

Occlusion-aware speed and road position planning for autonomous vehicles

Patiphon Narksri

Nagoya University

Doctoral Dissertation

**Occlusion-aware speed and road position planning for
autonomous vehicles**

by

Patiphon Narksri

Submitted to the Department of Intelligent Systems, Graduate School of Informatics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Informatics
August 2022

Abstract

Widespread utilization of autonomous vehicles is expected to bring multiple benefits to society, such as more efficient use of energy, less traffic congestion, and most importantly, fewer accidents, injuries and fatalities. As a result, there has been increasing research and investment in the development of these vehicles. Efforts to accelerate the mass adoption of autonomous driving technology have led to the rapid broadening of the operational domains of autonomous vehicles in recent years, i.e., autonomous vehicles are no longer expected to operate only in simplistic driving environments.

Surrounding structures and other objects often prevent autonomous vehicles from fully observing the driving environment, a phenomenon known as “occlusion”. In complex driving environments such as urban areas, partial occlusion of the driving environment by buildings, walls, foliage, etc. is inevitable, however, most autonomous driving systems do not explicitly take this limitation of their sensing modules into account during their motion planning stage, and only consider detected observable obstacles. This failure to consider occlusion while planning an autonomous vehicle’s motion can lead to catastrophic accidents at critical locations such as low visibility intersections, for example, where the vehicle can mistakenly assess that it is safe to enter the intersection as no obstacles are currently detected, while there is, in fact, a vehicle rapidly approaching from an occluded area of the intersection.

One approach commonly used in occlusion-aware motion planners to deal with the risk imposed by occlusions is to explicitly assume that there are virtual obstacles approaching at high speed from outside of the visible areas of the roadway, on a path that will intercept the ego vehicle’s trajectory. By assuming the worst-case scenario, and treating these virtual obstacles as if they are real, the motion of the ego vehicle can be efficiently and safely planned. However, approaches that are based on the assumed presence of virtual obstacles have two main drawbacks: excessively conservative movement due to the worst-case assumption, which can result in deadlock, i.e., the inability to move forward, and a lack of road position adjustment by the ego

vehicle, which prevents it from obtaining a better vantage point by moving to the side of its lane.

In the context of navigation, “deadlock” is sometimes used to refer to a situation where multiple traffic participants are waiting for the others to pass through a conflict zone, e.g., an intersection, a merging area, or a narrow one-lane road, such that none of them can actually proceed. However, the meaning of “deadlock” in the context of this dissertation is slightly different, as it does not refer to an encounter between actual vehicles but between the ego vehicle and virtual vehicles assumed to be approaching at high speed from occluded areas. By assuming that such hypothetical vehicles are always approaching at high speed, the ego vehicle is forced to stop and wait indefinitely when visibility is severely limited.

The lack of road position adjustment is a more straightforward problem. Most occlusion-aware motion planners mainly consider the longitudinal motion of the ego vehicle, i.e., its speed, and ignore lateral motion, i.e., a movement toward the vehicle’s lane boundaries. In other words, it is commonly assumed that the ego vehicle will always travel along the center of the lane. While this assumption simplifies the planning process, driving in the lane’s center does not always yield the best visibility of the driving environment.

In this dissertation, we attempt to solve these two main weaknesses of occlusion-aware motion planning approaches which are based on assumptions of approaching virtual vehicles. To solve the deadlock problem, we proposed using the hypothetical visibility of approaching, hidden vehicles, as well as potential changes in their behavior after observing the ego vehicle, to plan the ego vehicle’s speed. Our proposed speed planner first estimates visibility from point of view of the ego vehicle and potential hidden vehicles, using a 3D LiDAR scan and a road network map. The speed planner then predicts the states of occluded vehicles using a particle filter algorithm which supports possible changes in the behavior of approaching hidden vehicles based on their ability to see the ego vehicle. The behavior model for the hidden vehicles which is used to predict their states is based on an analysis of real driving data. An experiment confirms that our proposed speed planner can generate deadlock-free crossing motion

at a blind intersection of two narrow roads, a maneuver that a baseline planner was not able to execute since the ego vehicle stopped indefinitely before entering the intersection.

Regarding road position planning, we proposed predicting and quantifying the visibility conditions of driving environments using high-definition 3D point cloud maps and road network maps. The quantified visibility conditions are then used to plan road positions for the ego vehicle that will result in minimal occlusion (i.e., maximum visibility). Our proposed visibility estimation method first approximates a 3D scan of a specified viewpoint using a 3D point cloud map. The approximated scan and 3D points representing the area of the relevant surrounding lanes are then projected onto depth images, which are consequently compared to identify the visible and occluded regions of the relevant parts of the driving environment from the specified viewpoint. This visibility estimate is then quantified by calculating the ratio of the visible areas of the driving environment to the total area of the relevant driving environment, to determine an area’s visibility ratio. Candidate trajectories for the ego vehicle, with different lateral offsets from the reference path, are then generated, and visibility ratios along each candidate trajectory are calculated. Finally, a visibility cost derived from the visibility ratio is used along with other planning costs to determine the optimal output trajectory. The proposed planner was tested in various simulated traffic scenarios while using live localization and object detection results. Our experimental results show that the ego vehicle was able to effectively minimize occlusions, and consequently discover occluded vehicles earlier, in most cases, when the proposed road position planner was used, in comparison to a baseline planner.

Acknowledgments

I was told at the beginning of my doctoral studies that it would be a long and lonely journey, and in a way, it was. When I look back, however, I can clearly see that I was never truly alone. I have been helped, guided, inspired, and supported by many of the people around me throughout my journey. I would like to take this opportunity to express my sincere gratitude to those who helped make this achievement possible.

First and foremost, this research would not have been possible without the support of the Japan Science and Technology Agency (JST) Project on Open Innovation Platform with Enterprises, Research Institutes and Academia (OPERA), and the New Energy and Industrial Technology Development Organization (NEDO) on Next-Generation AI and Robot Core Technology Development Project, under Grant P15009.

I am also greatly indebted to my advisors, Prof. Eijiro Takeuchi and Prof. Kazuya Takeda. I could not have completed this journey without their continuous supervision, guidance, and encouragement. I would like to thank Prof. Katashi Nagao and Prof. Alexander Carballo for reviewing the contents of this dissertation. Moreover, I am extremely grateful to my former advisor, Prof. Nuksit Noomwongs, for introducing me to Nagoya University, and for encouraging me to apply there for my graduate studies.

I would also like to express my special thanks to all my friends and colleagues: Hatem Darweesh, Jacob Lambert, Kyle Sama, Atsushi Kuribayashi, Yuki Kitsukawa, Hiromi Takahashi, and many others not mentioned due to the lack of space, for their support and advice that helped me get through difficult times.

I would be remiss if I did not thank my parents, Sommung Narksri and Pilai Saelim, for their understanding and belief in me, which kept me going throughout this long and arduous journey.

Last but not least, I would like to thank Leslie Liang who has been by my side every step of the way. Thank you for enduring this long and stressful process with me, and for always offering your love and support.

Contents

Abstract	i
Acknowledgements	v
Contents	vii
List of Figures	xi
List of Tables	xxi
List of Acronyms	xxiii
List of Symbols	xxv
List of Author’s Publications	xxxv
1 Introduction	1
2 Related Work	9
2.1 Vehicle-to-everything (V2X) communication	10
2.2 Occlusion-aware motion planning	11
2.2.1 Machine learning	11
2.2.2 Partially Observable Markov Decision Process (POMDP) . . .	12
2.2.3 Virtual obstacles	14
2.3 Visibility estimation	16
2.4 Research gaps	17

3	Deadlock-free Speed Planning Using Estimated Visibility of Approaching Vehicles	21
3.1	Introduction	21
3.2	Blind intersection crossing	23
3.2.1	Safety constraint for intersection crossing	24
3.2.2	Visibility estimation	26
3.2.3	Particle filter	29
3.2.4	Decision making	30
3.3	Experiments	32
3.3.1	Experimental vehicle	32
3.3.2	Experimental environment	33
3.3.3	Closed-loop simulation based on recorded data	35
3.3.4	Speed profile comparison	36
3.4	Conclusion	39
4	Visibility-Dependent Behavior Model for Approaching Vehicles	41
4.1	Introduction	41
4.2	Proposed deadlock-free, blind intersection planner	44
4.2.1	Safe intersection crossing strategy	45
4.2.2	Visibility at blind intersections	49
4.2.3	Particle filter-based, occluded vehicle prediction	52
4.3	Proposed visibility-dependent behavior model for approaching vehicles	55
4.3.1	Collection of driving data	55
4.3.2	Driving data analysis	58
4.3.3	Proposed visibility-dependent behavior model	60
4.4	Experiments	62
4.4.1	Baseline comparison	63
4.4.2	Effects of perception inaccuracy	69
4.4.3	Effects of sensor mounting position	72
4.5	Conclusion	74

5	Visibility Prediction and Quantification Using High Definition Maps	77
5.1	Introduction	77
5.2	Visibility estimation using HD maps	79
5.2.1	Estimation of visible regions	80
5.2.2	3D scan approximation using a point cloud map	82
5.2.3	Calculation of visibility ratio	84
5.3	Experiments	86
5.3.1	Evaluation in a simulated environment	86
5.3.2	Evaluation in a real-world environment	91
5.4	Conclusion	96
6	Road Position Adjustment for Maximizing Visibility	99
6.1	Introduction	99
6.2	Proposed motion planner	102
6.2.1	Planner architecture	102
6.2.2	Visibility cost calculation	105
6.3	Visibility estimator	107
6.3.1	Estimation of visible regions	107
6.3.2	3D scan estimation	110
6.3.3	Calculation of visibility ratio	112
6.4	Experiments	114
6.4.1	Scenario I: T-junction crossing	116
6.4.2	Scenario II: Turning	130
6.4.3	Scenario III: Preparing to pass a parked vehicle	143
6.5	Conclusion	155
7	Conclusion and Future Work	157
7.1	Conclusion	157
7.2	Future Work	161
	Bibliography	163

List of Figures

1-1	Example of a low-visibility intersection in a residential area.	3
1-2	Low-visibility intersection where deadlock caused by excessively cautious motion planning can occur.	5
1-3	Low-visibility T-junction. Vehicle positioned on the left edge of its lane has better visibility of the merging road.	6
3-1	Low-visibility intersections in residential areas with mandatory stops.	22
3-2	Stages of crossing a severely occluded intersection using the proposed algorithm.	24
3-3	An example of a depth image created from a 3D LiDAR scan for the ego vehicle’s visibility estimation. For illustration purposes, the sample image only shows the region that corresponds to the frontal part of the 360-degree 3D scan. The brightness of each pixel indicates the depth, i.e., bright pixels are associated with object points far from the 3D LiDAR.	27
3-4	Examples of visibility estimation using real sensing data. Green rectangles represent area visible to the ego vehicle. Yellow rectangles represent locations where the ego vehicle can be observed from. Red rectangles represent occluded regions.	28
3-5	Experimental vehicle equipped with 3D LiDAR , GNSS receiver and wheel encoder.	33
3-6	High-definition point cloud map for localization using NDT scan-matching.	34

3-7	Road network map containing information about road networks in our experimental environment. The red line with an arrowhead shows one complete driving loop in the experiment.	34
3-8	Comparison of speed profiles of expert driver and proposed method at Intersection A.	37
3-9	Comparison of speed profiles of expert driver and proposed method at Intersection B.	38
4-1	Detection of the ego vehicle by an approaching, occluded vehicle. The occluded vehicle may change its behavior after observing the ego vehicle.	44
4-2	Simple, blind intersection model. Red shaded area represents the overlapping travel zone where a collision could potentially occur.	46
4-3	Example of visibility estimation results using real sensing data. The green, yellow and red squares indicate regions observable by the ego vehicle, regions from which other vehicles can observe the ego vehicle, and occluded regions, respectively. The visibility estimation approach proposed in the previous chapter can be directly incorporated into the planner proposed in this chapter for application in complex, real-world environments.	49
4-4	Visibility profiles of the ego vehicle and other vehicles at blind intersections of roadways of various widths. In each scenario, the ego vehicle's sensor is located 2 meters to the rear of its front bumper. The widths of the intersecting roadways are assumed to be equal.	51
4-5	Diagram of a particle filter algorithm	53
4-6	Experimental vehicle	57
4-7	Experimental environment	57
4-8	Speed profiles of all drivers when approaching and crossing blind intersections in a residential neighborhood (232 intersection crossings by 18 drivers).	58

4-9	Results of cluster analysis using K-Mean algorithm on collected speed profiles of drivers crossing all intersections. Red, yellow and green lines represent average speed profile of each cluster. Vertical bands show standard deviation of vehicle speed at each location.	59
4-10	Speed profiles of the three possible behaviors in proposed behavior model for vehicles approaching blind intersections.	60
4-11	Proposed behavior transition for vehicles approaching blind intersections.	61
4-12	Examples of simulation results for four stages of ego vehicle motion when crossing an occluded intersection.	64
4-12	Examples of simulation results for four stages of ego vehicle motion when crossing an occluded intersection. (cont.)	65
4-13	Comparison of motion generated by proposed and baseline planners at a 5 meter-wide, occluded intersection.	67
4-14	Comparison of proposed and baseline planners at a 15 meter-wide, occluded intersection.	68
4-15	Effects of perception noise on proposed planner at 5 meter-wide, occluded intersection.	70
4-16	Effects of perception noise on proposed planner at 15 meter-wide, occluded intersection.	71
4-17	Comparison of the effect of sensor mounting position on planner output at 5 meter-wide, occluded intersection.	73
4-18	Comparison of effect of sensor mounting position on planner output at 15 meter-wide, occluded intersection.	74
5-1	Street view of a section of hilly road close to Nagoya University. The occlusion in this scene is caused not only by the surrounding vertical structures, but also by the contour of road itself.	78

5-2	Projection of 3D points onto a 2D image plane. A set of 3D points representing the area of interest is projected onto a depth image. By comparing the resulting depth image with those of other objects in the scene, we can estimate which regions of the area of interest are visible from the specified viewing point.	80
5-3	Area of interest in the context of driving. The yellow shaded area indicates the range of the sensor ($\mathcal{S}_{\text{range}}$). A set of lane segments within the yellow area is defined as $\mathcal{L}_{\text{neighbor}}$, while $\mathcal{L}_{\text{reachable}}$ represents lane segments that share the direction of travel and reachable from p_{target} (highlighted in green). The segments highlighted in red represent $\mathcal{L}_{\text{intersect}}$, which are intersecting lane segments that could potentially be occupied by oncoming vehicles.	84
5-4	Visibility of the driving environment in CARLA’s Town01 as estimated using the proposed method. Lighter colors represent areas with higher visibility, while darker colors indicate lower visibility locations. . . .	87
5-5	Comparison of visibility at the original and modified reference intersections in CARLA’s Town01. The modified intersection has additional objects, i.e., kiosks and a bin, placed near the corner of the building to limit visibility.	88
5-6	Comparison of visibility at the original and modified intersections shown in Fig. 5-5.	89
5-7	Effects of lateral position shift on ego vehicle visibility at the original and modified intersections shown in Fig. 5-5.	90
5-8	Road network map of an area near Nagoya University, where colors indicate road elevation.	92
5-9	Visibility map of a residential area near Nagoya University (detail of area enclosed in dashed red rectangle in Fig. 5-8).	93
5-10	Visibility map of a hill near Nagoya University (detail of area enclosed in dashed blue rectangle in Fig. 5-8).	95

5-11	Visibility and elevation plot of the hill shown in Fig. 5-10. Left and right lanes are defined according to the driving scenario shown in Fig. 5-1.	96
6-1	Architecture of the proposed motion planner.	103
6-2	Generation of candidate trajectories.	104
6-3	Projection of 3D points onto a 2D image plane. A set of 3D points representing the area of interest is projected onto a depth image. By comparing the resulting depth image with those of other objects in the scene, regions visible from the specified viewing point can be estimated.	108
6-4	Area of interest $\mathcal{A}_{\text{interest}}$ for an ego vehicle at location p_{target} . The yellow shaded area represents sensing range $\mathcal{S}_{\text{range}}$. Lane segments located within the yellow area are defined as the set $\mathcal{L}_{\text{neighbor}}$. Lane segments within $\mathcal{L}_{\text{neighbor}}$ that have the same driving direction as the ego vehicle and are reachable from position p_{target} are designated as $\mathcal{L}_{\text{reachable}}$, and are represented by green lines. Lane segments with opposite traffic directions than the ego vehicle that intersect the $\mathcal{L}_{\text{reachable}}$ lanes are designated as $\mathcal{L}_{\text{intersect}}$, and are represented by red lines.	113
6-5	Town01 of the CARLA simulator. The orange rectangle indicates the location of the T-junction crossing scenario, the blue rectangle indicates the location of the right turn scenario, and the red rectangle indicates the location of the straight road where the ego vehicle has to prepare to pass a parked vehicle. The small, red rectangles are examples of the positions of hidden (or occluding) vehicles in each scenario.	115
6-6	T-junctions used in Scenario I.	117

6-7	Output trajectories generated by the baseline and proposed planners during the “clear intersection” T-junction traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	119
6-8	Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “clear intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	120
6-9	Distributions of the ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “clear intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	121
6-10	Output trajectories generated by the baseline and proposed planners during the “occluded intersection” T-junction traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	126
6-11	Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “occluded intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	127
6-12	Distributions of ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “occluded intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	128

6-13	Corners used for right turn in Scenario II.	131
6-14	Output trajectories generated by the baseline and proposed planners during the “clear corner” right turn traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	133
6-15	Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “clear corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	134
6-16	Distributions of the ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “clear corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	135
6-17	Output trajectories generated by the baseline and proposed planners during the “occluded corner” right turn traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	139
6-18	Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “occluded corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	140

6-19	Distributions of ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “occluded corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.	141
6-20	Straight road in Scenario III. Two trucks are parked in the ego vehicle’s lane, and a car is parked in the approaching lane.	144
6-21	Output trajectories generated by the baseline and proposed planners during the “preparing to pass” traffic scenario, with positions of the ego vehicle when the other vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.	145
6-22	Deviation from lane’s center and visibility ratio along the output trajectories of the baseline and proposed planners when traveling along a straight road toward the parked vehicles. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.	147
6-23	Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass” traffic scenario, where (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.	148

- 6-23 Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass” traffic scenario, (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane(cont.) 149
- 6-23 Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass traffic scenario, (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane(cont.) 150

List of Tables

2.1	Summary of advantages and disadvantages of popular occlusion-aware motion planning techniques proposed in previous studies.	18
3.1	Parameters used for speed profile generation.	36
4.1	Values of major parameters used in occluded intersection crossing experiments.	63
5.1	Parameters used for visibility estimation.	86
6.1	Parameters used by the proposed and baseline planners in our experiments.	116
6.2	Summary of statistics of ego vehicle positions when discovering hidden vehicles in the “T-junction crossing” traffic scenario (Scenario I), for both the “clear” and “occluded” variations of the intersection, for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials.	123
6.3	Summary of statistics of ego vehicle positions when discovering hidden vehicles in the “right turn” traffic scenario (Scenario II), for both the “clear” and “occluded” variations of the corners, for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials.	137

6.4 Summary of statistics of ego vehicle positions when discovering other vehicles during the “preparing to pass” traffic scenario (Scenario III), for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials. 153

List of Acronyms

CAN	Controller Area Network
CAV	Connected Automated Vehicles
DARPA	Defense Advanced Research Projects Agency
DQN	Deep Q-Network
ETA	Estimated Time of Arrival
FOV	Field of View
GNSS	Global Navigation Satellite System
HD	High Definition
HOMDP	Hierarchical Options for MDP
IDM	Intelligent Driver Model
IRL	Inverse Reinforcement Learning
LiDAR	Light Detection and Ranging
MDP	Markov Decision Process
MMS	Mobile Mapping System
MPC	Model Predictive Control

NDT Normal Distributions Transform

ODD Operational Design Domain

POMDP Partially Observable Markov Decision Process

RL Reinforcement Learning

SLAM Simultaneous Localization and Mapping

TAPIR Toolkit for Approximating and Adapting POMDP Solutions
In Realtime

TTC Time-to-Collision

V2I Vehicle-to-Infrastructure

V2V Vehicle-to-Vehicle

V2X Vehicle-to-Everything

List of Symbols

A_t^i	A Boolean variable that indicates whether a particle (i.e., an occluded vehicle) is aware of the ego vehicle
$\mathcal{A}_{\text{interest}}$	An area of interest
$\mathcal{A}_{\text{visible}}$	Total visible area of $\mathcal{A}_{\text{interest}}$
\mathcal{C}	A weighted sum of all of the costs associated with a candidate trajectory
\mathcal{D}	A depth image
$\mathcal{D}(\bar{u}, \bar{v})$	Pixel (\bar{u}, \bar{v}) of depth image \mathcal{D}
$\mathcal{D}_{\text{interest}}$	A depth image of an area of interest
D_{map}	HD map information
$\mathcal{D}_{\text{other}}$	A depth image of all surrounding objects
$\mathcal{D}_{\text{resolution}}^{\text{horizontal}}$	Width of a depth image
$\mathcal{D}_{\text{resolution}}^{\text{vertical}}$	Height of a depth image
D_{sensor}	Sensing data
Δt	Time step (i.e., a sampling interval)
$\mathcal{H}_{\text{object}}$	A z-coordinate representing vertical offset from a lane's surface

$\mathcal{H}_{\text{planning}}$	A planning horizon
$\mathcal{L}_{\text{neighbor}}$	Lane segments located within sensing range $\mathcal{S}_{\text{range}}$ of p_{target}
$\mathcal{L}_{\text{reachable}}$	Lane segments within $\mathcal{L}_{\text{neighbor}}$ that have the same driving direction and are reachable from p_{target}
\mathcal{M}	A set of rollouts
$\mathcal{N}_{\text{estimate}}$	The total number of 3D points in $\mathcal{P}_{\text{estimate}}$
$\mathcal{N}_{\text{interest}}$	The total number of 3D points representing the area above relevant lane segments
$\mathcal{N}_{\text{interest}}^{\text{visible}}$	The total number of 3D points in $\mathcal{P}_{\text{interest}}$ that are visible from p_{target}
$\mathcal{N}_{\text{lane}}$	The total number of 3D lane surface points
$\mathcal{N}_{\text{live}}$	The total number of 3D live scan points captured by a LiDAR unit
\mathcal{N}_{map}	The total number of points in a 3D point cloud map
$\mathcal{N}_{\text{object}}$	The total number of objects in a scene
$\mathcal{N}_{\text{particle}}$	The total number of particles
$\mathcal{N}_{\text{rollout}}$	The total number of rollouts (i.e., candidate trajectories)
$\mathcal{N}_{\text{sampling}}$	The total number of sampled locations along a rollout
\mathcal{O}	A 3D object
$\mathcal{O}'_{r,s}$	An occlusion ratio
$\mathcal{O}_{r,s}$	A squashed version of occlusion ratio $\mathcal{O}'_{r,s}$

$\mathcal{P}'_{\text{base}}$	A 3D scan estimated from a point cloud map, in sensor coordinates
$\mathcal{P}'_{\text{estimate}}$	A 3D scan estimated from a point cloud map and a live scan, in sensor coordinates
$\mathcal{P}_{\text{base}}$	A 3D scan estimated from a point cloud map, in point cloud map coordinates
$\mathcal{P}_{\text{estimate}}$	A 3D scan estimated from a point cloud map and a live scan, in point cloud map coordinates
$\mathcal{P}_{\text{interest}}$	A set of 3D points representing the area above relevant lane segments
$\mathcal{P}_{\text{lane}}$	A set of 3D points representing the surfaces of surrounding lanes
$\mathcal{P}_{\text{live}}$	A live 3D scan captured by a LiDAR unit, in point cloud map coordinates
\mathcal{P}_{map}	A 3D point cloud map
\mathcal{R}_r	A rollout
$\mathcal{S}_{\text{FOV}}^{\text{horizontal}}$	The horizontal field of view of a LiDAR unit
$\mathcal{S}_{\text{FOV}}^{\text{vertical}}$	The vertical field of view of a LiDAR unit
$\mathcal{S}_{\text{angle}}^{\text{lower}}$	An elevation angle from the specified viewing point to the lowest 3D point
$\mathcal{S}_{\text{angle}}^{\text{upper}}$	An elevation angle from the specified viewing point to the highest 3D point
$\mathcal{S}_{\text{range}}$	The sensing range of a LiDAR unit
T_{observe_i}	Length of time that a particle (i.e., an occluded vehicle) has observed the ego vehicle

T_{react}	Reaction time of an occluded vehicle after having detected the ego vehicle
\mathcal{T}_{lower}	Lower threshold of $\mathcal{O}'_{r,s}$
\mathcal{T}_{upper}	Upper threshold of $\mathcal{O}'_{r,s}$
\mathcal{V}	Visibility ratio
V_i	Visibility status of a particle (i.e., an occluded vehicle)
$\mathcal{V}_{r,s}$	Visibility ratio at sampled location $l_{r,s}$
V_{ego}	Estimated visibility of the ego vehicle
V_{other}	Estimated visibility of an occluded vehicle
W_{cross}	Width of an intersecting roadway
W_{ego}	Width of the ego vehicle's roadway
$\mathcal{W}_{rollout}$	Width of offset between parallel rollouts (i.e., candidate trajectories)
X_{ego}	Position of the ego vehicle's front end relative to an intersection's entrance
X_{sensor}	Position of the ego vehicle's sensor relative to the ego vehicle's front end
a	Acceleration of a vehicle
a_i, a_t^i	Acceleration of a particle (i.e., an occluded vehicle)
a_{ego}	Acceleration of the ego vehicle
a_{ego}^{cross}	Constant acceleration of the ego vehicle to cross an intersection

$a_{\text{ego}}^{\text{stop}}$	Constant deceleration of the ego vehicle to stop before an intersection
a_{enter}	Constant acceleration of the ego vehicle from a mandatory stop into an intersection to gain additional visibility
$a_{\text{freedrive}}$	The ego vehicle’s free road acceleration
$a_{\text{freedrive}_i}$	An occluded vehicle’s free road acceleration
a_{max}	Maximum acceleration in the Intelligent Driver Model (IDM)
$a_{\text{other}}^{\text{slow}}$	Constant deceleration of a particle (i.e., an occluded vehicle) during <i>Slowing down</i> behavior
$a_{\text{other}}^{\text{stop}}$	Constant deceleration needed for a particle (i.e., an occluded vehicle) to stop before an intersection
$a_{\text{other}}^{\text{yield}}$	Constant deceleration of a particle (i.e., an occluded vehicle) during <i>Yielding</i> behavior
α	Accuracy of the visible area classifier
b_{enter}	Constant deceleration of the ego vehicle to stop at s_{enter} from a mandatory stop
b_{stopping}	Deceleration required for the ego vehicle to stop at a mandatory stop
\bar{d}	Depth value of a depth image pixel
\bar{u}	A pixel position in a depth image along the horizontal axis
\bar{v}	A pixel position in a depth image along the vertical axis
χ_t	Estimated states of all occluded vehicles
$c'_{\text{visibility}_r}$	The discounted sum of squashed occlusion ratio $\mathcal{O}_{r,s}$ of all sampled locations along a rollout

c_{center}	Center cost of a trajectory
$c_{\text{collision}}^{\text{lat}}$	Lateral collision cost of a trajectory
$c_{\text{collision}}^{\text{lon}}$	Longitudinal collision cost of a trajectory
$c_{\text{transition}}$	Transition cost of a trajectory
$c_{\text{visibility}}^*$	Visibility cost of a trajectory
d	Travel distance of a vehicle
d_i	Displacement of an occluded vehicle from an intersection's center
d_{ego}	Minimum displacement necessary for the ego vehicle to completely cross an intersection
f	Acceleration function of a particle (i.e., an occluded vehicle)
γ	Discount factor
i	Index of a point, particle, or occluded vehicle
$l_{r,s}$	A sampled location on a rollout
l_{ego}	The ego vehicle's length
p	A 3D point
p'_{base_i}	An estimated 3D point in $\mathcal{P}'_{\text{base}}$
p'_{estimate_i}	An estimated 3D point in $\mathcal{P}'_{\text{estimate}}$
p_{current}	Current position of the LiDAR unit
p_{interest_i}	A representative point of the area above the surfaces of the relevant lanes

p_{lane_i}	A point on the surface of a lane
p_{live_i}	A 3D point in a live scan captured by a LiDAR unit in point cloud map coordinates
p_{map_i}	A 3D point in a point cloud map
p_{target}	A target position for visibility estimation
r	An index of a rollout
s	An index of a sampled location along a rollout
s_{accel}	Displacement of the ego vehicle during its acceleration from a mandatory stop into an intersection to gain additional visibility
s_{brake}	Displacement of the ego vehicle during its deceleration after entering an intersection to gain additional visibility
s_{center}	Distance between an intersection's center and a mandatory stop line
s_{ego}	Displacement of the ego vehicle from the mandatory stop line in its lane
s_{enter}	Total displacement of the ego vehicle from a mandatory stop into an intersection to gain additional visibility
t	Time
t_{ego}	Time it takes the ego vehicle to completely cross an intersection
$t_{\text{other}}, t_{\text{other}_i}$	Time a particle (i.e., an occluded vehicle) takes to arrive at an intersection

- $\mathbf{R}_{\text{current}}^{\text{map}}$ $\mathbf{R}_{3 \times 3}$ A 3D rotation matrix representing the current orientation of the LiDAR unit relative to the point cloud map coordinate system
- $\mathbf{T}_{\text{current}}^{\text{map}}$ $\mathbf{T}_{3 \times 1}$ A 3D translation matrix representing the current position of the LiDAR unit relative to the point cloud map coordinate system
- $\mathbf{R}_{\text{target}}^{\text{map}}$ $\mathbf{R}_{3 \times 3}$ A 3D rotation matrix representing the simulated orientation of the LiDAR unit relative to the point cloud map coordinate system
- $\mathbf{T}_{\text{target}}^{\text{map}}$ $\mathbf{T}_{3 \times 1}$ A 3D translation matrix representing the simulated position of the LiDAR unit relative to the point cloud map coordinate system
- v Velocity of a vehicle
- v_i, v_t^i Velocity of a particle (i.e., an occluded vehicle)
- v_{ego} Velocity of the ego vehicle
- $v_{\text{ego}}^{\text{allow}}$ Maximum velocity of the ego vehicle that will allow it to stop before an intersection
- v_{enter} Velocity of the ego vehicle while moving from a mandatory stop into an intersection to gain additional visibility
- v_{max} Maximum velocity in the Intelligent Driver Model (IDM)
- $v_{\text{other}}^{\text{cruise}}$ Constant velocity of a particle (i.e., an occluded vehicle) during *Cruising* behavior
- w_{center} A normalized weight for c_{center}
- $w_{\text{collision}}^{\text{lat}}$ A normalized weight for $c_{\text{collision}}^{\text{lat}}$

$w_{\text{collision}}^{\text{lon}}$	A normalized weight for $c_{\text{collision}}^{\text{lon}}$
w_{ego}	The ego vehicle's width
$w_{\text{other } i}$	An occluded vehicle's width
$w_{\text{transition}}$	A normalized weight for $c_{\text{transition}}$
$w_{\text{visibility}}$	A normalized weight for $w_{\text{visibility}}$
x	The x-coordinate of a 3D point
x_t^i	Estimated state of an occluded vehicle
$x'_{\text{live } i}$	The x-coordinate of a 3D point in a live scan
$x_{\text{base } i}$	The x-coordinate of a 3D point in $\mathcal{P}_{\text{base}}$
$x_{\text{estimate } i}$	The x-coordinate of a 3D point in $\mathcal{P}_{\text{estimate}}$
$x_{\text{interest } i}$	The x-coordinate of $p_{\text{interest } i}$
$x_{\text{lane } i}$	The x-coordinate of $p_{\text{lane } i}$
y	The y-coordinate of a 3D point
$y'_{\text{live } i}$	The y-coordinate of a 3D point in a live scan
$y_{\text{base } i}$	The y-coordinate of a 3D point in $\mathcal{P}_{\text{base}}$
$y_{\text{estimate } i}$	The y-coordinate of a 3D point in $\mathcal{P}_{\text{estimate}}$
$y_{\text{interest } i}$	The y-coordinate of $p_{\text{interest } i}$
$y_{\text{lane } i}$	The y-coordinate of $p_{\text{lane } i}$
z	The z-coordinate of a 3D point
$z'_{\text{live } i}$	The z-coordinate of a 3D point in a live scan

z_t An observation in a particle filter algorithm

$z_{\text{base}i}$ The z-coordinate of a 3D point in $\mathcal{P}_{\text{base}}$

$z_{\text{estimate}i}$ The z-coordinate of a 3D point in $\mathcal{P}_{\text{estimate}}$

$z_{\text{interest}i}$ The z-coordinate of $p_{\text{interest}i}$

$z_{\text{lane}i}$ The z-coordinate of $p_{\text{lane}i}$

List of Author's Publications

Journal Publications

1. Jacob Lambert, Leslie Liang, Luis Yoichi Morales, Naoki Akai, Alexander Carballo, Eijiro Takeuchi, **Patiphon Narksri**, Shunya Seiya, and Kazuya Takeda. 2018. "Tsukuba Challenge 2017 Dynamic Object Tracks Dataset for Pedestrian Behavior Analysis." *Journal of Robotics and Mechatronics* 30, 4 (2018), 598–612. DOI:<https://doi.org/10.20965/jrm.2018.p0598>
2. Alexander Carballo, Shunya Seiya, Jacob Lambert, Hatem Darweesh, **Patiphon Narksri**, Luis Yoichi Morales, Naoki Akai, Eijiro Takeuchi, and Kazuya Takeda. 2018. "End-to-End Autonomous Mobile Robot Navigation with Model-Based System Support." *Journal of Robotics and Mechatronics* 30, 4 (2018), 563–583. DOI:<https://doi.org/10.20965/jrm.2018.p0563>
3. Jacob Lambert, Alexander Carballo, Abraham Monrroy Cano, **Patiphon Narksri**, David Wong, Eijiro Takeuchi, and Kazuya Takeda. 2020. "Performance Analysis of 10 Models of 3D LiDARs for Automated Driving." *IEEE Access* 8, (2020), 131699–131722. DOI:<https://doi.org/10.1109/ACCESS.2020.3009680>
4. **Patiphon Narksri**, Eijiro Takeuchi, Yoshiki Ninomiya, and Kazuya Takeda. 2021. "Deadlock-Free Planner for Occluded Intersections Using Estimated Visibility of Hidden Vehicles." *Electronics* 10, 4 (2021). DOI:<https://doi.org/10.3390/electronics10040411>
5. **Patiphon Narksri**, Hatem Darweesh, Eijiro Takeuchi, Yoshiki Ninomiya, and

Kazuya Takeda. 2022. “Occlusion-aware Motion Planning with Visibility Maximization via Active Lateral Position Adjustment.” *IEEE Access* (2022), 1–1. DOI:<https://doi.org/10.1109/ACCESS.2022.3178729>

Conference Publications

1. **Patiphon Narksri**, Eijiro Takeuchi, Yoshiki Ninomiya, Yoichi Morales, Naoki Akai, and Nobuo Kawaguchi. 2018. “A Slope-robust Cascaded Ground Segmentation in 3D Point Cloud for Autonomous Vehicles.” In 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 497–504. DOI:<https://doi.org/10.1109/ITSC.2018.8569534>
2. **Patiphon Narksri**, Eijiro Takeuchi, Yoshiki Ninomiya, and Kazuya Takeda. 2019. “Crossing Blind Intersections from a Full Stop Using Estimated Visibility of Approaching Vehicles.” In 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 2427–2434. DOI:<https://doi.org/10.1109/ITSC.2019.8917323>
3. Alexander Carballo, Jacob Lambert, Abraham Monrroy, David Wong, **Patiphon Narksri**, Yuki Kitsukawa, Eijiro Takeuchi, Shinpei Kato, and Kazuya Takeda. 2020. “LIBRE: The Multiple 3D LiDAR Dataset.” In 2020 IEEE Intelligent Vehicles Symposium (IV), 1094–1101. DOI:<https://doi.org/10.1109/IV47402.2020.9304681>
4. Alexander Carballo, Abraham Monrroy, David Wong, **Patiphon Narksri**, Jacob Lambert, Yuki Kitsukawa, Eijiro Takeuchi, Shinpei Kato, and Kazuya Takeda. 2021. “Characterization of Multiple 3D LiDARs for Localization and Mapping Performance using the NDT Algorithm.” In 2021 IEEE Intelligent Vehicles Symposium Workshops (IV Workshops), 327–334. DOI:<https://doi.org/10.1109/IVWorkshops54471.2021.9669244>
5. **Patiphon Narksri**, Hatem Darweesh, Eijiro Takeuchi, Yoshiki Ninomiya, and

Kazuya Takeda. 2021. “Visibility Estimation in Complex, Real-World Driving Environments Using High Definition Maps.” In 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), 2847–2854. DOI:<https://doi.org/10.1109/ITSC48978.2021.9565003>

Chapter 1

Introduction

Autonomous vehicles, also known as driverless vehicles, have been gaining increasing attention from both the public and the research community due to their potential benefits. Since the majority of traffic accidents are caused by human error [1], traffic safety is expected to improve dramatically as conventional automobiles and trucks are replaced by autonomous vehicles. Apart from the safety aspect, these vehicles also offer other potential benefits, such as reduced fuel consumption, emissions and traffic congestion [2–4].

Initial development of autonomous vehicle dates as far back as the 1970s [5]. Two decades later, successful demonstrations of partially autonomous driving technology were presented to the public [6,7]. However, fully autonomous driving without human intervention has proved much more difficult. The first significant demonstration of successful, long-distance autonomous driving without human control was the DARPA Grand Challenge in 2005 [8,9], which was an off-road competition. In 2007, the DARPA Urban Challenge [10] was held in a simulated urban environment, and objectives and situations typically encountered in daily driving, such as avoiding collisions with static and dynamic objects while obeying traffic laws, negotiating intersections and merging into moving traffic, were introduced. Even though these tasks were not fully representative of real-world driving scenarios, they included the essential components of fully autonomous driving in an urban area. Only a few of the teams were able to navigate the course successfully [11–16], demonstrating that

autonomous driving in an urban area was possible, though not straightforward, with certain key issues remaining to be resolved.

The challenges faced when developing truly autonomous driving are not only due to the complexity of many driving scenarios, but also the result of uncertainty in the perception module of these vehicles. The perception module extracts and interprets information about the surrounding environment before passing that information on to other modules in the system. As a result, any perception uncertainty, if not handled properly, can lead to errors in the operation of the system [17, 18]. Uncertainty in the perception module comes from a variety of sources, such as intrinsic sensor error, adverse weather conditions or occlusions.

Uncertainty due to occlusion is frequently present and unavoidable in many driving scenarios. Various structures in the driving environment, such as buildings and walls, can prevent an autonomous vehicle from fully observing its surroundings. Hilly terrain, curves in the road, large trucks and foliage can also reduce the vehicle's visibility of the driving environment. While encountering occlusion is inevitable and routine, the majority of autonomous driving systems do not explicitly take this sensing limitation into account in their motion planning stage, since an autonomous vehicle's motion planning is typically based only on information about detected, observable objects in the local environment, resulting in potential occluded objects being entirely ignored [19–21]. Failure to consider occlusion during motion planning can lead to catastrophic accidents at critical locations, such as low visibility intersections in residential areas similar to the one shown in Fig. 1-1.

Due to walls and vegetation close to the intersection, the intersecting road on either side cannot be fully observed by the ego vehicle. If the ego vehicle ignores the occlusion completely, it will mistakenly assume that the intersection is safe to enter, as no obstacles or other traffic participants have been detected. However, there could actually be vehicles approaching from either side. In such a situation, an accident is likely to occur as it would be too late for the ego vehicle to initiate emergency braking after the approaching vehicle has become observable.

The occlusion problem could potentially be solved by utilizing additional data from



Figure 1-1: Example of a low-visibility intersection in a residential area.

external sources, for example via vehicle-to-everything (V2X) communication [22–26]. Such solutions are currently not feasible on a large scale because high-quality network devices with low latency are typically required to transmit the necessary information reliably in real-time. The hardware required also imposes additional infrastructure costs, and widespread deployment could take years or even decades. Therefore, relying solely on communication technologies to handle occlusion issues could delay the mass deployment of autonomous vehicles.

A feasible alternative solution to the occlusion problem is to incorporate visibility information during the motion planning stage, so that autonomous vehicles are aware of the possibility of hidden traffic participants, and can consider their sensing limitations before generating safe motion in areas with low visibility. Over the years, several different approaches for occlusion-aware motion planning have been proposed. These approaches can be broadly categorized into three categories; machine learning, Partially Observable Markov Decision Process (POMDP) and virtual obstacles.

Machine learning-based approaches utilize either human driving data or data from a simulation to train a function that maps the sensory inputs to control action outputs. While machine learning-based approaches have been shown to be effective for replicating the driving behavior of human drivers, especially those whose data they have been trained with, the relationships between sensor inputs and output actions are often not obvious, i.e., these approaches tend to lack interpretability.

POMDP-based approaches directly incorporate perception uncertainty and possible future developments in the surrounding traffic situation into the decision-making process via observation and transition models, making these approaches very versatile and applicable in many scenarios. Nevertheless, POMDP-based approaches are typically computationally expensive, and many cannot be used in real-time applications.

Approaches based on virtual obstacles explicitly assume that undetected vehicles are approaching from the occluded areas of the traffic environment, and the worst-case scenario is often assumed, i.e., these occluded objects are approaching at very high speed. By treating these virtual obstacles as if they are real, the motion of the ego vehicle can be efficiently planned, and risk due to occlusion, i.e., the risk of colliding with unobserved traffic participants, is acknowledged and clearly represented. Moreover, these virtual obstacle-based approaches tend to involve low computational complexity, therefore, they can be applied in various scenarios in real-time, making them suitable for real-world applications. Despite these advantages, virtual obstacle-based approaches are not without their drawbacks. These approaches have two main shortcomings: excessively conservative ego vehicle movement due to the worst-case assumption, which can lead to deadlock, and a lack of road position adjustment capability, which can result in less than optimal visibility.

In some locations, such as low-visibility intersections, being excessively cautious can lead to an indefinite deadlock of the ego vehicle. It is important to note that a deadlock is usually used to refer to an extreme traffic situation where vehicles on two or more roadways are all waiting for the others to pass a conflict area, such as an intersection or a narrow one-lane road, in order to free up the space for them, and thus none of them can successfully continue forward. However, the meaning of deadlock as used in this dissertation is slightly different as it does not refer to encounters between actual vehicles but between the ego vehicle and virtual vehicles assumed to be approaching at high speed from occluded areas. At low-visibility intersections, the ego vehicle can often see only a very limited portion of the intersecting roads before it enters the junction. Fig. 1-2a shows an example of such a situation. Since virtual vehicles are assumed to be approaching from the occluded areas at very high




(a) Visibility of intersecting roads before entering intersection.




(b) Visibility of intersecting roads after entering intersection.

Figure 1-2: Low-visibility intersection where deadlock caused by excessively cautious motion planning can occur.

speed, if there is an occluded area in close proximity to the ego vehicle, as is the case at this intersection, the ego vehicle is not allowed to proceed forward and has to come to a full stop, as it would collide with unobserved, virtual vehicles if it entered the intersection. Even if the ego vehicle advances all the way to the stop line, the view of the intersecting road is still limited, and high-speed vehicles could reach the intersection before the ego vehicle has transited. The only way the ego vehicle can increase its visibility is by moving forward and partially entering the intersection, as shown in Fig. 1-2b. Otherwise, the ego vehicle is stuck in a deadlock since it does not have sufficient visibility to move forward, yet cannot increase its visibility without doing so. This type of deadlock is one of the main shortcomings of occlusion-aware



(a) Visibility of intersecting road when ego vehicle is in the center of its lane.



(b) Visibility of intersecting road when ego vehicle is at the left edge of its lane.

Figure 1-3: Low-visibility T-junction. Vehicle positioned on the left edge of its lane has better visibility of the merging road.

planning approaches based on virtual vehicles.

The other main drawback of virtual obstacle-based motion planning is the lack of a road position adjustment feature. Most existing occlusion-aware planners only consider the longitudinal motion of the ego vehicle, while ignoring its lateral motion. Thus, it is often assumed that the ego vehicle will always travel along the center of its lane. Although staying in the center of the lane keeps the ego vehicle away from potential obstacles in adjacent lanes or on the roadside, it can result in inferior

visibility of the driving environment in some scenarios, as compared to being on the left or right edge of the lane. A sample scenario is shown in Fig. 1-3a, where the ego vehicle, which is shown in green, is in the center of its lane. From this position, the visibility of the ego vehicle, indicated by the green area, is poorer than when the ego vehicle keeps to the left, as shown in Fig. 1-3b. Note that the vehicle shown in Fig. 1-3, as well as other subsequent figures in this dissertation, is equipped with an omnidirectional sensor that covers 360 degrees around the vehicle. However, only the frontal visible region will be shown for the sake of a clearer illustration.

Our goal in this dissertation is to solve these two main problems with conventional occlusion-aware motion planning approaches based on virtual obstacles: deadlock due to the worst-case assumption and lack of road position planning. With the scope of this dissertation limited to only one type of occluded traffic participants, i.e., vehicles, our proposed solution for the deadlock problem is to utilize the visibility of the hidden vehicles to predict potential changes in their behavior based on this visibility, in order to better plan the speed of the ego vehicle. Specifically, instead of assuming hidden vehicles will always approach at a continuous, high speed, i.e., the worst-case assumption, we model interactions with approaching virtual vehicles under the assumption that they will likely change their behavior after observing the ego vehicle. This approach allows the ego vehicle to avoid deadlock by increasing its visibility of the occluded areas after sufficiently exposing itself to other traffic participants. Regarding road position planning, we propose predicting and quantifying visibility conditions of driving environments using high-definition (HD) maps consisting of 3D point clouds and road network maps. The quantified visibility condition will then be used for planning the road position of the ego vehicle in order to minimize occlusion.

The rest of this dissertation is organized as follows: In Chapter 2, previous work related to this research topic is introduced. Details of a speed planner designed to overcome deadlock at intersections with mandatory stops, by estimating the visibility of hidden vehicles, are then presented in Chapter 3. In Chapter 4, we propose a generic, deadlock-free speed planner that relies on a more accurate approaching vehicle behavior model, in order to more realistically predict the behavior of virtual

vehicles assumed to be approaching occluded intersections. The details of this approaching vehicle behavior model, which was derived through an analysis of real-world driving data, are also provided in Chapter 4. An approach for predicting and quantifying the visibility conditions of driving environments using HD maps is presented in Chapter 5. In Chapter 6, we develop and test a road position planning method that minimizes occlusions, which is based on the improved version of the visibility prediction and quantification approach described in Chapter 5. Finally, the conclusions of this dissertation are given in Chapter 7, along with a discussion of potential future work that could be conducted to expand on this research.

The main contributions of this dissertation are as follows:

- A deadlock-free speed planner for low-visibility intersections with mandatory stops.
- An approach for estimating visibility from the point of view of other vehicles, without using their actual sensing data.
- A visibility-dependent behavior model for vehicles approaching occluded intersections. based on an analysis of real driving data.
- A generic, deadlock-free speed planner for blind intersections that utilizes the visibility of both the ego vehicle and hidden vehicles.
- An approach for predicting the visibility conditions of driving environments using HD maps.
- The concept of a visibility ratio as a method which can be used to quantify visibility from a given location, as well as its calculation method.
- An extension of the proposed visibility prediction approach, to address occlusions caused by objects that are not included in HD maps.
- A road position planner that can adjust the ego vehicle's lateral position to minimize occlusions.

Chapter 2

Related Work

The planning module of an autonomous vehicle system relies upon the perception module to provide information about the driving environment [17, 18]. Most early motion planning research exclusively considered the risk of colliding with detected obstacles [19–21]. Although these approaches have proven to be sufficiently safe in many circumstances, they may not be safe in complex driving scenarios where portions of the environment are occluded. As the operational design domain (ODD) of autonomous vehicles continues to broaden, and encounters with occlusions become inevitable, the number of studies on how to achieve safe navigation in scenarios with occluded areas has increased.

One of the obvious solutions to the occlusion problem is to employ external aids via communication technologies. Therefore, previous investigations related to vehicle-to-everything (V2X) communication are presented first in Section 2.1, along with reasons why these approaches may not be feasible yet in the event of large-scale adoption of autonomous vehicles. In Section 2.2, research on occlusion-aware motion planning, which is a more viable alternative, and the details of commonly used techniques are discussed, along with their advantages and drawbacks. More specifically, previous research based on machine learning, Partially Observable Markov Decision Process (POMDP) and virtual obstacles are introduced in Sections 2.2.1, 2.2.2 and 2.2.3, respectively. Since visibility estimation is one of the most crucial components of an occlusion-aware motion planner, common visibility estimation methods are presented

in Section 2.3. In Section 2.4, both the strong points and shortcomings of different techniques used in existing occlusion-aware motion planners are summarized, and the research gaps we aim to fill with this dissertation are noted.

2.1 Vehicle-to-everything (V2X) communication

Several studies have proposed utilizing additional information obtained from external sources to ensure the safe navigation of autonomous vehicles in complex traffic scenarios. For example, in [26] additional sensing data from road junction infrastructure, transmitted via a low-latency mobile network, is used in parallel with the onboard perception module to plan longitudinal motion that is safe and comfortable for passengers. Zhao et al. [22] proposed an algorithm for safely scheduling connected automated vehicles (CAVs) traveling through an uncontrolled intersection that involves sharing information among the vehicles. In [25], an approach for increasing the situational awareness of an automated bus by fusing information from other connected vehicles, as well as from infrastructure with local sensors, was introduced.

While leveraging additional information from other sources can improve overall sensing coverage of targeted environments, a high-quality communications network with low latency is typically required in order to transmit the necessary information reliably in real-time. Additionally, in the case of vehicle-to-infrastructure (V2I) communication, where a vehicle receives data from fixed installations, multiple sensors per installation are often required to provide adequate local coverage, as demonstrated in [27] and [28]. These additional hardware requirements translate into extra time and investment that will be needed before autonomous vehicles can operate safely. As a result, the wide-scale adoption of approaches that depend on information from external sources will likely be slow.

2.2 Occlusion-aware motion planning

Instead of relying on additional data from external sources to reduce occlusion-related hazards, many researchers have focused on ensuring that autonomous vehicles can navigate safely even without complete observation of the driving environment by identifying occluded regions and considering them during motion planning. Several different approaches for occlusion-aware planning have been proposed in recent years. These approaches can be broadly classified into three categories: machine learning, Partially Observable Markov Decision Process (POMDP) and virtual obstacles.

2.2.1 Machine learning

Machine learning-based approaches have been widely used for planning autonomous vehicle motion when encountering occlusions. Occluded intersection navigation tasks are often modeled as reinforcement learning (RL) problems [29–31]. In [29], the navigation problem was modeled as a Markov Decision Process (MDP), where the optimal MDP policy was obtained indirectly by learning a state-action value function, i.e., a Q-function, using a Deep Q-Network (DQN). In [32], the problem was formulated as an MDP with hierarchical options (HOMDP), in which the planner first assesses the status of the driving environment before generating an output action in the form of continuous acceleration or deceleration, accordingly. While these studies have shown that reinforcement learning has the potential to be used for decision-making when encountering occluded areas, training an RL agent outside a simulation can be a challenging task. It is important to note that with recent advancements in RL algorithms, it is possible to train an RL agent directly on raw sensor input obtained from the real world. One of the most notable examples of such recent RL algorithms is the soft actor-critic algorithm introduced in [33], which allows an RL actor to learn to control a 3-finger dexterous robotic hand to manipulate an object from raw RGB images of the object. This type of approach will likely become more applicable to other real-world tasks, including navigation in occluded areas, in the future.

Instead of utilizing data from simulations, other researchers have focused on de-

veloping approaches that can learn from real-world driving data. As reported by Yoshihara et al. in [34], expert drivers, i.e., driving school instructors with years of experience, take proactive action to avoid possible collisions by slowing down as they approach intersections with poor visibility. In [35], proactive braking prior to entering occluded intersections without traffic signals was modeled using a potential risk function, with expert driving data being used to estimate the parameters of the function. Similarly, in [36] the driving data of experts was used to determine the parameters of a speed planner for blind intersections. In order to train a planner to learn proactive driving strategies for navigating blind intersections directly from driving data, Morales et al. [37] used a set of trajectory features, e.g., acceleration, speed and jerk, to describe driving behavior. The weights of such features were then obtained through Inverse Reinforcement Learning (IRL). In [38], reliance on hand-crafted features was avoided by using Deep Autoencoders.

While machine learning-based methods have proven to be very effective at replicating the driving behavior of experts, the inexplicability of their output is a major drawback. Another limitation of ML approaches is their inability to generalize, i.e., their weakness when applied in driving scenarios other than the ones which they were specifically trained for.

2.2.2 Partially Observable Markov Decision Process (POMDP)

Partially Observable Markov Decision Processes (POMDPs) are often used for planning the motion of autonomous vehicles in scenarios where there is incomplete knowledge of the driving environment, e.g., when encountering occlusions. By framing motion planning as a POMDP, sensing limitations and uncertainty are directly reflected by the observation model. The states of hidden traffic participants can be estimated by utilizing both current and past observations via belief state updates. Due to the generality of the POMDP framework, it has been used in several studies on occlusion-aware planning.

In [39], the optimal level of acceleration/deceleration in a merging scenario at an occluded intersection was formulated as a continuous POMDP problem. An initial

assumption regarding the pose and velocity of a possible occluded vehicle was predefined, then the POMDP model was solved offline using a Monte Carlo Value Iteration (MCVI) algorithm. Likewise, POMDP was used for occluded intersection crossing tasks in [40], where the worst-case assumption regarding occluded vehicles was used, i.e., occluded vehicles were assumed to be right at the boundaries of the visible area and to be traveling at the maximum allowable speed. The routes of the occluded vehicles were determined using belief state updates. In order to reduce the level of computational complexity, the authors proposed a custom POMDP solver designed specifically for their model. Moreover, they utilized an action prioritization method which makes the ego vehicle favor acceleration over maintaining a constant speed or deceleration, in order to avoid deadlock. In [41], a hidden vehicle model included an associated “existence probability”, which allows more flexibility than simply assuming the worst-case scenario. The Toolkit for approximating and Adapting POMDP solutions In Realtime (TAPIR) [42] was used to solve their POMDP model. The possible presence of hypothetical hidden vehicles was handled using a similar probabilistic approach in [43], where the probability of their existence was modeled to be dependent on the traffic density of each lane. In addition to the current visibility of the ego vehicle, simulated future visibility along the planning horizon was also incorporated into planning, allowing the ego vehicle to actively position itself to minimize occlusion. An optimized policy was obtained by solving the POMDP using TAPIR [42].

While a POMDP is very versatile and can be applied in various scenarios, solving its optimal policy, which maps the current belief state to an output action, is often computationally expensive. Some prior studies, such as [39, 44, 45], have relied on pre-computing the policy offline. As a result of continuous improvement of POMDP solvers and increasing hardware performance over the years, several recent studies have successfully planned the output motion of an autonomous vehicle during execution time, as in [40, 41, 43, 46]. However, these approaches utilize a coarsely discretized action space to reduce computation complexity.

2.2.3 Virtual obstacles

One of the most common approaches used to reduce the risk of collision with unobserved traffic participants, which has been adopted in various autonomous driving systems, is to explicitly assume that there are always virtual obstacles approaching at high speed from the boundaries of the visible region, on a course that will intercept the ego vehicle’s trajectory. By treating these approaching, virtual traffic participants as real moving obstacles, it is possible to regulate the speed of the ego vehicle so that collisions with these potential occluded objects can be avoided.

In [47], a safe ego vehicle speed when encountering a blind intersection was calculated based on the time-to-collision (TTC) of the ego vehicle and a virtual occluded vehicle at a collision point located within the intersection. A more conservative calculation of TTC was utilized in [48], which included reaction time as a delay before the autonomous vehicle begins braking. In [49, 50], instead of collision points, potential collision zones were used in order to represent a wider range of possible collision scenarios. While these studies propose effective approaches for modeling risk from unobserved, moving objects, they may suffer from uncertainty in perception, i.e., inaccurate prediction, as most rely on strict assumptions regarding the position and speed of virtual traffic participants. One way to handle the uncertainty associated with the states of virtual obstacles is to utilize a probabilistic approach. In [51], a particle filter-based approach for predicting the motion of occluded traffic participants was used. Possible unobserved objects were represented using particles whose states were continuously updated as new information from the vehicle’s sensors became available. A similar particle filter-based approach for probabilistic risk assessment in areas containing occlusions was also proposed in [52]. As these methods explicitly deal with uncertainty during hidden object prediction, they are likely to be more robust in scenarios where sensing information is noisy. Nevertheless, these methods are often criticized for their lack of a safety guarantee, as virtual traffic participants are often treated as points or particles that lack dimensionality.

In order to ensure the safety of vehicles navigating in areas with occlusions, Orze-

chowski et al. [53] proposed extending the set-based approach used for predicting the future occupancy of detected traffic participants, introduced in [54], to handle potential road users that are occluded. The future states of virtual objects traveling from the edges of the visible area were over-approximated and represented as reachable sets. These reachable sets cover all possible future states of the virtual objects, based on a set of assumptions regarding their initial speed, position, orientation and local traffic rules. Since the reachable sets include the worst-case scenario, the output motion of the ego vehicle was deemed safe if its future trajectory did not intersect with any of the reachable sets within the planning horizon. Reachable set prediction has subsequently been utilized in numerous studies [55–63], since using reachable sets to represent the possible future states of occluded traffic participants provides a guarantee of safety. However, since these reachable sets include the worst-case scenario, they often result in overly conservative ego vehicle control outputs.

One of the main advantages of approaches that are based on virtual obstacles is a clear relationship between potential risks and the output actions of the ego vehicle. Another strong point of these approaches is that they are typically computationally efficient and real-time capable, as the assumption regarding the states of virtual obstacles, namely that they are approaching from just outside the visible area at a high and constant speed, simplifies the planning problem. While this worst-case assumption simplifies motion planning when encountering occlusions, it can result in excessively cautious ego vehicle movement, and in locations where visibility is severely limited, it can lead to indefinite deadlock. Vulnerability to deadlock is a significant drawback of virtual obstacle-based motion planning approaches, as the ego vehicle sometimes does not have sufficient visibility to proceed from its current position, but unless it moves forward it cannot improve its visibility of the driving environment. Another shortcoming of these approaches is that they often lack road position adjustment capability. Many approaches simply assume that the ego vehicle will always travel along the center of its lane, and therefore lateral position planning is not supported. While driving in the center of a lane creates buffer zones between the ego vehicle and potential obstacles on the roadside or in adjacent lanes, in some

situations, e.g. at highly occluded intersections, it may result in poorer visibility of the driving environment compared to moving to the edge of the lane.

2.3 Visibility estimation

Identifying both the visible and occluded regions of a vehicle’s surroundings given its position, a process known as visibility estimation, is an essential component of safe navigation when encountering occluded areas, and various approaches have been proposed.

In studies that focus on motion planning in a specific scenario, such as intersection crossing [47, 64] or parked car passing [65], an environmental model is used to estimate either the visible or occluded areas of the current driving environment. These approaches have low computational costs, as the relationship between the location of the vehicle and its visibility can often be expressed in analytic form. However, these approaches tend to be limited to specific scenarios, and thus are not generally applicable to other driving situations.

Another common approach for estimating visible or occluded regions of an environment is 2D ray-tracing. First, structures or objects in the scene are represented using 2D geometric shapes, then rays originating from the sensor’s position and hitting the edges of these shapes, along with the sensor’s range, are used to define the boundaries of the visible area. This type of approach is often used in conjunction with road network maps in order to obtain road boundaries and the positions of stationary objects in the local environment [39–41, 43–46, 48–50, 52, 53, 55, 63, 66, 67]. While 2D ray-tracing is more versatile than using a model of the environment for visibility estimation, its accuracy may suffer in areas with complex occlusions, e.g., areas where a hill causes the occlusion.

Zhang et al. [61] used Octomap [68] to estimate both current and future visibility of the ego vehicle in 3D. The importance of performing visibility estimation in 3D was demonstrated in one of their experiments, in which an autonomous vehicle driving uphill is faced with an occlusion caused by the hill’s summit. In [69], a point cloud

map obtained from a 3D light detection and ranging (LiDAR) unit was used along with a road network map to estimate the ego vehicle’s visibility of the surrounding lanes. This approach converts the captured 3D scan and known positions of the surrounding lanes into depth images. These two depth images are then compared in order to determine the visible and occluded regions of the lanes. This approach was shown to work well in real-world environments with hilly road segments. The main drawback of this approach, however, is the requirement that a 3D scan be captured at the location where visibility is to be estimated. This limits the usability of this approach to locations where 3D scans are available. Moreover, when using this method, the 3D scans can only be used to estimate visibility under one specific sensor configuration, which is the configuration used to acquire the scans.

2.4 Research gaps

As presented in Sections 2.2.1, 2.2.2, and 2.2.3, each commonly used technique for planning ego vehicle motion in areas with occlusions has its own strengths and weaknesses. These points are summarized in Table 2.1 to facilitate comparison.

The ability of virtual obstacle-based methods to plan output motion in real-time in various driving scenarios makes these approaches the most suitable for real-world applications. Nevertheless, most virtual obstacle-based approaches have two key drawbacks which need to be addressed: excessively conservative ego vehicle movement due to the worst-case assumption which can result in deadlock, and inability to adjust the position of the ego vehicle within its lane to improve visibility.

In order to alleviate the problem of excessively cautious motion, in [59] and [63] the reachable sets were repeatedly updated as new observations became available. As a result, more accurate estimates of the states of hidden traffic participants could be obtained and utilized to improve driving efficiency. In [60] and [61], overly cautious motion was mitigated by considering estimated future visibility, and by encouraging the ego vehicle to actively improve its visibility of the driving environment. In [70], a heuristic strategy was used to escape deadlock in an intersection merging scenario.

Table 2.1: Summary of advantages and disadvantages of popular occlusion-aware motion planning techniques proposed in previous studies.

	Machine learning	POMDP	Virtual obstacles
Advantages	Human driving styles are directly modeled	Perception uncertainty is considered	Risk is clearly represented
		Future development of driving situation taken into account	Applicable in various scenarios
			Low computational complexity
			Real-time capable
Disadvantages	Relationship between output action and occlusion is unclear	Computationally intensive	Excessively conservative movement and deadlock risk due to the worst-case assumption
		Usually not real-time capable	Cannot adjust position of the ego vehicle within lane to improve visibility
		Scenario specific	

The ego vehicle was allowed to slowly advance into the intersection for a certain distance after its velocity reached zero at the edge of the intersection, allowing it to obtain additional visibility and escape deadlock. Although these previous studies offer countermeasures for avoiding deadlock in some driving scenarios, the underlying cause of the deadlock problem, i.e., an unrealistically strict assumption regarding the behavior of occluded vehicles, remains unaddressed.

The lack of road position planning in virtual obstacle-based motion planning approaches has been addressed in a few recent studies. In [47], road positions that were not safe from potential collisions with occluded obstacles were discarded by the planner. Although unsafe road positions could be eliminated using this approach, visibility gain by adjusting the ego vehicle’s lateral position was not directly considered. Andersen et al. [65] proposed a framework for passing parked vehicles that takes the

size of the blind spot into account during trajectory optimization via Model Predictive Control (MPC). While this approach has been shown to generate efficient driving behavior when encountering occlusions, it was designed specifically for passing parked vehicles. In [60] and [61], estimated future visibility was used to actively adjust the ego vehicle’s road position in order to increase visibility of the driving environment and reduce the risk of collision with potential occluded vehicles. Although the approach proposed in [60] was shown to be applicable in various situations, it lacks highly accurate 3D visibility estimation. Zhang et al. [61] addressed this problem by utilizing Octomap [68] to estimate both current and future visibility, however, the future visibility was estimated using a pre-constructed Octomap and predicted occupancy of other obstacles. As a result, estimations of future occlusions resulting from objects that were not present during the offline generation of the Octomap may not be accurate. Additionally, ground truth information regarding the state of objects was obtained directly from the simulator. Therefore, performance of the planner with a realistic level of uncertainty in the perception module was not tested in [61].

The research presented in this dissertation aims to fill gaps left unaddressed in previous studies regarding the two main shortcomings of virtual obstacle-based techniques. It will do so by proposing:

- A deadlock-free speed planner that directly addresses the underlying cause of deadlock when using virtual obstacle-based motion planning techniques, i.e., their strict assumptions regarding the behavior of occluded vehicles, by taking into account potential changes in the behavior of hidden vehicles after they have detected the ego vehicle.
- An occlusion-aware motion planning approach that is capable of minimizing occlusion by strategically adjusting the ego vehicle’s position within its lane, based on highly accurate, 3D estimates of current and future visibility conditions, which is applicable in various traffic scenarios and able to operate with a realistic level of perception uncertainty.

Chapter 3

Deadlock-free Speed Planning Using Estimated Visibility of Approaching Vehicles

3.1 Introduction

Most autonomous vehicles plan their trajectories based on the environmental information they observe, such as the positions of detected objects, but in order to achieve comfortable and safe navigation unobservable risks, such as the risk caused by hidden vehicles, need to be appropriately incorporated into motion planning.

In some critical locations, such as intersections with low visibility, failing to take occlusions into account during planning could lead to a serious accident. Intersections like the ones shown in Fig. 3-1 are dangerous because oncoming traffic traveling along the intersecting roads cannot be completely observed due to the occlusions caused by the walls and buildings on either side of the ego vehicle. In order to navigate such blind intersections safely, autonomous vehicles must be capable of predicting possible oncoming vehicles approaching from the occluded areas and plan their motion accordingly.

When encountering uncontrolled, blind intersections, a simple driving strategy



(a) Intersection A.



(b) Intersection B.

Figure 3-1: Low-visibility intersections in residential areas with mandatory stops.

that autonomous vehicles can employ to avoid collisions with occluded vehicles is to reduce their speed until they gain sufficient visibility of the intersection, and then accelerate across if it is safe to do so. However, when encountering blind intersections with stop signs or red traffic lights, they must first come to a full stop and then begin crossing from the stop line. Since visibility from the stop line at some intersections is insufficient to ensure safe intersection crossing, it can lead to a deadlock situation because the ego vehicle cannot move forward and cannot complete the cross. One possible motion the ego vehicle could perform is to move forward from the stop line and enter the intersection slightly in order to increase its visibility. Nevertheless, blindly entering the intersection without considering if possible oncoming vehicles can observe and have enough time to respond to the ego vehicle's movement into the intersection, can be dangerous.

In this chapter, we describe a method for crossing low-visibility intersections with mandatory stops when visibility from the stopping point is insufficient to continue across. We utilize the estimated visibility of possible approaching vehicles to decide whether it is safe to slowly proceed forward after stopping. To validate the proposed method, driving data was collected from an expert driver when encountering occluded intersections in residential areas with mandatory stops. Speed profiles generated when using the proposed method were similar to those observed when these types of intersections were encountered by the expert driver.

The contributions of the proposed method are:

- A speed planner which allows safe crossing of low-visibility intersections with mandatory stops, even when there is insufficient visibility to advance from the stopping point.
- Estimation of visibility of the ego vehicle from the perspective of approaching vehicles, without obtaining their actual sensing data.

The remainder of this chapter is organized as follows: In Section 3.2, the proposed method is described in detail. Comparisons between the speed profiles of an expert driver and those generated by the proposed method are then presented in Section 3.3. In Section 3.4, our experimental results are discussed and the conclusions of this investigation are presented.

3.2 Blind intersection crossing

The proposed method can be understood as consisting of two parts, movement planning before, and after, making a mandatory stop at a stop line. As it approaches the stop line, the ego vehicle reduces its speed to prepare to come to a full stop. This stage is shown in Fig. 3-2a. After the full stop (Fig. 3-2b), if the ego vehicle determines that it can reach the other side of the intersection safely before any possible oncoming vehicles arrive, it begins crossing. However, if the ego vehicle's visibility is insufficient to safely cross the intersection from its stopping point, it slowly proceeds

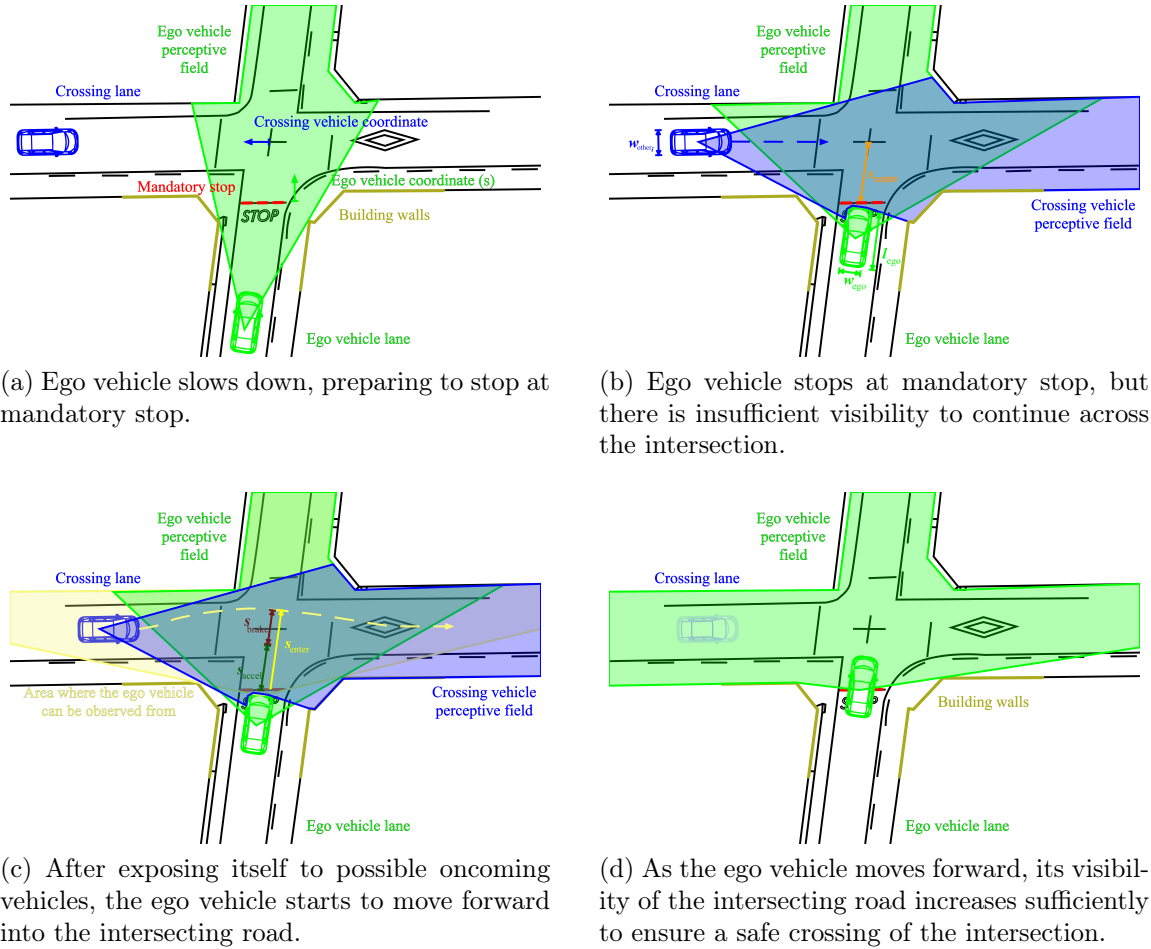


Figure 3-2: Stages of crossing a severely occluded intersection using the proposed algorithm.

into the intersection to increase its visibility. After the ego vehicle has exposed itself to possible approaching vehicles and given them enough time to respond (Fig. 3-2c), the ego vehicle moves forward. As it does so, its visibility will finally be sufficient to cross the intersection, as depicted in Fig. 3-2d. The details of the proposed method are given in the following subsections.

3.2.1 Safety constraint for intersection crossing

The structure of a typical intersection in a residential area is depicted in Fig. 3-2a. As can be seen in the figure, there is a mandatory stop ahead in the lane on which the ego vehicle is traveling, indicating that users of the intersecting road may

have higher priority. Given that the ego vehicle is not on a priority road, it is clear that it can safely cross the intersection if, and only if, it will completely clear the intersection before any vehicles approaching along the intersecting roadway arrive at the intersection.

In order to facilitate the formulation of a safety constraint for the crossing of an intersection, vehicle positions are represented with respect to a system of coordinates based on their displacement from a key point along their current lane, which is shown in Fig. 3-2a. The ego vehicle's position is represented using its displacement (s_{ego}) from the mandatory stop line in its lane, such that (s_{ego}) remains negative until the ego vehicle passes the mandatory stop line. The positions of vehicles moving along the intersecting lanes are defined as their displacement from the center of the intersection (d_i).

Having defined the coordinates and assuming that the acceleration of the ego vehicle is constant, the displacement the ego vehicle achieves in time interval t_{ego} can be calculated as follows:

$$d_{\text{ego}} = v_{\text{ego}}t_{\text{ego}} + \frac{1}{2}a_{\text{ego}}t_{\text{ego}}^2 \quad (3.1)$$

where d_{ego} , v_{ego} , and a_{ego} represent the displacement, velocity, and acceleration of the ego vehicle respectively. Therefore, the time it will take the ego vehicle to completely cross the intersection (t_{ego}) can be calculated using the following equation:

$$t_{\text{ego}} = \begin{cases} \frac{\sqrt{v_{\text{ego}}^2 + 2a_{\text{ego}}d_{\text{ego}}} - v_{\text{ego}}}{a_{\text{ego}}} & , \text{ if } a_{\text{ego}} \neq 0 \\ \frac{d_{\text{ego}}}{v_{\text{ego}}} & , \text{ if } a_{\text{ego}} = 0 \end{cases} \quad (3.2)$$

where d_{ego} is the minimum displacement the ego vehicle needs to achieve in order to completely avoid a collision with an approaching vehicle traveling on the center of its lane, which is given by:

$$d_{\text{ego}} = l_{\text{ego}} + \frac{w_{\text{other}i}}{2} + s_{\text{center}} - s_{\text{ego}} \quad (3.3)$$

where l_{ego} is the length of the ego vehicle, w_{other_i} is the width of a vehicle traveling in the intersecting lane and s_{center} is the distance between the center of the intersection and the mandatory stop line, as shown in Fig. 3-2b. In a similar manner, the time any vehicle traveling along the intersecting road takes to arrive at the intersection (t_{other_i}) is given by the following equation:

$$t_{\text{other}_i} = \begin{cases} \frac{\sqrt{v_i^2 + 2a_i(|d_i| - \frac{w_{\text{ego}}}{2})} - v_i}{a_i}, & \text{if } a_i \neq 0 \\ \frac{|d_i| - \frac{w_{\text{ego}}}{2}}{v_i}, & \text{if } a_i = 0 \end{cases} \quad (3.4)$$

where w_{ego} is the width of the ego vehicle and d_i , v_i , and a_i represent the displacement, velocity and acceleration of the vehicle traveling along the intersecting road, respectively. Thus, the ego vehicle can safely cross the intersection if the following inequality is satisfied with regard to all of the vehicles on the intersecting road:

$$t_{\text{ego}} < t_{\text{other}_i} \quad (3.5)$$

3.2.2 Visibility estimation

In order to produce a safe and reliable control output for the crossing of an intersection with limited visibility, it is crucial to be able to accurately estimate the current visibility of the ego vehicle.

The proposed planning method adopts a visibility estimation approach proposed in [69]. As this estimation method utilizes real sensing data, it is capable of correctly identifying which parts of the driving environment are occluded from the point of view of the ego vehicle, even in a complex scenario. The visibility estimation method consists of the following steps:

1. Create a depth image from a 3D LiDAR scan.
2. Select a pixel from a region of interest in the depth image.




Figure 3-3: An example of a depth image created from a 3D LiDAR scan for the ego vehicle’s visibility estimation. For illustration purposes, the sample image only shows the region that corresponds to the frontal part of the 360-degree 3D scan. The brightness of each pixel indicates the depth, i.e., bright pixels are associated with object points far from the 3D LiDAR.

3. Compare the actual, measured depth of the pixel’s location with the depth of the region of interest in the depth image.
4. Mark the region of interest as occluded if its depth is greater than the depth of the actual measurement.


An example of a depth image created from a 3D LiDAR scan for the ego vehicle’s visibility estimation is shown in Fig. 3-3.

Since our planning method requires information about whether vehicles traveling along the intersecting road are able to see the ego vehicle, visibility estimates from the points of view of these vehicles are also needed. Although the estimated visibility from the vantage points of the other vehicles cannot be achieved directly using the same estimation method described above, it is possible to estimate which regions the ego vehicle can be observed from without the actual sensing data of those vehicles.

As shown by the shaded blue area in Fig. 3-2c, if a vehicle on the intersecting road has a clear line of sight from its sensor, or from the driver’s seat in the case of a human-driven car, to the front bumper of the ego vehicle, then the ego vehicle can be observed. Therefore, the ego vehicle can also be seen from any position where a clear line of sight originating from the ego vehicle’s front bumper can reach. Based on this assumption, the areas that the ego vehicle is visible from can be approximated by shifting the origin of a 3D LiDAR scan to the front bumper of the ego vehicle and performing visibility estimation. Therefore, the yellow shaded area in Fig. 3-2c shows the areas from which the ego vehicle can be observed. Note that since the LiDAR sensor is installed slightly behind the leading edge of the ego vehicle’s front bumper,



(a) Both the area visible to the ego vehicle (represented by the green rectangles) and the area where the ego vehicle can be observed from (represented by the yellow rectangles) are small as the ego vehicle approaches the stop line.



(b) Area visible to the ego vehicle (represented by the green rectangles) remains small, while area where the ego vehicle can be observed from (represented by the yellow rectangles) grows significantly larger as the ego vehicle stopping the stop line.

Figure 3-4: Examples of visibility estimation using real sensing data. Green rectangles represent area visible to the ego vehicle. Yellow rectangles represent locations where the ego vehicle can be observed from. Red rectangles represent occluded regions.

oncoming vehicles can usually observe the ego vehicle before the ego vehicle can detect them, as shown in both Figs. 3-2b and 3-2c. Examples of visibility estimation near an intersection using real sensing data are shown in Fig. 3-4.

3.2.3 Particle filter

While the simplest approach to modeling the risk of collision with occluded vehicles is to assume that there are always virtual vehicles traveling at a certain speed toward the intersection from the boundaries of the perceptive field, in this study prediction of the states of oncoming vehicles approaching along the intersecting road is performed using a particle filter, in order to handle uncertainty that might arise from the perception module.

Each particle in our model represents a hypothetical oncoming vehicle. The state of a particle is described by a set of parameters, $\{d_i, v_i, a_i, V_i, T_{\text{observe}_i}\}$, where d_i , v_i , and a_i are the distance from the intersection's center to the particle, the particle's velocity and the particle's acceleration, respectively. The visibility status of each particle, denoted by V_i , includes three possible states: visible by the ego vehicle, occluded from the ego vehicle, or occluded from the ego vehicle but able to see the ego vehicle. Finally, T_{observe_i} is the length of time the ego vehicle has been visible to a particle.

The prediction process starts by generating a number of particles traveling at random speeds toward the intersection from locations. Their states are then updated at every sampling interval using the following motion model:

$$\begin{aligned} a_i &= a_{\text{freedrive}_i} \\ v_i &= v_i t + a_i t \\ d_i &= d_i + v_i t + \frac{1}{2} a_i t^2 \end{aligned} \tag{3.6}$$

where $a_{\text{freedrive}_i}$ is the free road acceleration component in the Intelligent Driver Model (IDM) [71], given by the following equation:

$$a_{\text{freedrive}_i} = a_{\text{max}} \left[1 - \left(\frac{v_i}{v_{\text{max}}} \right)^2 \right]. \tag{3.7}$$

Lastly, every particle has a weight associated with it which decreases as the particle becomes visible to the ego vehicle. A visible particle is depicted in Fig. 3-2d as a

faded vehicle. The remaining particles with non-zero weights are treated as actual vehicles during the ego vehicle’s motion planning.

3.2.4 Decision making

Controlling ego vehicle motion is equivalent to deciding proper ego vehicle acceleration at each time step given its current state, $\{s_{\text{ego}}, v_{\text{ego}}, \text{ and } a_{\text{ego}}\}$ and current sensory input. Crossing a blind intersection with a mandatory stop mainly involves two types of motion; braking, to come to a full stop at the stop line, followed by acceleration, to cross the intersection and continue forward. The required deceleration for the ego vehicle to brake to a stop (b_{stopping}) can be calculated as follows:

$$b_{\text{stopping}} = \frac{v_{\text{ego}}^2}{2s_{\text{ego}}}. \quad (3.8)$$

Note that the b_{stopping} is always negative, indicating that it is deceleration, and that s_{ego} remains negative until the ego vehicle reaches the stop line. For acceleration, free road acceleration in Eq. 3.7 is used. Unlike braking, the ego vehicle can safely accelerate at $a_{\text{freedrive}}$ if, and only if, inequality 3.5 holds for all possible oncoming vehicles on the intersecting road. However, at intersections where visibility is severely limited, the safety constraint may never be satisfied while the ego vehicle is at a complete stop at the mandatory stop line. Consequently, this could result in a deadlock problem, where the ego cannot move forward from the stop line.

In order to prevent a deadlock situation, it is assumed that if a vehicle approaching along the intersecting road has observed the ego vehicle waiting at the stop line, it will decelerate and slightly change its lateral position, as shown in Fig. 3-2c, after some reaction time, T_{react} . Given such an assumption, the ego vehicle is consequently allowed to perform a displacement of s_{enter} into the intersecting road if it has been observed by possible oncoming vehicles for a certain period of time (T_{react}) and still cannot cross the entire intersection safely from the stop line.

Moving from the mandatory stop to a full stop again after performing displacement s_{enter} can be achieved by constantly accelerating at a_{enter} until the ego vehicle has

Algorithm 1 Algorithm for crossing intersections with mandatory stops.

Input: s_{ego} , v_{ego} , a_{ego} **Output:** a_{ego}

```
1: if  $s_{ego} \geq 0$  then
2:    $a_{ego} \leftarrow a_{freedrive}$ 
3:   for each particle  $i$  do
4:     if  $(t_{ego} < t_{other\ i})$  then
5:        $a_{ego} = a_{freedrive}$ 
6:     else
7:       if  $(s_{ego} < s_{accel})$  then
8:         if  $(T_{observe\ i} \geq T_{react})$  then
9:            $a_{ego} = a_{enter}$ 
10:        else
11:           $a_{ego} = 0$ 
12:        end if
13:       else
14:          $a_{ego} = b_{enter}$ 
15:       end if
16:       break
17:     end if
18:   end for
19: else
20:    $a_{ego} = b_{stopping}$ 
21: end if
22: return  $a_{ego}$ 
```

achieved displacement s_{accel} from the stop line, then continuously decelerating at b_{enter} to come to a complete stop after achieving an additional displacement of s_{brake} .

Values of s_{accel} and s_{brake} can be calculated using the following equations:

$$v_{enter}^2 = 2a_{enter}s_{accel} \quad (3.9)$$

$$0 = v_{enter}^2 + 2b_{enter}s_{brake}$$

$$s_{enter} = s_{accel} + s_{brake} \quad (3.10)$$

$$s_{accel} = s_{enter} \left(\frac{|b_{enter}|}{a_{enter} + b_{enter}} \right) \quad (3.11)$$
$$s_{brake} = s_{enter} \left(\frac{a_{enter}}{a_{enter} + |b_{enter}|} \right)$$

By moving forward, i.e., closer to the intersection, the ego vehicle can increase its

visibility sufficiently so that it can finally cross. The complete decision-making process is summarized in Algorithm 1.

3.3 Experiments

In order to evaluate the proposed method of crossing highly occluded intersections with mandatory stops, the proposed planner’s output driving behaviors were compared to those of an expert human driver, an instructor at a driving school in Japan who has over 15 years of safe driving experience. The driving behavior of the expert driver was used as a reference for safe human driving behavior. An experimental vehicle equipped with various sensors for data collection was driven through several intersections in a residential area by the expert driver, and speed profiles of the vehicle were recorded. To fairly evaluate the proposed automated intersection crossing method, a simple yet realistic closed-loop simulation was created using the same driving data recorded in the real-world residential area. The proposed method was then used in the closed-loop simulation to generate speed profiles when crossing the selected intersections.

Details of the experimental conditions, along with a comparison of the speed profiles of the expert driver and those generated by the proposed, automated crossing method at low-visibility intersections with a mandatory stop, are presented in the following subsections.

3.3.1 Experimental vehicle

The experimental vehicle, a Toyota Prius (shown in Fig. 3-5), was equipped with a 3D LiDAR (Velodyne HDL64E-S2) on its roof along with a Global Navigation Satellite System (GNSS) receiver for localization. A wheel encoder was also installed for speed measurement. Positioning or localization of the experimental vehicle was performed in real-time at 10 Hz, primarily using the Normal Distributions Transform (NDT) scan-matching method proposed in [72]. After initialization using a position from the GNSS receiver, NDT scan-matching was used to match the current 3D scan of the



Figure 3-5: Experimental vehicle equipped with 3D LiDAR , GNSS receiver and wheel encoder.

driving environment with a high-definition point cloud map of the same environment, as shown in Fig. 3-6. The vehicle's localization system then returns the current position of the vehicle in the point cloud map.

3.3.2 Experimental environment

The experiment was conducted on a public road in a residential area in Nagoya, Japan. As this chapter focuses on crossing a low-visibility intersection from a mandatory stop, two intersections that meet these criteria were selected. Street views of the selected intersections are shown in Fig. 3-1. Intersections A (Fig. 3-1a) and B (Fig. 3-1b) both have a clear stop line, indicating that vehicles need to come to a full stop before entering the intersection. There are also buildings or walls on both sides of the road that severely limit the visibility of the intersecting road, a type of intersection that is very common in residential areas in Japan.

Two types of maps that complementarily represent the experimental environment were created. The first type of map is a high-definition point cloud map, an example of which is shown in Fig. 3-6. These maps contain 3D renderings of the environment in the form of multiple scan points captured by a Mobile Mapping System (MMS) using a LiDAR . The point cloud map is mainly used for localization, as described




Figure 3-6: High-definition point cloud map for localization using NDT scan-matching.

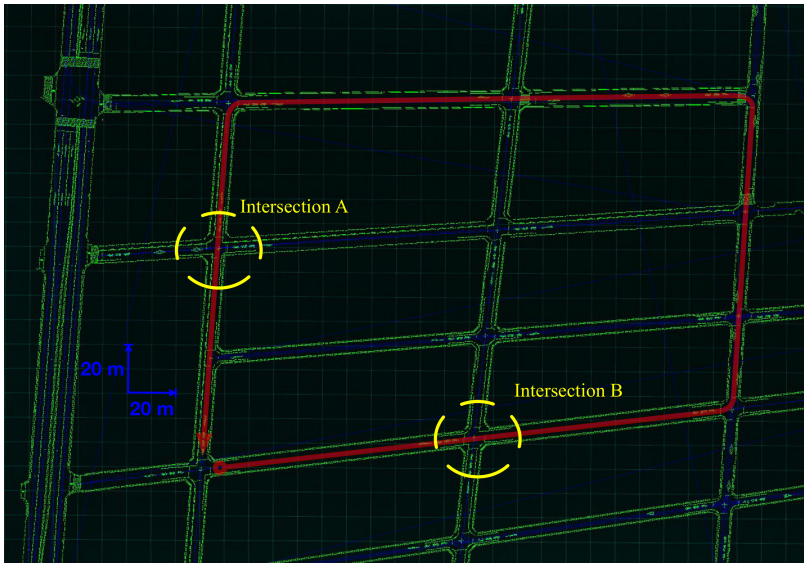


Figure 3-7: Road network map containing information about road networks in our experimental environment. The red line with an arrowhead shows one complete driving loop in the experiment.

in Section 3.3.1. The other type of map is a road network map an example of which is shown in Fig. 3-7. Unlike the point cloud map, the road network map does not contain an exact representation of the driving environment but is a higher level of representation that shows the positions of lane boundaries, traffic signs, pedestrian crossings and intersections.

In order to record the driving data efficiently, the driving route was designed to

be a loop. The driving direction is indicated by the red line with an arrowhead in Fig. 3-7. During each completed loop, the expert driver crossed Intersections A and B once. A total of eight laps were completed, resulting in eight crossings of each intersection.

3.3.3 Closed-loop simulation based on recorded data

As mentioned in Section 3.2.2, the visibility estimation approach used in the proposed method utilizes sensing data captured from the current position of the ego vehicle. As sensory input is used for visibility estimation, it consequently affects the planner's output as well. Therefore, in order to fairly compare the speed profiles produced when using the proposed method with those of the expert driver, input sensing data for the proposed planner must be similar to that encountered when the expert was driving. A simple simulator was therefore created to satisfy this requirement under the following assumptions:

- The expert driver was driving in a straight line.
- The sampling rate of the 3D LiDAR is exactly 10 Hz.
- There were no dynamic objects in the scenes.

The working principle of the simulator is straightforward; when given a position, the simulator outputs the corresponding 3D scan from the recorded data. The proposed planner takes this 3D scan and the position of the scan as inputs from the simulator. The planner then estimates visibility, computes output acceleration or deceleration, and sends the longitudinal displacement that will occur by the next data sampling, which takes place every 0.1 seconds, given the current velocity and the output acceleration, back to the simulator. The simulator then computes the new position of the ego vehicle using the received displacement. After the simulator compares the new position with all of the available positions in the recorded data, it returns the 3D scan associated with the closest position to the planner. The simulation loop is therefore complete. It is important to note that even though the ego

Table 3.1: Parameters used for speed profile generation.

Parameter	Value	Parameter	Value
l_{ego}	4.5 m	$w_{\text{other}_i}, w_{\text{ego}}$	1.0 m
a_{max}	2.5 m/s ²	v_{max}	10.0 m/s
T_{react}	1.2 s	s_{enter}	6.0 m
a_{enter}	1.2 m/s ²	b_{enter}	-0.2 m/s ²


vehicle’s movement is entirely simulated, the sensory input provided by the simulator is realistic as it uses recorded, real-world data.

Using the closed-loop simulation and the parameters shown in Table 3.1, speed profiles at Intersections A and B were generated. Parameters l_{ego} , w_{ego} , and w_{other_i} were defined using the dimensions of the Toyota Prius. Maximum acceleration a_{max} was determined based on the data from [73], and maximum velocity v_{max} was set to approximately the local speed limit. Reaction time T_{react} was selected based information presented in [74]. The allowable displacement from the stop line, s_{enter} , was selected to ensure that the ego vehicle could obtain sufficient visibility of the intersecting road while leaving enough space for approaching vehicles to cross the intersection without a collision. Lastly, a_{enter} and b_{enter} were determined such that together they represent a careful advance from the stop line into the intersecting road.

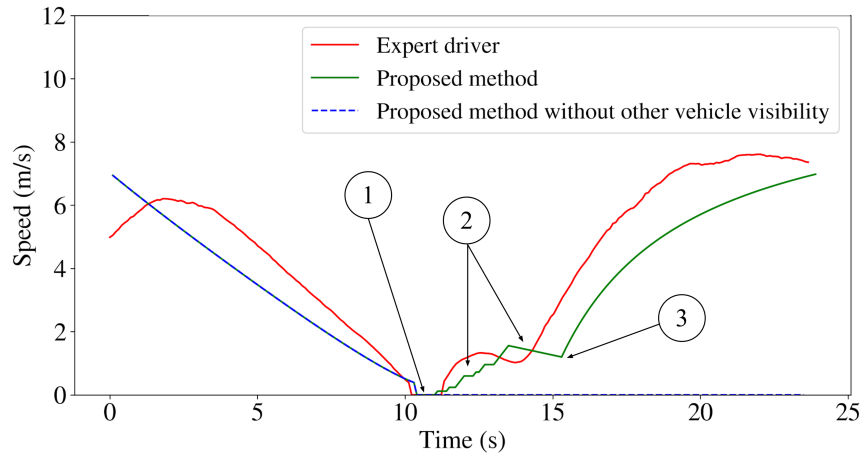
The simulation started 50 m before the intersection, with the ego vehicle traveling at a velocity of 6.94 m/s (25 km/h), and ended 50 m after exiting the intersection. In order to validate the effect of utilizing approximations of the visibility of the approaching vehicles for motion planning, speed profiles were created with and without this information.

3.3.4 Speed profile comparison

As can be seen by comparing Figs. 3-8 and 3-9, the speed profiles generated when using the proposed method are similar to those of the expert driver in both the time and displacement dimensions, especially after stopping. In addition, when the proposed method was used without estimation of the visibility of approaching vehicles,



(a) Variation in speed in relation to displacement from center of intersection.

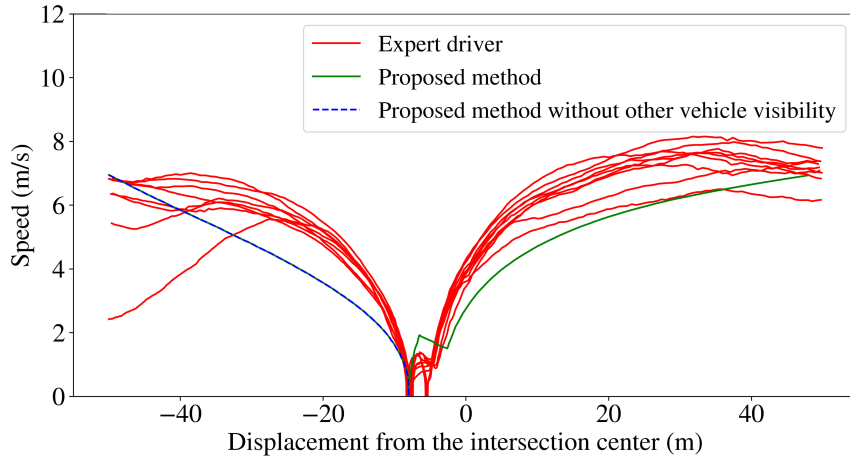


(b) Variation in speed over time.


Figure 3-8: Comparison of speed profiles of expert driver and proposed method at Intersection A.

indicated by the blue dotted lines in Figs. 3-8 and 3-9, the ego vehicle could not move forward to cross the intersection after stopping, as visibility from the ego vehicle's point of view was very limited.

The resulting speed profiles for both intersections consist of a similar sequence of behaviors, labeled by numbers in Figs. 3-8b and 3-9b. As indicated by ①, the ego vehicle remained at the stop line for a short period of time at both intersections, in order to expose itself to possible oncoming traffic, when using the proposed method and when being driven by the expert driver. It then began to move forward slowly



(a) Variation in speed in relation to displacement from center of intersection.



(b) Variation in speed over time.

Figure 3-9: Comparison of speed profiles of expert driver and proposed method at Intersection B.

into the intersecting road ②. As the ego vehicle moved forward, its visibility of the intersection became sufficient to ensure safe crossing. As a result, it began accelerating again at ③. However, a slight difference between the speed profiles at intersections A and B can be observed. As shown in Fig. 3-9b, the ego vehicle's acceleration when slowly moving forward, ②, was steady at Intersection B, while this acceleration fluctuated slightly when using the proposed method at intersection A, as shown in Fig. 3-8b.

3.4 Conclusion

In this chapter, a method for crossing low-visibility intersections with a mandatory stop was proposed. The proposed method utilizes the estimated visibility of possible approaching vehicles to decide whether it is safe to proceed forward from the mandatory stopping point when visibility is insufficient for intersection crossing.

Speed profiles of an expert driver when crossing the chosen blind intersections were used for comparison with those produced when using the proposed method. Our results showed that the proposed method could generate driving behavior comparable to that of an expert driver at low-visibility intersections with mandatory stops. The resulting speed profiles revealed similar motion when the vehicle was being driven by an expert driver and when the vehicle was autonomously driven using the proposed method, at both of the selected intersections. This included pausing at the stop line in order to expose the vehicle to possible oncoming traffic, moving slowly into the intersecting road to gain additional visibility, and accelerating to cross the intersection after having obtained sufficient visibility. However, when using the proposed motion planning method, the speed of the ego vehicle as it began to move slowly forward from the stop line was slightly different at Intersections A and B, with smoother acceleration observed at Intersection B. This may be due to more limited visibility at Intersection A. Furthermore, it was found that when the proposed method was used without estimating the visibility of approaching vehicles, the ego vehicle could not successfully cross the intersections due to severely limited visibility. This suggests that estimation of the visibility of possible approaching vehicles is crucial information for an automated crossing of such intersections when using a virtual obstacle-based motion planner.

Lastly, it should be noted that the proposed method relies on several parameters which were manually selected, e.g., s_{enter} . The optimal values of these parameters may differ at different intersections.

Chapter 4

Visibility-Dependent Behavior Model for Approaching Vehicles

4.1 Introduction

A common strategy used to ensure safe motion by the ego vehicle when the environment is not fully observable, especially at intersections, is to assume that there are virtual vehicles approaching at a constant, high speed from immediately outside the visible area toward a conflict point where a collision with the ego vehicle could occur. Under this assumption, it is possible to limit the ego vehicle's speed such that it will not collide with any dynamic objects that could potentially be approaching from the occluded areas. While this approach has been shown to be effective for generating a safe speed profile for the ego vehicle, it can also result in a deadlock situation, where the ego vehicle does not have sufficient visibility to safely proceed forward into an intersection, but cannot gain additional visibility without doing so. As a result, the ego vehicle remains stopped indefinitely at the edge of the intersection. This is especially likely to occur at an intersection with severely restricted visibility since the distance the ego vehicle can observe along the intersecting road is very limited. If it is assumed that there is a hypothetical vehicle traveling towards the intersection at a high and constant speed from just beyond this observable distance, this hypothetical vehicle will be able to reach the conflict point very quickly, making it impossible for

the ego vehicle to cross the intersection without potentially colliding with this virtual, approaching vehicle. In order to escape this situation, a simple heuristic strategy can be implemented such as having the ego vehicle advance slowly into the blind intersection for a limited distance after its speed reaches zero at the edge of the intersection. However, driving too far into the intersection without knowing whether or not other traffic participants are expecting such an event could cause an accident.

In Chapter 3, we investigated a deadlock scenario at blind intersections with mandatory stops. Visibility from the points of view of approaching vehicles was estimated using road network map information and real sensing data from the ego vehicle’s 3D LiDAR unit. This estimated visibility was then used to regulate the ego vehicle’s forward movement out of the deadlock position. It was found that the planner proposed in the previous chapter allowed the ego vehicle to proceed forward from a stop line to gain additional visibility before starting to cross an intersection, in a manner similar to that of an experienced driver. However, the planner proposed in the previous chapter relied on a special action (the ego vehicle creeping forward) that was triggered after the vehicle was stopped long enough at a mandatory stop to be sufficiently observed by other traffic participants. Moreover, that investigation lacked a detailed analysis of changes in the behavior of approaching vehicles after the ego vehicle became observable to them, i.e., an analysis of the visibility-dependent behavior of approaching vehicles.

In this chapter, we extend and generalize our proposed occlusion-aware motion planner that utilizes the visibility of both the ego vehicle and approaching vehicles to formulate safe crossing maneuvers at blind intersections by proposing a visibility-dependent behavior model for vehicles approaching occluded intersections which is based on our analysis of collected, real-world driving data. The proposed visibility-dependent behavior model will allow our planner to predict changes in the behavior of approaching vehicles in response to detecting the ego vehicle waiting to enter the intersection. Our proposed generic, deadlock-free planner utilizes a particle filter algorithm for occluded object prediction, which allows the planner to handle perception uncertainty and incorporate changes in the behavior of occluded traffic participants

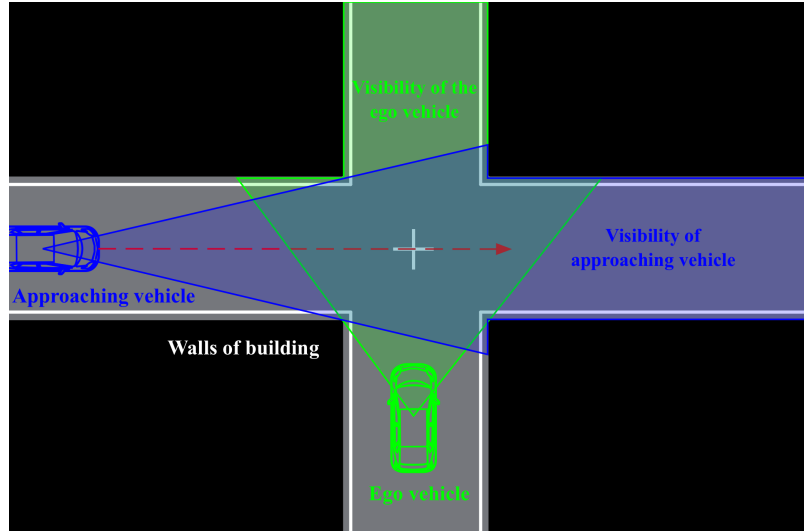
after they have observed the ego vehicle.

In order to verify the effectiveness of our expanded motion planner and visibility-dependent behavior model, experiments were carried out using a simulation of the crossing of a blind intersection. We found that the proposed planner could avoid a deadlock situation even at intersections with severely limited visibility. Moreover, our experimental results revealed that both the level of perception accuracy assumed in the occluded object prediction algorithm, as well as the mounting position of the sensor used for detecting surrounding obstacles, had significant impacts on the speed at which the ego vehicle could safely cross occluded intersections.

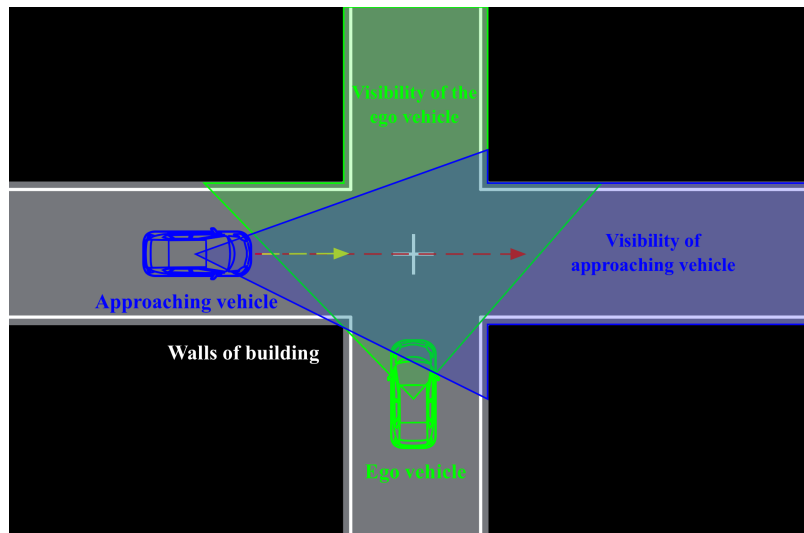
The work in this chapter is a direct extension of the work in Chapter 3. The new contributions of the work described in this chapter are as follows:

- A generic, deadlock-free motion planner for blind intersections which utilizes the visibility of both the ego vehicle and approaching vehicles.
- A visibility-dependent behavior model for vehicles approaching occluded intersections, based on an analysis of real-world driving data.

The remainder of this chapter is organized as follows: In Section 4.2, details of the expanded planner are provided. In Section 4.3, our driving data collection method is described, an analysis of the data is presented, and the proposed visibility-dependent approaching vehicle behavior model for blind intersections is described. In Section 4.4, details of the experimental procedures used to evaluate the performance of the proposed planner, the results of those experiments, including a baseline comparison to verify the planner’s ability to achieve better, deadlock-free motion generation performance, and the effects of perception uncertainty and sensor mounting position are presented in Sections 4.4.1, 4.4.2, and 4.4.3, respectively. Finally, the conclusions of this chapter are discussed in Section 4.5.



(a) An occluded vehicle travels toward the intersection at a constant speed. The ego vehicle is not yet observable by the occluded vehicle.



(b) As the occluded vehicle moves closer to the intersection, the front bumper of the ego vehicle becomes visible, but the occluded vehicle is still hidden from the perspective of the ego vehicle. After detecting the ego vehicle, the occluded vehicle may change its behavior.

Figure 4-1: Detection of the ego vehicle by an approaching, occluded vehicle. The occluded vehicle may change its behavior after observing the ego vehicle.

4.2 Proposed deadlock-free, blind intersection planner

In order to cross an intersection safely, the ego vehicle must pass through the area where the different roadways intersect before any other object, visible or occluded,

arrives at the same location and intercepts it. The proposed planner predicts occluded vehicles using a particle filter, and then determines whether the ego vehicle can cross an intersection without colliding with any of these predicted vehicles based on this safe crossing condition. To handle the deadlock problem, a particle filter-based, occluded vehicle prediction module in the proposed planner utilizes the visibility of both the ego vehicle and hidden vehicles. Visibility from the ego vehicle’s perspective is used to identify regions where occluded vehicles might exist, while the visibility of hidden vehicles is used to determine their behavior. An illustration of an occluded vehicle before and after observing the ego vehicle is given in Fig. 4-1. By considering possible changes in the behavior of approaching vehicles after detecting the ego vehicle, the proposed planner can avoid becoming trapped in a deadlock situation.

The details of each of the proposed planner’s components are given in this section in the following order. First, the conditions required for safe intersection crossing are explained, and a safe crossing strategy is derived based on these conditions. Next, as the main focus of this work is navigating intersections with limited visibility, characteristics of the visibility of both the ego and approaching vehicles at blind intersections are then illustrated. In the third and final subsection, the prediction of occluded objects based on a particle filter algorithm is described.

4.2.1 Safe intersection crossing strategy

An intersection, by definition, is the area where two or more roadways intersect, and this is where collisions between vehicles traveling along the intersecting roadways occur. In the case of a simple intersection involving only two roadways crossing at a 90-degree angle, as shown in Fig. 4-2, the overlapping area is represented by the red rectangle. In order for the ego vehicle to cross the intersection safely, it has to enter and exit the overlapping area before any vehicle traveling on the other roadway arrives. If the ego vehicle cannot ensure that it is safe to cross, it must remain outside the overlapping zone by stopping before the intersection to avoid a potential collision.

This analysis leads us to a strategy for choosing safe actions for the ego vehicle as




Figure 4-2: Simple, blind intersection model. Red shaded area represents the overlapping travel zone where a collision could potentially occur.

it approaches an intersection:

$$Safe\ Action = \begin{cases} Cross, & \text{if } t_{ego} < t_{other} \\ Stop, & \text{if } t_{ego} \geq t_{other} \end{cases}, \quad (4.1)$$

where t_{ego} is the time required by the ego vehicle to completely pass through the overlapping area from its current position, while t_{other} is the time an approaching vehicle will take to reach the overlapping area. In the event of multiple approaching vehicles, the minimum time until one of the other vehicles on the crossing roadway arrives at the overlapping zone is used as t_{other} .

In the proposed planner, the action *Cross* is defined as applying a constant acceleration of a_{ego}^{cross} to completely cross the intersection. On the other hand, the action *Stop* represents a constant reduction of speed at the rate of a_{ego}^{stop} in order to stop before the overlapping zone if the ego vehicle is traveling at the maximum allowable speed, v_{ego}^{allow} , which is given by:

$$v_{ego}^{allow} = \sqrt{-2a_{ego}^{stop} X_{ego}}, \quad (4.2)$$

where X_{ego} is the ego vehicle's current position.

By assuming constant acceleration, we can calculate time t , the time required by any vehicle to travel a total distance of d , using the following kinematic equation:

$$t = \begin{cases} \frac{\sqrt{v^2 + 2ad} - v}{a} & , \text{if } a \neq 0 \\ \frac{d}{v} & , \text{if } a = 0 \end{cases}, \quad (4.3)$$

where a and v represent the instantaneous acceleration and velocity of the vehicle, respectively.

In the case of the ego vehicle, despite uncertainty in the state estimation process, the calculation of t_{ego} is relatively straightforward as the relevant information, i.e., v_{ego} , $a_{\text{ego}}^{\text{cross}}$, and, X_{ego} , are all available. Additionally, the total distance the ego vehicle needs to travel in order to pass through the overlapping zone completely can be calculated as follows:

$$d_{\text{ego}} = X_{\text{ego}} + l_{\text{ego}} + W_{\text{cross}}, \quad (4.4)$$

where l_{ego} is the length of the ego vehicle and W_{cross} is the width of the intersecting roadway.

In contrast, the state of approaching vehicles cannot be fully observed at occluded intersections, therefore, a prediction of the existence of occluded vehicles and their current state, χ_t , has to be made in order to estimate their time of arrival at the overlapping zone, i.e., t_{other} . One common approach for estimating t_{other} at occluded intersections is to assume that there is always a vehicle traveling toward the intersection at high speed, from outside the visible zone. While such an assumption simplifies the estimation of t_{other} , it may also lead to a deadlock situation when the intersecting roadway is heavily occluded.

In order to solve the deadlock problem, a more elaborate behavior model for approaching vehicles, one that captures possible changes in their motion after they have observed the ego vehicle, i.e., a visibility-dependent behavior model, is proposed in this chapter. Therefore, a particle filter-based occluded object prediction method that allows the use of a complex motion model is utilized in the proposed planner to predict the current state of approaching occluded vehicles, χ_t . The proposed occluded

Algorithm 2 Algorithm of proposed generic, deadlock-free planner.

Input: $X_{\text{ego}}, v_{\text{ego}}, a_{\text{ego}}^{\text{cross}}, a_{\text{ego}}^{\text{stop}}, X_{\text{sensor}}, l_{\text{ego}}, W_{\text{ego}}, W_{\text{cross}}, \chi_{t-1}, D_{\text{sensor}}, D_{\text{map}}$ **Output:** $a_{\text{ego}}^{\text{cross}}, a_{\text{ego}}^{\text{stop}}$, or 0

```
1:  $v_{\text{ego}}^{\text{allow}} \leftarrow \text{CalcAllowableSpeed}(X_{\text{ego}}, a_{\text{ego}}^{\text{stop}})$ 
2:  $d_{\text{ego}} \leftarrow \text{CalcCrossingDistance}(X_{\text{ego}}, l_{\text{ego}}, W_{\text{cross}})$ 
3:  $t_{\text{ego}} \leftarrow \text{CalcTravelTime}(d_{\text{ego}}, v_{\text{ego}}, a_{\text{ego}}^{\text{cross}})$ 
4:  $V_{\text{ego}} \leftarrow \text{EstEgoVisibility}(X_{\text{ego}}, X_{\text{sensor}}, D_{\text{sensor}}, D_{\text{map}})$ 
5:  $V_{\text{other}} \leftarrow \text{EstOtherVisibility}(X_{\text{ego}}, X_{\text{sensor}}, D_{\text{sensor}}, D_{\text{map}})$ 
6:  $\chi_t \leftarrow \text{PredOccludedVehicle}(\chi_{t-1}, V_{\text{ego}}, V_{\text{other}})$ 
7:  $t_{\text{other}} \leftarrow \text{CalcTravelTime}(\chi_t)$ 
8: if  $t_{\text{ego}} < t_{\text{other}}$  then
9:   return  $a_{\text{ego}}^{\text{cross}}$ 
10: else
11:   if  $v_{\text{ego}} \geq v_{\text{ego}}^{\text{allow}}$  then
12:     return  $a_{\text{ego}}^{\text{stop}}$ 
13:   else
14:     return 0
15:   end if
16: end if
```

vehicle prediction method uses estimated visibility of both the ego vehicle, V_{ego} , and of the approaching vehicles, V_{other} . In most complex scenarios, both V_{ego} and V_{other} can be estimated using map information D_{map} , sensing data D_{sensor} , and the position of the ego vehicle and its sensor, X_{ego} and X_{sensor} , respectively, as demonstrated in Chapter 3. However, in some situations, e.g., at simple occluded intersections, visibility estimation can also be performed without using real sensing data. Visibility estimation will be described in more detail in Section 4.2.2.

By using the occluded vehicle prediction output t_{other} and the time required by the ego vehicle to completely cross the intersection t_{ego} , the appropriate output action for the ego vehicle can be determined. Therefore, the proposed planner can be summarized as shown in Algorithm 2, where `CalcAllowableSpeed`, `CalcCrossingDistance`, and `CalcTravelTime` refer to Eq. 4.2, Eq. 4.4, and Eq. 4.3, respectively. The visibility estimation of both the ego vehicle (`EstEgoVisibility`) and approaching vehicles (`EstOtherVisibility`) will be explained in Section 4.2.2 as previously mentioned. The function `PredOccludedVehicle` represents a particle filter-based occluded object prediction method, which will be described in Section 4.2.3.

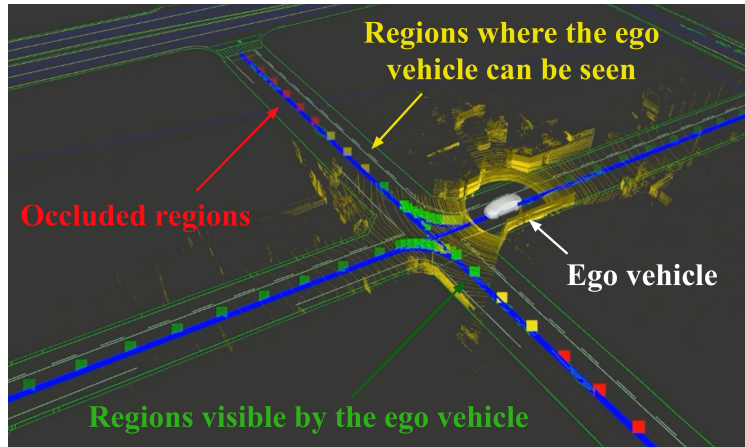


Figure 4-3: Example of visibility estimation results using real sensing data. The green, yellow and red squares indicate regions observable by the ego vehicle, regions from which other vehicles can observe the ego vehicle, and occluded regions, respectively. The visibility estimation approach proposed in the previous chapter can be directly incorporated into the planner proposed in this chapter for application in complex, real-world environments.

4.2.2 Visibility at blind intersections

The first step toward the prediction of occluded dynamic objects and their current state is to understand the characteristics of visibility at blind intersections. Visibility from the ego vehicle’s point of view at an intersection depends on various factors, some of which are static, i.e., sensor position, sensor coverage, the intersection’s geometry, the presence of stationary objects such as buildings and walls, etc. On the other hand, there are components affecting visibility that can change over time, such as dynamic surrounding objects, parked vehicles, trees that shed their leaves, etc. Therefore, in order to correctly determine the optimal motion of the ego vehicle in areas with occlusions, visibility should be estimated in real-time using sensing data. However, in the absence of dynamic occlusions, i.e., occlusions caused by moving objects, visibility can also be estimated with adequate accuracy without using current sensing data. In order to demonstrate how visibility changes at an occluded intersection, in this investigation, visibility estimation is performed using a closed-form expression based on the intersection’s geometry and the position of the ego vehicle and its sensor.

However, it is important to emphasize that the planner proposed in this work

can be used in more complex, real-world environments by replacing the geometry-based visibility estimation approach used here with an estimation approach that relies on actual sensing data. For example, the visibility estimation method employed in Chapter 3 utilizes 3D LiDAR data and HD maps to accurately identify the occluded and visible regions of a complex environment, as shown in Fig. 4-3.

In order to estimate visibility without sensing data, a simple model of a blind intersection such as the one depicted in Fig. 4-2 can be used. As can be seen in the figure, the ego vehicle cannot fully observe the intersecting roadway due to the occluding walls of the buildings at the intersection's corners. The solid, green line in Fig. 4-2 represents a clear line of sight from the sensor of the ego vehicle, defining the maximum visible distance along the intersecting roadway in that direction. Based on the geometry of the intersection, the furthest distance that is visible from the ego vehicle, i.e., the ego vehicle's visibility, V_{ego} , can be estimated using the following equation:

$$V_{\text{ego}} = \frac{(X_{\text{ego}} + X_{\text{sensor}} + \frac{W_{\text{cross}}}{2}) \times (\frac{W_{\text{ego}}}{2})}{(X_{\text{ego}} + X_{\text{sensor}})}, \quad (4.5)$$

where X_{ego} is the distance from the intersection's entrance to the front bumper of the ego vehicle. The sensor's position relative to the ego vehicle's front bumper is represented by X_{sensor} , and the width of the ego vehicle's roadway is denoted by W_{ego} . Therefore, it is evident from Eq. 4.5 that the ego vehicle's visibility depends on its distance from the entrance to the intersection, the intersection's geometry and the mounting position of its sensor.

Using Eq. 4.5, it is possible to determine how visibility from the point of view of the ego vehicle changes as the ego vehicle approaches an intersection, i.e., we can generate a visibility profile. As can be seen from the visibility profiles shown in Fig. 4-4, the ego vehicle's visibility remains virtually constant when the ego vehicle is located at a position distant from the intersection, but its visibility starts to increase exponentially as the ego vehicle advances closer to the intersection. The effects of the intersection's geometry can also be observed from the visibility profile since visibility begins to increase sooner the greater the width of the ego vehicle's roadway. Also note

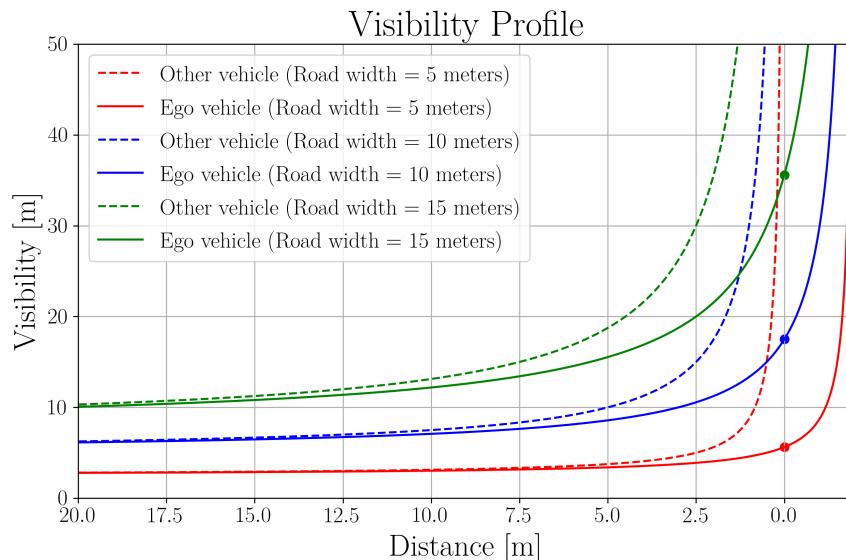


Figure 4-4: Visibility profiles of the ego vehicle and other vehicles at blind intersections of roadways of various widths. In each scenario, the ego vehicle’s sensor is located 2 meters to the rear of its front bumper. The widths of the intersecting roadways are assumed to be equal.

that the visibility of the ego vehicle when it reaches the intersection’s entrance, i.e., $X_{\text{ego}} = 0$, is significantly different depending on the width of the intersecting road, as indicated in Fig. 4-4 with (•) markers along the solid lines at the 0-meter position, where the color of each line represents an intersecting roadway of different width. Although the ego vehicle’s visibility becomes infinitely large as the sensor reaches the intersection’s entrance regardless of the width of the intersecting roadway, when the intersecting road is narrower there is a steeper and more delayed increase in the ego vehicle’s visibility.

While modern autonomous vehicles are generally equipped with multiple sensors mounted at various locations, the main sensors for perception tasks are usually mounted close to the center of the vehicle, e.g., on the roof, to increase overall sensing coverage. When the sensor is mounted in this location, the ego vehicle can be observed by occluded vehicles traveling along the intersecting roadway before they can be detected by the ego vehicle, as depicted in Fig. 4-2, where the solid yellow and green lines represent the distance the ego vehicle can be observed from the intersecting roadway, V_{other} , and visibility from the ego vehicle’s sensor, V_{ego} , respectively.

It is important to note that V_{other} is not the actual visibility of any particular vehicle traveling along the intersecting roadway, from the perspective of the “other vehicle”, but the furthest distance from the intersection that the ego vehicle can be seen from. However, for the sake of conciseness, will be used to refer to the visibility of other vehicles. In a manner similar to that used in Eq. 4.5 to define visibility from the perspective of the ego vehicle, the relationship between visibility from other vehicles and the position of the ego vehicle can be expressed as follows:

$$V_{\text{other}} = \frac{(X_{\text{ego}} + \frac{W_{\text{cross}}}{2}) \times (\frac{W_{\text{ego}}}{2})}{(X_{\text{ego}})}. \quad (4.6)$$

In Fig. 4-4, similar characteristics can be observed between the visibility of the ego vehicle and that of the other vehicles, but with an earlier increase in the visibility of the other vehicles, as illustrated by the dashed lines. As shown by the red lines in Fig. 4-4, when the ego vehicle arrives at the intersection its visibility is still very limited, while approaching vehicles traveling on the intersecting road can observe it from a much more distant location. Thus, the difference between V_{other} and V_{ego} results in the ego vehicle becoming visible to other vehicles before the ego vehicle can observe them. In a deadlock situation, even though the ego vehicle does not have sufficient visibility to cross the intersection from its current position, it is likely that it can be observed by vehicles approaching from the occluded area. This is important because this difference in visibility is utilized in our proposed method to solve the deadlock problem, by modeling approaching vehicle behavior to be dependent on whether or not these vehicles can observe the ego vehicle attempting to cross the intersection.

4.2.3 Particle filter-based, occluded vehicle prediction

As pointed out in Section 4.2.1, the prediction of occluded vehicles is a crucial component of planning safe ego vehicle motion through blind intersections. In order to incorporate perception uncertainty and allow flexibility in the occluded vehicle behavior model during prediction, we used a particle filter algorithm in our model. The schematic diagram shown in Fig. 4-5 illustrates the main steps of a basic particle

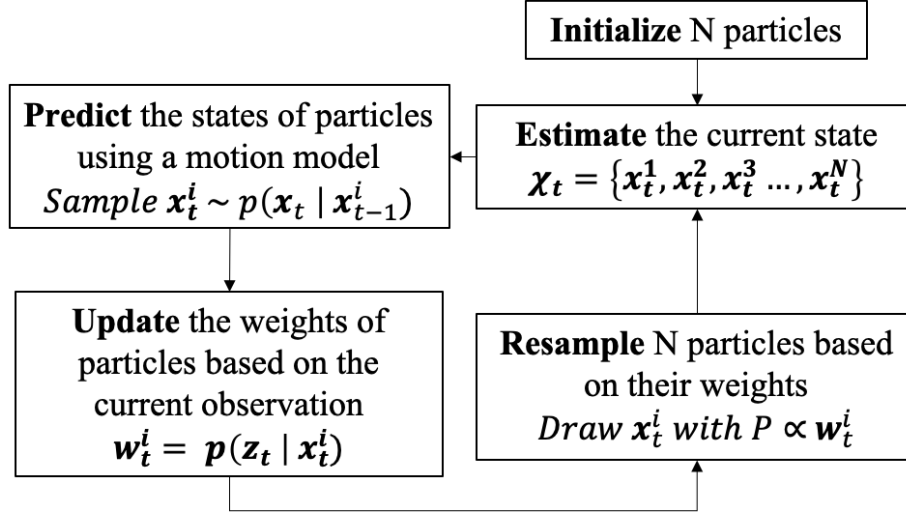


Figure 4-5: Diagram of a particle filter algorithm

filter algorithm.

The algorithm begins by generating a finite number of particles $\mathcal{N}_{\text{particle}}$, where a set of $\mathcal{N}_{\text{particle}}$ particles at time t can be denoted as follows:

$$\chi_t = \{x_t^1, x_t^2, x_t^3, \dots, x_t^{\mathcal{N}_{\text{particle}}}\}. \quad (4.7)$$

Each particle x_t^i , when $1 \leq i \leq \mathcal{N}_{\text{particle}}$, contains a state. In this study, a particle's state represents a hypothesis of the state of an occluded vehicle at time t . Specifically, the state of each vehicle traveling on the intersecting roadway is represented by the state of a particle, which can be represented as:

$$x_t^i = \{d_t^i, v_t^i, a_t^i, A_t^i\}, \quad (4.8)$$

where d_t^i represents displacement from the current position of the particle to the intersection's center, while v_t^i and a_t^i are the particle's speed and acceleration, respectively. The last state variable, A_t^i , is a Boolean variable indicating whether or not the particle is aware of the ego vehicle, which will be explained in more detail below.

After the particles are initialized, the current estimation of each particle's (i.e., each occluded vehicle's) state, χ_t , is recursively updated at every time step Δt , as a new observation becomes available. Starting from the previously estimated state,

χ_{t-1} , the current state of every particle is predicted using a motion model. The motion model used in this study is as follows:

$$\begin{aligned}
 d_t^i &= d_{t-1}^i + v_t^i \Delta t + \frac{1}{2} a_t^i \Delta t^2 \\
 v_t^i &= v_{t-1}^i + a_t^i \Delta t \\
 a_t^i &= f(d_{t-1}^i, v_{t-1}^i, A_t^i) \quad , \\
 A_t^i &= \begin{cases} True & , \text{if } T_{\text{observe}_i} \geq T_{\text{react}} \\ False & , \text{if } T_{\text{observe}_i} < T_{\text{react}} \end{cases}
 \end{aligned} \tag{4.9}$$

where T_{observe_i} is the length of time the ego vehicle has been observable by the particle, which is determined by $d_t^i < V_{\text{other}}$. The parameter T_{react} represents a recognition delay, i.e., the time it takes the approaching vehicle to become aware of the ego vehicle after it has become visible. The instantaneous acceleration of a particle, a_t^i , is determined by a function f , which in this work is designed to be dependent on the particle's visibility of the ego vehicle. This f function represents the visibility-dependent behavior of approaching vehicles, which will be used to avoid a deadlock situation, as will be explained in more detail in Section 4.3.

It is important to note that function f is not limited to the behavior model proposed in this work, nor it is restricted to only one specific type of traffic participant. It can also be used to represent completely different behavior models, including the conventional approaching vehicle behavior model commonly used in blind intersection planners, i.e., a vehicle traveling toward the intersection at a constant speed can be expressed by a constant function, which outputs zero acceleration regardless of the state of the approaching vehicle. Moreover, the occluded object behavior function f can be applied to other types of traffic participants, e.g., pedestrians and bicycles.

After predicting the states of all of the particles using the motion model, these particles represent samples drawn from a state transition distribution, $p(x_t|x_{t-1}^i)$. As observation z_t becomes available, each particle is assigned an importance weight based on an observation model, $p(z_t|x_t^i)$. In this work, the Bernoulli distribution is used as

our observation model:

$$p(z_t|x_t^i) = \begin{cases} 1 - \alpha, & d_t^i < V_{\text{ego}} \\ \alpha, & d_t^i \geq V_{\text{ego}} \end{cases}, \quad (4.10)$$

where α represents the accuracy of the visible area classifier.

Lastly, a set of $\mathcal{N}_{\text{particle}}$ particles is re-sampled, based on their importance weights. This set of $\mathcal{N}_{\text{particle}}$ particles now represents the current state estimate, χ_t . All of the processes described above are then repeated at every time step.

4.3 Proposed visibility-dependent behavior model for approaching vehicles

In order to address the occluded intersection deadlock problem, a more elaborate model for the behavior of approaching vehicles is proposed in this section. The proposed behavior model aims to capture possible changes in the behavior of approaching vehicles conditional on their visibility of the ego vehicle, i.e., a visibility-dependent behavior model. The proposed behavior model is incorporated into the particle filter-based occluded vehicle prediction algorithm introduced in Section 4.2.3, via function f .

As the proposed behavior model is based on the analysis of collected real-world driving data, the details of the collection of this data are given below in the first subsection. An analysis of this driving data is then presented in the following subsection. Lastly, the proposed visibility-dependent behavior model is described.

4.3.1 Collection of driving data

It is essential to model the behavior of approaching vehicles in a way that represents the real behavior of actual drivers. In order to understand how drivers behave when they encounter an intersection with limited visibility, driving data collected in such an environment is required. Therefore, an experiment was carried out to collect the

required driving data, the details of which are described below.

Driver types

Studies have shown that experienced and expert drivers tend to take proactive action to avoid collisions, e.g., by slowing down as they approach an uncontrolled intersection with poor visibility [34,75]. However, it is neither safe nor reasonable to assume that every driver approaching such an intersection will behave in this manner. Therefore, our data collection experiment included three different types of drivers in order to capture a wide range of driving behavior at blind intersections:

- Expert drivers: Driving school instructors
- Elderly drivers: Drivers 65 years old or older
- Typical drivers: Drivers who do not belong to the other two groups

A total of 18 drivers from these three groups participated in our data collection experiment. The participants included five male expert drivers, two elderly female drivers, two elderly male drivers, four typical female drivers, and five typical male drivers. Informed consent was obtained from all participants involved in the study. All personal information was handled carefully to protect the privacy of the participants.

Experimental vehicle

The vehicle used for data collection was the Toyota Prius shown in Fig. 4-6. It was equipped with a 3D LiDAR sensor (Velodyne VLP-16) and a GPS receiver, which allowed real-time localization via the 3D Normal Distributions Transform (3D-NDT) scan-matching method proposed in [72]. Along with data from the previously mentioned sensors, CAN-bus data, which included speed and brake pedal pressure, was also collected. By fusing data from these different sensors, the position, orientation and speed of the vehicle at each location along its route could be obtained.



Figure 4-6: Experimental vehicle



Figure 4-7: Experimental environment

Experimental environment

The environment selected for data collection was a residential area of Nagoya, Japan containing several intersections with notably low visibility and no traffic signals. The driving route used for data collection is shown in Fig. 4-7. From the beginning to the end of the route, each driver crossed a total of four blind intersections. Two of the intersections required drivers to make a mandatory stop (Intersections A and B),




Figure 4-8: Speed profiles of all drivers when approaching and crossing blind intersections in a residential neighborhood (232 intersection crossings by 18 drivers).

while the other two intersections (Intersections C and D), did not require drivers to stop prior to crossing.¹

4.3.2 Driving data analysis

A total of 232 intersection crossings by 18 drivers were collected. In order to visualize all of the crossing data from the various intersections in one graph, the distance from the center of each intersection to the experimental vehicle was used to represent the location of the vehicle. Profiles of crossing speeds at the selected blind intersections are shown in Fig. 4-8. We can see that most of the drivers initially approached the intersections at a speed of approximately 8 m/s, which is just below the local speed limit of 8.33 m/s, i.e., 30 km/h, and that the majority of drivers slowed down as they approached the intersections. However, in some cases, drivers maintained a relatively high speed as they crossed the intersections. The data suggests there is a variety of possible driving behavior when drivers encounter blind intersections.

In order to further analyze the behavior of the approaching vehicles, cluster anal-

¹A satellite view of the experimental environment and driving route can be found at https://drive.google.com/open?id=1CbrP0BYfK0tPR_Qs60Qkjov-PJDk3pEK&usp=sharing

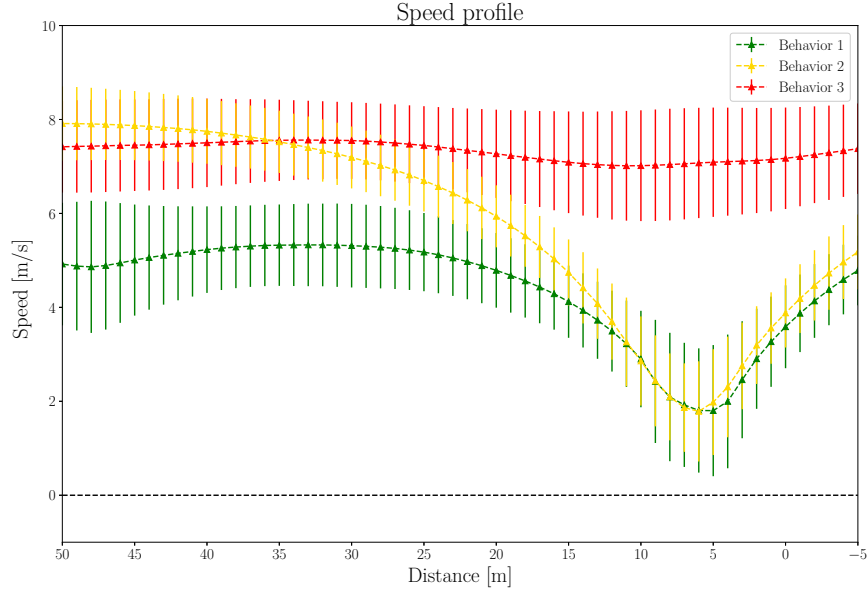


Figure 4-9: Results of cluster analysis using K-Mean algorithm on collected speed profiles of drivers crossing all intersections. Red, yellow and green lines represent average speed profile of each cluster. Vertical bands show standard deviation of vehicle speed at each location.

ysis was performed. The K-Mean clustering algorithm [76] was used to group similar speed profiles into clusters. The optimal number of clusters, which is the main parameter of the K-Mean clustering algorithm, was determined based on the silhouette coefficient [77]. As a result, the speed profiles were divided into three groups. The average speed at each location in relation to the intersection, i.e., the average speed profile, for each of the three clusters is represented by the red, yellow, and green lines in Fig. 4-9. The standard deviation of vehicle speed at each location for each cluster is represented by a vertical band.

As indicated by the red speed profile in Fig. 4-9, cruising through intersections at a constant speed appears to be one of common the behaviors of approaching vehicles. The yellow and green speed profiles both show approaching vehicles slowing down prior to a blind intersection, however, the green speed profile indicates a more abrupt deceleration compared to the yellow speed profile.

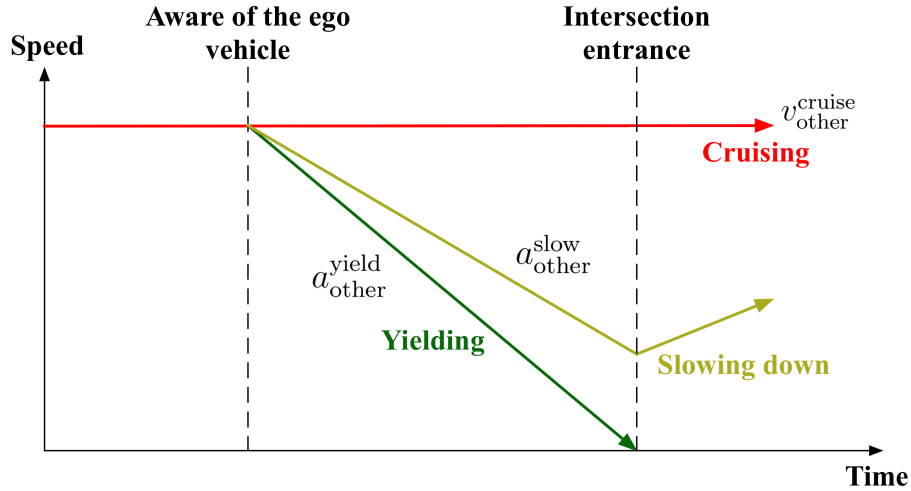


Figure 4-10: Speed profiles of the three possible behaviors in proposed behavior model for vehicles approaching blind intersections.

4.3.3 Proposed visibility-dependent behavior model

Based on our analysis of the collected data, we created a behavior model for vehicles approaching blind intersections. The proposed model consists of three possible behaviors, *Cruising*, *Slowing down* and *Yielding*, which correspond to the red, yellow and green speed profiles shown in Fig. 4-9, respectively. Note that these speed profiles are not linked directly to particular drivers, but represent behaviors, i.e., sometimes the same driver may fit a different speed profile, based on the driving environment at the time they encountered the intersection. In other words, the driving behavior of particular drivers may vary based on the visibility of the intersection, and whether or not approaching vehicles are observed along the intersecting road.

Our ultimate goal when designing the approaching vehicle behavior model was to use it for defining function f in our proposed particle filter-based occluded vehicle prediction method. As function f is used for estimating the instantaneous acceleration of occluded vehicles, a speed profile for each behavior in the time domain is required. Therefore, we next modeled a speed profile for each intersection approaching behavior, and these speed profiles are shown in Fig.4-10. For the *Cruising* profile, the vehicle approaches and crosses the intersection at a constant speed of $v_{\text{other}}^{\text{cruise}}$, resulting in zero acceleration the entire time. As for the *Slowing down* profile, the vehicle also travels

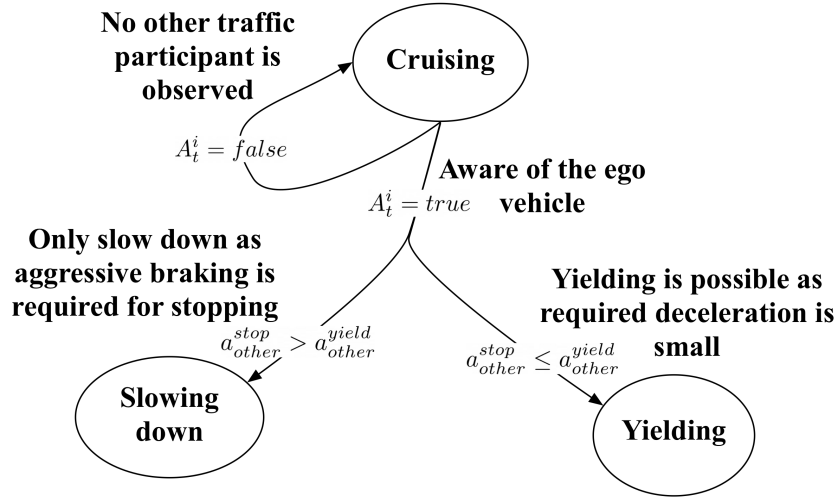


Figure 4-11: Proposed behavior transition for vehicles approaching blind intersections.

toward the intersection at a speed of $v_{\text{other}}^{\text{cruise}}$ initially, however, at some point, it starts to decelerate at a rate of $a_{\text{other}}^{\text{slow}}$, before accelerating again to cross the intersection. Lastly, the *Yielding* profile represents behavior similar to *Slowing down*, but the vehicle slows down to a full stop prior to the intersection, from an initial speed of $v_{\text{other}}^{\text{cruise}}$. Apart from the minimum speed before crossing the intersection, another difference between the *Yielding* and *Slowing down* profiles is the rate at which the vehicles decelerate. For the *Yielding* profile, the approaching vehicle applies a relatively sharper deceleration of $a_{\text{other}}^{\text{yield}}$ in order to come to a full stop before the intersection.

We propose that these behaviors are connected, and that transitions between them can be explained by whether or not the driver detects another traffic participant on the intersecting road and the deceleration required to yield after detecting another vehicle. These transitions are represented in the form of a diagram in Fig. 4-11, which shows an approaching vehicle starts in the *Cruising* state, i.e., traveling toward an intersection at a constant speed of $v_{\text{other}}^{\text{cruise}}$. When no other traffic participant is observed on the intersecting roadway, the approaching vehicle remains in the same state. The transition from *Cruising* behavior occurs when the driver of the approaching vehicle becomes aware of the ego vehicle on the intersecting roadway, i.e., when A_t^i is true, since the ego vehicle has become observable for a period of at least T_{react} . This

transition results in *Yielding* if the amount of deceleration required to come to a full stop, $a_{\text{other}}^{\text{stop}}$, does not exceed a limit of $a_{\text{other}}^{\text{yield}}$, as the approaching vehicle can easily come to a full stop in order to yield to the ego vehicle in this case. Otherwise, the driver behavior switches from *Cruising* to *Slowing down* since the approaching vehicle would have to decelerate too abruptly to stop, and therefore is likely to only slow down and pass the intersection before the ego vehicle.

This proposed behavior model can be viewed as a more elaborate extension of the worst-case scenario behavior model commonly used in related studies on virtual obstacle-based motion planning, where occluded vehicles are assumed to maintain a constant speed regardless of the situation, since it provides a behavior transition from worst-case scenario behavior, i.e., *Cruising*, to other behaviors which can be observed in real traffic. Even though the proposed behavior transitions are deterministic and only conditioned on two factors, namely, observation of other traffic participants and the amount of deceleration required for stopping, it can be further extended to probabilistic transitions which are conditioned on other factors, e.g., the probability of an approaching vehicle changing its state from *Cruising* to *Yield* could be set higher in situations where it is traveling toward an intersection with a mandatory stop

4.4 Experiments

In order to validate the performance of the proposed intersection crossing planner at occluded intersections, a closed-loop intersection crossing simulator was developed, which allowed us to test the proposed planner under various visibility conditions and sensor configurations while leaving other factors that might affect the output motion unchanged. Fig. 4-12 shows the outputs of the closed-loop simulator, illustrating the operation of the proposed planner during an intersection crossing event. The simulation begins with the ego vehicle 50 meters from the intersection’s entrance, traveling toward the intersection at the maximum speed of 8.3 m/s. At each time step, the simulator uses the output acceleration from the proposed planner to compute the ego vehicle’s speed and position at the next time step. These processes are then

Table 4.1: Values of major parameters used in occluded intersection crossing experiments.

Parameter	Value	Parameter	Value
l_{ego}	4.5 m	w_{ego}	1.7 m
W_{ego}	5.0 m, 15.0 m	W_{cross}	5.0 m, 15.0 m
$a_{\text{ego}}^{\text{cross}}$	3.0 m/s ²	$a_{\text{ego}}^{\text{stop}}$	-3.0 m/s ²
$a_{\text{other}}^{\text{yield}}$	-1.5 m/s ²	$a_{\text{other}}^{\text{slow}}$	-0.8 m/s ²
$v_{\text{other}}^{\text{cruise}}$	8.3 m/s	T_{react}	2.3 s

repeated every 0.1 seconds until the simulation times out after 20 seconds.

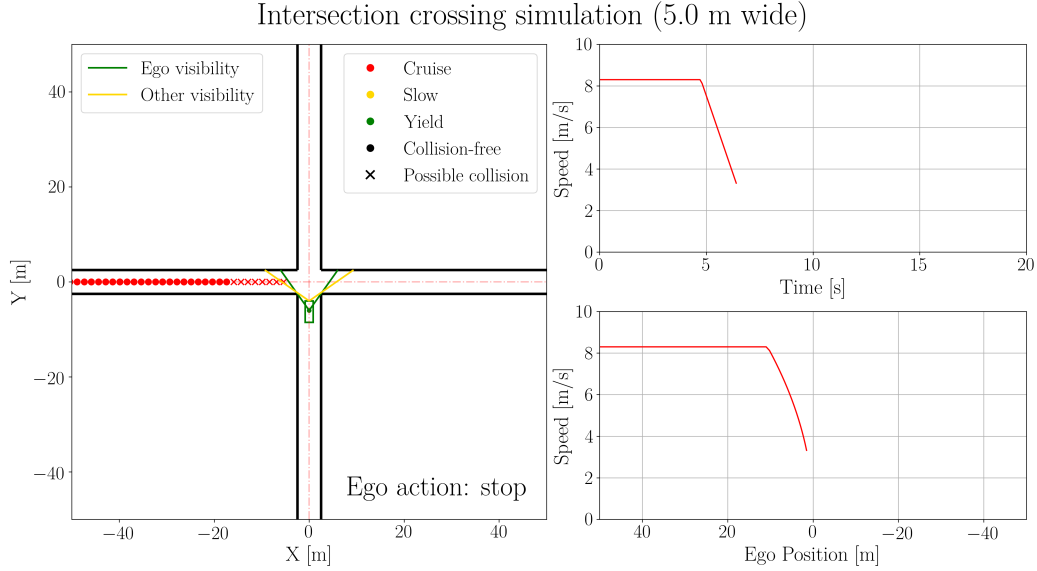
Three different experiments were carried out in order to highlight various characteristics of the proposed planner.² In each experiment, the planner was tested at two intersections with different visibility conditions; one where the intersecting roads were both 5 meters wide, and one where the intersecting roads were both 15 meters wide. The major parameters used in these simulation experiments are summarized in Table 4.1.

4.4.1 Baseline comparison


An initial experiment was carried out to verify the effectiveness of the visibility-dependent behavior model for approaching vehicles at solving the deadlock problem. The proposed planner was compared with a baseline model which assumed that an occluded vehicle was approaching the intersection at a constant speed from outside the area visible to the ego vehicle. In this experiment, the sensor was mounted 2 meters from the front end of the ego vehicle, i.e., $X_{\text{sensor}} = 2\text{ m}$. Moreover, perception accuracy was set to be 100%, i.e., $\alpha = 1.0$.

Fig. 4-12 shows a sequence of ego vehicle motions generated by the proposed planner at a 5-meter-wide, blind intersection. The hypothetical occluded vehicles (or “particles”) are depicted with either an \times (representing a possible collision) or a \bullet (representing a collision-free crossing) depending on whether or not they will reach

²Videos of our experiment can be found at <https://youtu.be/Bic2QL2RQps>



(a) Ego vehicle slows down and prepares to stop at the intersection's entrance as its visibility, shown by the green lines, is insufficient for crossing.

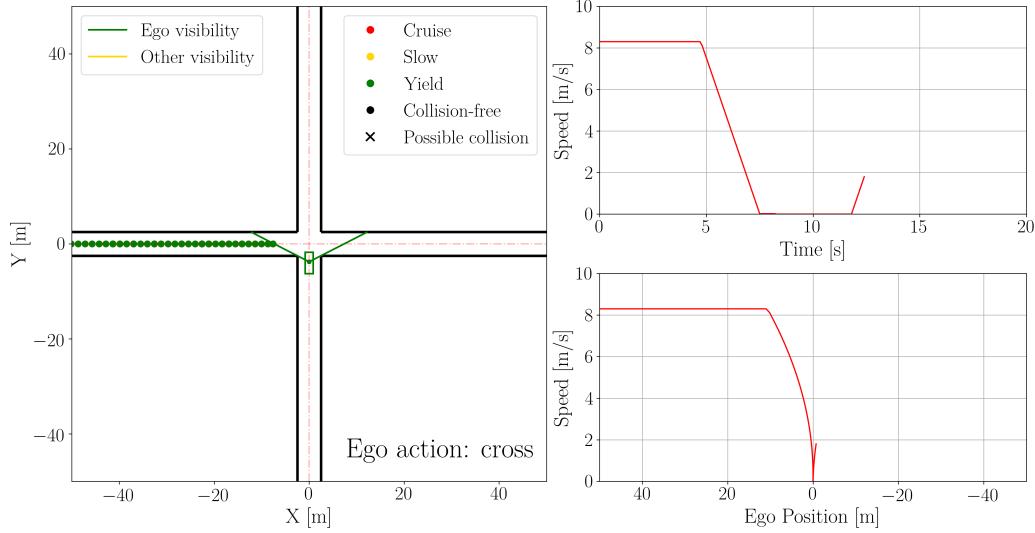


(b) Ego vehicle has come to a full stop at the intersection's entrance. Even though its visibility is still limited, it can be observed by other vehicles from some distance since its front bumper is aligned with the edge of the intersection.

Figure 4-12: Examples of simulation results for four stages of ego vehicle motion when crossing an occluded intersection.

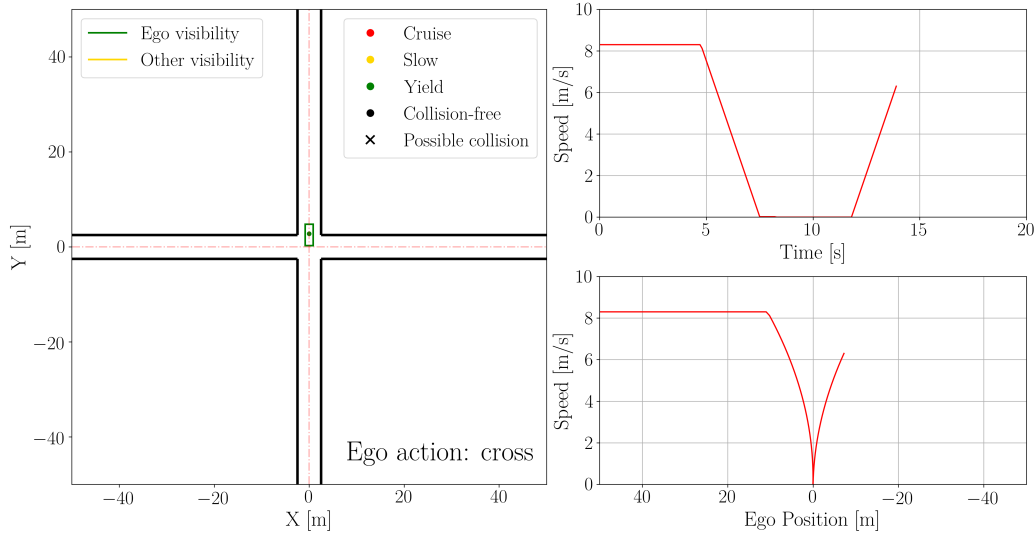
the overlapping zone of the intersection before the ego vehicle finishes crossing, i.e., $t_{\text{ego}} \geq t_{\text{other}}$ or $t_{\text{ego}} < t_{\text{other}}$, respectively. The colors of these particles indicate their current behavior; green represents *Yield* behavior, yellow represents *Slowing down* behavior and red represents *Cruise* behavior.

Intersection crossing simulation (5.0 m wide)



(c) Ego vehicle begins to move into the intersection after it has stopped and estimated that other, hidden vehicles have seen it waiting to cross. As a result of moving forward, the ego vehicle gains additional visibility of the intersecting roadway.

Intersection crossing simulation (5.0 m wide)



(d) Ego vehicle crosses intersection as the visibility has become sufficient.


Figure 4-12: Examples of simulation results for four stages of ego vehicle motion when crossing an occluded intersection. (cont.)

As can be seen in Fig. 4-12a, the ego vehicle slowed down as it approached the intersection since its visibility, shown by the green lines, was very limited. Similarly, the ego vehicle was observable by vehicles approaching on the intersecting roadway only after they reached a position very close to the intersection, i.e., after entering

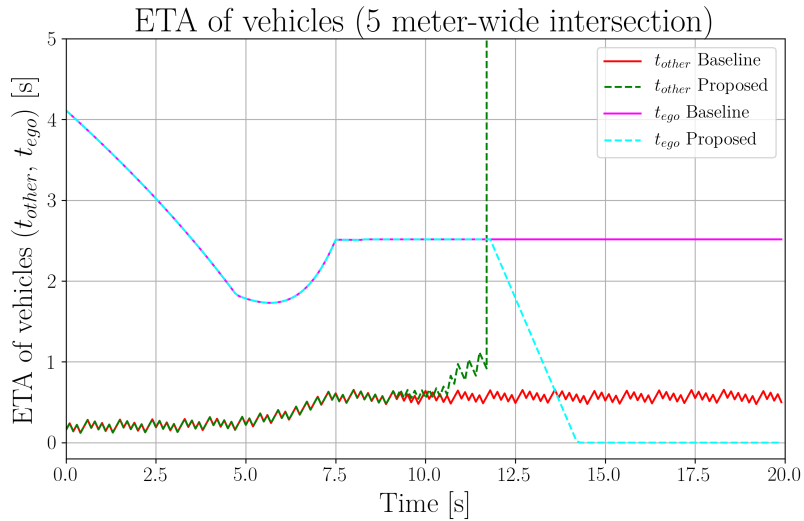
the area between the two yellow lines. As depicted in Fig. 4-12b, the ego vehicle then came to a full stop at the entrance to the intersection. From this position, the ego vehicle’s visibility was still insufficient for crossing, however, the ego vehicle could now be observed by approaching vehicles from a more distant location, as its front end was now aligned with the corner. Therefore, some of the particles started to change their behavior from *Cruising* to *Slowing down* or *Yield*, depending on their location. As shown in Fig. 4-12c, once the ego vehicle had stopped at the intersection for some time, all of the particles close to the intersection changed their behavior to *Yield*. Consequently, the ego vehicle began crossing the intersection, as the safe crossing condition defined in Eq. 4.1 was satisfied, as shown in Fig. 4-12d. As a result of proceeding forward, the ego vehicle gained more visibility along the intersecting roadway, causing the hypothetical occluded vehicles to be eliminated.

The proposed planner’s speed profile when crossing a 5-meter-wide intersection is represented by the dashed green line in Fig. 4-13a. We can see that the ego vehicle reduced its speed continuously until reaching a full stop. It then remained at rest for a few seconds before starting to accelerate again. As for the baseline planner, its speed profile is depicted by the solid red line in Fig. 4-13a. When approaching the intersection, the baseline and proposed planners generated identical motion, therefore, both speed profiles overlap perfectly. However, after the ego vehicle came to a full stop, the baseline planner did not generate any crossing motion command due to insufficient visibility. Consequently, the ego vehicle remained in a deadlocked situation until the simulation timed out.

The estimated time of arrival (ETA) for the ego vehicle, t_{ego} , and the predicted occluded vehicle, t_{other} , for both the baseline and proposed planners, are shown in Fig. 4-13b. The green and cyan dashed lines represent t_{ego} and t_{other} , respectively, as estimated by the proposed planner. Initially, t_{other} was extremely low and it remained almost unchanged. Meanwhile, t_{ego} continuously decreased as the ego vehicle approached the intersection at a constant speed. Since t_{ego} remained larger than t_{other} , the ego vehicle began braking. As the ego vehicle’s speed fell, t_{ego} decreased at a slower rate before beginning to increase again as the vehicle slowed to a full



(a) Ego vehicle speed profiles at a 5 meter-wide, occluded intersection.




(b) Estimated time of arrival of ego and other vehicles at a 5 meter-wide, occluded intersection.

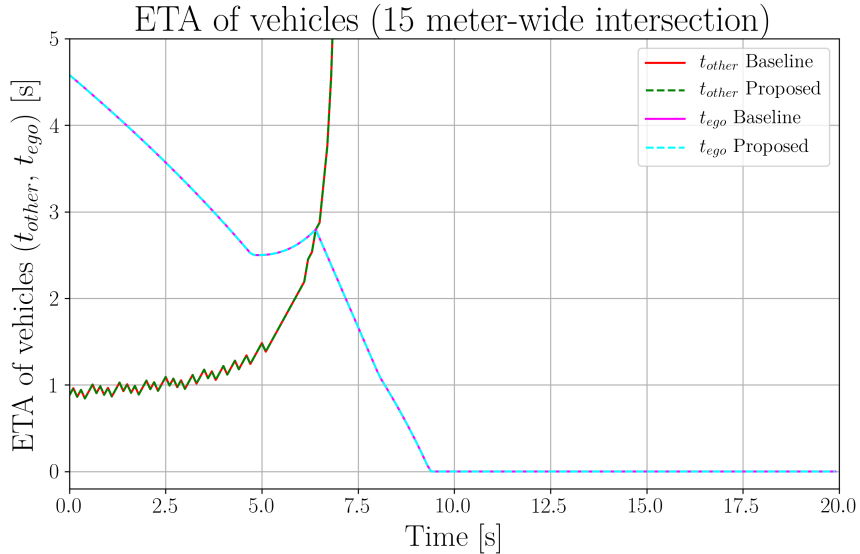
Figure 4-13: Comparison of motion generated by proposed and baseline planners at a 5 meter-wide, occluded intersection.

stop. After coming to a full stop, t_{ego} remained constant, while t_{other} began increasing after a few seconds as the particles began to adjust their speed. The increasing t_{other} eventually surpassed the constant t_{ego} , allowing the ego vehicle to begin crossing the intersection. As the ego vehicle crossed the intersection, t_{ego} began to drop until it finally reached zero once the crossing of the intersection was completed.

For the baseline planner, plots of t_{ego} and t_{other} are identical to those for the



(a) Ego vehicle speed profiles at 15 meter-wide, occluded intersection.



(b) Estimated time of arrival of ego and other vehicles at 15 meter-wide, occluded intersection.

Figure 4-14: Comparison of proposed and baseline planners at a 15 meter-wide, occluded intersection.

proposed planner prior to the full stop, as shown in Fig. 4-13b. In contrast to the proposed planner, however, t_{other} and t_{ego} remained almost unchanged after the stop, causing the gap between them to remain nearly constant. As t_{other} never exceeded t_{ego} , the ego vehicle continued to be trapped in a deadlock until the end of the simulation.

We also compared the two planners at an intersection with relatively better visi-

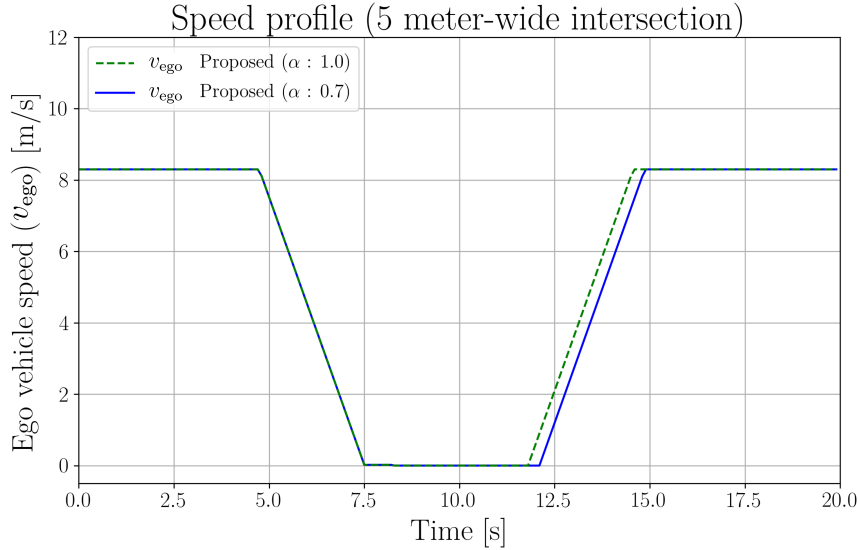
bility, i.e., at an intersection of two roads that were both 15 meters wide. As shown in Fig. 4-14a, the output speed profiles of both planners are identical. Initially, the ego vehicle approached the intersection at a constant speed, then it started to slow down in order to prepare for a stop prior to entering the intersection. However, as it approached the intersection, the ego vehicle started to accelerate and crossed the intersection without stopping.

Similarly, the estimated t_{other} and t_{ego} of the proposed and baseline planners are indistinguishable, as can be seen in Fig. 4-14b. As the ego vehicle progressed toward the intersection at a constant speed, t_{ego} gradually declined, while t_{other} increased at an accelerating rate. Even though t_{ego} started to rise due to braking, the ego vehicle’s visibility increased exponentially as the ego vehicle reached the intersection, and t_{other} eventually exceeded t_{ego} . This was the moment the ego vehicle started to accelerate to cross the intersection.

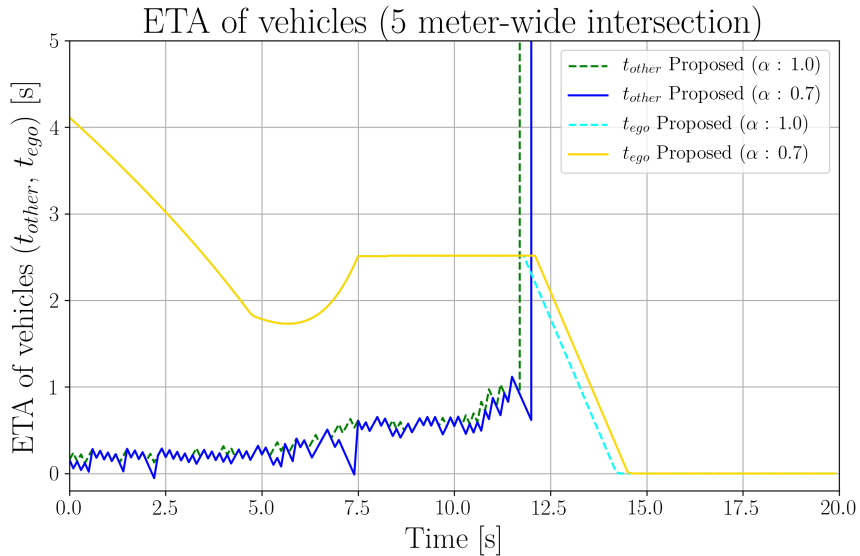
4.4.2 Effects of perception inaccuracy

As mentioned in Section 4.2.3, the proposed planner utilizes a particle filter algorithm for occluded vehicle prediction in order to account for uncertainty from the perception module. This experiment aims to investigate how perception inaccuracy affects the output motion of the proposed planner. The proposed planner was tested using occluded intersections of both 5 and 15 meter-wide roadways using two different perception accuracy values ($\alpha = 1.0$ and $\alpha = 0.7$), while X_{sensor} was set to 2 meters, representing the location of a rooftop-mounted LiDAR unit.

When encountering the intersection of the 5 meter-wide roads, the speed profiles shown in Fig. 4-15a indicate that, overall, the output motions of the proposed planner are similar no matter which α value is used. More specifically, when assuming perfect perception ($\alpha = 1.0$), the ego vehicle slowed down and came to a full stop prior to the intersection. After a few seconds of waiting at the entrance to the intersection, the ego vehicle finally began crossing it. However, in the case of imperfect perception, i.e., when $\alpha = 0.7$, the ego vehicle remained at rest slightly longer before it started crossing the intersection.



(a) Ego vehicle speed profiles at 5 meter-wide, occluded intersection.

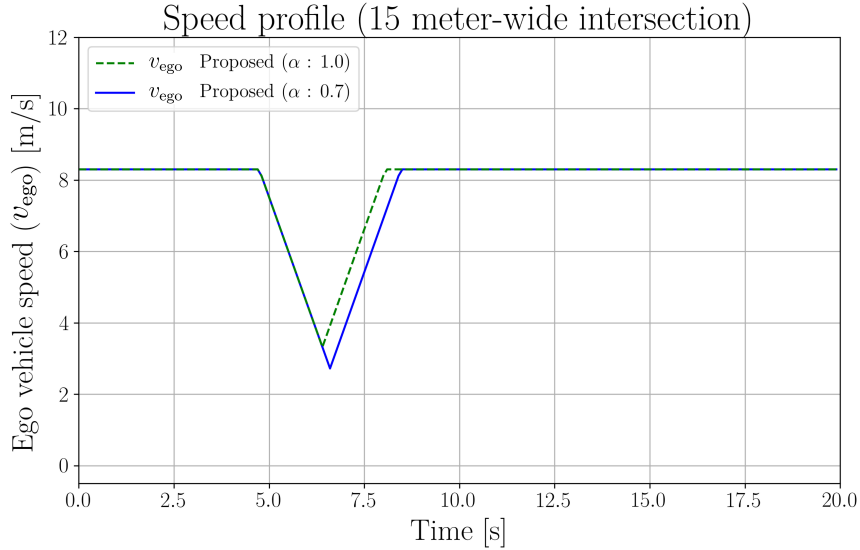


(b) Estimated time of arrival of ego and other vehicles at 5 meter-wide, occluded intersection.


Figure 4-15: Effects of perception noise on proposed planner at 5 meter-wide, occluded intersection.

The effect of noisy perception is evident in Fig. 4-15b. The estimated time of arrival of occluded vehicle, t_{other} , appears to fluctuate more when $\alpha = 0.7$. Moreover, t_{other} began exceeding t_{ego} later than when perception was ideal.

Fig. 4-16a shows the resulting speed profiles for both perception accuracy settings at an occluded intersection of two roads with a width of 15 meters. The output



(a) Ego vehicle speed profiles at 15 meter-wide, occluded intersection.



(b) Estimated time of arrival of ego and other vehicles at 15 meter-wide, occluded intersection.

Figure 4-16: Effects of perception noise on proposed planner at 15 meter-wide, occluded intersection.

motions at both settings are similar, i.e., the ego vehicle slowed down before crossing the intersection without stopping. While both settings output similar motions, when α was set to 0.7 the proposed planner generated a lower minimum speed and had a more delayed crossing action than when $\alpha = 1.0$. Furthermore, as shown in Fig. 4-16b, the estimate of t_{other} is noisier when perception is imperfect.

4.4.3 Effects of sensor mounting position


As described in Section 4.2.2, one factor that significantly affects the ego vehicle’s visibility at blind intersections is the sensor mounting position, X_{sensor} . The purpose of this experiment was to investigate the effects of the sensor mounting position on the motion outputs of the proposed planner.

Two sensor positions were tested in this experiment. The sensor was either mounted close to the vehicle’s center, i.e., $X_{\text{sensor}} = 2.0$ meters, which is consistent with a rooftop-mounted LiDAR unit, or the sensor was simulated to be mounted at the front of the vehicle, i.e., $X_{\text{sensor}} = 0.0$ meters, consistent with a front bumper-mounted sensor. In both cases, α was set to 1.0, representing perfect perception.


As can be seen in Fig. 4-17a, the output speed profiles for the two sensor locations are clearly different when the intersecting roads are 5 meters wide. The green dashed line shows that the ego vehicle came to a full stop and waited before crossing the intersection when the sensor was mounted close to the vehicle’s center, $X_{\text{sensor}} = 2.0$. On the contrary, when the sensor was mounted at the front of the ego vehicle, as illustrated by the solid blue line, the ego vehicle slowed down until it reached a minimum speed of 1.82 m/s before crossing the intersection without stopping.

Variation in the outputs of the proposed planner with different sensor mounting positions can also be observed in Fig. 4-17b. When the sensor was front-mounted, t_{other} increased rapidly as the ego vehicle was about to reach the entrance of the intersection, as depicted by the blue line. Meanwhile, t_{other} remained nearly unchanged when using the center-mounted configuration.

At an intersection of two 15 meter-wide roads, an intersection with relatively better visibility, the output speed profiles of the two sensor configurations are more similar, as can be seen in Fig. 4-18a. In both cases, the ego vehicle slowed down before accelerating again and continuing through the intersection. However, when the sensor was mounted close to the center of the ego vehicle, the vehicle decelerated until it reached a minimum speed of 3.32 m/s, which is 1.50 m/s slower than when the front-mounted sensor configuration was used.




(a) Ego vehicle speed profiles at 5 meter-wide, occluded intersection with sensor at 0 and 2 meters from front bumper.




(b) Estimated time of arrival of ego and other vehicles at 5 meter-wide, occluded intersection with sensor at 0 and 2 meters from front bumper.

Figure 4-17: Comparison of the effect of sensor mounting position on planner output at 5 meter-wide, occluded intersection.

Despite slight variations in timing between the increase in t_{other} and t_{ego} , depending on the location of the sensor unit, the motion planning outputs for both sensor positions were very similar at the intersection of the 15 meter-wide roads, as shown in Fig. 4-18b.



(a) Ego vehicle speed profiles at 15 meter-wide, occluded intersection with sensor at 0 and 2 meters from front bumper.



(b) Estimated time of arrival of ego and other vehicles at 15 meter-wide, occluded intersection with sensor at 0 and 2 meters from front bumper.

Figure 4-18: Comparison of effect of sensor mounting position on planner output at 15 meter-wide, occluded intersection.

4.5 Conclusion

In this chapter, we proposed a generic, deadlock-free motion planner that utilizes the visibility of both the ego and approaching vehicles to generate safe crossing motion

at blind intersections. In order to support possible changes in the behavior of approaching traffic participants, based on their ability to see the ego vehicle, and to account for uncertainty in perception accuracy, the proposed planner utilizes a particle filter algorithm for occluded vehicle prediction. To model the visibility-dependent behavior of occluded vehicles at intersections, real driving data was collected from multiple drivers when crossing blind intersections in a residential area, which was then analyzed. Based on our analysis of the behavior of these vehicles when approaching low-visibility intersections, an approaching vehicle behavior model, dependent on their visibility of the ego vehicle, was introduced.

To validate the ability of our proposed method to overcome the deadlock problem at blind intersections, the proposed planner was compared with a baseline planner that simply assumed other vehicles were approaching from the occluded area at a constant speed. Our comparison of the performance of the proposed and baseline planners in a simulation experiment showed that the proposed planner could generate deadlock-free crossing motion at a blind intersection of two narrow roads (each 5 meters wide), while the baseline planner could not. The effects of perception accuracy and sensor position on the output motion of the planner were also investigated. As for the effects of noisy perception, it was found that inaccuracy in perception generally caused the proposed planner to slightly delay its intersection crossing action. Furthermore, it was found that the sensor mounting position significantly affected the motion output of the proposed planner at intersections of narrow roadways with poor visibility. When the sensor was mounted at the front of the ego vehicle, the vehicle slowed down but did not stop prior to crossing the intersection. In contrast, when a center-mounted sensor configuration was used, the ego vehicle slowed to a complete stop before it started crossing the intersection.

It is important to note that while closed-form expressions were used in this investigation to estimate visibility at blind intersections, a more sophisticated visibility estimation approach, such as the one used in our previous work described in Chapter 3, could be used instead, enabling the proposed planner to deal with complex, real-world environments.

Furthermore, even though this work only considers potential occluded vehicles, other types of traffic participants, such as pedestrians and cyclists, could also be integrated into the occluded object prediction by properly defining their behavior models. Moreover, as the proposed planner offers an approach for considering potential changes in the behavior of occluded traffic participants, it could be extended to cover other situations where more than one behavior by the occluded dynamic objects is possible, e.g., potential deceleration by an occluded cyclist who, after becoming visible, is warned by the ego vehicle's horn. Note also that the concept of changes in the behavior of occluded dynamic objects can be utilized within the POMDP framework as well, to solve for the optimal action in more complex scenarios.

Chapter 5

Visibility Prediction and Quantification Using High Definition Maps

5.1 Introduction

As autonomous driving technology advances, the scope of operation of autonomous vehicles also expands, and they are increasingly expected to navigate not only along relatively consistent, well-defined roadways like expressways, but also in more complex environments such as cities. One of the problems faced when navigating in such complex environments is occlusion which prevents autonomous vehicles from fully observing the traffic environment. With incomplete knowledge of its surroundings, safe motion planning becomes increasingly difficult for the vehicle.

In order to achieve safe, occlusion-aware motion planning, the ability to identify the visible and occluded areas of any given location, based on the vehicle's sensor configuration, is indispensable, not only to estimate visibility at its current location but also to predict visibility along the road ahead. In addition, a method of quantifying these visibility conditions is equally essential for planning the ego vehicle's motion when encountering occlusions, especially its motion in the lateral direction, i.e., road




Figure 5-1: Street view of a section of hilly road close to Nagoya University. The occlusion in this scene is caused not only by the surrounding vertical structures, but also by the contour of road itself.

position planning. Such a quantified approach would allow the ego vehicle to compare two or more locations in terms of visibility and plan its motion accordingly.

Several related studies have used simple 2D shapes to represent objects in driving scenes, in order to use geometry to estimate the visible area of the driving environment from the vehicle's point of view. While computationally inexpensive, such an approach may not be applicable in complex terrains, such as on the hill shown in Fig. 5-1, where occlusions are caused not only by surrounding objects but also by the undulating contour of the road itself.

In this chapter, a method of estimating visibility in such complex driving environments is proposed. In addition, a visibility ratio, which is a number that represents the degree of visibility from a particular location, is introduced. The proposed method uses high-definition (HD) maps, which are a combination of 3D point clouds and road network maps, to estimate visibility at any viewing point within the maps. The proposed approach first approximates a 3D scan at the specified viewing point using a 3D point cloud map and a sensor configuration. The scan is then projected onto an image plane to create a depth image. A similar projection is applied to a set of 3D points representing the area above the surrounding road's surface, which are obtained from a road network map. Subsequently, a depth buffering algorithm is used to determine

which regions of the road network are visible from a specified viewing point. Finally, the visibility ratio is calculated by dividing the visible area of interest by the total driving environment relevant to the location of the viewing point.

The proposed method was tested in both simple, simulated environments and in complex, real-world settings with irregular terrain, and our experimental results show that the proposed method is able to estimate visibility in both settings. Moreover, the visibility ratio was found to be indicative of changes in visibility under several conditions, such as when the vehicle's lateral position was varied at an intersection, or when driving up or down a hill. These experimental results suggest that the proposed method could be used to predict and quantify the visibility conditions of any location contained within the HD maps.

The main contributions of this work are as follows:

- An approach to estimate visible and occluded regions of complex driving environments using HD maps.
- The concept of a visibility ratio that represents the degree of visibility from any location with a numerical value, as well as a practical method of calculating such a value.

The remainder of this chapter is organized as follows: In Section 5.2, the proposed visibility estimation method is described in detail. Experimental conditions and results, in both simulated and real-world driving environments, are presented in Section 5.3, along with a discussion of these results. Finally, in Section 5.4, our conclusions and the limitations of the proposed method are provided.

5.2 Visibility estimation using HD maps

In this section, the proposed visibility estimation algorithm is explained in detail. First, the concept of how to estimate the visible regions of a driving environment using a road network map and a 3D scan is described. Next, an approach to approximate




Figure 5-2: Projection of 3D points onto a 2D image plane. A set of 3D points representing the area of interest is projected onto a depth image. By comparing the resulting depth image with those of other objects in the scene, we can estimate which regions of the area of interest are visible from the specified viewing point.

the required 3D scan from a point cloud map is presented. Finally, the method for calculating a visibility ratio is explained.

5.2.1 Estimation of visible regions

The proposed algorithm depends heavily on the ability to determine which regions of the driving environment are observable from a specified location. The task of identifying whether an object in a scene is visible from a given viewpoint has been well-studied in the field of computer graphics. One of the standard algorithms for solving such problems is the depth buffering or z-buffering algorithm.

The depth buffering algorithm works by first projecting 3D points representing an object's surface, onto a 2D image, \mathcal{D} . The width and height of image \mathcal{D} are fixed and denoted by $\mathcal{D}_{\text{resolution}}^{\text{horizontal}}$ and $\mathcal{D}_{\text{resolution}}^{\text{vertical}}$, respectively. As illustrated in Fig. 5-2, the 3D coordinates $[x, y, z]^T$ of each representative point p on the object's surface are converted into a corresponding pixel position on an image plane $[\bar{u}, \bar{v}]^T$. Each pixel

of image $\mathcal{D}(\bar{u}, \bar{v})$ contains the distance from the viewing point to the object's surface or depth, \bar{d} .

This process is repeated for all of the objects in the scene. Therefore, each 3D object \mathcal{O} has an associated depth image \mathcal{D} . In order to determine whether or not the object of interest, denoted by $\mathcal{O}_{\text{interest}}$, is visible from a particular viewing point, depth image $\mathcal{D}_{\text{interest}}$ is compared with the depth images of other objects, $\mathcal{D}_1, \dots, \mathcal{D}_{\mathcal{N}_{\text{object}}}$, where $\mathcal{N}_{\text{object}}$ is the number of other objects in the scene, a comparison which is performed pixel by pixel. Suppose the depth value of $\mathcal{D}_{\text{interest}}$ at pixel $[\bar{u}, \bar{v}]^\top$ is smaller than those of the other objects at the same pixel position. In that case, the object of interest is visible at coordinates $[x, y, z]^\top$ which corresponds to pixel $[\bar{u}, \bar{v}]^\top$.

In the case of driving environments, the object of interest is the space above the road surface that may be occupied by other vehicles or obstacles. A set of 3D points representing the surface of the lane's center, $\mathcal{P}_{\text{lane}} = \{p_{\text{lane}_1}, \dots, p_{\text{lane}_{\mathcal{N}_{\text{lane}}}}\}$, where $\mathcal{N}_{\text{lane}}$ is the total number of surface points, are first extracted from a road network map. Each point on the lane's surface can be defined as follows:

$$p_{\text{lane}_i} = \begin{bmatrix} x_{\text{lane}_i} \\ y_{\text{lane}_i} \\ z_{\text{lane}_i} \end{bmatrix}. \quad (5.1)$$

A representative point in the area above the lane is then obtained by shifting the z-coordinate of the extracted lane surface point upward by fixed value $\mathcal{H}_{\text{object}}$:

$$p_{\text{interest}_i} = \begin{bmatrix} x_{\text{interest}_i} \\ y_{\text{interest}_i} \\ z_{\text{interest}_i} \end{bmatrix} = \begin{bmatrix} x_{\text{lane}_i} \\ y_{\text{lane}_i} \\ z_{\text{lane}_i} + \mathcal{H}_{\text{object}} \end{bmatrix}. \quad (5.2)$$

By projecting all of the resulting points, $\mathcal{P}_{\text{interest}} = \{p_{\text{interest}_1}, \dots, p_{\text{interest}_{\mathcal{N}_{\text{interest}}}}\}$, onto an image plane, depth image $\mathcal{D}_{\text{interest}}$ is generated. In order to generate depth images of other objects in the environment, $\mathcal{D}_1, \dots, \mathcal{D}_{\mathcal{N}_{\text{object}}}$, surface points of each of those objects are required. However, it is not necessary to create a separate depth image

for each individual object, as the goal is only to determine whether the object of interest, i.e., the area above the lane’s surface, is visible or occluded by any other object in the scene. Consequently, surface points of all the surrounding objects can be collectively projected onto a single depth image designated $\mathcal{D}_{\text{other}}$. Depending on the density of the 3D scan, output depth image $\mathcal{D}_{\text{other}}$ may be sparse. In order to address the sparsity problem, box filtering is applied to $\mathcal{D}_{\text{other}}$. As a result, many of the empty pixels of $\mathcal{D}_{\text{other}}$ are filled.

Filled depth image $\mathcal{D}_{\text{other}}$ is then directly compared with $\mathcal{D}_{\text{interest}}$ to determine the visible regions of the area of interest. Using this approach, if a 3D scan of the environment captured by a LiDAR unit is available, the scan points can be used directly as the surface points of the objects in the scene. Therefore, it is possible to determine which regions of the driving environment are visible by the sensor from the location where the scan was captured from.

5.2.2 3D scan approximation using a point cloud map

As described in 5.2.1, if a road network map and a 3D scan of the local surroundings acquired by a LiDAR sensor are both available, visibility of the driving environment from the sensor’s point of view can be estimated. However, this estimated visibility is the visibility from a single position, i.e., the position of the sensor when the scan was obtained. Therefore, 3D scans from different points of view are needed to estimate visibility for the entire environment.

In order to approximate 3D scans acquired from different locations, a point cloud map is used. Let $\mathcal{P}_{\text{map}} = \{p_{\text{map}_1}, \dots, p_{\text{map}_{\mathcal{N}_{\text{map}}}}\}$ be the available 3D point cloud map, where \mathcal{N}_{map} is the number of points in the map. To simulate a 3D scan captured at a target position p_{target} by a LiDAR unit with the sensing range of $\mathcal{S}_{\text{range}}$, points neighboring p_{target} in \mathcal{P}_{map} within the radius of $\mathcal{S}_{\text{range}}$ are selected to construct a base scan as follows:

$$\mathcal{P}_{\text{base}} = \{p_{\text{map}} \in \mathcal{P}_{\text{map}} \mid \|p_{\text{map}} - p_{\text{target}}\| < \mathcal{S}_{\text{range}}\}. \quad (5.3)$$

Each of the selected points is then transformed to a local coordinate of the simulated sensor as follows:

$$\begin{bmatrix} \mathcal{P}'_{\text{base}i} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{T}_{3 \times 1} \\ \mathbf{0} & 1 \end{bmatrix}^{-1} \begin{bmatrix} x_{\text{base}i} \\ y_{\text{base}i} \\ z_{\text{base}i} \\ 1 \end{bmatrix} \quad (5.4)$$

where $\mathbf{T}_{3 \times 1}$ and $\mathbf{R}_{3 \times 3}$ are a 3D translation vector and a 3D rotation matrix, respectively. $\mathbf{T}_{3 \times 1}$ and $\mathbf{R}_{3 \times 3}$ represent the position and orientation of the LiDAR sensor with respect to the coordinate system of point cloud map \mathcal{P}_{map} , respectively. The transformed point cloud $\mathcal{P}'_{\text{base}}$ contains all of the points within the sensing range of the simulated LiDAR unit. However, some of the points in $\mathcal{P}'_{\text{base}}$ would not have been captured if an actual sensor had been to be placed at the specified viewing point, as some of them would have been outside of the sensor's field of view (FOV) or occluded. These points need to be removed in order to obtain a realistic, simulated 3D scan from the specified location. The removal of these points is done by projecting all of the points in $\mathcal{P}'_{\text{base}}$ onto depth image $\mathcal{D}_{\text{other}}$. The pixel position (\bar{u}, \bar{v}) and the depth value \bar{d} of the projected points can be calculated as follows:

$$\bar{u} = \left\lfloor \arctan\left(\frac{y}{x}\right) \times \frac{\mathcal{D}_{\text{resolution}}^{\text{horizontal}}}{2\pi} \right\rfloor \quad (5.5)$$

$$\bar{v} = \left\lfloor \left(\arctan\left(\frac{z}{\sqrt{x^2 + y^2}}\right) \times \frac{180}{\pi} - \mathcal{S}_{\text{angle}}^{\text{lower}} \right) \times \frac{\mathcal{D}_{\text{resolution}}^{\text{vertical}}}{\mathcal{S}_{\text{FOV}}^{\text{vertical}}} \right\rfloor \quad (5.6)$$

$$\bar{d} = \sqrt{x^2 + y^2 + z^2} \quad (5.7)$$

where $\mathcal{S}_{\text{FOV}}^{\text{vertical}}$ is the vertical field of view of the sensor in degrees, which is equivalent to:

$$\mathcal{S}_{\text{FOV}}^{\text{vertical}} = \mathcal{S}_{\text{angle}}^{\text{upper}} - \mathcal{S}_{\text{angle}}^{\text{lower}}. \quad (5.8)$$

After the projection, only pixels that fall within the depth image, i.e., $\{(\bar{u}, \bar{v}) : 0 \leq \bar{u} < \mathcal{D}_{\text{resolution}}^{\text{horizontal}}, 0 \leq \bar{v} < \mathcal{D}_{\text{resolution}}^{\text{vertical}}\}$, are kept, as they correspond to points within the sensor's field of view. Furthermore, if multiple points fall onto the same pixel




Figure 5-3: Area of interest in the context of driving. The yellow shaded area indicates the range of the sensor (S_{range}). A set of lane segments within the yellow area is defined as $\mathcal{L}_{\text{neighbor}}$, while $\mathcal{L}_{\text{reachable}}$ represents lane segments that share the direction of travel and reachable from p_{target} (highlighted in green). The segments highlighted in red represent $\mathcal{L}_{\text{intersect}}$, which are intersecting lane segments that could potentially be occupied by oncoming vehicles.

position, the point with the minimum depth value is kept since it is not occluded by other points.

The resulting depth image $\mathcal{D}_{\text{other}}$ can then be reverted into a 3D scan. However, for the estimation of visible regions, $\mathcal{D}_{\text{other}}$ can be up-sampled using box filtering and then compared directly with $\mathcal{D}_{\text{interest}}$.

5.2.3 Calculation of visibility ratio

By using the approach described in Sections 5.2.1 and 5.2.2, it is possible to determine which regions are visible from a specified viewing point. However, in some circumstances, it is useful to encapsulate the visibility information for a specified location in a single numerical value, e.g., when comparing the visibility of two locations. Therefore, the concept of a visibility ratio is introduced. A visibility ratio, denoted by \mathcal{V} , is a number representing the degree of visibility from a particular location. It is defined as a ratio between the area visible from the specified location and the total area of interest. The value of a visibility ratio can range from 0 to 1, where 0 indicates the

area of interest is entirely occluded, while 1 indicates the area of interest is entirely visible.

Suppose a visibility ratio for position p_{target} , shown as a blue arrow in Fig. 5-3, is to be calculated. First, the area of interest needs to be defined. Let $\mathcal{L}_{\text{neighbor}}$ be a set of lane segments within the sensor's range ($\mathcal{S}_{\text{range}}$) measured from p_{target} . Under the assumption that traffic participants will not travel against a lane's intended direction, the area of interest, i.e., the relevant area in the context of driving, at position p_{target} , consists of the space above subsets of $\mathcal{L}_{\text{neighbor}}$ which are located on the following lanes:

- Lanes with the same driving direction that are reachable from position p_{target} , denoted by $\mathcal{L}_{\text{reachable}}$.
- Lanes with the opposite driving direction from position p_{target} that intersect any of the $\mathcal{L}_{\text{reachable}}$ lanes, denoted by $\mathcal{L}_{\text{intersect}}$.

Therefore, $\mathcal{L}_{\text{reachable}} \cup \mathcal{L}_{\text{intersect}}$ represents all of the relevant lane segments which are potentially observable from position p_{target} . After the relevant lane segments are identified, a set of 3D points representing the area above those segments, i.e., $\mathcal{P}_{\text{interest}} = \{p_{\text{interest}_1}, \dots, p_{\text{interest}_{\mathcal{N}_{\text{interest}}}}\}$, and a corresponding depth image, $\mathcal{D}_{\text{interest}}$, can be obtained from a road network map as described in Section 5.2.1. Next, a depth image of the environment, $\mathcal{D}_{\text{other}}$, is generated using the approach explained in Section 5.2.2. Let $\mathcal{N}_{\text{interest}}^{\text{visible}}$ be the number of points in $\mathcal{P}_{\text{interest}}$ that are visible from p_{target} . By comparing $\mathcal{D}_{\text{other}}$ with $\mathcal{D}_{\text{interest}}$, $\mathcal{N}_{\text{interest}}^{\text{visible}}$ can be determined. Finally, assuming that the points in $\mathcal{P}_{\text{interest}}$ are uniformly distributed across the surface of the lanes, a visibility ratio \mathcal{V} can be calculated as follows:

$$\mathcal{V} = \frac{\mathcal{N}_{\text{interest}}^{\text{visible}}}{\mathcal{N}_{\text{interest}}}, \quad (5.9)$$

where $\mathcal{N}_{\text{interest}}$ is the total number of points in $\mathcal{P}_{\text{interest}}$.

Table 5.1: Parameters used for visibility estimation.

Parameter	Value	Parameter	Value
$\mathcal{H}_{\text{object}}$	1.5 m	$\mathcal{S}_{\text{range}}$	200.0 m
$\mathcal{D}_{\text{resolution}}^{\text{horizontal}}$	720 pixels	$\mathcal{D}_{\text{resolution}}^{\text{vertical}}$	64 pixels
$\mathcal{S}_{\text{angle}}^{\text{upper}}$	2.0 deg	$\mathcal{S}_{\text{angle}}^{\text{lower}}$	-26.8 deg

5.3 Experiments

The proposed method was tested in both simulated and real-world driving environments. The goal of the first experiment was to verify that the proposed algorithm can determine visibility in relatively simple driving environments, therefore, a simulated, urban area was used. The goal of the second experiment was to confirm that the proposed approach is also applicable to complex, real-world environments, therefore, HD maps of a real-world location were created and used for visibility estimation. The same set of parameters, which are listed in Table 5.1, was used in both experiments.

5.3.1 Evaluation in a simulated environment

As the objective of the first experiment was to validate the proposed method in a simplistic landscape, Town01 of the CARLA simulator was selected as the test environment [78]. Town01 is a simulated urban area with entirely flat terrain. It has a right-hand traffic system that consists of two-lane roads and intersections. The HD maps for Town01, which include both road network and point cloud maps, were obtained from a publicly available source.¹

Visibility along Town01 roadways was estimated using the proposed method, under the assumption that the sensor was located at the center of the lane, parallel to the ground, and pointed directly forward in alignment with the lane’s travel direction. The resulting visibility estimations are shown in Fig. 5-4. The color of a lane indicates visibility at each location, where lighter colors are associated with higher visibility and darker colors with lower visibility. As can be seen in Fig. 5-4, straight

¹HD maps of Town01 can be downloaded from <https://bitbucket.org/carla-simulator/autoware-contents/src/master/maps/>





Figure 5-4: Visibility of the driving environment in CARLA’s Town01 as estimated using the proposed method. Lighter colors represent areas with higher visibility, while darker colors indicate lower visibility locations.

road segments tend to have higher visibility than areas close to intersections or corners. Apart from the distance to junctions or curves, a lane’s direction of travel also affects how visibility changes. As indicated by the color gradient in the figure, visibility gradually decreases when a lane approaches a junction or corner, until reaching the area immediately preceding an intersection or corner, where it starts to increase again. As a vehicle travels away from a junction or a curve towards a straight segment of the road, visibility increases rapidly.

This is an expected result, as the terrain is completely flat. In such an environment, occlusions are likely to be caused only by vertical objects such as buildings or walls located near intersections or corners. The more cluttered the area, the lower its



(a) Original intersection.



(b) Modified (occluded) intersection.

Figure 5-5: Comparison of visibility at the original and modified reference intersections in CARLA’s Town01. The modified intersection has additional objects, i.e., kiosks and a bin, placed near the corner of the building to limit visibility.

visibility tends to be.

In order to investigate the effects of occluding objects on visibility near intersections, an additional experiment was conducted. The intersection highlighted by the red rectangle in Fig. 5-4 was selected as a reference location. The selected intersection was then modified by adding three small kiosks and a bin close to the intersection to further obstruct the view. Fig. 5-5a shows the reference intersection from the original Town01, and the modified intersection is shown in Fig. 5-5b.

The point cloud map of the modified Town01 was created using an open-source




Figure 5-6: Comparison of visibility at the original and modified intersections shown in Fig. 5-5.


tool.² Visibility estimation was then performed using the proposed algorithm under the same assumptions as for the previous experiment. Fig. 5-6 shows the visibility of the lane leading to the reference intersection, i.e., the right-hand lane shown in the red rectangle in Fig. 5-4, in both the original and modified Town01 intersections. The dashed blue line in Fig. 5-6 represents visibility at the modified intersection, revealing that visibility remains significantly lower prior to the modified intersection, i.e., at $y = -215\text{ m}$ to $y = -205\text{ m}$, compared to visibility prior to the original intersection, represented by the solid blue line. Note that the x and y values given in the rest of this chapter, are in reference to a Cartesian grid used to identify locations in Town01 as a whole, and do not represent distances from corners or intersections, etc.

In the previous experiments, the sensor was assumed to be located at the exact center of the lane. Consequently, the visibility estimated in those experiments represents the visibility of the driving environment a vehicle will have at each location

²The tool for point cloud map creation is available at https://github.com/carla-simulator/ros-bridge/tree/master/pcl_recorder



(a) Visibility at original intersection.



(b) Visibility at modified intersection.

Figure 5-7: Effects of lateral position shift on ego vehicle visibility at the original and modified intersections shown in Fig. 5-5.

if it stays at the center of the lane, assuming its sensor is mounted at the center of the vehicle. However, if a vehicle deviates from the lane’s center, its visibility will be altered. To examine the effects of lateral shift from the lane’s center on ego vehicle visibility, the visibility of the selected intersection was also estimated under the condition that there is a 2-meter deviation from the center of the lane, to either the left or right.

Fig. 5-7a shows visibility from the perspective of the ego vehicle at the original reference intersection from three different lateral positions (lane center, 2 meters to the left of center and 2 meters to the right of center). As can be seen from the figure, visibility tends to be lowest when the vehicle is located on the right side of the lane, especially in the area close to the entrance of the intersection. Moreover, keeping left appears to result in slightly better visibility than staying in the middle of the lane at this intersection. Similar patterns were also found at the modified (occluded) intersection, as shown in Fig. 5-7b. In the case of the modified intersection, where more occluding objects are present, the visibility gain from keeping left is most apparent right before the intersection, i.e., at $y = -207\text{ m}$ to $y = -205\text{ m}$.

This improvement in visibility when driving on the left side of the lane is likely caused by the increased line-of-sight angle, which is the angle measured horizontally from the forward direction of the sensor to the edge of the closest occluding object. A larger line-of-sight angle generally translates into a broader field of view. As the occluding objects at the reference intersection were on the right side, moving to the left increased the line-of-sight angle and, consequently, visibility along the intersecting roadway.

5.3.2 Evaluation in a real-world environment

In order to verify that the proposed algorithm is also applicable to complex, real-world environments, a second experiment was conducted using a location in a real neighborhood near Nagoya University, in Nagoya, Japan. Point cloud and road network maps of the area were created using the Mobile Mapping System (MMS) developed by Mitsubishi Electric Corporation.




Figure 5-8: Road network map of an area near Nagoya University, where colors indicate road elevation.

Although an MMS was used to create point cloud and road network maps in this experiment, it is worth noting there are other ways to generate HD maps. Point cloud maps could be generated by using Simultaneous localization and mapping (SLAM) algorithms [79, 80]. SLAM algorithms mainly utilize LiDAR scans captured from the environment to estimate the mapping platform’s trajectory during the data collection. With the estimated trajectory, captured scans can be transformed to the correct coordinates and then combined to create a point cloud map. Road network maps can be generated automatically by extracting information such as road and lane boundaries from the point cloud maps of the same environment [81, 82]. However, more often than not, road network maps are created manually using specialized software, which allows users to use the point cloud maps of the same driving environment as




Figure 5-9: Visibility map of a residential area near Nagoya University (detail of area enclosed in dashed red rectangle in Fig. 5-8).

references.³

The road network map is shown in Fig. 5-8, where colors are used to indicate the elevation of the road. Unlike the simulated urban environment in the previous experiment, it is apparent from this map that the roads in this area are not all level, as a significant change in elevation can be observed in several locations.

A section of this real-world neighborhood with some resemblance to the environment used in the previous experiment was first investigated. The targeted region, indicated by the dashed red rectangle in Fig. 5-8, is a residential area that consists of narrow roads, intersections and buildings.⁴ As can be seen from the road network

³Commonly used free software for creating road network maps include VectorMap-Builder (<https://tools.tier4.jp>) and ASSURE mapping tools (<https://github.com/hatem-darweesh/assuremappingtools>)

⁴A satellite view of the targeted residential area can be found at <https://www.google.com/maps/@35.1617009,136.957285,235m/data=!3m1!1e3>

map, the selected residential area is relatively flat, similar to Town01 of the CARLA simulator.

Fig. 5-9 shows visibility in the selected area as estimated using the proposed method. The arrows in the figure indicate a lane’s direction of travel, while the colors of the arrows indicate the visibility from that location. Similar to the results shown in Fig. 5-4, visibility slowly decreases as a lane approaches intersections and quickly increases as the lane exits intersections and continues along straight segments of the roads. Moreover, overall visibility in the selected area appears to be low, as was expected since the area is cluttered with buildings, while the straight segments of the road are relatively short.

Additionally, in order to investigate the effects of uneven terrain on visibility, a hilly segment of the road, located in the area enclosed by the dashed blue rectangle in Fig. 5-8, was chosen as another region of interest. Fig. 5-1 is a photograph of the selected road segment. In contrast to CARLA Town01, Japan uses a left-hand driving traffic system, therefore, the photograph illustrates the surroundings from the vehicle’s point of view as it travels uphill in the left lane.⁵ As can be seen in the photograph in Fig. 5-1, visibility further along the road is obscured by the undulating contour of the road itself, with zero visibility of the other side of the hill.

Visibility at the selected hill was estimated using the proposed algorithm, and is shown in Fig. 5-10, where it can be observed that visibility gradually decreases as the road approaches the top of the hill, but as the road begins to go downhill visibility increases sharply. The visibility profile of both lanes, along with the hill’s elevation profile, are shown in Fig. 5-11. The visibility profiles reveal that there also appears to be a significant drop in visibility around $y = -94150\text{ m}$ to $y = -94100\text{ m}$ in both lanes.

These changes in visibility were mainly caused by the change in the road’s slope. If a vehicle is traveling uphill in the left lane, as illustrated in Fig. 5-1, its visibility profile will be similar to the blue line in Fig. 5-11. As the vehicle approaches the

⁵The Google Streetview of the selected hill can be found at <https://goo.gl/maps/eW1V2uzJEW3GP2Gt5>




Figure 5-10: Visibility map of a hill near Nagoya University (detail of area enclosed in dashed blue rectangle in Fig. 5-8).

bottom of the hill, at $y = -94200\text{ m}$, its sensor will remain almost horizontal as the area is relatively flat. The farthest position it can observe from this location will be the bottom of the hill, i.e., $y = -94150\text{ m}$. Since the farthest position that is visible will not change until the vehicle begins to go uphill and its sensor starts to point upward, its visibility will decrease as it gets closer to the bottom of the hill. As soon as the vehicle starts to go uphill, its visibility will increase sharply, as the farthest position it can now observe has become the top of the hill. After that, the vehicle's visibility will gradually decrease again until it reaches the top of the hill,




Figure 5-11: Visibility and elevation plot of the hill shown in Fig. 5-10. Left and right lanes are defined according to the driving scenario shown in Fig. 5-1.

but will finally increase once more as it starts descending the other side of the hill. A similar but opposite sequence of events will happen to vehicles traveling in the opposite direction, as indicated by the orange line in Fig. 5-11.

5.4 Conclusion

In this chapter, we have proposed a visibility estimation method that can be used in complex, real-world driving environments, as well as a method of quantifying visibility at a particular location in the driving environment. The proposed visibility estimation method utilizes HD maps that include both 3D point clouds and road network maps. The 3D point cloud map is used to approximate a 3D scan taken from a specified viewpoint. The approximated scan and 3D points representing the area of relevant surrounding lanes are projected onto depth images. The resulting depth images are then compared in order to determine the visible and occluded regions of the driving environment from that particular viewpoint. The proposed approach

was experimentally verified to be effective in both simulated and real-world driving environments, including a real driving environment with complex, hilly terrain.

The proposed visibility ratio, which is a numerical value representing the degree of visibility from a particular location, is computed by dividing the visible area of interest by the total relevant driving area. Our experimental results suggest that the visibility ratio calculated using the proposed method is representative of actual road visibility from a specified location.

As the proposed method is able to identify the areas of a driving environment that are visible and occluded from a specified viewpoint, it could be used in the motion planning stage to determine occluded regions of the surrounding roads from the ego vehicle’s current point of view. The information regarding occluded areas could then be used for predicting unobserved traffic participants that may pose a collision risk to the ego vehicle. Additionally, since the visibility ratio is representative of the degree of visibility from a particular location, it could be used as one of the costs of an optimization-based motion planner, so that the vehicle would not only avoid collision with observable objects but also reduce its potential risk of colliding with hidden objects by actively gaining more visibility of the traffic environment via active road position adjustment.

Since the proposed method relies on HD maps, its accuracy is highly dependent on the quality of the information available in those maps. Furthermore, it cannot detect occlusions caused by dynamic objects which are not included in the point cloud maps. Moreover, there are other sources of occlusion apart from objects present in a driving scene, such as adverse weather conditions, which have been reported to cause degradation in sensor performance in several studies [83–86]. The dynamic object problem could potentially be mitigated by leveraging detection results from the perception module. As for limitations due to adverse weather, sensor fusion seems to be one of the more promising solutions. However, both problems remain to be investigated and resolved.

Chapter 6

Road Position Adjustment for Maximizing Visibility

6.1 Introduction

In complex driving scenarios, ranging from urban streets full of obstacles and low-visibility intersections, to hilly and winding roads in rural areas, it is almost impossible for autonomous vehicles to obtain complete information about the driving environment at all times due to occlusions, i.e., sensor blind spots caused by obstructions. In order to navigate safely in areas with occlusions, the visibility conditions of the driving environment must be considered during the motion planning stage.

The topic of visibility-aware, or occlusion-aware, motion planning has gained increasing attention in recent years. However, despite the steadily growing body of research on motion planning in areas with limited visibility, the majority of existing studies have focused almost exclusively on longitudinal motion planning, in which the autonomous vehicle is either assumed to be driving exclusively in the center of a lane or along a pre-generated trajectory. In many situations, by strategically making the autonomous vehicle deviate from the center of its lane, occlusion can be significantly reduced, i.e., visibility of the driving environment can be increased, increasing the ego vehicle's chances of discovering hidden traffic participants. Therefore, ignoring road position adjustment during motion planning can result in inefficient output motion.

A few studies have considered optimizing the ego vehicle’s lateral position, or road position, in order to proactively improve visibility. However these approaches tend to be scenario-specific, e.g., they are designed specifically for intersection crossing [34, 43] or passing parked vehicles [65]. Another component missing from existing approaches is the use of highly accurate visibility estimation. In most of these studies, visible region estimation is typically performed by representing surrounding obstacles with simple 2D geometric shapes and then performing ray-casting. While using 2D representation is computationally efficient, it results in inaccurate estimation of visibility in areas with complexly-shaped obstacles or undulated terrain. Lastly, uncertainty within the perception modules used for localization and object detection is rarely taken into consideration.

In this chapter, we propose a motion planner capable of actively adjusting the ego vehicle’s lateral position within its lane, i.e., its road position, to minimize occlusions, which is achieved by estimating current and future visibility. The planner proposed in this chapter is applicable in various scenarios and tolerates realistic uncertainty from perception modules. Additionally, this work extends the 3D visibility estimation approach described in Chapter 5, which allows the proposed system to handle occlusions caused by objects that are not present in high-definition (HD) maps by incorporating live sensing data from the onboard LiDAR unit during the estimation process.

The road position adjustment-capable motion planner proposed in this chapter first generates several physically feasible trajectories for the ego vehicle from its current position, using different lateral offsets from the reference path. Visibility along each candidate trajectory is then estimated in 3D using HD maps and live sensing data, in order to achieve accurate estimation in complex environments. Subsequently, the visibility conditions of each trajectory are quantified and converted into a visibility cost. Finally, the optimal trajectory is selected from the generated candidates based on visibility and other planning costs.

The proposed planner was tested in the CARLA simulator [78] in three different scenarios where occlusions were present. These test scenarios consisted of crossing a

T-junction, making a right turn, and preparing to pass a parked vehicle not present in the HD maps. In order to verify that the proposed planner is practical under conditions of perception uncertainty, actual localization and object detection results obtained in real-time were used. Our experimental results showed that the proposed planner could generate a trajectory that allowed the ego vehicle to minimize occlusions and consequently detect occluded traffic participants earlier than the baseline planner in most cases. Furthermore, our improved visibility estimation approach was able to accurately estimate occlusions caused by a parked vehicle, even though it was not present in the HD maps.

The main contributions of this work are as follows:

- A motion planner capable of actively adjusting the ego vehicle’s lateral position to minimize occlusions, which is applicable in various scenarios and under conditions where perception uncertainty is present.
- Extension of our previously proposed visibility estimation approach for handling occlusions caused by objects which do not appear in HD maps, by incorporating live sensing data.

The remainder of this chapter is organized as follows: In Section 6.2, details of our proposed motion planner are provided, including conversion of visibility conditions into trajectory costs, as well as optimal trajectory selection. In Section 6.3, details of our approach for estimation of visibility at specified locations using 3D HD maps and live sensing data are provided. In Section 6.4, the details of the experimental procedure used to assess the proposed planner are given. Our experimental results when using the proposed and baseline planners in multiple variations of the three scenarios described above, i.e., T-junction crossing, turning right, and preparing to pass a parked vehicle, are presented in Sections 6.4.1, 6.4.2, and 6.4.3, respectively. Finally, the conclusions of this investigation are provided in Section 6.5.

6.2 Proposed motion planner

This section provides the operational details of our proposed motion planner, which is based on OpenPlanner [21], an open-source, integrated planner used for mobile robot navigation in highly dynamic environments. OpenPlanner includes various components such as a global planner, a trajectory generator and a behavior state generator.¹ The difference between our proposed planner and OpenPlanner is mainly in the local planning stage, where our planner generates a set of candidate trajectories and then selects the best trajectory. Additionally, the proposed planner includes a visibility estimator that can quantify the visibility conditions of a specified location. Unlike the local planner in OpenPlanner, the proposed planner also considers visibility along the candidate trajectories, in addition to other standard planning criteria such as the risk of collision with detected obstacles, lane center deviation and stability.

The overall architecture of our proposed planner is described in Section 6.2.1. In Section 6.2.2, integration of the visibility cost of the candidate trajectories, which is the primary improvement over the original OpenPlanner, is presented in detail.

6.2.1 Planner architecture

Fig. 6-1 shows the overall architecture of our proposed motion planner, which consists of three main components: a global planner, a local planner, and a visibility estimator. Since the proposed planner shares a similar design with OpenPlanner, some of its components, namely the global planner and the behavior state generator inside the local planner, are adopted into the planner proposed in this work. These components are shown with dashed boundaries in Fig. 6-1.

The objective of the global planner is to find the optimal route from the vehicle's starting position to its destination. A road network map, which contains information on lanes, intersections, traffic directions, speed limits, and traffic signs, is used by the global planner to construct a routing graph. Consequently, dynamic programming is

¹The latest source code of OpenPlanner can be found at <https://github.com/hatem-darweesh/autoware.ai.openplanner>

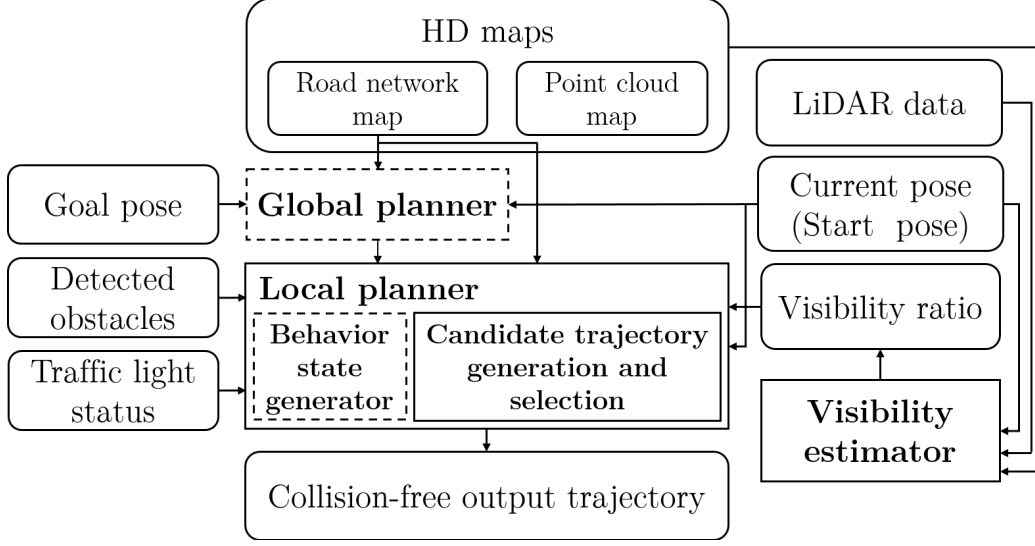


Figure 6-1: Architecture of the proposed motion planner.

applied to calculate the best route from the starting vertex to the goal vertex in the routing graph.

After the global route is determined, it is utilized by the local planner as a reference for creating candidate trajectories, i.e., rollout generation. The local planner generates a total of $\mathcal{N}_{\text{rollout}}$ rollouts along the global reference path. The generated rollouts are denoted by $\{\mathcal{R}_r\}_{r=1}^{\mathcal{N}_{\text{rollout}}}$. All rollouts are parallel to the global path and are spaced evenly to each side of the reference path, with each rollout being $\mathcal{W}_{\text{rollout}}$ apart from the adjacent candidate trajectories. The rollouts originate from the ego vehicle’s current position and extend forward longitudinally to the end of the planning horizon ($\mathcal{H}_{\text{planning}}$). Fig. 6-2 shows an example of rollouts generated by the local planner.

Each rollout has associated costs, including a center cost (c_{center}), transition cost ($c_{\text{transition}}$), longitudinal collision cost ($c_{\text{collision}}^{\text{lon}}$), lateral collision cost ($c_{\text{collision}}^{\text{lat}}$), and visibility cost ($c_{\text{visibility}}$). Each cost indicates how good or bad a candidate trajectory is relative to other candidate trajectories regarding that particular aspect of the trajectory. A summation of the costs for all of the generated rollouts in a particular cost category is equal to one. The calculation of common planning costs, i.e., c_{center} , $c_{\text{transition}}$, $c_{\text{collision}}^{\text{lon}}$, $c_{\text{collision}}^{\text{lat}}$, is the same as described in [21]. Visibility cost $c_{\text{visibility}}$




Figure 6-2: Generation of candidate trajectories.

is calculated using visibility ratio \mathcal{V} , a numerical value that quantifies the visibility condition at a specified location from the perspective of the ego vehicle, as proposed in Chapter 5 and improved upon in this work. As shown in Fig. 6-1, the visibility estimator takes a road network map, a point cloud map, the ego vehicle's current pose, and live LiDAR data as inputs, and outputs visibility ratio \mathcal{V} , which is then utilized by the local planner. More information about the visibility estimator is provided in Section 6.3. A weighted sum of all of the costs associated with each rollout (\mathcal{C}) is calculated as follows:

$$\mathcal{C} = w_{\text{center}}c_{\text{center}} + w_{\text{transition}}c_{\text{transition}} + w_{\text{collision}}^{\text{lon}}c_{\text{collision}}^{\text{lon}} + w_{\text{collision}}^{\text{lat}}c_{\text{collision}}^{\text{lat}} + w_{\text{visibility}}c_{\text{visibility}} \quad (6.1)$$

where w_{center} , $w_{\text{transition}}$, $w_{\text{collision}}^{\text{lon}}$, $w_{\text{collision}}^{\text{lat}}$, and, $w_{\text{visibility}}$ are normalized weights. Finally, the rollout with the lowest \mathcal{C} is selected as the output trajectory.

After the output trajectory is selected, the behavior state generator, adopted from [21], determines the output action, i.e., going forward along the same trajectory, stopping or changing the trajectory, based on information about surrounding obstacles and traffic light status.

6.2.2 Visibility cost calculation

The visibility cost of each rollout ($c_{\text{visibility}}$) is calculated from the visibility ratios \mathcal{V} of positions sampled evenly along the rollout. Let a total of $\mathcal{N}_{\text{sampling}}$ sampled locations along a rollout \mathcal{R}_r be $\{l_{r,s}\}_{s=1}^{\mathcal{N}_{\text{sampling}}}$, where $l_{r,1}$ represents the sampled location closest to the current position. The visibility ratio of $l_{r,s}$ is denoted by $\mathcal{V}_{r,s}$, and can be calculated using the approach detailed in Section 6.3. This visibility ratio is positively correlated with visibility conditions, i.e., the higher the value, the better the ego vehicle’s visibility at location $l_{r,s}$. However, visibility cost $c_{\text{visibility}}$ has an inverse relationship with the visibility condition. Therefore, visibility ratio $\mathcal{V}_{r,s}$ is first converted into occlusion ratio $\mathcal{O}'_{r,s}$ using the following equation:

$$\mathcal{O}'_{r,s} = 1 - \mathcal{V}_{r,s}. \quad (6.2)$$

The value of $\mathcal{O}'_{r,s}$ also ranges from zero to one, similar to $\mathcal{V}_{r,s}$. However, higher values of $\mathcal{O}'_{r,s}$ are associated with a higher degree of occlusion.

Theoretically, if $l_{r,s}$ is located on a straight, level segment of road without any obstacles, the value of $\mathcal{O}'_{r,s}$ should be zero. However, since $\mathcal{O}'_{r,s}$ is estimated using live sensing data and point cloud maps, which can be noisy, $\mathcal{O}'_{r,s}$ is sometimes a very small number close to zero. Similarly, there are also cases where $\mathcal{O}'_{r,s}$ is very close to one, despite $l_{r,s}$ being completely occluded. In order to mitigate the noise issue in practice, extreme values of $\mathcal{O}'_{r,s}$ are squashed to zero or one using the following

equation:

$$\mathcal{O}_{r,s} = \begin{cases} 0, & \mathcal{O}'_{r,s} < \mathcal{T}_{\text{lower}} \\ \mathcal{O}'_{r,s}, & \mathcal{T}_{\text{lower}} \leq \mathcal{O}'_{r,s} \leq \mathcal{T}_{\text{upper}} \\ 1, & \mathcal{O}'_{r,s} > \mathcal{T}_{\text{upper}} \end{cases} \quad (6.3)$$

where $\mathcal{T}_{\text{lower}}$ and $\mathcal{T}_{\text{upper}}$ are the lower and upper thresholds of $\mathcal{O}'_{r,s}$, respectively. Squashed $\mathcal{O}'_{r,s}$ values are denoted by $\mathcal{O}_{r,s}$.

In order to calculate the visibility cost of an entire rollout \mathcal{R}_r , a discounted sum of $\mathcal{O}_{r,s}$ for all of the sampled locations along the rollout ($c'_{\text{visibility}_r}$) is first calculated as follows:

$$c'_{\text{visibility}_r} = \sum_{s=1}^{\mathcal{N}_{\text{sampling}}} \gamma^{s-1} \mathcal{O}_{r,s}, \quad (6.4)$$

where $\gamma \in (0, 1)$ is a discount factor used to account for an increasing level of uncertainty in the calculated value of $\mathcal{O}_{r,s}$ along rollout \mathcal{R}_r . The further the sampled position is from the current position, i.e., the higher the s index, the more delay there is before the ego vehicle will actually reach that position. During the delay, circumstances within the surrounding traffic environment may change, e.g., occluding obstacles may move, or new obstacles may be detected. Therefore, the significance of $\mathcal{O}_{r,s}$ is discounted by γ . The closer to zero the value of γ , the lower the impact of $\mathcal{O}_{r,s}$ at distant positions on overall rollout visibility cost $c'_{\text{visibility}_r}$. This overall rollout visibility cost $c'_{\text{visibility}_r}$ is further normalized to obtain the final visibility cost of a rollout ($c_{\text{visibility}_r}$), such that the sum of $c_{\text{visibility}_r}$ over all generated rollouts is equal to one. The normalization process begins by identifying a set of rollouts with the lowest $c'_{\text{visibility}_r}$. This set of rollouts is denoted by \mathcal{M} , and can be formally defined as follows:

$$\mathcal{M} = \{c'_{\text{visibility}_r} \mid c'_{\text{visibility}_r} = c^*_{\text{visibility}}\}, \quad (6.5)$$

where

$$c^*_{\text{visibility}} = \min\{c'_{\text{visibility}_r}\}_{r=1}^{\mathcal{N}_{\text{rollout}}}. \quad (6.6)$$

The value of visibility cost $c_{\text{visibility}_r}$ can then be calculated using the following equa-

tion:

$$c_{\text{visibility}_r} = \begin{cases} 0, & c'_{\text{visibility}_r} \in \mathcal{M} \\ (\mathcal{N}_{\text{rollout}} - |\mathcal{M}|)^{-1}, & c'_{\text{visibility}_r} \notin \mathcal{M} \end{cases}. \quad (6.7)$$

Finally, $c_{\text{visibility}_r}$ is used in Eq. 6.1 to calculate the overall cost of a rollout. In special cases where every rollout has an indistinguishable visibility condition, i.e., where $c'_{\text{visibility}_r}$ is equal for all r , all of the outputs of Eq. 6.7 automatically become zero. Thus, the visibility cost does not have any impact on the overall rollout cost calculation in Eq. 6.1.

6.3 Visibility estimator

In this section, our visibility estimator is explained in detail. The goal of the visibility estimator is to calculate a numerical value that is indicative of the visibility condition of an input location. The visibility estimation method used in this chapter is based on an algorithm previously presented in Chapter 5 that leverages the information available in HD maps, i.e., a combination of a road network map and a point cloud map. The main improvement introduced here, over the method proposed in Chapter 5, is the ability to take into account occlusions caused by objects that are not present in the point cloud map, which is achieved by incorporating a 3D scan captured in real time by a LiDAR sensor.

Our approach for identifying the visible regions of surrounding roads using a 3D scan and a road network map is first explained in Section 6.3.1. In Section 6.3.2, a method for estimating the required local 3D scan from a point cloud map and a live 3D scan is then described. Finally, a method of calculating a visibility ratio from the estimated visible region is presented in Section 6.3.3.

6.3.1 Estimation of visible regions

Regions visible from a specified location are estimated using the z-buffering (depth buffering) algorithm introduced in Chapter 5. As shown in Fig. 6-3, the algorithm first projects a set of 3D points that are representative of the surfaces of surrounding




Figure 6-3: Projection of 3D points onto a 2D image plane. A set of 3D points representing the area of interest is projected onto a depth image. By comparing the resulting depth image with those of other objects in the scene, regions visible from the specified viewing point can be estimated.

objects onto a 2D image plane. The 3D coordinates of each surface point $[x, y, z]^T$ are then converted into a pixel position $[\bar{u}, \bar{v}]$ and a pixel depth value \bar{d} of a depth image \mathcal{D} that has a fixed pixel width and height of $\mathcal{D}_{\text{resolution}}^{\text{horizontal}}$ and $\mathcal{D}_{\text{resolution}}^{\text{vertical}}$, respectively. Pixel position (\bar{u}, \bar{v}) and depth value \bar{d} of the projected point can be calculated as follows:

$$\bar{u} = \left\lceil \arctan\left(\frac{y}{x}\right) \times \frac{\mathcal{D}_{\text{resolution}}^{\text{horizontal}}}{\mathcal{S}_{\text{FOV}}^{\text{horizontal}}} \right\rceil \quad (6.8)$$

$$\bar{v} = \left\lceil \left(\arctan\left(\frac{z}{\sqrt{x^2 + y^2}}\right) - \mathcal{S}_{\text{angle}}^{\text{lower}} \right) \times \frac{\mathcal{D}_{\text{resolution}}^{\text{vertical}}}{\mathcal{S}_{\text{FOV}}^{\text{vertical}}} \right\rceil \quad (6.9)$$

$$\bar{d} = \sqrt{x^2 + y^2 + z^2}, \quad (6.10)$$

where $\mathcal{S}_{\text{FOV}}^{\text{horizontal}}$ and $\mathcal{S}_{\text{FOV}}^{\text{vertical}}$ are the horizontal and vertical fields of view, respectively, in radians. The variable $\mathcal{S}_{\text{angle}}^{\text{lower}}$ denotes the elevation angle with respect to the ground plane, from the specified viewing point to the lowest 3D point. During the projection, only the points that fall within the boundaries of the depth image, i.e., $\{(\bar{u}, \bar{v}) \mid$

$\bar{u} \in [0, \mathcal{D}_{\text{resolution}}^{\text{horizontal}}), \bar{v} \in [0, \mathcal{D}_{\text{resolution}}^{\text{vertical}})$, are kept. Moreover, if more than one point is projected onto the same pixel position, only the point with the lowest depth value is preserved, as it is not occluded by other points.

Visible areas of an object of interest with an associated depth image $\mathcal{D}_{\text{interest}}$ can subsequently be identified by comparing $\mathcal{D}_{\text{interest}}$ with a depth image of other surrounding objects $\mathcal{D}_{\text{other}}$. This comparison is performed pixel-wise. If the depth value \bar{d} of $\mathcal{D}_{\text{interest}}$ in pixel $[\bar{u}, \bar{v}]^\top$ is smaller than that of the other depth images located at the same pixel position, surface point $[x, y, z]^\top$ of the object of interest, which corresponds to pixel $[\bar{u}, \bar{v}]^\top$, is visible from the specified viewing location.

In driving environments, the space above the surfaces of the surrounding lanes that other traffic participants or obstacles may occupy is considered the object of interest for visible region estimation. Therefore, a set of 3D points that represents the surfaces of those lanes is first extracted from a road network map. This set of points is denoted by $\mathcal{P}_{\text{lane}} = \{p_{\text{lane}_i}\}_{i=1}^{\mathcal{N}_{\text{lane}}}$ where $\mathcal{N}_{\text{lane}}$ is the total number of lane surface points. Each point in the set can be defined as follows:

$$p_{\text{lane}_i} = \begin{bmatrix} x_{\text{lane}_i} \\ y_{\text{lane}_i} \\ z_{\text{lane}_i} \end{bmatrix}. \quad (6.11)$$

Representative points of the area above the surfaces of the lanes can then be obtained by offsetting the z-coordinate of each lane surface point upward by fixed value $\mathcal{H}_{\text{object}}$:

$$p_{\text{interest}_i} = \begin{bmatrix} x_{\text{interest}_i} \\ y_{\text{interest}_i} \\ z_{\text{interest}_i} \end{bmatrix} = \begin{bmatrix} x_{\text{lane}_i} \\ y_{\text{lane}_i} \\ z_{\text{lane}_i} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \mathcal{H}_{\text{object}} \end{bmatrix}. \quad (6.12)$$

The resulting points $\mathcal{P}_{\text{interest}} = \{p_{\text{interest}_i}\}_{i=1}^{\mathcal{N}_{\text{lane}}}$ are subsequently projected onto an image plane to create depth image $\mathcal{D}_{\text{interest}}$. Similarly, surface points of all surrounding objects are projected onto a depth image $\mathcal{D}_{\text{other}}$. Since $\mathcal{D}_{\text{other}}$ can be sparse, depending on its source, a box filter is utilized to fill the empty pixels of $\mathcal{D}_{\text{other}}$, based on the

values of their neighboring pixels. Finally, the filled $\mathcal{D}_{\text{other}}$ is compared with $\mathcal{D}_{\text{interest}}$ to determine which parts of the surrounding lanes are visible from the specified viewing location.

If a 3D scan of the driving environment, taken by a LiDAR unit from the specified viewing point, is available, the captured scan points can be used as the surface points of the objects in the scene. In that case, depth image $\mathcal{D}_{\text{other}}$ can then be generated directly. As a result, this approach can be used to estimate the visible regions of the surrounding lanes from any viewing location within the road network map if a 3D scan captured from that location is available.

6.3.2 3D scan estimation

Visibility from any position within a road network map can be estimated if a 3D scan of the surrounding area captured at that position is available. Estimating regions that are visible from the ego vehicle’s current position is also straightforward, as the 3D scan can be directly acquired from the onboard LiDAR unit. However, as described in Section 6.2.2, the calculation of the visibility cost of a trajectory candidate requires information regarding visibility conditions at future ego vehicle positions. Therefore, the ability to predict 3D scans at those future positions is essential.

In order to estimate 3D scans at future ego vehicle positions, a point cloud map and a scan obtained from the LiDAR unit at the current position are used. The incorporation of the live scan acquired at the vehicle’s current position is the main distinction between the approach proposed here and the algorithm proposed in Chapter 5. The live scan is used to account for objects that were not present during point cloud map generation.

First, the point cloud map is used to obtain a base scan at a specified future location. Let $\mathcal{P}_{\text{map}} = \{p_{\text{map}_i}\}_{i=1}^{\mathcal{N}_{\text{map}}}$ be the input 3D point cloud map, where \mathcal{N}_{map} represents the total number of points in the map. In order to simulate a 3D scan obtained at a targeted future position p_{target} by a LiDAR unit with the sensing range of $\mathcal{S}_{\text{range}}$, the 3D points in \mathcal{P}_{map} that are within range $\mathcal{S}_{\text{range}}$ from p_{target} are denoted

as base scan points and can be defined as follows:

$$\mathcal{P}_{\text{base}} = \{p_{\text{map}} \in \mathcal{P}_{\text{map}} \mid \|p_{\text{map}} - p_{\text{target}}\| < \mathcal{S}_{\text{range}}\}. \quad (6.13)$$

The resulting base scan $\mathcal{P}_{\text{base}}$ cannot be used directly for visible region estimation as it only contains points from objects that were present during map generation.

In order to add objects that do not appear in the point cloud map to the estimated scan, a live scan captured by the LiDAR unit at its current position p_{current} is used. Each point in the live scan, $[x'_{\text{live}i}, y'_{\text{live}i}, z'_{\text{live}i}]^\top$, is first transformed to point cloud map coordinates using the following equation:

$$\begin{bmatrix} p_{\text{live}i} \\ 1 \end{bmatrix} = \begin{bmatrix} {}^{\text{map}}_{\text{current}}\mathbf{R}_{3 \times 3} & {}^{\text{map}}_{\text{current}}\mathbf{T}_{3 \times 1} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} x'_{\text{live}i} \\ y'_{\text{live}i} \\ z'_{\text{live}i} \\ 1 \end{bmatrix}, \quad (6.14)$$

where ${}^{\text{map}}_{\text{current}}\mathbf{T}_{3 \times 1}$ is a 3D translation vector, and ${}^{\text{map}}_{\text{current}}\mathbf{R}_{3 \times 3}$ is a 3D rotation matrix, representing the current position and orientation of the LiDAR sensor, respectively, relative to the coordinate system of point cloud map \mathcal{P}_{map} . The resulting transformed live scan is denoted as:

$$\mathcal{P}_{\text{live}} = \{p_{\text{live}i}\}_{i=1}^{\mathcal{N}_{\text{live}}}, \quad (6.15)$$

where $\mathcal{N}_{\text{live}}$ is the number of live scan points. The scan $\mathcal{P}_{\text{live}}$ is now in the same coordinate system as $\mathcal{P}_{\text{base}}$ and therefore can be combined directly as follows:

$$\mathcal{P}_{\text{estimate}} = \mathcal{P}_{\text{base}} + \mathcal{P}_{\text{live}}. \quad (6.16)$$

All of the points in the resulting $\mathcal{P}_{\text{estimate}}$ are in point cloud map coordinates. Therefore, they need to be further transformed to the local coordinates of the simulated LiDAR scanner, whose origin is located at p_{target} . The transformation of each point

$[x_{\text{estimate}_i}, y_{\text{estimate}_i}, z_{\text{estimate}_i}]^\top$ in $\mathcal{P}_{\text{estimate}}$ is performed using the following equation:

$$\begin{bmatrix} p'_{\text{estimate}_i} \\ 1 \end{bmatrix} = \begin{bmatrix} \text{map}_{\text{target}} \mathbf{R}_{3 \times 3} & \text{map}_{\text{target}} \mathbf{T}_{3 \times 1} \\ \mathbf{0} & 1 \end{bmatrix}^{-1} \begin{bmatrix} x_{\text{estimate}_i} \\ y_{\text{estimate}_i} \\ z_{\text{estimate}_i} \\ 1 \end{bmatrix}, \quad (6.17)$$

where the target 3D position and orientation of the simulated LiDAR sensor with respect to the coordinate system of the point cloud map \mathcal{P}_{map} are represented by $\text{map}_{\text{target}} \mathbf{T}_{3 \times 1}$ and $\text{map}_{\text{target}} \mathbf{R}_{3 \times 3}$, respectively.

Finally, resulting estimated scan $\mathcal{P}'_{\text{estimate}} = \{p'_{\text{estimate}_i}\}_{i=1}^{N_{\text{estimate}}}$ is projected onto a 2D image plane using the method described in Section 6.3.1, in order to generate depth image $\mathcal{D}_{\text{other}}$.

6.3.3 Calculation of visibility ratio

The approach described in Sections 6.3.1 and 6.3.2 can be used to identify which regions of the surfaces of surrounding lanes are visible from a specified viewing point. However, the proposed motion planner cannot easily use the resulting visible regions for candidate trajectory selection as they are not quantified. Therefore, this section explains an approach for calculating a quantitative value representing visibility conditions at a particular location from visible region estimation results.

Visibility ratio \mathcal{V} is a numerical value that represents the degree of visibility from a particular location, a concept which was first introduced in Chapter 5. Its value is equal to a ratio of the area visible from the specified viewing location in relation to the total area of interest:

$$\mathcal{V} = \frac{\mathcal{A}_{\text{visible}}}{\mathcal{A}_{\text{interest}}}, \quad (6.18)$$

where $\mathcal{A}_{\text{interest}}$ is the total area of interest, and $\mathcal{A}_{\text{visible}}$ is the portion of the total area of interest which is visible from a particular viewing location. As $\mathcal{A}_{\text{visible}} \subseteq \mathcal{A}_{\text{interest}}$, the value of a visibility ratio ranges from 0 to 1, where 0 indicates that the area of interest is completely occluded from the viewing position, while 1 indicates the area




Figure 6-4: Area of interest $\mathcal{A}_{\text{interest}}$ for an ego vehicle at location p_{target} . The yellow shaded area represents sensing range $\mathcal{S}_{\text{range}}$. Lane segments located within the yellow area are defined as the set $\mathcal{L}_{\text{neighbor}}$. Lane segments within $\mathcal{L}_{\text{neighbor}}$ that have the same driving direction as the ego vehicle and are reachable from position p_{target} are designated as $\mathcal{L}_{\text{reachable}}$, and are represented by green lines. Lane segments with opposite traffic directions than the ego vehicle that intersect the $\mathcal{L}_{\text{reachable}}$ lanes are designated as $\mathcal{L}_{\text{intersect}}$, and are represented by red lines.

of interest is entirely visible.

In order to calculate the visibility ratio of a given position p_{target} , the area of interest $\mathcal{A}_{\text{interest}}$ first needs to be defined. Let $\mathcal{L}_{\text{neighbor}}$ be a set of lane segments that are within sensing range $\mathcal{S}_{\text{range}}$ from location p_{target} . Assuming that no road users will travel in the wrong, i.e., opposite, direction, areas that are relevant in the context of driving from position p_{target} , i.e., $\mathcal{A}_{\text{interest}}$, are comprised of the space above the following lane segments within $\mathcal{L}_{\text{neighbor}}$:

- Lanes with the same driving direction as the ego vehicle's lane that are reachable from position p_{target} , which are denoted by $\mathcal{L}_{\text{reachable}}$.
- Lanes with the opposite driving direction (in relation to the ego vehicle) from position p_{target} that intersect any of the $\mathcal{L}_{\text{reachable}}$ lanes, which are denoted by $\mathcal{L}_{\text{intersect}}$.

Examples of lane segments $\mathcal{L}_{\text{reachable}}$ and $\mathcal{L}_{\text{intersect}}$ are shown in Fig. 6-4 as green and red lines, respectively.

Once the relevant lane segments are identified, a set of 3D points representing the area above those segments, i.e., $\mathcal{P}_{\text{interest}}$, is obtained from the road network map and used to generate corresponding depth image $\mathcal{D}_{\text{interest}}$ using the approach described in Section 6.3.1. The resulting depth image $\mathcal{D}_{\text{interest}}$ is then compared with $\mathcal{D}_{\text{other}}$, which is estimated using the method explained in Section 6.3.2. Finally, assuming that the points in $\mathcal{P}_{\text{interest}}$ are uniformly distributed across the surface of the lanes, $\mathcal{A}_{\text{interest}}$ in Eq. 6.18 can be replaced by the total number of points in $\mathcal{P}_{\text{interest}}$. Similarly, the number of points in $\mathcal{P}_{\text{interest}}$ that are visible from p_{target} can be used as $\mathcal{A}_{\text{visible}}$.

6.4 Experiments

In order to verify the effectiveness and applicability of the proposed planner in reducing the risk of collision with potentially occluded obstacles in various situations, experiments were carried out in three different traffic scenarios: T-junction crossing, turning right, and preparing to pass a parked vehicle. One or more target vehicles were placed within an occluded area in each scenario. The position of the ego vehicle when these hidden vehicles were first detected, measured by the distance from its starting position, was used for evaluation. In other words, shorter distances from the ego vehicle’s starting position represent better performance when detecting occluded targets, as this represents earlier, more effective, and thus safer detection. In the T-junction crossing and right turn scenarios, the primary sources of occlusion are static objects such as buildings, kiosks, and walls. In order to verify the ability of the proposed system to take into account occlusions caused by objects that are not included in the HD maps, in the third scenario, the occlusion is caused by a parked vehicle that was not present when the HD map was generated.

Town01 in the CARLA simulator [78], which represents a flat, urban area, was chosen as the experimental environment. All of the roads in the town are two-lane roads, and the traffic travels in a right-side direction. A map of Town01 is shown in Fig. 6-5, where the orange, blue, and red rectangles represent the locations of the T-junction crossing, right turn, and preparing to pass a parked vehicle traffic scenarios,

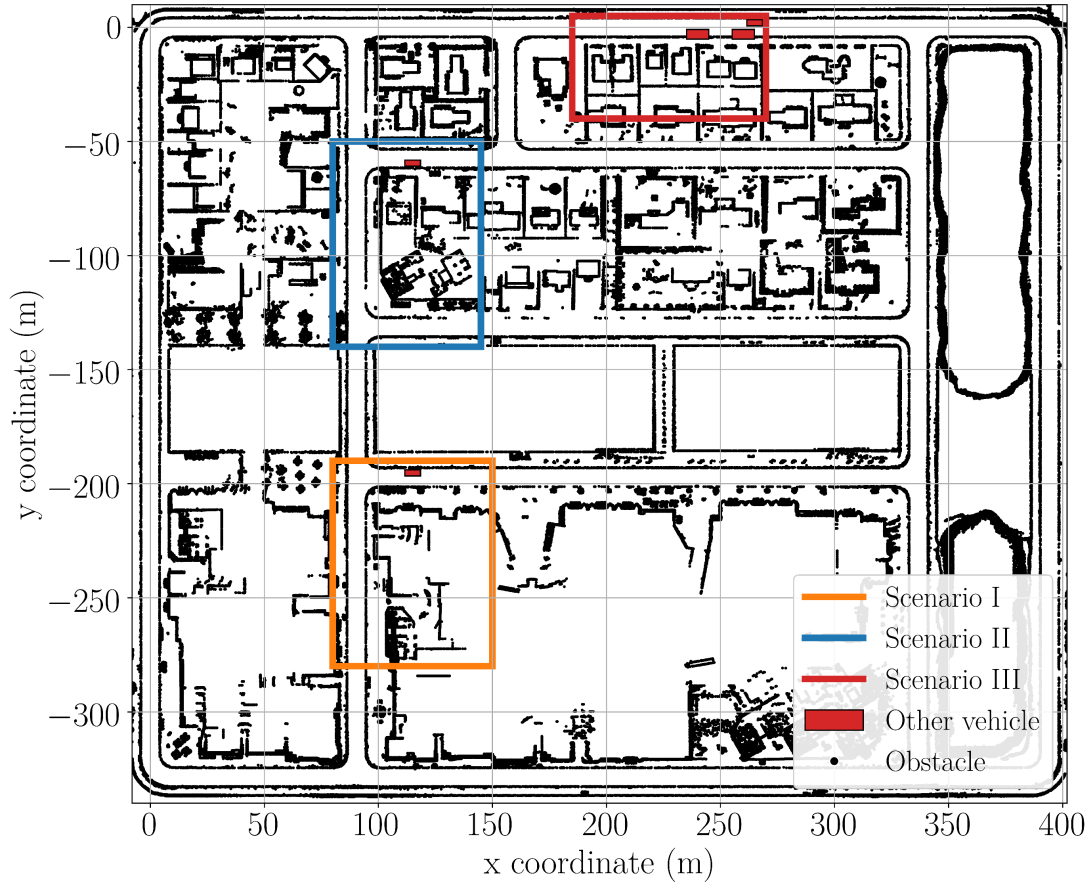


Figure 6-5: Town01 of the CARLA simulator. The orange rectangle indicates the location of the T-junction crossing scenario, the blue rectangle indicates the location of the right turn scenario, and the red rectangle indicates the location of the straight road where the ego vehicle has to prepare to pass a parked vehicle. The small, red rectangles are examples of the positions of hidden (or occluding) vehicles in each scenario.

respectively.

The performance of the proposed planner was compared to that of a baseline planner (OpenPlanner [21]), which does not consider visibility when planning the vehicle’s output trajectory. Real-time object detection and localization results were achieved using “lidar_euclidean_cluster_detect” and “ndt_matching” modules available in Autoware [87], respectively. The maximum lateral deviation of the ego vehicle was implicitly limited so that it would never go out of its lane via $\mathcal{W}_{\text{rollout}}$ and $\mathcal{N}_{\text{rollout}}$. In this experiment, $\mathcal{N}_{\text{rollout}}$ was set to 3 to keep the computational load manageable for real-time planning. The value of $\mathcal{W}_{\text{rollout}}$ was set to 0.5 m resulting in the expected

Table 6.1: Parameters used by the proposed and baseline planners in our experiments.

Parameter	Value	Parameter	Value
$\mathcal{T}_{\text{lower}}$	0.1	$\mathcal{T}_{\text{upper}}$	0.9
$\mathcal{N}_{\text{sampling}}$	10	$\mathcal{N}_{\text{rollout}}$	3
γ	0.85	$\mathcal{H}_{\text{planning}}$	50 m
$\mathcal{D}_{\text{resolution}}^{\text{horizontal}}$	720 pixels	$\mathcal{D}_{\text{resolution}}^{\text{vertical}}$	64 pixels
$\mathcal{S}_{\text{angle}}^{\text{lower}}$	-0.47 rad	$\mathcal{S}_{\text{range}}$	50.0 m
$\mathcal{S}_{\text{FOV}}^{\text{horizontal}}$	2π rad	$\mathcal{S}_{\text{FOV}}^{\text{vertical}}$	0.50 rad
$\mathcal{H}_{\text{object}}$	1.5 m	$\mathcal{W}_{\text{rollout}}$	0.5 m
Baseline planner-specific parameters:			
w_{center}	0.25	$w_{\text{transition}}$	0.25
$w_{\text{collision}}^{\text{lon}}$	0.25	$w_{\text{collision}}^{\text{lat}}$	0.25
Proposed planner-specific parameters:			
w_{center}	0.1	$w_{\text{transition}}$	0.3
$w_{\text{collision}}^{\text{lon}}$	0.1	$w_{\text{collision}}^{\text{lat}}$	0.1
$w_{\text{visibility}}$	0.4		

maximum lateral shift of 1 m in both left and right directions. The same parameters were used in both planners, except for those which were specific to each planner, as shown in Table 6.1. All parameters remained the same in each test scenario.

6.4.1 Scenario I: T-junction crossing

In this scenario, the ego vehicle must drive straight through an uncontrolled T-junction, where the merging roadway is on its right. The ego vehicle starting position in this scenario is at the bottom edge of the orange rectangle in Fig. 6-5. From its starting position, the ego vehicle has to go straight, pass through the T-junction, and reach its goal at the top edge of the orange rectangle. We tested two variations of this scenario, the first of which involves occlusion of the intersection, primarily by a large building set back from the road, which we called the “clear intersection” variation. The second variation included additional structures located very close to the corner,



(a) Slightly occluded T-junction in the original Town01 (clear intersection).



(b) Heavily occluded T-junction in the modified Town01 (occluded intersection).

Figure 6-6: T-junctions used in Scenario I.

which we called the “occluded intersection” variation.

The challenge in this scenario mainly stems from the occlusions caused by the structures near the intersection, as shown in Fig. 6-6a, which prevent the ego vehicle from fully observing the roadway merging from the right. Therefore, the merging road was chosen as the location for the target vehicle. Three different location configurations were used for the target vehicle. In the first configuration, the vehicle was placed at $x = 105\text{ m}$, which is approximately 11.5 m from the intersection’s entrance. In the second and third configurations, the vehicle was placed at $x = 110\text{ m}$ and $x = 115\text{ m}$, respectively. In all three configurations, the target vehicle was positioned

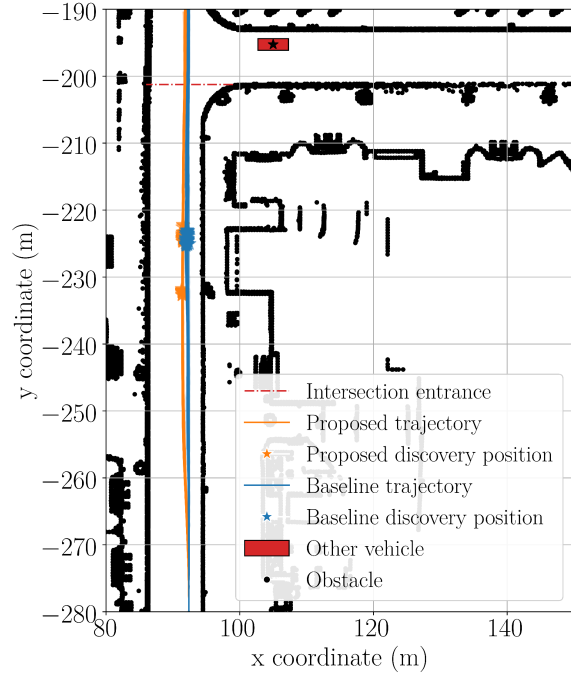
at the center of the right lane of the merging road, i.e., $y = -195.25 m$. Both planners were tested 50 times in each of the three configurations, resulting in 150 junction crossings per planner.

Clear intersection variation

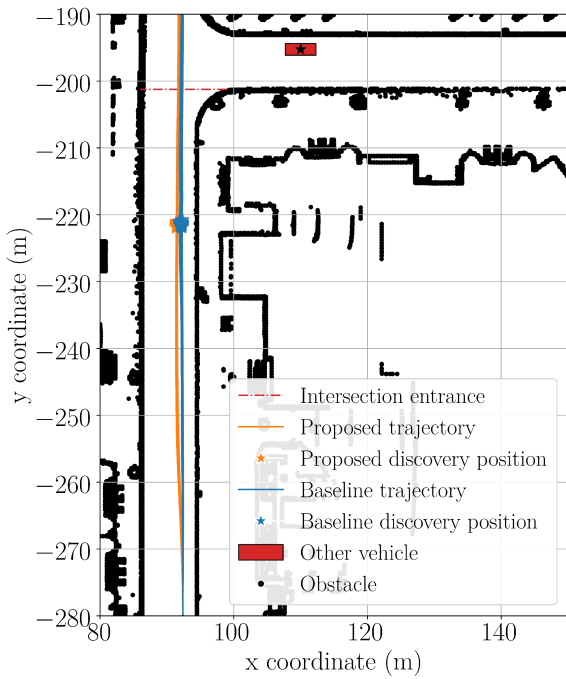
First, we will discuss the “clear intersection” variation of the T-junction scenario. The output trajectories for the proposed and baseline planners in all three of the target vehicle configurations tested ($x = 105 m$, $x = 110 m$, and $x = 115, m$) are shown in Fig. 6-7 as orange and blue lines, respectively. We can see that the output trajectories for each planner are similar for all three target vehicle configurations. When using the baseline planner, the ego vehicle stayed close to the center of its lane most of the time. On the other hand, when using the proposed planner, the ego vehicle started moving to the left side of the lane soon after the experiment began and kept left until finally moving back closer to the lane’s center as it crossed the intersection.

As can be observed in Fig. 6-8 (bottom), although the lateral position shift differed slightly between each run for both planners, their output trajectories were consistent overall. Fig. 6-8 (top) also shows the change in visibility ratio \mathcal{V} along the ego vehicle’s trajectory, averaged over 50 runs. At the beginning of the experiment, the visibility ratio remained constant at approximately 1.0 for both planners, which indicates that all relevant areas of the driving environment could be fully observed during that period. As the ego vehicle approached the intersection and the intersection came within sensor range \mathcal{S}_{range} , the visibility ratio started to drop, signifying incomplete observation of the merging roadway from that point onward. As shown by the shaded area in Fig. 6-8 (top), there is a slight difference in the drop in the visibility ratio between the two planners, with the visibility ratio along the proposed planner’s trajectories remaining relatively higher than those of the baseline planner, suggesting that the proposed planner was able to choose an output trajectory with better visibility conditions compared to the baseline planner.

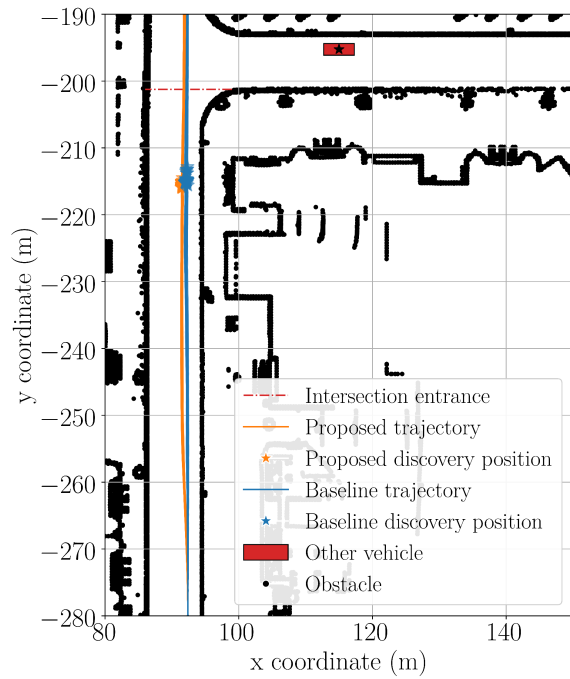
Although the trajectories output by both planners were consistent for all three target vehicle location configurations, the locations where the ego vehicle first detected



(a) $x = 105$ m




(b) $x = 110$ m

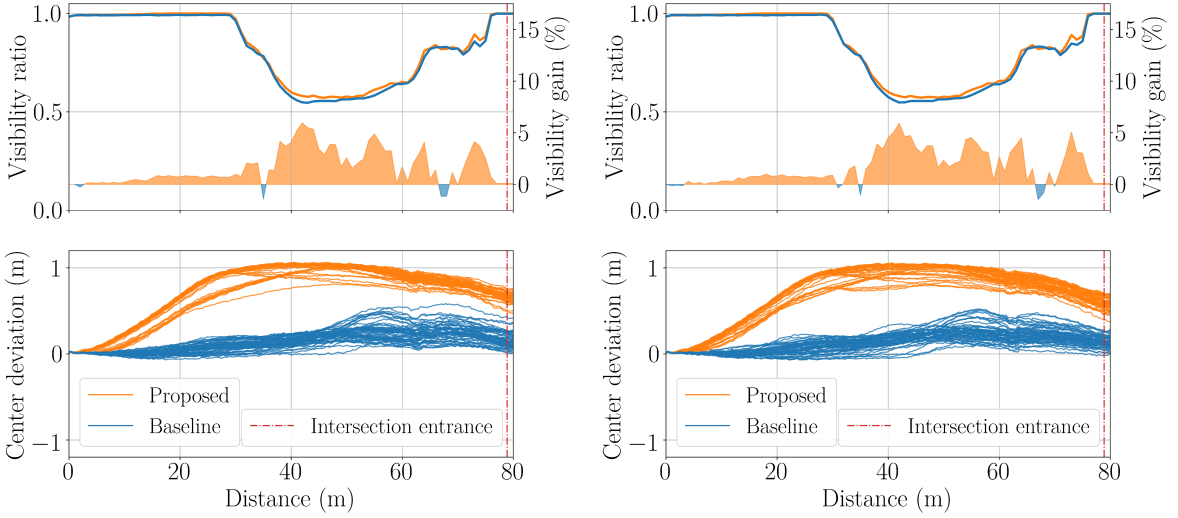


(c) $x = 115$ m

Figure 6-7: Output trajectories generated by the baseline and proposed planners during the “clear intersection” T-junction traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.



(a) $x = 105 m$




(b) $x = 110 m$


(c) $x = 115 m$

Figure 6-8: Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “clear intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

the target vehicle, i.e., the discovery positions, in each hidden vehicle configuration are clearly different, as shown by the orange and blue stars in Fig. 6-7, where a red rectangle depicts the hidden vehicle in each configuration. Differences in the discovery positions when using the various target vehicle locations can also be observed in Fig. 6-9, in which distributions of the discovery positions for both planners in the tested configurations are shown using histograms and cumulative histograms. The x-axes in



(a) $x = 105\text{ m}$



(b) $x = 110\text{ m}$

(c) $x = 115\text{ m}$

Figure 6-9: Distributions of the ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “clear intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

Fig. 6-9 show the distance traveled along the lane’s center from the starting position of the ego vehicle. Note that the distance traveled along the lane’s center is exactly the same as the actual distance the ego vehicle has traveled if, and only if, the ego vehicle remains at the center of the lane without any lateral deviation during the entire run.

However, since lateral motion within the ego vehicle’s lane has a negligible effect on distance traveled, for the sake of conciseness distance traveled along the ego vehicle’s lane from its starting position will be referred to as distance traveled, travel distance, or distance in this chapter unless stated otherwise.

As can be observed in Fig. 6-9a, a distribution of the proposed planner’s discovery position when the target vehicle was placed close to the intersection, at $x = 105\text{ m}$, is clearly bimodal, with the first and second peaks located at approximately 48 m and 56 m from the ego vehicle’s starting position, respectively. The discovery position distribution of the baseline planner appears to be bimodal as well. However, the two peaks are very close to each other, with all of the discovery positions clustered around 55.5 m from the ego vehicle’s starting position.

Fig. 6-9b shows distributions of the discovery positions in the second configuration, where the target vehicle was located at $x = 110\text{ m}$. The resulting discovery positions are very consistent for both planners in this configuration. For the baseline planner, the discovery positions appear to be normally distributed with their center at approximately 58.5 m . On the other hand, in the case of the proposed planner, the majority of the first detections of the hidden vehicle occurred slightly earlier, at 58.2 m , causing the distribution to skew to the right.

Discovery position distributions for the last configuration, where the hidden vehicle was placed furthest from the intersection at $x = 115\text{ m}$, are shown in Fig. 6-9c. The distribution results when using the proposed planner show two visible peaks that are relatively close to each other. The first peak is located at roughly 64.5 m from the ego vehicle’s starting position, while the second peak occurs slightly later at approximately 65 m . For the baseline planner, the resulting discovery position distribution generally resembles the second mode of the distribution for the proposed planner, with its peak located at 65 m . However, as can be observed in Fig. 6-9c, in a few trials, the ego vehicle detected the hidden vehicle significantly later when using the baseline planner.

A summary of the statistics of the ego vehicle positions when discovering the hidden vehicles in each hidden vehicle configuration during the “clear intersection”

Table 6.2: Summary of statistics of ego vehicle positions when discovering hidden vehicles in the “T-junction crossing” traffic scenario (Scenario I), for both the “clear” and “occluded” variations of the intersection, for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials.

Clear T-junction						
	$x = 105\ m$		$x = 110\ m$		$x = 115\ m$	
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	50	50
mean	52.17	55.73	58.41	58.58	64.59	65.22
std	4.20	0.70	0.21	0.27	0.36	0.61
min	47.27	54.67	58.03	58.10	64.02	64.36
median	55.27	55.60	58.37	58.57	64.54	65.02
max	57.82	57.00	58.95	59.24	65.52	67.15
Occluded T-junction						
	$x = 105\ m$		$x = 110\ m$		$x = 115\ m$	
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	50	50
mean	63.10	62.92	68.74	69.61	72.74	72.87
std	0.80	0.55	0.23	0.36	0.20	0.25
min	61.69	61.17	68.26	68.77	72.31	72.27
median	63.14	62.95	68.76	69.69	72.79	72.88
max	64.61	64.22	69.27	70.34	73.07	73.34

variation of the T-junction traffic scenario (mean, standard deviation, minimum, median and maximum distances from the starting point until detection) as well as the total number of experiments where the ego vehicle successfully detected the target vehicle (count), are given in the top part of Table 6.2. As can be seen from the table, the ego vehicle was able to detect the hidden vehicles in all configurations with both planners. Moreover, the table shows that both the mean and median distances from the ego vehicle’s starting position when the hidden vehicle was detected were smaller when using the proposed planner than when using the baseline planner in every target vehicle configuration tested.

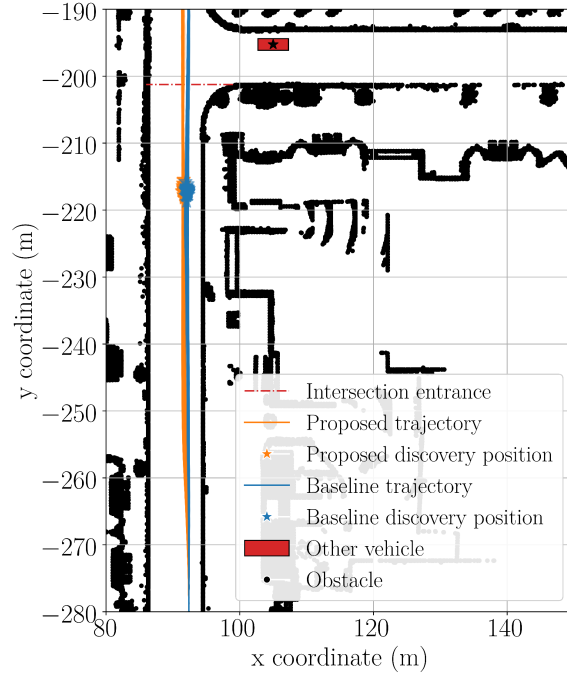
From the results shown in Figs. 6-7, 6-8, and 6-9, it is clear that the output trajectories and discovery positions during each run differed slightly, even when using the same planner and hidden target configuration. This slight deviation was expected since an actual localization and detection module were used in the experiment instead of directly utilizing the ground truth information from the simulation. The effects of using live perception modules seem to be most evident in the first configuration, as shown in Fig. 6-7a and Fig. 6-9a, where the discovery position distribution appears to be bimodal, with the two modes being significantly far apart in the case of the proposed planner. In this configuration, the hidden vehicle was initially occluded from the ego vehicle’s point of view by the corner of the building and a street light pole, located at $(x = 101\text{ m}, y = -202\text{ m})$ in Fig. 6-7a. Although the size of the pole is small, when it is far from the ego vehicle it can effectively prevent the front end of the target vehicle from being detected by the LiDAR-based detector used in the experiments, as point cloud sparsity positively correlates with range. Therefore, there are two possible ways the ego vehicle can detect the target vehicle. The first option, which allows the target vehicle to be discovered earlier, is to detect it through the gap between the building corner and the street light pole. This early detection through the opening between two occluding objects could explain the first peak of the proposed planner’s discovery position distribution, shown in Fig. 6-9a, which occurs at approximately 48 m from the ego vehicle’s starting position. By keeping the ego vehicle on the left side of the lane, the proposed planner widened the gap between the building and the pole, thus allowing the ego vehicle to briefly detect the target vehicle earlier. While early detection was physically possible, the target vehicle was not discovered in some runs due to the imperfect performance of the perception module. Another relatively more straightforward way to observe the hidden vehicle is to get close enough to the intersection that both the pole and the corner of the building no longer obstruct the view. The second peak of the proposed planner’s discovery position distribution, and the entire position distribution for the baseline planner, reflect this later but simpler method of detecting the hidden vehicle. Similar situations also occurred in the third hidden vehicle configuration, as can be observed

in Fig. 6-7c and 6-9c. In the third configuration, the gap occurred between the corner of the building and a tree located at $(x = 107\text{ m}, y = -202\text{ m})$ in Fig. 6-7c. Despite some minor variation, the results are generally consistent, suggesting that the proposed planner could cope with noisy perception to some degree in this scenario. The smaller mean and median discovery distances when using the proposed planner to detect the hidden vehicles in all three configurations, as shown in Table 6.2, suggest that the ego vehicle would likely detect a hidden vehicle earlier in this scenario when using the proposed planner, as compared to the baseline planner.

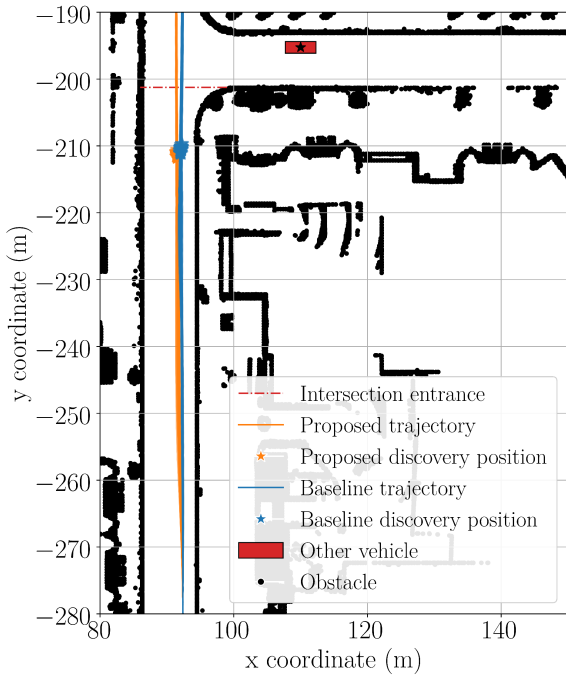
Occluded intersection variation

In order to investigate how different levels of occlusion can affect the output trajectories produced by the proposed planner, additional experiments were conducted using the modified (occluded) T-junction. Several additional obstructions were placed close to the corner of the building, as can be seen in Fig. 6-6b. Except for the modifications to the nearest corner of the T-junction, the experiments were set up in exactly the same way as before, i.e., 50 runs for each planner in each of the three target vehicle location configurations, resulting in a total of 150 trials per planner.

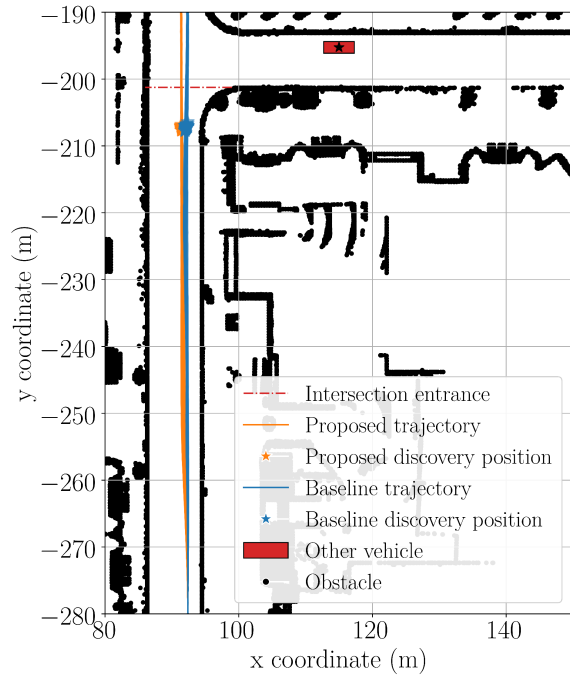
Fig. 6-10 shows the output trajectories for both planners at the modified intersection for all three hidden target vehicle placements. Similar to the results at the original T-junction, shown in Fig. 6-7, the ego vehicle generally stayed close to the lane's center when using the baseline planner, while the ego vehicle gradually moved to the left side of the lane soon after leaving its starting point when using the proposed planner, as happened in the previous clear intersection scenario. However, the ego vehicle did not move back to the center of its lane as it approached the occluded intersection when using the proposed planner; instead, it deviated further to the left of its lane before entering and eventually crossing the intersection. Differences in the lateral deviation profiles of the ego vehicle when using the proposed planner at the two experimental T-junctions can be further observed by comparing Fig. 6-11 with Fig. 6-8. As can be seen, when comparing these two figures, apart from the lateral deviation, the visibility ratio along the output trajectory at the heavily occluded T-



(a) $x = 105$ m




(b) $x = 110$ m

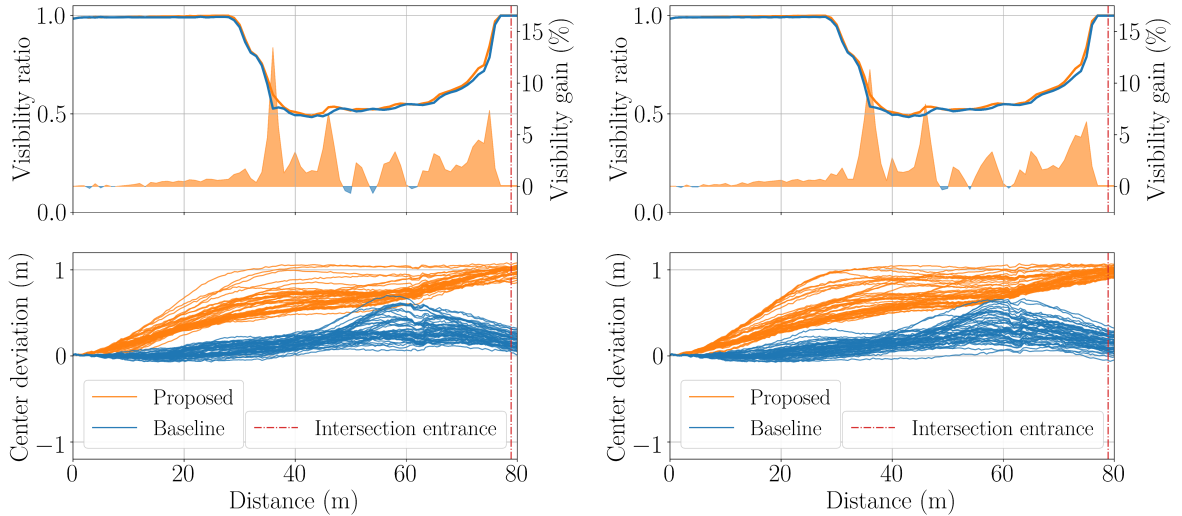


(c) $x = 115$ m

Figure 6-10: Output trajectories generated by the baseline and proposed planners during the “occluded intersection” T-junction traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.



(a) $x = 105\text{ m}$




(b) $x = 110\text{ m}$


(c) $x = 115\text{ m}$

Figure 6-11: Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “occluded intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

junction is also dissimilar to that of the original intersection, and this was also true when using the baseline planner, with the drop in the visibility ratios falling more sharply at the modified intersection. Moreover, the visibility ratios dropped to a lower level than when the planners encountered the less cluttered T-junction. Another distinction between the two tested intersections is the recovery of the visibility ratio as the ego vehicle approaches the intersection. The increase in the visibility ratios at



(a) $x = 105 m$



(b) $x = 110 m$

(c) $x = 115 m$

Figure 6-12: Distributions of ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “occluded intersection” T-junction traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

the modified intersection happened later than at the original intersection, although the increase was relatively smoother and more steady, as can be seen in Fig. 6-11.

The stable gain of visibility at the modified intersection also directly affected the distributions of the discovery positions, as the distributions at this heavily occluded

junction were all unimodal, as can be seen in Fig. 6-12. As shown in Fig. 6-12a, in the first configuration, where the hidden vehicle was placed closest to the intersection, the ego vehicle usually discovered the target vehicle after having traveled approximately 63 m from its starting position, regardless of which planner was used, therefore, both planners achieved similar discovery positions. However, when the target vehicle was placed at $x = 110\text{ m}$ in the second configuration, the difference in the discovery position distributions when using the two planners is apparent. As can be observed in Fig. 6-12b, the majority of initial detections occurred at 68.75 m and 69.75 m when using the proposed and baseline planners, respectively, indicating that the ego vehicle could detect the occluded target vehicle earlier in this hidden vehicle configuration when the proposed planner was used. Fig. 6-12c shows the resulting distributions of discovery positions for both planners in the last configuration, where the hidden vehicle was furthest from the intersection, i.e., at $x = 115\text{ m}$. The resulting distributions for both planners appear to be very similar, with discovery of the hidden vehicle occurring at around 72.9 m , whether the proposed or baseline planner was used.

The summary statistics of the discovery positions in the experiments carried out at the modified (occluded) T-junction are given in the bottom half of Table 6.2, which shows that the proposed planner achieved lower mean and median detection distances from the ego vehicle starting position when the hidden target vehicles were detected in the second and third configurations. However, both mean and median detection distances were lower in the first configuration when using the baseline planner.

The experimental results, which are shown in Fig. 6-10 and 6-11, indicate that the ego vehicle also moved to the left side of the lane prior to crossing the heavily occluded T-junction when using the proposed planner. This movement was anticipated because being on the left side of the lane provides the ego vehicle with a better view of the occluded, merging road to the right of the ego vehicle, even with the additional occlusions. However, unlike at the original, less-occluded T-junction, where the road to the right becomes minimally obstructed as soon as the ego vehicle passes the corner of the building, the view of the merging roadway is blocked right up to the corner

of the modified T-junction. Therefore, when using the proposed planner, the ego vehicle continued traveling on the left side of its lane and entered the occluded T-junction without reverting to the center of the lane, in order to maintain additional visibility of the intersection. A sharper and deeper decline in the visibility ratio and its later but smoother recovery, can be observed at the modified T-junction compared to the original T-junction. These are the effects of a higher degree of occlusion caused by the additional obstacles. The more severe occlusion also resulted in a smoother and steadier increase in the visibility ratio, as there was no gap between these closely placed obstructing objects where the ego vehicle could “peek” through. Therefore, the occlusion is relatively more consistent. This consistent occlusion near the modified junction is reflected by the unimodal distributions of the discovery positions in all of the tested hidden vehicle configurations, as shown in Fig. 6-12. It is evident from Table 6.2 that, in every configuration, hidden vehicle detections at the occluded intersection generally happened later than at the original intersection. The table also shows that the proposed planner achieved earlier detections compared to the baseline planner in two of the three hidden vehicle configurations, with the first configuration, where the hidden target was placed very close to the heavily occluded intersection, being the exception.

6.4.2 Scenario II: Turning

In this scenario, the ego vehicle initially drives straight along a short segment of road, then turns right at an intersection. This scenario occurs in the region indicated by the blue rectangle in Fig. 6-5. As in the T-junction crossing scenario, two variations of this scenario were tested, the first of which, referred to as the “clear corner”, involves occlusion mainly from a roadside wall. In the second variation, which we called the “occluded corner”, additional structures were added close to the corner to further obstruct the view of the intersection from the ego vehicle.



(a) Slightly occluded corner in the original Town01 (clear corner).



(b) Heavily occluded corner in the modified Town01 (occluded corner).

Figure 6-13: Corners used for right turn in Scenario II.

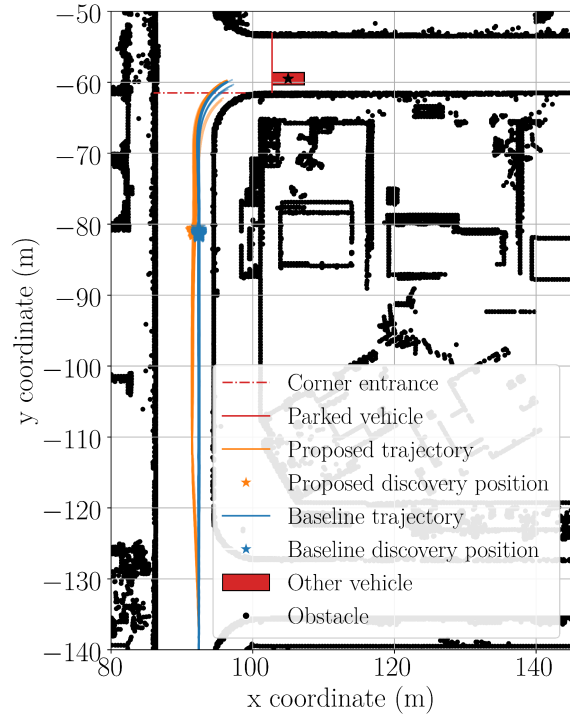
Clear corner variation

In this variation of the right turn scenario, after driving along a straight road the ego vehicle makes a right turn at a corner which is partially occluded by an opaque wall that obstructs the ego vehicle's line of sight of the intersecting road, as shown in Fig. 6-13a, a corner in the original CARLA Town01. Although the ego vehicle does not need to cross into other lanes in this scenario, there is still the risk of collision with vehicles stopped in the ego vehicle's lane, especially if the road to the right is not fully observable. Therefore, the hidden target vehicle was placed in various locations in the same lane the ego vehicle will be traveling in after rounding the corner. In the


first, second and third configurations, the target vehicles were placed in the middle of the ego vehicle's lane $x = 105\text{ m}$, $x = 110\text{ m}$, and $x = 115\text{ m}$, respectively, where x is a coordinate of a Cartesian system used to map CARLA Town01. The proposed and baseline planners were tested 50 times each in each configuration.

Fig. 6-14 shows the output trajectories generated in all three hidden vehicle configurations by the proposed and baseline planners, represented by the orange and blue lines, respectively. The initial output trajectories appear to be similar for all three configurations. However, when using the proposed planner, the ego vehicle moved towards the left side of its lane to increase its visibility of the intersecting road when approaching the intersection. In contrast, when using the baseline planner, the ego vehicle stayed close to the lane's center right up to the entrance of the intersection. As shown by the ends of the orange and blue lines in Fig. 6-14, during or just after the right turn maneuver, the ego vehicle stopped at a different position in each hidden vehicle configuration to avoid colliding with the hidden vehicles, approximately 5 m to 6 m from the rear end of the parked target vehicles.

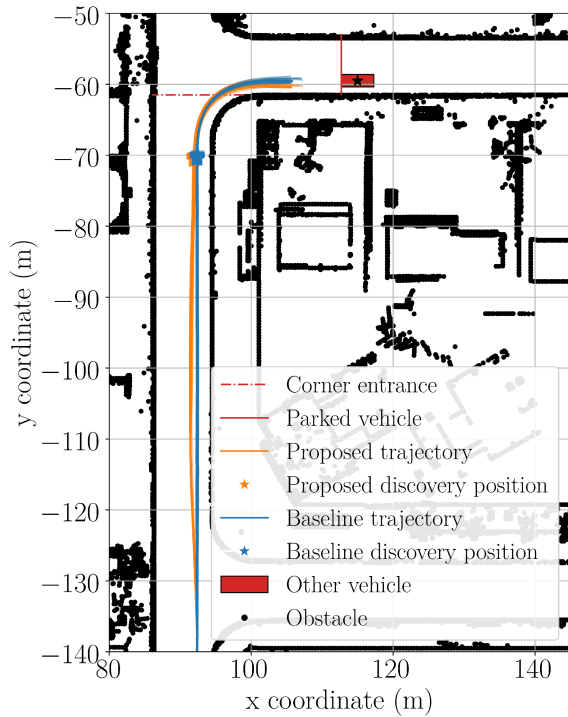
Differences in the stopping position of the ego vehicle in each configuration can also be observed in the lateral deviation plot in Fig. 6-15, where a dashed red line indicates the entrance of the intersection. As can be seen from the lateral deviation profiles of the tested configurations, when using the proposed planner, the ego vehicle initially moved to the left of the lane, then generally moved back to the center of the lane before it started turning right at the corner. Nonetheless, there appears to be a slight difference in the output trajectories produced by the two planners during the first hidden vehicle configuration, however, the ego vehicle briefly inched to the left immediately before turning, regardless of which planner was used. Fig. 6-15 also shows similar visibility ratio profiles for the two planners in the tested configurations. When the ego vehicle was close to the corner, the visibility ratio dropped, indicating that the area around the corner was not fully observable from these positions. The visibility ratio increased again once the ego vehicle almost reached the intersection's entrance and started to turn. Of the two tested planners, the proposed planner seems to have maintained a higher visibility ratio as the ego vehicle approached the entrance



(a) $x = 105\text{ m}$




(b) $x = 110\text{ m}$

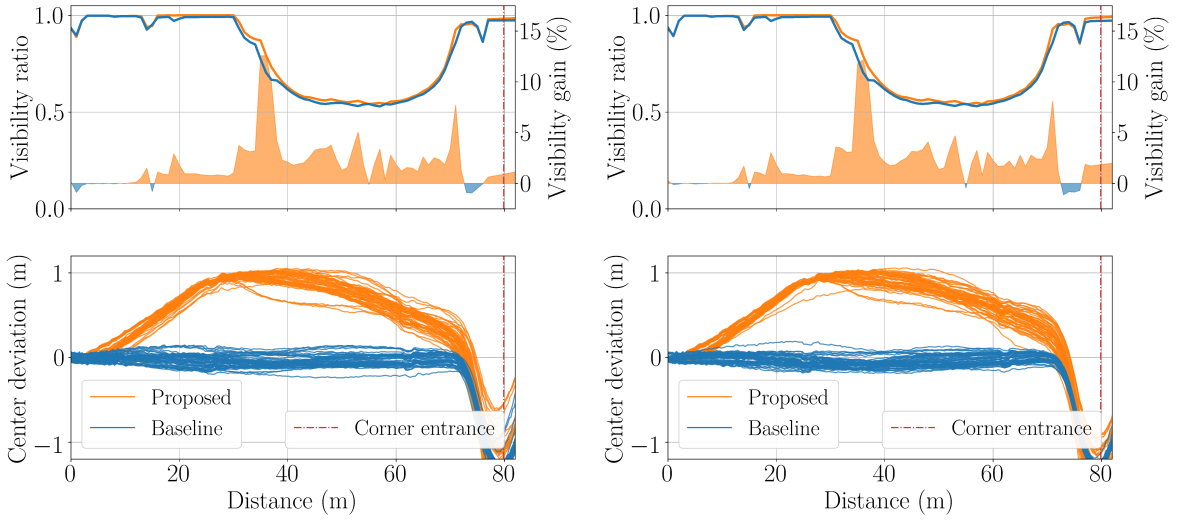


(c) $x = 115\text{ m}$

Figure 6-14: Output trajectories generated by the baseline and proposed planners during the “clear corner” right turn traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.



(a) $x = 105 m$




(b) $x = 110 m$

(c) $x = 115 m$


Figure 6-15: Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “clear corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

of the intersection.

As indicated by the orange and blue markers in Fig. 6-14, the locations where the ego vehicle first discovered the hidden target vehicles in this scenario are very consistent for all configurations under both planners. The distributions of the discovery positions, which are shown in Fig. 6-16, also confirm this consistency, as all of the distributions are unimodal. In the configuration where the target vehicle was



(a) $x = 105 m$



(b) $x = 110 m$

(c) $x = 115 m$

Figure 6-16: Distributions of the ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “clear corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

placed closest to the corner, i.e., the first configuration, the ego vehicle using the baseline planner made the majority of its detections at around $58.8 m$ from its starting position, whereas these detections mainly occurred at approximately $59 m$ when using the proposed planner, as shown in Fig. 6-16a, therefore, the baseline planner

yielded earlier detection in the first hidden vehicle configuration. However, as shown in Fig. 6-16b, in the second configuration where the target vehicle was placed at $x = 110\text{ m}$, the hidden vehicle was detected later when using the baseline planner, compared to the proposed planner. The peak of the discovery distribution in the case of the baseline planner is at 66.9 m , but at 66.7 m for the proposed planner. In the last configuration, where the target vehicle is farthest from the corner, the peaks of the distributions for the proposed and baseline planners are very close to each other at around 69.8 m , as can be observed in Fig. 6-16c, although there were a number of trials in this last hidden vehicle configuration where the ego vehicle using the proposed planner was able to detect the target vehicle relatively earlier, as indicated by the portion of the orange distribution on the left that is not overlapped by the blue distribution.

The distributions of ego vehicle positions when discovering the hidden vehicles using each planner under the three hidden vehicle placement configurations in the “clear corner” variation of Scenario II are also summarized using the number of successful detections (count), as well as mean, median, standard deviation, minimum value, and maximum value of the discovery positions in Table 6.3 (top). As can be observed from the table, using the proposed planner resulted in lower mean and median distances from the ego vehicle’s starting position when first detecting the hidden vehicle in the second and third configurations but in higher values in the first configuration.

As the results in Figs. 6-14 and 6-15 show, when using the proposed planner, the ego vehicle initially moved to the left and then reverted back to the center of its lane before turning right at the corner. The initial movement to the left was expected. Since the ego vehicle’s view of the corner on its right is blocked by the wall near the intersection, being on the left side of the lane increases the visible area of the road around the corner. This increased visibility is also reflected by a relatively higher visibility ratio along the trajectories generated by the proposed planner compared to those output by the baseline planner, as shown in Fig. 6-15. This figure also shows that the output lateral deviation profiles are similar for all of the hidden vehicle configurations, with a slight difference in the first configuration in which the hidden

Table 6.3: Summary of statistics of ego vehicle positions when discovering hidden vehicles in the “right turn” traffic scenario (Scenario II), for both the “clear” and “occluded” variations of the corners, for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials.

	Clear corner					
	$x = 105\ m$		$x = 110\ m$		$x = 115\ m$	
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	50	50
mean	58.95	58.77	66.56	66.82	69.59	69.76
std	0.24	0.24	0.21	0.29	0.22	0.19
min	58.51	58.24	66.02	66.06	69.10	69.35
median	58.93	58.75	66.56	66.89	69.63	69.75
max	59.76	59.38	66.99	67.40	69.98	70.28
	Occluded corner					
	$x = 105\ m$		$x = 110\ m$		$x = 115\ m$	
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	50	50
mean	64.95	66.41	71.09	72.44	74.31	74.25
std	0.74	0.33	0.56	0.25	0.37	0.25
min	63.87	65.84	70.43	71.54	73.73	73.49
median	64.72	66.36	70.94	72.42	74.22	74.28
max	66.48	67.19	72.69	72.90	75.64	74.88

vehicle was placed closest to the corner. The difference shown in Fig. 6-15a is likely caused by the ego vehicle detecting the target vehicle earlier, while still on the initial, straight segment of the road, and trying to avoid the parked vehicle by moving to the left. However, when the ego vehicle arrived at the intersection, it had to abort its diversion and start turning right. In the other configurations, the ego vehicle detected the target vehicles later, i.e., closer to the turning point, and therefore no brief diversion to the left occurred. As can be observed in Fig. 6-16, the difference in the resulting discovery positions for both planners is modest at this corner. The reason for this marginal difference is believed to be an adequate level of visibility at

this corner. Although the wall indeed blocks the view of the intersection to the right of the ego vehicle, it is not located very close to the corner. As a result, visibility increases rapidly as the ego vehicle moves sufficiently close to the intersection.

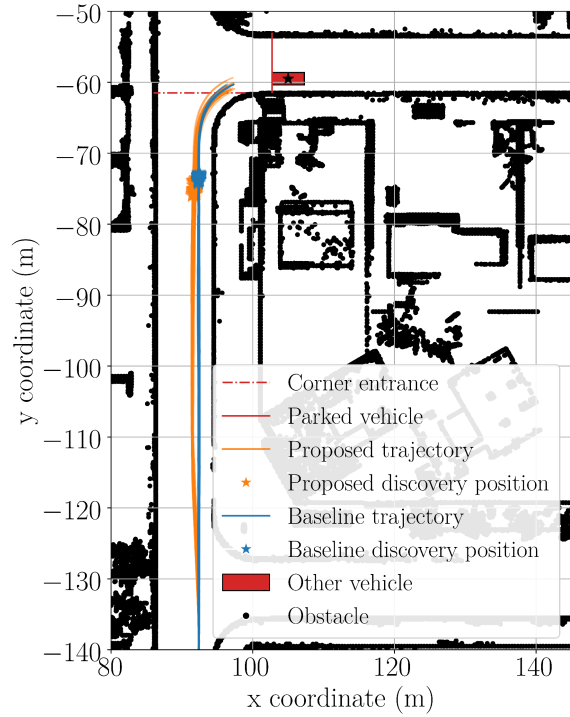
Occluded corner variation

In order to investigate the effects of severe occlusion on motion planner output trajectories and discovery positions during this turning scenario, modifications were made to the original corner in Town01 to reduce ego vehicle visibility at the corner. Two kiosks were added, very close to the intersection, adjacent to the existing wall, as shown in Fig. 6-13b. All other experimental conditions were identical to those used in the previous experiment, i.e., three different placements of the hidden vehicle and a total of 50 trials per planner in each placement configuration.


The output trajectories of both planners at the highly occluded corner are shown in Fig. 6-17. Similar to the output trajectories at the original corner, when using the proposed planner the ego vehicle moved to the left as it approached the intersection, before reverting to the center immediately before turning. After turning, the ego vehicle came to a complete stop 5 *m* to 6 *m* away from the rear end of the parked target vehicle in all three of the hidden vehicle configurations.

Unlike at the original, slightly occluded corner, Fig. 6-18 shows that at the modified, heavily occluded corner the lane deviation profiles were nearly indistinguishable among the three hidden vehicle configurations for each planner, i.e., although the deviation profiles differed between the two planners, the results remained very consistent for each planner. When using the proposed planner, the ego vehicle did not divert to the left immediately before turning at the occluded corner in the first hidden vehicle configuration, as can be observed in Fig. 6-18a. Moreover, the ego vehicle stayed on the left side of the lane longer before moving back to the center when turning, as compared to its trajectory the less occluded corner, i.e., the “clear corner.” Fig. 6-18 also shows that the visibility ratio dropped to a lower level and then increased later, in comparison to the right turn at the less occluded corner.

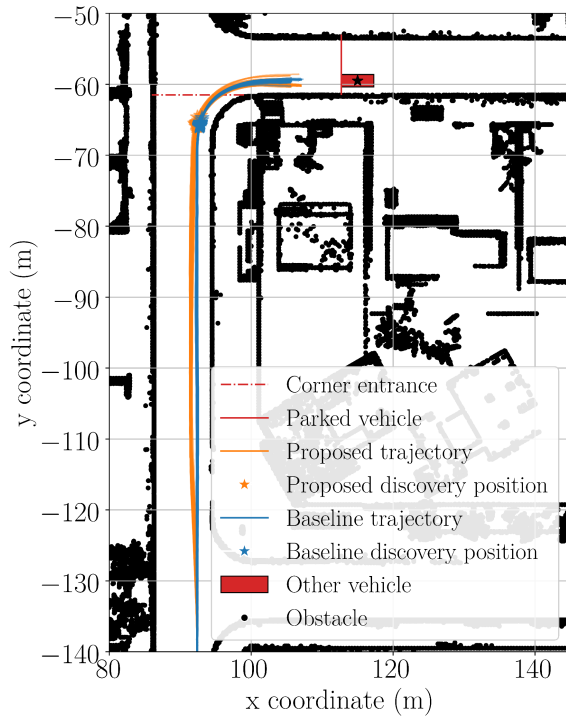
The distributions of the discovery positions appear to be the major distinction



(a) $x = 105\text{ m}$

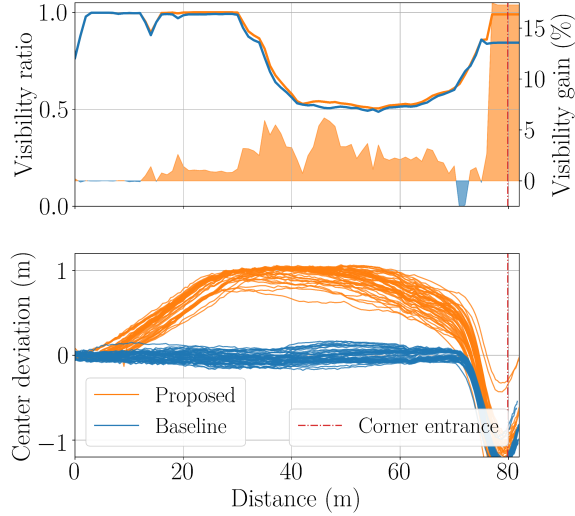


(b) $x = 110\text{ m}$

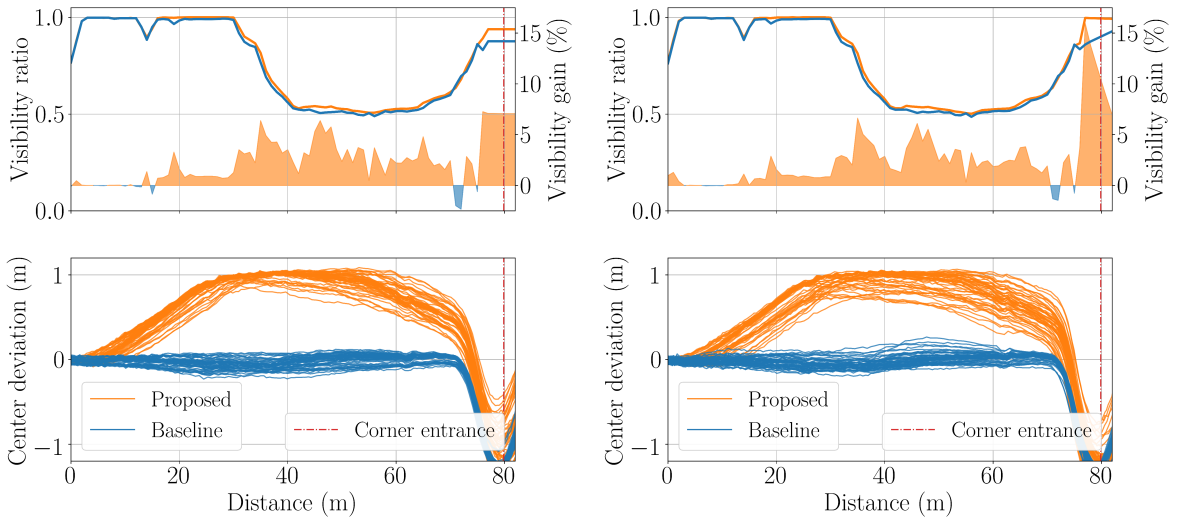


(c) $x = 115\text{ m}$

Figure 6-17: Output trajectories generated by the baseline and proposed planners during the “occluded corner” right turn traffic scenario, with positions of the ego vehicle when the occluded vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.



(a) $x = 105 m$




(b) $x = 110 m$


(c) $x = 115 m$

Figure 6-18: Average visibility ratio (top) and lane center deviation (bottom) along the output trajectories of the baseline and proposed planners during the “occluded corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

between the planner results at the original and the modified corners, as can be observed in Fig. 6-19. Except for the third hidden vehicle configuration, the difference between the ego vehicle’s discovery positions when using the proposed and baseline planners is apparent. In the first configuration, when using the proposed planner, there was a bimodal distribution of the discovery positions, with the first and second peaks occurring at approximately $64.5 m$ and $65.75 m$, respectively. On the other



(a) $x = 105 m$



(b) $x = 110 m$

(c) $x = 115 m$

Figure 6-19: Distributions of ego vehicle discovery positions (i.e., distance of the ego vehicle from its starting point when occluded vehicles were detected) during the “occluded corner” right turn traffic scenario. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the corner.

hand, when using the baseline planner, there was a unimodal distribution in the discovery position results, with its peak at roughly $66.75 m$. When the target vehicle was placed at $x = 110 m$, i.e., during the second hidden vehicle configuration, results for both planners show a unimodal distribution of the discovery positions, with the

proposed planner achieving earlier detection overall, as the ego vehicle was generally able to spot the hidden vehicle at around 70.8 m from its starting position when using the proposed planner compared to 72.5 m when using the baseline planner. However, in the last configuration, where the target was placed furthest from the corner, the distribution results for both planners were very similar, with both planners having their discovery peaks at around 74.25 m .

The difference in discovery positions when using each planner can be further observed in the lower half of Table 6.3. During the first two hidden vehicle configurations, the proposed planner achieved significantly lower mean and median distances from the ego vehicle's starting position to where the target vehicles were discovered. Nonetheless, no difference in the discovery position is apparent in the last configuration, as the proposed planner yielded a lower median but higher mean distance compared to the baseline planner.

Fig. 6-17 shows that the overall movement of the ego vehicle when using the proposed planner at the heavily occluded intersection is similar to that at the original intersection, i.e., moving to the left side of the lane soon after the experiment begins, then returning to the center of the lane just before turning right. However, by comparing the results shown in Fig. 6-18, we can see that the ego vehicle stayed on the left side of its lane longer in the scenario with the occluded corner, as compared to the original, less occluded corner. The extended deviation to the left is the direct effect of the higher degree of occlusion at the modified corner. Unlike at the original corner, the additional obstacles at the modified corner obstructed the ego vehicle's view of the road ahead all the way to the intersection's entrance. Therefore, the proposed planner kept the ego vehicle to the left side of its lane to gain additional visibility for as long as possible before beginning its turn. The increased occlusion is also responsible for the absence of the brief diversion to the left that occurred at the original corner in the first configuration, as shown in Fig. 6-15a. The ego vehicle did not detect the target vehicle in the first hidden vehicle configuration as early when encountering the heavily occluded corner as it did at the original, less obstructed corner, therefore, the ego vehicle did not initiate an avoidance maneuver before turning.

We can see from these results that the proposed planner offers a clear advantage over the baseline planner when detecting the hidden target vehicles behind the heavily occluded corner. As shown in Fig. 6-19, the majority of detections occurred significantly earlier when traveling along the output trajectory of the proposed planner, except during the last hidden vehicle configuration, during which both planners achieved similar performance. When using the proposed planner in the last hidden vehicle configuration, the target vehicle was not in detection range until the ego vehicle had already reverted to the center of its lane for turning. Therefore, the proposed planner did not offer any advantage over the baseline planner in detecting the target vehicle.

6.4.3 Scenario III: Preparing to pass a parked vehicle

In this scenario, the ego vehicle travels along a straight road until it encounters a truck parked in its lane. The location in Town01 where this scenario takes place is indicated by the red rectangle in Fig. 6-5. The ego vehicle has to come to a complete stop behind the parked truck to prepare to pass it. This scenario is different from the previously tested scenarios as the occlusion is not caused by static structures present in the point cloud map, e.g., buildings or walls, but instead, the road ahead is obstructed by a parked truck, which is not a part of the mapped environment. Experiments in this scenario aim to verify improvements in our previously proposed visibility estimation method for handling occlusions caused by objects not included in HD maps. As safe passing requires additional computation and other considerations that are not currently supported by the proposed planner, the experiments in this traffic scenario focus exclusively on the output motion of the ego vehicle prior to passing.

In this scenario, the truck parked in the ego vehicle’s lane is an obstacle occluding the ego vehicle’s view of the road ahead. In addition to this occluding obstacle, there are two detection targets; a second truck parked in the same lane as the occluding parked truck, located at various distances ahead of the first truck in each configuration of the scenario, and a normal passenger car parked in the opposite lane, with its

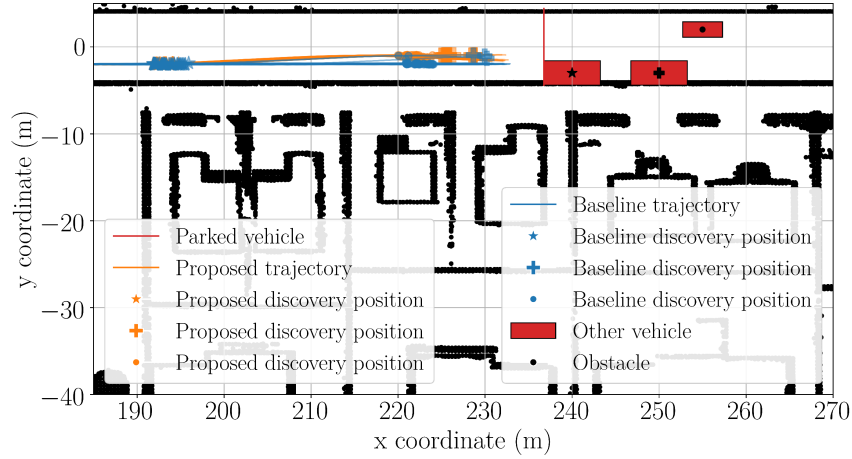


Figure 6-20: Straight road in Scenario III. Two trucks are parked in the ego vehicle’s lane, and a car is parked in the approaching lane.

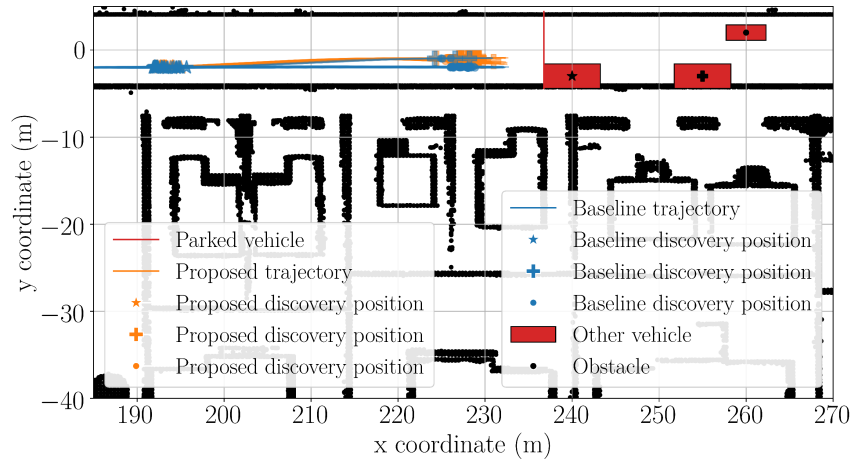
front end is facing the approaching ego vehicle. The passenger car is also placed at various distances from the occluding obstacle in each configuration of the scenario. An example of the arrangement of the three vehicles in this scenario (besides the ego vehicle) is shown in Fig. 6-20.

Three target vehicle location configurations were used in this experiment. In all three configurations, the occluding vehicle (the first truck) is parked at $x = 240\text{ m}$. In the first configuration, the targeted truck and passenger car are placed at $x = 250\text{ m}$ and $x = 255\text{ m}$, respectively. In the second configuration, both target vehicles are moved further from the occluding vehicle, i.e., $x = 255\text{ m}$ for the targeted truck and $x = 260\text{ m}$ for the targeted car. In the third configuration, the targeted truck and passenger car are placed at $x = 260\text{ m}$ and $x = 265\text{ m}$, respectively. A total of 50 experimental runs were carried out per planner (baseline and proposed) in each of the three configurations.

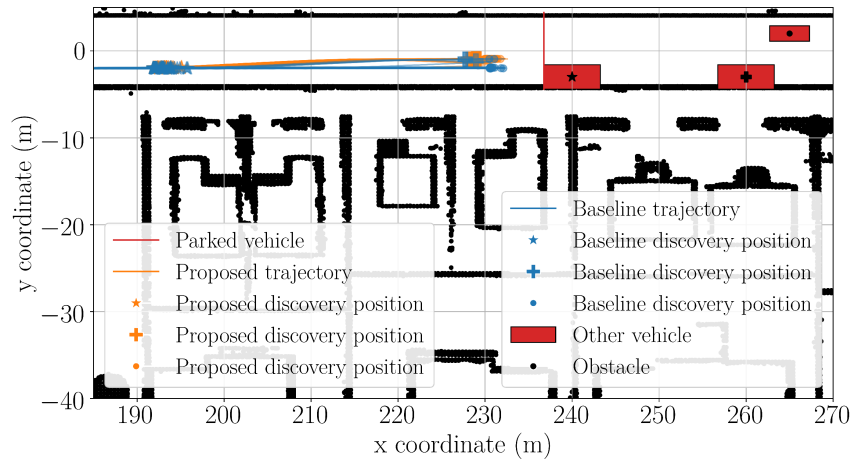
Fig. 6-21 shows the output trajectories of both planners in all three of the tested configurations. Ego vehicle trajectories generated by the proposed planner are shown in orange, while the blue lines represent the trajectories generated by the baseline planner. The output motion for each planner appears to be very consistent over all three target vehicle configurations. When using the baseline planner, the ego vehicle traveled straight ahead from its starting position and remained in the center of its



(a) $x = 240\text{ m}$, $x = 250\text{ m}$, $x = 255\text{ m}$



(b) $x = 240\text{ m}$, $x = 255\text{ m}$, $x = 260\text{ m}$




(c) $x = 240\text{ m}$, $x = 260\text{ m}$, $x = 265\text{ m}$

Figure 6-21: Output trajectories generated by the baseline and proposed planners during the “preparing to pass” traffic scenario, with positions of the ego vehicle when the other vehicles were discovered. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.

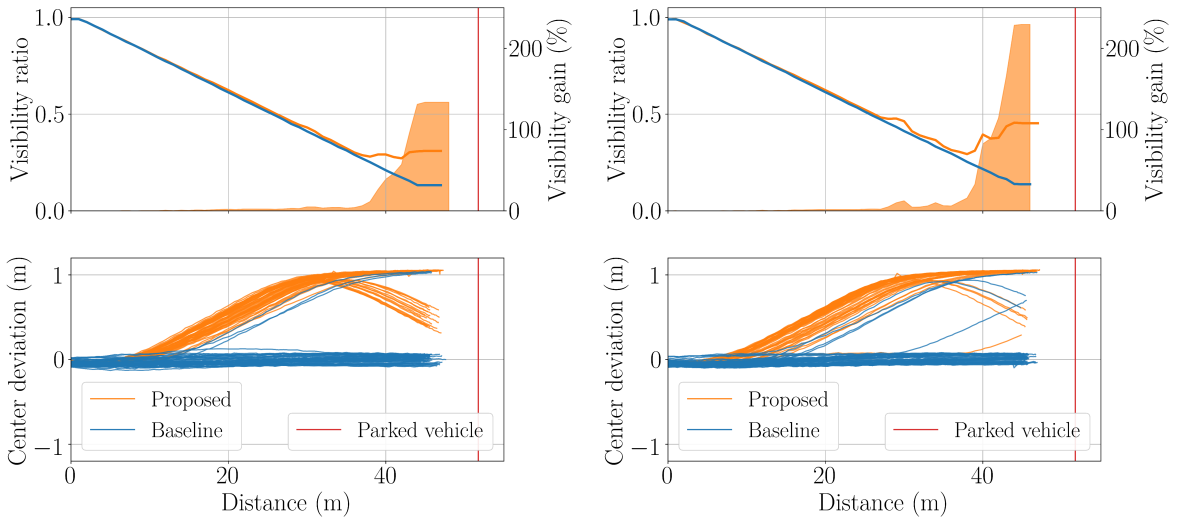
lane until coming to a complete stop behind the occluding truck, although there were a small number of trials in which the ego vehicle made a slight deviation to the left before stopping. In contrast, when using the proposed planner, the ego vehicle always moved to the left side of its lane after leaving its starting location. Moreover, when the ego vehicle came to a complete stop behind the parked truck, it was always located on the left side of its lane.

Differences in the lateral deviation profiles when using each of the two planners can be clearly seen at the bottom of Fig. 6-22. The baseline planner generally kept the ego vehicle at the center of its lane, as indicated by the blue lines, apart from a few trials in which the ego vehicle eventually drifted towards the left side of the lane. When using the proposed planner, the orange lines show that the ego vehicle consistently moved to the left border of its lane and remained there until it came to a complete stop behind the occluding truck. The top of Fig. 6-22 also shows a stark difference in the visibility ratios for the ego vehicle as it moved along the trajectories produced by the proposed and baseline planners, especially near the ego vehicle's stopping position. Initially, the visibility ratios declined linearly for both planners. However, in the case of the proposed planner, the visibility ratio started to increase again as soon as the ego vehicle reached the left border of its lane. When stopping behind the occluding vehicle, the proposed planner generally achieved a higher visibility ratio compared to the baseline planner.

The positions where the three other vehicles in this scenario (the occluding truck, the occluded truck, and the occluded car) were detected are indicated by markers in Fig. 6-21. The distributions of these discovery positions are shown in Fig. 6-23. By comparing Figs. 6-23a, 6-23b, and 6-23c, we can see that the discovery positions for the parked occluding truck, i.e., the first truck, are similar for both planners in all vehicle placement configurations. The ego vehicle generally detected the occluding truck after traveling approximately 7 m from its starting position, regardless of the planner used. However, Figs. 6-23g, 6-23h, and 6-23i show a significant difference in the discovery distributions between the two planners when detecting the target truck hidden behind the first parked truck. As can be seen in Fig. 6-23g, when the



(a) $x = 240\text{ m}, x = 250\text{ m}, x = 255\text{ m}$




(b) $x = 240\text{ m}, x = 255\text{ m}, x = 260\text{ m}$


(c) $x = 240\text{ m}, x = 260\text{ m}, x = 265\text{ m}$

Figure 6-22: Deviation from lane’s center and visibility ratio along the output trajectories of the baseline and proposed planners when traveling along a straight road toward the parked vehicles. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.

target truck was placed closest to the occluding truck, i.e., only 10 m ahead of it, and the proposed planner was used, the hidden truck was generally discovered by the ego vehicle at 40 m or 43.5 m from its starting position. As shown by the cumulative histograms in Fig. 6-23g, the proposed planner detected the hidden truck in 17 out of the 50 trials, whereas the same hidden truck was only detected in 2 of the 50 trials when using the baseline planner. Similarly, Fig. 6-23h shows that the baseline planner




(a) $x = 240\text{ m}$




(b) $x = 240\text{ m}$

(c) $x = 240\text{ m}$


Figure 6-23: Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass” traffic scenario, where (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane. The x distances represent coordinates of a Cartesian system for CARLA Town01, not distances from the starting point.



(d) $x = 255 m$




(e) $x = 260 m$

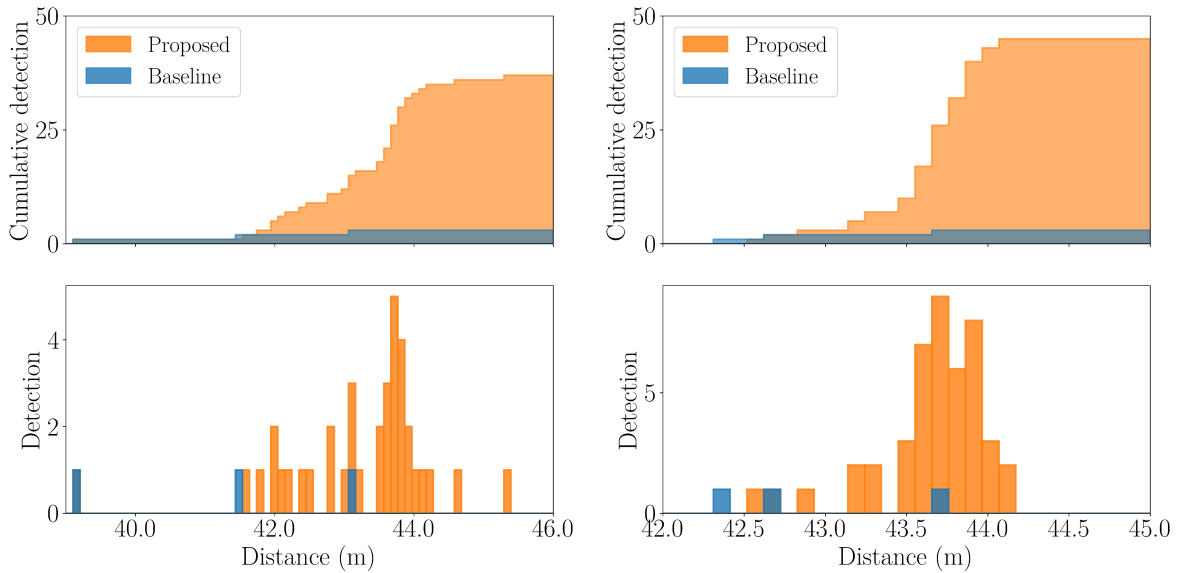


(f) $x = 265 m$

Figure 6-23: Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass” traffic scenario, (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane(cont.)



(g) $x = 250\text{ m}$



(h) $x = 255\text{ m}$

(i) $x = 260\text{ m}$

Figure 6-23: Distributions of the ego vehicle discovery positions (i.e., distance of ego vehicle from its starting point when occluded vehicles were detected) during the “preparing to pass traffic scenario, (a, b, c) = discovery of the occluding truck in each target vehicle configuration, (d, e, f) = discovery of hidden truck parked in front of the occluding truck, and (g, h, i) = discovery of car parked in the opposite lane(cont.)

only made six successful detections of the target truck when it was parked 15 *m* ahead of the first truck, while the proposed planner yielded 39 detections, most of which occurred at approximately 43.75 *m* from its starting position, a significantly higher number of detections. The same pattern can also be observed in the last configuration, as shown in Fig. 6-23i. The distribution peak for detection distances when using the proposed planner is roughly at 43.75 *m* when the targeted truck was parked 20 *m* ahead of the first truck, and the proposed planner resulted in 45 successful detections of the target truck out of the 50 trials. On the other hand, the second truck was only detected during 5 of 50 trials when using the baseline planner.

The distributions of the ego vehicle’s positions when detecting the targeted passenger vehicle in the opposite lane, in all three configurations are shown in Figs. 6-23d, 6-23e, and 6-23f. As shown in Figs. 6-23d and 6-23e, the resulting distributions for both planners appear to be very wide, with no prominent peak in the first and second vehicle location configurations. In both of these configurations, the proposed planner achieved slightly earlier discovery as can be seen in Figs. 6-23d and 6-23e. In both the first and second target vehicle configurations, the car was detected in every trial by both planners. In the last configuration, where the targeted passenger vehicle in the opposite lane was located a considerable distance from the occluding truck, most detections occurred at 45 *m* and 45.5 *m* when using the proposed and baseline planners, respectively. Nevertheless, the proposed planner achieved a noticeably higher detection rate (32/50) compared to the baseline planner (18/50), as indicated by the cumulative histograms in Fig. 6-23f.

Table 6.4 shows a summary of the statistics of the ego vehicle discovery success rates and positions when detecting the other vehicles in this scenario in each of the three vehicle configurations tested. As can be seen from the table, the occluding truck was detected by the ego vehicle in every trial, regardless of the planner used. The proposed planner achieved lower mean and median travel distances before detecting the occluding truck in the first and third configurations. However, it yielded a lower mean but higher median travel distance in the second configuration. In the case of the hidden truck parked in front of the first truck, the proposed planner achieved a

significantly higher number of successful detections by the ego vehicle. Although the baseline planner achieved lower mean and median travel distances before discovery, it is worth noting that the ego vehicle only detected the hidden truck during six trials in the second vehicle configuration and during five trials in the third configuration, out of 50 total trials in each configuration. As shown in the lower part of Table 6.4, the second target vehicle, i.e., the car parked on the other side of the road, was discovered in all experimental trials in the first and second hidden vehicle configurations when using both planners, but not in the third configuration, where the proposed planner achieved significantly better detection performance. The proposed planner yielded lower mean and median travel distances prior to the detection of the parked car in all three vehicle configurations.

As can be seen in Figs. 6-21 and 6-22, in the majority of the trials when using the proposed planner, the ego vehicle deviated to the left of its lane and eventually came to a complete stop behind the occluding vehicle on the left side of the lane, because moving to the left allowed the ego vehicle to achieve a better view of the road ahead. Since there was not enough space for the ego vehicle to pass the parked truck without briefly entering the opposite lane, it had to come to a stop behind the truck and prepare for a passing maneuver. Although the ego vehicle was not able to simply pass the parked truck, coming to a stop on the left side of the lane benefited the ego vehicle by enabling it to observe both the area beyond the truck and obstacles in the opposite lane, as indicated by its higher visibility ratio at its stopping position in Fig. 6-22. Moreover, Fig. 6-23 and Table 6.4 show that the ego vehicle was more likely to discover target vehicles placed behind the occluding truck and in the opposite lane, compared to the baseline planner. Information regarding obstacles in these areas is crucial in determining whether a safe passing maneuver can be performed, therefore, the proposed planner put the ego vehicle in a relatively safer position for passing in this scenario, compared to the baseline planner. Note that the deviation to the left of the lane by the ego vehicle when using the proposed planner did not occur immediately after the beginning of the experiment. As can be observed in Figs. 6-21 and 6-22, there is a brief segment after departing from its

Table 6.4: Summary of statistics of ego vehicle positions when discovering other vehicles during the “preparing to pass” traffic scenario (Scenario III), for both the proposed and baseline methods. Distances are measured from the ego vehicle’s starting point. “Count” represents the number of successful detections per 50 trials.

Occluding truck						
$x = 240\ m$		$x = 240\ m$		$x = 240\ m$		
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	50	50
mean	7.72	7.99	7.55	7.80	7.54	7.86
std	0.72	1.05	0.72	1.07	0.77	1.15
min	6.74	6.55	6.48	6.59	6.52	6.45
median	7.60	7.79	7.64	7.34	7.40	7.58
max	9.26	10.45	9.50	10.36	9.76	10.59
Target truck						
$x = 250\ m$		$x = 255\ m$		$x = 260\ m$		
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	17	2	39	6	45	5
mean	41.51	43.41	43.05	42.36	43.69	43.32
std	1.67	1.72	1.02	1.79	0.32	0.76
min	39.71	42.19	39.12	39.16	42.60	42.38
median	40.66	43.41	43.47	43.03	43.72	43.71
max	43.83	44.62	44.59	44.12	44.17	44.10
Target vehicle in the opposite lane						
$x = 255\ m$		$x = 260\ m$		$x = 265\ m$		
	Proposed	Baseline	Proposed	Baseline	Proposed	Baseline
count	50	50	50	50	32	18
mean	36.31	37.03	41.72	41.89	45.18	45.58
std	0.98	1.01	1.07	1.00	0.42	0.39
min	34.72	35.55	39.84	39.85	44.36	45.19
median	36.07	37.01	41.59	41.95	45.06	45.51
max	38.88	38.86	43.33	43.47	46.08	46.62

starting position where the ego vehicle was traveling straight along the center of its lane. This behavior was expected because the parked truck was initially outside of

the ego vehicle’s field of perception, i.e., its sensing range. Therefore, the proposed planner calculated the output motion as if the road ahead was clear, and thus there was no benefit in moving away from the lane’s center. However, the planner shifted the ego vehicle to the left as soon as the parked truck was within detection range.

It is important to emphasize that the parked truck was not included in the HD maps, so the ego vehicle had to rely on its live detection results to evaluate the unexpected occlusion it encountered in the surrounding traffic environment. Nevertheless, the detection results in these experiments were not perfect. The shape and orientation of several obstacles in the vicinity of the ego vehicle were sometimes incorrectly estimated, and these occasional, incorrect detection results caused the baseline planner to veer the ego vehicle to the left in a few trials, as can be observed in Figs. 6-21 and 6-22. In those incidents, the obstacles were mistakenly judged to be very close to the right side of the ego vehicle’s lane, as a result, the baseline planner tried to avoid them. These sporadic deviations to the left account for why the ego vehicle was sometimes able to discover the hidden target truck when using the baseline planner, as shown in Figs. 6-23g, 6-23h, and 6-23i, and in Table 6.4. In most of the trials where the erroneous deviation in lane position did not occur, the target truck could not be detected when using the baseline planner. Similarly, as can be observed in Fig. 6-23f and Table 6.4, the number of successful detections of the target vehicle in the opposite lane when using the baseline planner in the third vehicle configuration was inflated by these irregular swerves. These deviations increased the chance of successful detections because the visibility of the target vehicle in the opposite lane was better when the ego vehicle was on the left side of its lane, compared to the lane’s center. Furthermore, when the target vehicles were placed very far away from the occluding truck in the third configuration, almost at the edge of the sensing field of the ego vehicle from its stopping position, reducing the distance to the target vehicle even marginally may have significantly increased the detection rate.

6.5 Conclusion

In this chapter, we have proposed a motion planner capable of actively adjusting the ego vehicle’s lateral position in order to minimize occlusions in various traffic scenarios under uncertainty from the perception modules. We have also extended the 3D visibility estimation approach proposed in Chapter 5 for handling occlusions caused by objects not present in HD maps.

The proposed planner first generates candidate trajectories with different lateral offsets from the reference path. Current and future visibility along each trajectory is then estimated using live sensing data from the LiDAR unit and HD maps, which are then converted into a visibility cost. Finally, the visibility cost and other planning costs are used to determine the optimal output trajectory.

Experiments were conducted in the CARLA simulator [78] to evaluate our proposed and improved methods. Live localization and object detection results were used in three traffic scenarios where occlusions were present; crossing an occluded T-junction, making a right turn at an occluded corner, and preparing to pass a large, parked vehicle not present in the HD maps. Our results showed that the ego vehicle was able to effectively minimize occlusions and consequently discover occluded vehicles earlier in most cases when the proposed planner was used, in comparison to a baseline planner. Moreover, the occlusions caused by the parked vehicle that was not present in the HD maps were correctly estimated when using our extended visibility estimation approach.

During our review of related research in Chapter 2, we noted that there is a common assumption, adopted in several of the approaches that have been proposed for regulating the ego vehicle’s speed when encountering occluded areas in the driving environment, that the ego vehicle will generally travel along the lane’s center. Instead of simply using the center of the lane as a reference, the planner proposed in this chapter can be used to plan a reference trajectory that will minimize occlusions by altering the lateral position, i.e., road position, of the ego vehicle within its lane. Therefore, the use of trajectories generated by our proposed motion planner, in con-

junction with existing speed planners, should result in less conservative and more efficient driving in traffic scenarios involving occlusions.

This work has focused only on occlusions caused by surrounding objects. However, occlusions can also be caused by several other factors, such as sensor failure or adverse weather conditions. These types of occlusions should also be considered in future research in order to ensure the safe navigation of autonomous vehicles in broader operational domains. One potential solution is to extend the visibility estimation method improved upon in this work to support multiple sources of sensing information. Although only one sensing modality, namely LiDAR, was considered in this paper, the concept of quantifying estimated visibility conditions along candidate trajectories and converting them into visibility costs for trajectory selection could also be applied when using multiple sensing modalities.

While deviating from the lane center can increase driving efficiency in occluded areas, it is essential to note that being close to the lane boundaries could also have a negative outcome, as the ego vehicle would be closer to other traffic participants on the adjacent lanes than it would be from the center. Although the proposed planner already takes into account a collision with other traffic participants when planning the output trajectory, other road users may not understand the ego vehicle's change in direction toward them and react in an unsafe way. This potential issue could be mitigated by utilizing multiple sensors to increase the total perception coverage of the ego vehicle so that only minimal deviation is sufficient to enhance the visibility significantly.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

In this dissertation, we have attempted to solve the two major drawbacks of occlusion-aware motion planning approaches based on virtual vehicles, namely, deadlock due to their worst-case assumptions, and their lack of road position planning functionality. In order to solve the deadlock problem, we proposed using the estimated visibility of hidden vehicles, and potential changes in their behavior after having observed the ego vehicle, to plan the ego vehicle's speed. For road position planning, we proposed first predicting and quantifying the visibility conditions of driving environments using HD maps consisting of 3D point cloud and road network maps and then using the quantified visibility conditions to plan the road position of the ego vehicle so that occlusions are minimized.

In Chapter 3, we first proposed a deadlock-free method for crossing low-visibility intersections with a mandatory stop, as deadlock is prone to occur at these types of intersections. The proposed method uses the estimated visibility of possible approaching vehicles to decide whether it is safe to trigger a predefined, special action for escaping deadlock, which involves proceeding slowly forward from the mandatory stop, where visibility is insufficient for a complete intersection crossing, in order to gain additional visibility. A comparison of speed profiles generated when using the proposed method to cross such intersections, and when an expert driver crossed them,

showed that the proposed planner could generate driving behavior similar to that of the expert at low-visibility intersections with a mandatory stop. Additionally, our experimental results showed that when the estimated visibility of approaching vehicles was not used for motion planning, the proposed method could not successfully cross the intersections due to deadlocks. This finding indicates that estimation of the visibility of possible approaching vehicles is crucial information for crossing such intersections without deadlock.

While the planner proposed in Chapter 3 could escape deadlocks at blind intersections with a mandatory stop, it relied on forward movement from the stop line which was triggered using estimated visibility of the ego vehicle by approaching vehicles. However, the proposed planner lacked a detailed analysis of the actual behavior of approaching vehicles after they observed the ego vehicle, i.e., the visibility-dependent behavior of the approaching vehicles. Therefore, in Chapter 4, we proposed a generic, deadlock-free motion planner that uses the visibility of both the ego vehicle and approaching vehicles to generate safe crossing motion at blind intersections. The proposed planner utilizes a particle filter algorithm for occluded vehicle prediction in order to support possible changes in the behavior of approaching traffic participants based on their ability to see the ego vehicle and to account for uncertainty in perception accuracy. Real driving data from multiple drivers crossing blind intersections in a residential area was also collected and analyzed in order to model the visibility-dependent behavior of occluded vehicles at intersections. We then compared the proposed generic, deadlock-free planner with a baseline planner that simply assumes the worst-case scenario, i.e., hidden vehicles are approaching from the occluded areas at a constant, high speed. The results of our comparison under various conditions showed that our proposed planner could generate deadlock-free crossing motion at a blind intersection of two narrow roads (each 5 meters wide), while the baseline planner could not. The effects of perception accuracy and sensor position on the output motion of the planner were also investigated. It was found that noisy perception generally caused the proposed planner to slightly delay intersection crossing action. Regarding the sensor mounting position, a front-mounted configuration resulted in a

less reduction in the ego vehicle speed and the elimination of a complete stop before crossing uncontrolled intersections, as compared to a sensor mounted at the center of the vehicle.

As a first step toward road position planning when encountering occluded areas, in Chapter 5, we proposed a visibility estimation method using HD maps, i.e., a combination of 3D point cloud and road network maps, that is applicable in complex driving environments. We also introduced an approach for quantifying the visibility conditions at a particular location with a numerical value, namely a visibility ratio. A 3D scan of a specified viewpoint is first approximated using a 3D point cloud map. The approximated scan and 3D points representing the area of the relevant surrounding lanes are then projected onto depth images. To identify the visible and occluded regions of the driving environment from the specified viewpoint, the resulting depth images are then compared. Finally, the visibility ratio is computed by dividing the visible area of interest by the total relevant driving area. The proposed visibility estimation approach was tested in both simulated and real-world terrain. Our results suggest that the visibility ratio calculated using the proposed method was representative of actual road visibility from a specified location within the HD maps. Furthermore, using the visibility estimation method proposed in Chapter 5, it is also possible to predict and quantify the visibility conditions of the road ahead. Consequently, this predicted visibility can be used for planning the optimal road position which will result in the least occlusion.

In Chapter 6, we proposed a motion planner capable of actively adjusting the ego vehicle’s lateral position, i.e., road position, in order to minimize occlusions in various traffic scenarios under perception uncertainty. We also extended an approach for 3D visibility estimation proposed in Chapter 5 to also include the use of live sensing data, so that it can handle occlusions caused by objects that are not present in the HD maps. First, candidate trajectories with different lateral offsets from the reference path are generated. The current and future visibility conditions along each trajectory are then estimated, allowing their visibility ratios to be calculated using the proposed extended visibility estimation approach. Finally, a visibility cost, computed

from the visibility ratio, is used along with other planning costs to determine the optimal output trajectory. The proposed planner was tested in various simulated traffic scenarios using both live localization and object detection results. Our experimental results showed that the ego vehicle was able to effectively minimize occlusions and consequently discover occluded vehicles earlier when the proposed planner was used, in most cases, in comparison to a baseline planner. Moreover, occlusions caused by a parked vehicle that was not present in the HD maps were correctly estimated when using our extended visibility estimation approach.

To unify the approaches proposed in this dissertation into a single system, one could use the road position planner described in Chapter 6 to first plan the reference trajectory that will result in maximum visibility of the driving environment, based on an approach for visibility prediction and quantification introduced in Chapter 5. Once a reference trajectory is obtained, the deadlock-free speed planner presented in Chapters 3 and 4 could then be used to regulate the ego vehicle’s motion along the reference trajectory. By planning the motion of the ego vehicle in this manner, the resulting motion should be efficient. The ego vehicle would not need to reduce its speed as much to safely traverse through occluded areas, due to the better visibility achieved by strategically adjusting its road position. As a result, deadlocks would also be less likely. In a situation where a deadlock is inevitable, the ego vehicle would still be able to escape an indefinite standstill by taking the visibility of hidden vehicles, and potential changes in their behavior after detecting the ego vehicle, into account.

Being able to ensure the safe navigation of autonomous vehicles in areas that contain occlusions will certainly broaden their operational domains, making the mass adoption of autonomous vehicles one step closer to reality. This will eventually translate into natural resources being used more efficiently and millions of invaluable lives saved — the potential societal benefits envisioned decades ago by many great minds when autonomous driving technology was first being developed.

7.2 Future Work

The speed planner proposed in Chapters 3 and 4 only considers potential occluded vehicles. In future research, other types of traffic participants such as pedestrians and cyclists should also be considered. The addition of other types of traffic participants should be as straightforward as modifying the motion model of the particle filter-based occluded object prediction module used in our proposed speed planner.

Moreover, in Chapter 4, the behavior of hidden traffic participants is modeled to be affected mainly by their visibility of the ego vehicle. Future improvement of the proposed behavior model should take into account other actions or events that are common in driving scenarios and can significantly influence the motion of traffic participants. For example, possible deceleration by an approaching cyclist or vehicle after the ego vehicle gives an audible warning with the horn should be considered. Similarly, other factors such, as the right-of-way, can also be incorporated into the behavior model of hidden traffic participants.

Although motion planning of the ego vehicle at an intersection with a mandatory stop is already discussed in Chapter 3, a mandatory stop is only one of many traffic rules that directly affect how the ego vehicle should navigate in the environment. Other traffic regulations, such as traffic lights which are very common at intersections, should also be considered in future research to increase feasibility in real-world applications of the motion planner.

The visibility estimation method proposed in this dissertation still has much room for improvement, since our approach exclusively utilizes sensing data from a 3D LiDAR to identify the visible and occluded regions of driving environments. Modern autonomous vehicles typically have multimodal sensing capabilities, i.e., they are often equipped with several types of sensors in order to provide redundancy or complementarity to the system, therefore, visibility estimation methods should be able to make use of data from various types of sensors. One possible way to achieve such fusion is to use Bayesian inference. Regions of interest in the driving environment can first be identified as either visible or occluded using data from an individual sensing

modality. Then, individual results from all of the sensors can be fused together, taking into consideration the accuracy of each sensor for estimating visibility, in order to achieve reliable, final estimation results. This fusion can also be applied to other sources of information, such as information from other vehicles via V2V communication, or indirect observations of the driving environment via the convex mirrors which are typically placed near occluded intersections to improve visibility [88].

Additionally, the visibility estimation approach in this dissertation focuses only on occlusions caused by surrounding objects. Nevertheless, many other factors such as defective sensors or adverse weather conditions have been reported in several studies [83–86] to have caused degradation of visibility or occlusions in driving environments. These potential sources of occlusion should be considered in future investigations as well.

The proposed road position planner could be further enhanced by taking into account predicted future trajectories of surrounding traffic participants, both visible and occluded, as they directly influence the future visibility of the ego vehicle. Finally, instead of calculating visibility ratios entirely online, it would be more computationally efficient to precalculate them using only point cloud and road network maps, embed the results in the road network map, and then dynamically update the ratios again online using live 3D scans to reflect the changes in visibility caused by objects not present in the point cloud map.

Bibliography

- [1] S. Singh, “Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey,” National Center for Statistics and Analysis, Washington, DC, Tech. Rep., 2018. [Online]. Available: <https://trid.trb.org/view/1507603>
- [2] D. J. Fagnant and K. Kockelman, “Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations,” *Transportation Research Part A: Policy and Practice*, vol. 77, pp. 167–181, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.tra.2015.04.003>
- [3] Z. Wadud, D. MacKenzie, and P. Leiby, “Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles,” *Transportation Research Part A: Policy and Practice*, vol. 86, pp. 1–18, 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0965856415002694>
- [4] C. Y. Chan, “Advancements, prospects, and impacts of automated driving systems,” *International Journal of Transportation Science and Technology*, vol. 6, no. 3, pp. 208–216, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2046043017300035>
- [5] S. Tsugawa, T. Yatabe, T. Hirose, and S. Matsumoto, “An automobile with artificial intelligence,” in *International Joint Conference on Artificial Intelligence*, Aug. 1979, pp. 893–895.
- [6] E. D. Dickmanns, “Vehicles capable of dynamic vision,” in *International Joint Conference on Artificial Intelligence*, Aug. 1997, pp. 1577–1592. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1622270.1622386>
- [7] P. Dean and J. Todd, “No Hands Across America Journal,” 1995, Accessed: 2022/10/13. [Online]. Available: https://www.cs.cmu.edu/~tjochem/nhaa/nhaa_home_page.html
- [8] M. Buehler, K. Iagnemma, and S. Singh, Eds., *The 2005 DARPA Grand Challenge*, ser. Springer Tracts in Advanced Robotics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, vol. 36. [Online]. Available: <http://link.springer.com/10.1007/978-3-540-73429-1>

- [9] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L. E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekerk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney, “Stanley: The robot that won the DARPA Grand Challenge,” *Springer Tracts in Advanced Robotics*, vol. 36, pp. 1–43, 2007.
- [10] M. Buehler, K. Iagnemma, and S. Singh, Eds., *The DARPA Urban Challenge*, ser. Springer Tracts in Advanced Robotics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, vol. 56. [Online]. Available: <http://link.springer.com/10.1007/978-3-642-03991-1>
- [11] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y. W. Seo, S. Singh, J. Snider, A. Stentz, W. Whittaker, Z. Wolkowicki, J. Zigar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson, “Autonomous driving in Urban environments: Boss and the Urban Challenge,” in *Springer Tracts in Advanced Robotics*, 2009, pp. 1–59.
- [12] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun, “Junior: The stanford entry in the urban challenge,” in *Springer Tracts in Advanced Robotics*, vol. 56, 2009, pp. 91–123.
- [13] C. Reinholtz, D. Hong, A. Wicks, A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, T. Alberi, D. Anderson, S. Cacciola, P. Currier, A. Dalton, J. Farmer, J. Hurdus, S. Kimmel, P. King, A. Taylor, D. Van Covern, and M. Webster, “Odin: Team VictorTango’s entry in the DARPA Urban Challenge,” in *Springer Tracts in Advanced Robotics*, 2009, pp. 125–162.
- [14] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams, “A perception-driven autonomous urban vehicle,” in *Springer Tracts in Advanced Robotics*, 2009, pp. 163–230.
- [15] J. Bohren, T. Foote, J. Keller, A. Kushleyev, D. Lee, A. Stewart, P. Vernaza, J. Derenick, J. Spletzer, and B. Satterfield, “Little Ben: The Ben Franklin Racing Team’s entry in the 2007 DARPA Urban Challenge,” in *Springer Tracts in Advanced Robotics*, 2009, pp. 231–255.

- [16] I. Miller, M. Campbell, D. Huttenlocher, A. Nathan, F. R. Kline, P. Moran, N. Zych, B. Schimpf, S. Lupashin, E. Garcia, J. Catlin, M. Kurdziel, and H. Fujishima, “Team Cornell’s Skynet: Robust perception and planning in an urban environment,” in *Springer Tracts in Advanced Robotics*, 2009, pp. 257–304.
- [17] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, “A Survey of Autonomous Driving: Common Practices and Emerging Technologies,” *IEEE Access*, vol. 8, pp. 58 443–58 469, 2020.
- [18] D. González, J. Pérez, V. Milanés, and F. Nashashibi, “A Review of Motion Planning Techniques for Automated Vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 4, pp. 1135–1145, 2016.
- [19] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots,” in *1985 IEEE International Conference on Robotics and Automation (ICRA)*, Mar. 1985, pp. 500–505.
- [20] D. Fox, W. Burgard, and S. Thrun, “The dynamic window approach to collision avoidance,” *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, Mar. 1997. [Online]. Available: <http://ieeexplore.ieee.org/document/580977/>
- [21] H. Darweesh, E. Takeuchi, K. Takeda, Y. Ninomiya, A. Sujiwo, L. Y. Morales, N. Akai, T. Tomizawa, and S. Kato, “Open Source Integrated Planner for Autonomous Navigation in Highly Dynamic Environments,” *Journal of Robotics and Mechatronics*, vol. 29, no. 4, pp. 668–684, Aug. 2017. [Online]. Available: <https://www.fujipress.jp/jrm/rb/robot002900040668/>
- [22] X. Zhao, J. Wang, G. Yin, and K. Zhang, “Cooperative Driving for Connected and Automated Vehicles at Non-signalized Intersection based on Model Predictive Control,” in *2019 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2019, pp. 2121–2126.
- [23] M. R. Hafner, D. Cunningham, L. Caminiti, and D. Del Vecchio, “Cooperative collision avoidance at intersections: Algorithms and experiments,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1162–1175, 2013.
- [24] V. Milanés, J. Villagrà, J. Godoy, J. Simó, J. Pérez, and E. Onieva, “An intelligent V2I-based traffic management system,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 1, pp. 49–58, 2012.
- [25] V. Narri, A. Alanwar, J. Mårtensson, C. Norén, L. Dal Col, and K. H. Johansson, “Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios,” in *2021 IEEE Intelligent Vehicles Symposium (IV)*, Jul. 2021, pp. 385–392.
- [26] J. Müller, J. Strohbeck, M. Herrmann, and M. Buchholz, “Motion Planning for Connected Automated Vehicles at Occluded Intersections With Infrastructure

- Sensors,” *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–12, 2022.
- [27] L. Kloeker, C. Geller, A. Kloeker, and L. Eckstein, “High-Precision Digital Traffic Recording with Multi-LiDAR Infrastructure Sensor Setups,” in *2020 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Sep. 2020, pp. 1–8.
- [28] F. Geissler and R. Gräfe, “Optimized sensor placement for dependable roadside infrastructures,” in *2019 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2019, pp. 2408–2413.
- [29] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, “Navigating Occluded Intersections with Autonomous Vehicles Using Deep Reinforcement Learning,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 2034–2039.
- [30] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, “Safe Reinforcement Learning with Scene Decomposition for Navigating Complex Urban Environments,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 1469–1476.
- [31] D. Kamran, C. F. Lopez, M. Lauer, and C. Stiller, “Risk-Aware High-level Decisions for Automated Driving at Occluded Intersections with Reinforcement Learning,” in *2020 IEEE Intelligent Vehicles Symposium (IV)*, Oct. 2020, pp. 1205–1212.
- [32] Z. Qiao, K. Muelling, J. Dolan, P. Palanisamy, and P. Mudalige, “POMDP and Hierarchical Options MDP with Continuous Actions for Autonomous Driving at Intersections,” in *2018 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Nov. 2018, pp. 2377–2382.
- [33] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, “Soft actor-critic algorithms and applications,” 2018, Accessed: 2022/10/13. [Online]. Available: <https://arxiv.org/abs/1812.05905>
- [34] Y. Yoshihara, E. Takeuchi, Y. Ninomiya, and T. R. Board, “Accurate Analysis of Expert and Elderly Driving at Blind Corners for Proactive Advanced Driving Assistance Systems,” in *Transportation Research Board 95th Annual Meeting*, Jan. 2016, p. 19. [Online]. Available: <https://trid.trb.org/view/1392723>
- [35] Y. Akagi and P. Raksincharoensak, “Stochastic driver speed control behavior modeling in urban intersections using risk potential-based motion planning framework,” in *2015 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2015, pp. 368–373.

- [36] Y. Morales, Y. Yoshihara, N. Akai, E. Takeuchi, and Y. Ninomiya, “Proactive driving modeling in blind intersections based on expert driver data,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 901–907.
- [37] L. Y. Morales, A. Naoki, Y. Yoshihara, and H. Murase, “Towards Predictive Driving through Blind Intersections,” in *2018 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Nov. 2018, pp. 716–722.
- [38] K. Sama, Y. Morales, H. Liu, N. Akai, A. Carballo, E. Takeuchi, and K. Takeda, “Extracting Human-Like Driving Behaviors From Expert Driver Data Using Deep Learning,” *IEEE Transactions on Vehicular Technology*, vol. 69, no. 9, pp. 9315–9329, Sep. 2020.
- [39] S. Brechtel, T. Gindele, and R. Dillmann, “Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs,” in *2014 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2014, pp. 392–399.
- [40] X. Lin, J. Zhang, J. Shang, Y. Wang, H. Yu, and X. Zhang, “Decision Making through Occluded Intersections for Autonomous Driving,” in *2019 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2019, pp. 2449–2455.
- [41] P. Schörner, L. Töttel, J. Doll, and J. M. Zöllner, “Predictive Trajectory Planning in Situations with Hidden Road Users Using Partially Observable Markov Decision Processes,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 2299–2306.
- [42] D. Klimenko, J. Song, and H. Kurniawati, “TAPIR: A software Toolkit for approximating and adapting POMDP solutions online,” in *Australasian Conference on Robotics and Automation (ACRA)*, Dec. 2014.
- [43] C. Hubmann, N. Quetschlich, J. Schulz, J. Bernhard, D. Althoff, and C. Stiller, “A POMDP Maneuver Planner For Occlusions in Urban Scenarios,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 2172–2179.
- [44] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, “Scalable Decision Making with Sensor Occlusions for Autonomous Driving,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 2076–2081.
- [45] M. Schratte, M. Bouton, M. J. Kochenderfer, and D. Watzenig, “Pedestrian Collision Avoidance System for Scenarios with Occlusions,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 1054–1060.
- [46] K. H. Wray, B. Lange, A. Jamgochian, S. J. Witwicki, A. Kobashi, S. Hagari-bommanahalli, and D. Ilstrup, “POMDPs for Safe Visibility Reasoning in Autonomous Vehicles,” in *2021 IEEE International Conference on Intelligence and Safety for Robotics (ISR)*, Mar. 2021, pp. 191–195.

- [47] Y. Yoshihara, Y. Morales, N. Akai, E. Takeuchi, and Y. Ninomiya, “Autonomous predictive driving for blind intersections,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 3452–3459.
- [48] O. S. Tas and C. Stiller, “Limited Visibility and Uncertainty Aware Motion Planning for Automated Driving,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2018, pp. 1171–1178.
- [49] M. Lee, K. Jo, and M. Sunwoo, “Collision risk assessment for possible collision vehicle in occluded area based on precise map,” in *2017 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Oct. 2017, pp. 1–6.
- [50] R. Poncelet, A. Verroust-Blondet, and F. Nashashibi, “Safe Geometric Speed Planning Approach for Autonomous Driving through Occluded Intersections,” in *2020 International Conference on Control, Automation, Robotics and Vision (ICARCV)*, Dec. 2020, pp. 393–399.
- [51] E. Takeuchi, Y. Yoshihara, and N. Yoshiki, “Blind Area Traffic Prediction Using High Definition Maps and LiDAR for Safe Driving Assist,” in *2015 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Sep. 2015, pp. 2311–2316.
- [52] M. Yu, R. Vasudevan, and M. Johnson-Roberson, “Occlusion-Aware Risk Assessment for Autonomous Driving in Urban Environments,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2235–2241, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8645673>
- [53] P. F. Orzechowski, A. Meyer, and M. Lauer, “Tackling Occlusions & Limited Sensor Range with Set-based Safety Verification,” in *2018 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Nov. 2018, pp. 1729–1736. [Online]. Available: <https://ieeexplore.ieee.org/document/8569332/>
- [54] M. Althoff and S. Magdici, “Set-Based Prediction of Traffic Participants on Arbitrary Road Networks,” *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 2, pp. 187–202, Jun. 2016.
- [55] M. Naumann, H. Konigshof, M. Lauer, and C. Stiller, “Safe but not Overcautious Motion Planning under Occlusions and Limited Sensor Range,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 140–145. [Online]. Available: <https://ieeexplore.ieee.org/document/8814251/>
- [56] S. Lee, W. Lim, M. Sunwoo, and K. Jo, “Limited Visibility Aware Motion Planning for Autonomous Valet Parking Using Reachable Set Estimation,” *Sensors*, vol. 21, no. 4, p. 1520, Jan. 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/4/1520>
- [57] E. Debada, A. Ung, and D. Gillet, “Occlusion-Aware Motion Planning at Roundabouts,” *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 2, pp. 276–287, Jun. 2021.

- [58] M. Yu, R. Vasudevan, and M. Johnson-Roberson, “Risk Assessment and Planning with Bidirectional Reachability for Autonomous Driving,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 5363–5369.
- [59] G. Neel and S. Saripalli, “Improving Bounds on Occluded Vehicle States for Use in Safe Motion Planning,” in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, Nov. 2020, pp. 268–275.
- [60] L. Wang, C. F. Lopez, and C. Stiller, “Generating Efficient Behaviour with Predictive Visibility Risk for Scenarios with Occlusions,” in *2020 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, Sep. 2020, pp. 1–7. [Online]. Available: <https://ieeexplore.ieee.org/document/9294403/>
- [61] Z. Zhang and J. F. Fisac, “Safe Occlusion-aware Autonomous Driving via Game-Theoretic Active Perception,” in *Proceedings of Robotics: Science and Systems*, Jul. 2021. [Online]. Available: <http://arxiv.org/abs/2105.08169>
- [62] M. Koschi and M. Althoff, “Set-Based Prediction of Traffic Participants Considering Occlusions and Traffic Rules,” *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 2, pp. 249–265, Jun. 2021.
- [63] L. Wang, C. Burger, and C. Stiller, “Reasoning about Potential Hidden Traffic Participants by Tracking Occluded Areas,” in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Sep. 2021, pp. 157–163.
- [64] Y. Fujinami, P. Raksincharoensak, D. Ulbricht, and R. Adomat, “Risk Predictive Driver Assistance System for Collision Avoidance in Intersection Right Turns,” *Journal of Robotics and Mechatronics*, vol. 30, no. 1, pp. 15–23, Feb. 2018. [Online]. Available: <https://www.fujipress.jp/jrm/rb/robot003000010015/>
- [65] H. Andersen, J. Alonso-Mora, Y. H. Eng, D. Rus, and M. H. Ang, “Trajectory Optimization and Situational Analysis Framework for Autonomous Overtaking With Visibility Maximization,” *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 1, pp. 7–20, Mar. 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/8911227/>
- [66] M. Koschi and M. Althoff, “Set-based Prediction of Traffic Participants Considering Occlusions and Traffic Rules,” *IEEE Transactions on Intelligent Vehicles*, pp. 1–1, 2020.
- [67] F. Damerow, T. Puphal, Y. Li, and J. Eggert, “Risk-based driver assistance for approaching intersections of limited visibility,” in *2017 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, Jun. 2017, pp. 178–184.

- [68] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, “OctoMap: an efficient probabilistic 3D mapping framework based on octrees,” *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, Apr. 2013. [Online]. Available: <https://doi.org/10.1007/s10514-012-9321-0>
- [69] E. Takeuchi, Y. Ninomiya, and S. Kato, “Lane visibility check methods based on high precision maps and 3D LiDAR for traffic prediction,” in *FAST-zero’15: 3rd International Symposium on Future Active Safety Technology Toward zero traffic accidents*, 2015, pp. pp 61–66. [Online]. Available: <https://trid.trb.org/view/1412203>
- [70] S. Hoermann, F. Kunz, D. Nuss, S. Renter, and K. Dietmayer, “Entering cross-roads with blind corners. A safe strategy for autonomous vehicles,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 727–732.
- [71] M. Treiber and D. Helbing, “Explanation of Observed Features of Self-Organization in Traffic Flow,” Jan. 1999, Accessed: 2022/10/13. [Online]. Available: <http://arxiv.org/abs/cond-mat/9901239>
- [72] E. Takeuchi and T. Tsubouchi, “A 3-D Scan Matching using Improved 3-D Normal Distributions Transform for Mobile Robotic Mapping,” in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2006, pp. 3068–3073.
- [73] P. Bokare and A. Maurya, “Acceleration-Deceleration Behaviour of Various Vehicle Types,” *Transportation Research Procedia*, vol. 25, pp. 4733–4749, 2017. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S2352146517307937>
- [74] H. Summala, “Brake Reaction Times and Driver Behavior Analysis,” *Transportation Human Factors*, vol. 2, no. 3, pp. 217–226, Sep. 2000. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1207/STHF0203_2
- [75] Y. Akagi and P. Raksincharoensak, “An analysis of an elderly driver behaviour in urban intersections based on a risk potential model,” in *2015 Annual Conference of the IEEE Industrial Electronics Society (IECON)*, Nov. 2015, pp. 1627–1632. [Online]. Available: <http://ieeexplore.ieee.org/document/7392334/>
- [76] D. Arthur and S. Vassilvitskii, “K-means++: The advantages of careful seeding,” in *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, Jan. 2007, pp. 1027–1035.
- [77] P. J. Rousseeuw, “Silhouettes: A graphical aid to the interpretation and validation of cluster analysis,” *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987.
- [78] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, “CARLA: An Open Urban Driving Simulator,” in *Proceedings of the 1st Annual*

Conference on Robot Learning, Oct. 2017, pp. 1–16. [Online]. Available: <https://proceedings.mlr.press/v78/dosovitskiy17a.html>

- [79] T. Shan and B. Englot, “LeGO-LOAM: Lightweight and Ground-Optimized Lidar Odometry and Mapping on Variable Terrain,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 4758–4765.
- [80] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, “LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 5135–5142.
- [81] M. Billah, A. Maskooki, F. Rahman, and J. A. Farrell, “Roadway feature mapping from point cloud data: A graph-based clustering approach,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 475–480. [Online]. Available: <http://ieeexplore.ieee.org/document/7995763/>
- [82] D. Zai, J. Li, Y. Guo, M. Cheng, Y. Lin, H. Luo, and C. Wang, “3-D Road Boundary Extraction From Mobile Laser Scanning Data via Supervoxels and Graph Cuts,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 802–813, Mar. 2018. [Online]. Available: <http://ieeexplore.ieee.org/document/7938403/>
- [83] K. Yoneda, N. Suganuma, R. Yanase, and M. Aldibaja, “Automated driving recognition technologies for adverse weather conditions,” *IATSS Research*, vol. 43, no. 4, pp. 253–262, Dec. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0386111219301463>
- [84] A. Carballo, J. Lambert, A. Monrroy, D. Wong, P. Narksri, Y. Kitsukawa, E. Takeuchi, S. Kato, and K. Takeda, “LIBRE: The Multiple 3D LiDAR Dataset,” in *2020 IEEE Intelligent Vehicles Symposium (IV)*, Oct. 2020, pp. 1094–1101.
- [85] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, “Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020, pp. 11 679–11 689. [Online]. Available: <https://ieeexplore.ieee.org/document/9157107/>
- [86] R. Heinzler, P. Schindler, J. Seekircher, W. Ritter, and W. Stork, “Weather Influence and Classification with Automotive Lidar Sensors,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 1527–1534.
- [87] S. Kato, E. Takeuchi, Y. Ishiguro, Y. Ninomiya, K. Takeda, and T. Hamada, “An Open Approach to Autonomous Vehicles,” *IEEE Micro*, vol. 35, no. 6, pp. 60–68, Nov. 2015.

- [88] C. Zhang, F. Steinhauser, G. Hinz, and A. Knoll, “Traffic Mirror-Aware POMDP Behavior Planning for Autonomous Urban Driving,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2022, pp. 323–330. [Online]. Available: <https://ieeexplore.ieee.org/document/9827139/>