technologie. Karlsruhe, Germany: KIT Scientific Publishing, 2011, no. 34, doctoral dissertation, DOI: <http://dx.doi.org/10.5445/KSP/1000021738>.
- [WZKT10] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, “Optimal Trajectory Generation for Dynamic Street Scenarios in a Frenét Frame,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Anchorage, AK, USA, May 2010, pp. 987–993, DOI: <http://dx.doi.org/10.1109/ROBOT.2010.5509799>.
- [XPWD14] W. Xu, J. Pan, J. Wei, and J. M. Dolan, “Motion Planning under Uncertainty for On-Road Autonomous Driving,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Hong Kong, China, May 2014, pp. 2507–2512, DOI: <http://dx.doi.org/10.1109/ICRA.2014.6907209>.
- [ZBDS14] J. Ziegler, P. Bender, T. Dang, and C. Stiller, “Trajectory Planning for Bertha — A Local, Continuous Method,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Dearborn, MI, USA, Jun. 2014, pp. 450–457, DOI: <http://dx.doi.org/10.1109/IVS.2014.6856581>.
- [ZBS⁺14] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. G. Keller *et al.*,

- “Making Bertha Drive—An Autonomous Journey on a Historic Route,” in *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 2, pp. 8–20, 2014, DOI: <http://dx.doi.org/10.1109/MITS.2014.2306552>.
- [ZCC⁺17] W. Zhan, J. Chen, C. Chan, C. Liu, and M. Tomizuka, “Spatially-Partitioned Environmental Representation and Planning Architecture for on-Road Autonomous Driving,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, USA, Jun. 2017, pp. 632–639, DOI: <http://dx.doi.org/10.1109/IVS.2017.7995789>.
- [Zie17] J. Ziegler, *Optimale Trajektorienplanung für Automobile*, ser. Schriftenreihe des Instituts für Mess- und Regelungstechnik, Karlsruher Institut für Technologie. Karlsruhe, Germany: KIT Scientific Publishing, 2017, no. 35, doctoral dissertation, DOI: <http://dx.doi.org/10.5445/KSP/1000056530>.
- [ZLCT16] W. Zhan, C. Liu, C.-Y. Chan, and M. Tomizuka, “A Non-Conservatively Defensive Strategy for Urban Autonomous Driving,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, Nov. 2016, pp. 459–464, DOI: <http://dx.doi.org/10.1109/ITSC.2016.7795595>.
- [ZS09] J. Ziegler and C. Stiller, “Spatiotemporal State Lattices for Fast Trajectory Planning in Dynamic On-Road Driving Scenarios,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, St. Louis, MO, USA, Oct. 2009, pp. 1879–1884, DOI: <http://dx.doi.org/10.1109/IROS.2009.5354448>.

Publications by the Author

- [1] M. Naumann and A.-M. Hellmund, “Multi-Drive Road Map Generation on Standardized High-Velocity Roads using Low-Cost Sensor Data,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, Nov. 2016, pp. 113–120, DOI: <http://dx.doi.org/10.1109/ITSC.2016.7795540>.
- [2] M. Naumann, P. Orzechowski, C. Burger, Ö. Ş. Taş, and C. Stiller, “Herausforderungen für die Verhaltensplanung kooperativer automatischer Fahrzeuge,” in *AAET Automatisiertes und vernetztes Fahren*, ser. Beiträge zum gleichnamigen 18. Braunschweiger Symposium vom 8. und 9. Februar 2017. Braunschweig, Germany: ITS automotive nord e.V., Feb. 2017, pp. 287–307, ISBN: 978-3-937655-41-3.
- [3] M. Naumann and C. Stiller, “Towards Cooperative Motion Planning for Automated Vehicles in Mixed Traffic,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 9th Workshop on Planning, Perception and Navigation for Intelligent Vehicles*, Vancouver, Canada, Sep. 2017, also available via: <https://arxiv.org/abs/1708.06962v1>, last retrieved 2020-05-28.
- [4] E. Rehder, M. Naumann, N. O. Salscheider, and C. Stiller, *Cooperative Motion Planning for Non-Holonomic Agents with Value Iteration Networks*, Sep. 2017, *arXiv:1709.05273v1*, <https://arxiv.org/abs/1709.05273v1>, last retrieved 2020-05-28.
- [5] M. Naumann, F. Poggenhans, M. Lauer, and C. Stiller, “CoInCar-Sim: An Open-Source Simulation Framework for Cooperatively Interacting Automobiles,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Changshu, China, Jun. 2018, pp. 1879–1884, DOI: <http://dx.doi.org/10.1109/IVS.2018.8500405>.
- [6] M. Naumann, M. Lauer, and C. Stiller, “Generating Comfortable, Safe and Comprehensible Trajectories for Automated Vehicles in Mixed Traffic,” in *Proceedings of the IEEE International Conference on Intelligent*

- Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 575–582, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569658>.
- [7] F. Poggenschans, J.-H. Pauls, J. Janosovits, S. Orf, M. Naumann, F. Kuhnt, and M. Mayr, “Lanelet2: A High-Definition Map Framework for the Future of Automated Driving,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 1672–1679, DOI: <http://dx.doi.org/10.1109/ITSC.2018.8569929>.
- [8] M. Naumann, H. Königshof, M. Lauer, and C. Stiller, “Safe but not Overcautious Motion Planning under Occlusions and Limited Sensor Range,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, Jun. 2019, pp. 140–145, DOI: <http://dx.doi.org/10.1109/IVS.2019.8814251>.
- [9] M. Naumann, H. Königshof, and C. Stiller, “Provably Safe and Smooth Lane Changes in Mixed Traffic,” in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Auckland, NZ, Oct. 2019, pp. 1832–1837, DOI: <http://dx.doi.org/10.1109/ITSC.2019.8917461>.
- [10] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kümmerle, H. Königshof, C. Stiller, A. de La Fortelle *et al.*, *INTERACTION Dataset: An INTERnational, Adversarial and Cooperative motion Dataset in Interactive Driving Scenarios with Semantic Maps*, Sep. 2019, *arXiv:1910.03088v1*, <https://arxiv.org/abs/1910.03088v1>, last retrieved 2020-05-28.
- [11] M. Naumann, L. Sun, W. Zhan, and M. Tomizuka, “Analyzing the Suitability of Cost Functions for Explaining and Imitating Human Driving Behavior based on Inverse Reinforcement Learning,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, France, Jun. 2020, accepted.

Supervised Theses

Chenyang Zhou, “Decentralized Planning of Macro-Actions for Cooperative Automated Vehicles with Hierarchical Monte Carlo Tree Search,” Master’s Thesis, Karlsruhe Institute of Technology, Institute of Measurement and Control, Karlsruhe, Germany, Feb. 2018. Supervised jointly with Karl Kurzer.

Florian Engelhorn, “Dezentrale kooperative Planung autonomer Fahrzeuge auf Basis einer kontinuierlichen Monte Carlo Tree Search,” Master’s Thesis, Karlsruhe Institute of Technology, Institute of Measurement and Control, Karlsruhe, Germany, Mar. 2018. Supervised jointly with Karl Kurzer.

Bernhard Trünkle, “Entwurf und Parameteroptimierung eines Kostenfunktionalns für die Trajektorienplanung von automatisierten Fahrzeugen,” Master’s Thesis, Karlsruhe Institute of Technology, Institute of Measurement and Control, Karlsruhe, Germany, Aug. 2018.

Rudolf Ehle, “Entwurf eines Sampling-Verfahrens für die Trajektorienplanung automatischer Fahrzeuge mittels Simulated Annealing,” Master’s Thesis, Karlsruhe Institute of Technology, Institute of Measurement and Control, Karlsruhe, Germany, Nov. 2018.

Gunnar Diestmann, “Trajektorienplanung für einen kooperativen Fahrstreifenwechsel hochautomatisierter Fahrzeuge,” Master’s Thesis, Karlsruhe Institute of Technology, Institute of Measurement and Control, Karlsruhe, Germany, May 2019.

Schriftenreihe

Institut für Mess- und Regelungstechnik

Karlsruher Institut für Technologie

(1613-4214)

- Band 001** Hans, Annegret
Entwicklung eines Inline-Viskosimeters
auf Basis eines magnetisch-induktiven
Durchflussmessers. 2004
ISBN 3-937300-02-3
- Band 002** Heizmann, Michael
Auswertung von forensischen Riefenspuren
mittels automatischer Sichtprüfung. 2004
ISBN 3-937300-05-8
- Band 003** Herbst, Jürgen
Zerstörungsfreie Prüfung von Abwasserkanälen
mit Klopferschall. 2004
ISBN 3-937300-23-6
- Band 004** Kammel, Sören
Deflektometrische Untersuchung spiegelnd
reflektierender Freiformflächen. 2005
ISBN 3-937300-28-7
- Band 005** Geistler, Alexander
Bordautonome Ortung von Schienenfahrzeugen
mit Wirbelstrom-Sensoren. 2007
ISBN 978-3-86644-123-1
- Band 006** Horn, Jan
Zweidimensionale Geschwindigkeitsmessung
texturierter Oberflächen mit flächenhaften
bildgebenden Sensoren. 2007
ISBN 978-3-86644-076-0

- Band 007** Hoffmann, Christian
Fahrzeuginnenraumdetektion durch Fusion monoskopischer Videomerkmale. 2007
ISBN 978-3-86644-139-2
- Band 008** Dang, Thao
Kontinuierliche Selbstkalibrierung von Stereokameras. 2007
ISBN 978-3-86644-164-4
- Band 009** Kapp, Andreas
Ein Beitrag zur Verbesserung und Erweiterung der Lidar-Signalverarbeitung für Fahrzeuge. 2007
ISBN 978-3-86644-174-3
- Band 010** Horbach, Jan
Verfahren zur optischen 3D-Vermessung spiegelnder Oberflächen. 2008
ISBN 978-3-86644-202-3
- Band 011** Böhringer, Frank
Gleiselektive Ortung von Schienenfahrzeugen mit bordautonomer Sensorik. 2008
ISBN 978-3-86644-196-5
- Band 012** Xin, Binjian
Auswertung und Charakterisierung dreidimensionaler Messdaten technischer Oberflächen mit Riefentexturen. 2009
ISBN 978-3-86644-326-6
- Band 013** Cech, Markus
Fahrspurschätzung aus monokularen Bildfolgen für innerstädtische Fahrerassistentenanwendungen. 2009
ISBN 978-3-86644-351-8
- Band 014** Speck, Christoph
Automatisierte Auswertung forensischer Spuren auf Patronenhülsen. 2009
ISBN 978-3-86644-365-5

- Band 015** Bachmann, Alexander
Dichte Objektsegmentierung in Stereobildfolgen. 2010
ISBN 978-3-86644-541-3
- Band 016** Duchow, Christian
Videobasierte Wahrnehmung markierter Kreuzungen mit lokalem Markierungstest und Bayes'scher Modellierung. 2011
ISBN 978-3-86644-630-4
- Band 017** Pink, Oliver
Bildbasierte Selbstlokalisierung von Straßenfahrzeugen. 2011
ISBN 978-3-86644-708-0
- Band 018** Hensel, Stefan
Wirbelstromsensorbasierte Lokalisierung von Schienenfahrzeugen in topologischen Karten. 2011
ISBN 978-3-86644-749-3
- Band 019** Carsten Hasberg
Simultane Lokalisierung und Kartierung spurgeführter Systeme. 2012
ISBN 978-3-86644-831-5
- Band 020** Pitzer, Benjamin
Automatic Reconstruction of Textured 3D Models. 2012
ISBN 978-3-86644-805-6
- Band 021** Roser, Martin
Modellbasierte und positionsgenaue Erkennung von Regentropfen in Bildfolgen zur Verbesserung von videobasierten Fahrerassistenzfunktionen. 2012
ISBN 978-3-86644-926-8

- Band 022** Loose, Heidi
Dreidimensionale Straßenmodelle für Fahrerassistenzsysteme auf Landstraßen. 2013
ISBN 978-3-86644-942-8
- Band 023** Rapp, Holger
Reconstruction of Specular Reflective Surfaces using Auto-Calibrating Deflectometry. 2013
ISBN 978-3-86644-966-4
- Band 024** Moosmann, Frank
Interlacing Self-Localization, Moving Object Tracking and Mapping for 3D Range Sensors. 2013
ISBN 978-3-86644-977-0
- Band 025** Geiger, Andreas
Probabilistic Models for 3D Urban Scene Understanding from Movable Platforms. 2013
ISBN 978-3-7315-0081-0
- Band 026** Hörter, Marko
Entwicklung und vergleichende Bewertung einer bildbasierten Markierungslichtsteuerung für Kraftfahrzeuge. 2013
ISBN 978-3-7315-0091-9
- Band 027** Kitt, Bernd
Effiziente Schätzung dichter Bewegungsvektorfelder unter Berücksichtigung der Epipolarometrie zwischen unterschiedlichen Ansichten einer Szene. 2013
ISBN 978-3-7315-0105-3
- Band 028** Lategahn, Henning
Mapping and Localization in Urban Environments Using Cameras. 2013
ISBN 978-3-7315-0135-0

- Band 029** Tischler, Karin
**Informationsfusion für die kooperative
Umfeldwahrnehmung vernetzter Fahrzeuge.** 2014
ISBN 978-3-7315-0166-4
- Band 030** Schmidt, Christian
**Fahrstrategien zur Unfallvermeidung im
Straßenverkehr für Einzel- und
Mehrobjektszenarien.** 2014
ISBN 978-3-7315-0198-5
- Band 031** Firl, Jonas
**Probabilistic Maneuver Recognition
in Traffic Scenarios.** 2014
ISBN 978-3-7315-0287-6
- Band 032** Schönbein, Miriam
**Omnidirectional Stereo Vision
for Autonomous Vehicles.** 2015
ISBN 978-3-7315-0357-6
- Band 033** Nicht erschienen
- Band 034** Liebner, Martin
**Fahrerabsichtserkennung und Risikobewertung
für warnende Fahrerassistenzsysteme.** 2016
ISBN 978-3-7315-0508-2
- Band 035** Ziegler, Julius
Optimale Trajektorienplanung für Automobile. 2017
ISBN 978-3-7315-0553-2
- Band 036** Harms, Hannes
**Genauigkeitsuntersuchung von
binokularen Normalenvektoren für
die Umfeldwahrnehmung.** 2017
ISBN 978-3-7315-0628-7

- Band 037** Ruhhammer, Christian
Inferenz von Kreuzungsinformationen aus Flottendaten. 2017
ISBN 978-3-7315-0721-5
- Band 038** Stein, Denis
Mobile laser scanning based determination of railway network topology and branching direction on turnouts. 2018
ISBN 978-3-7315-0743-7
- Band 039** Yi, Boliang
Integrated Planning and Control for Collision Avoidance Systems. 2018
ISBN 978-3-7315-0785-7
- Band 040** Schwarze, Tobias
Compact Environment Modelling from Unconstrained Camera Platforms. 2018
ISBN 978-3-7315-0801-4
- Band 041** Knorr, Moritz
Self-Calibration of Multi-Camera Systems for Vehicle Surround Sensing. 2018
ISBN 978-3-7315-0765-9
- Band 042** Rabe, Johannes
Lane-Precise Localization with Production Vehicle Sensors and Application to Augmented Reality Navigation. 2018
ISBN 978-3-7315-0854-0
- Band 043** Weiser, Andreas
Probabilistische Vorhersage von Fahrstreifenwechseln für hochautomatisiertes Fahren auf Autobahnen. 2019
ISBN 978-3-7315-0794-9

- Band 044** Tian, Wei
Novel Aggregated Solutions for Robust Visual Tracking in Traffic Scenarios. 2019
ISBN 978-3-7315-0915-8
- Band 045** Gräter, Johannes
Monokulare Visuelle Odometrie auf Multisensorplattformen für autonome Fahrzeuge. 2019
ISBN 978-3-7315-0935-6
- Band 046** Spindler, Max
Ferromagnetische Inhomogenitäten zur berührungslosen Bestimmung der Geschwindigkeit und gleiselektiven Position von Schienenfahrzeugen. 2020
ISBN 978-3-7315-1000-0
- Band 047** Sons, Marc
Automatische Erzeugung langzeitverfügbarer Punktmerkmalskarten zur robusten Lokalisierung mit Multi-Kamera-Systemen für automatisierte Fahrzeuge. 2020
ISBN 978-3-7315-1029-1
- Band 048** Hubmann, Constantin
Belief State Planning for Autonomous Driving: Planning with Interaction, Uncertain Prediction and Uncertain Perception. 2020
ISBN 978-3-7315-1039-0
- Band 049** Naumann, Maximilian
Probabilistic Motion Planning for Automated Vehicles. 2020
ISBN 978-3-7315-1070-3



This work targets the problem of motion planning for automated vehicles. As a prerequisite for their on-road deployment, automated vehicles must show an appropriate and reliable driving behavior in mixed traffic, i.e. alongside human drivers. Besides the uncertainties resulting from imperfect perception, occlusions and limited sensor range, also the uncertainties in the behavior of other traffic participants have to be considered.

Related approaches often employ a deterministic problem formulation, where large uncertainties lead to conservative and over-cautious behavior. On the other hand, due to the need for real-time capability, a holistic uncertainty treatment in probabilistic approaches often induces a strong limitation of the action space of automated vehicles. Moreover, safety and traffic rule compliance are often not considered.

Thus, in this work, three motion planning approaches and a scenario-based safety approach are presented. The safety approach enhances an existing concept - which targets the guarantee that automated vehicles will never cause accidents - by the consideration of traffic rules for crossing and merging traffic, occlusions, limited sensor range and lane changes. The three presented motion planning approaches are targeted towards the different predominant uncertainties in different scenarios, while operating in a continuous action space. The performance of the presented approaches is shown in various scenarios.

ISSN 1613-4214

ISBN 978-3-7315-1070-3

Gedruckt auf FSC-zertifiziertem Papier

