



HAL
open science

Estimation du risque aux intersections pour applications sécuritaires avec véhicules communicants

Stéphanie Lefèvre

► **To cite this version:**

Stéphanie Lefèvre. Estimation du risque aux intersections pour applications sécuritaires avec véhicules communicants. Autre [cs.OH]. Université de Grenoble, 2012. Français. NNT : 2012GRENM070 . tel-00858906

HAL Id: tel-00858906

<https://theses.hal.science/tel-00858906>

Submitted on 6 Sep 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE

Spécialité : **Mathématiques-Informatique**

Arrêté ministériel : 7 août 2006

Présentée par

Stéphanie Lefèvre

Thèse dirigée par **Christian Laugier**
et codirigée par **Javier Ibañez-Guzmán**

préparée au sein de **Inria Grenoble Rhône-Alpes**
et de **Ecole Doctorale Mathématiques, Sciences et Technologies de l'Information, Informatique**

Risk Estimation at Road Intersections for Connected Vehicle Safety Applications

Thèse soutenue publiquement le **22 octobre 2012**,
devant le jury composé de :

M. Augustin Lux

Professeur, Institut National Polytechnique Grenoble, France, Président

M. Roland Chapuis

Professeur, Université Blaise Pascal, France, Rapporteur

M. Christoph Stiller

Professeur, Karlsruher Institut für Technologie, Allemagne, Rapporteur

M. Urbano Nunes

Professeur, Universidade de Coimbra, Portugal, Examineur

M. Christian Laugier

Directeur de Recherche, Inria Grenoble, France, Directeur de thèse

M. Javier Ibañez-Guzmán

Ingénieur de Recherche, Renault S.A.S., France, Co-Directeur de thèse



Abstract

Intersections are the most complex and dangerous areas of the road network. Statistics show that most road intersection accidents are caused by driver error and that many of them could be avoided through the use of Advanced Driver Assistance Systems. In this context, vehicular communications are a very promising technology. The sharing of information between vehicles over wireless links allows vehicles to perceive their environment beyond the field-of-view of their on-board sensors. Thanks to this enlarged representation of the environment in time and space, situation assessment is improved and dangerous situations can be detected earlier.

This thesis tackles the risk estimation problem from a new perspective: a framework is proposed for reasoning about traffic situations and collision risk at a semantic level, while classic approaches typically reason at a trajectory level. Risk is assessed by estimating the intentions of drivers and detecting conflicts between them, rather than by predicting the future trajectories of the vehicles and detecting intersections between them. More specifically, dangerous situations are identified by comparing *what drivers intend to do* with *what drivers are expected to do according to the traffic rules*. The reasoning about *intentions* and *expectations* is performed in a probabilistic manner, in order to take into account sensor uncertainties and interpretation ambiguities.

This framework can in theory be applied to any type of traffic situation; here we present its application to the specific case of road intersections. The proposed motion model takes into account the mutual influences between the maneuvers performed by vehicles at an intersection. It also incorporates information about the influence of the geometry and topology of the intersection on the behavior of a vehicle. The approach was validated with field trials using passenger vehicles equipped with Vehicle-to-Vehicle wireless communication modems, and in simulation. The results demonstrate that the algorithm is able to detect dangerous situations early and complies with real-time constraints.

Résumé

Les intersections sont les zones les plus dangereuses du réseau routier. Les études d'accidentologie montrent que la plupart des accidents aux intersections sont causés par des erreurs des conducteurs, et qu'une majorité d'accidents pourraient être évités grâce à l'utilisation de systèmes d'aide à la conduite. Dans ce cadre, les communications inter-véhiculaire sont une technologie particulièrement prometteuse. Le partage d'informations entre les véhicules via des liens sans fil permet à chaque véhicule de percevoir son environnement au-delà des limites de champ de vision de ses capteurs embarqués. Il en résulte une représentation de l'environnement plus étendue dans l'espace et dans le temps, ce qui améliore la compréhension de situation et permet d'anticiper le danger.

Cette thèse aborde le problème de l'estimation du risque sous un angle nouveau : elle propose une structure de raisonnement pour analyser les scènes routières et le risque de collision à un niveau sémantique, contrairement aux approches classiques qui raisonnent au niveau des trajectoires. Le risque est calculé en estimant les intentions des conducteurs et en détectant les conflits, sans avoir à prédire les trajectoires futures des véhicules. Plus précisément, la détection des situations dangereuses est basée sur la comparaison entre *ce que les conducteurs ont l'intention de faire* et *ce que les conducteurs devraient faire d'après les règles de la circulation*. Ce raisonnement est réalisé de manière probabiliste afin de prendre en compte les incertitudes sur les mesures capteur et les ambiguïtés sur l'interprétation de la scène.

En théorie ce raisonnement peut être appliqué à tout type de scène routière ; dans cette thèse nous présentons son application aux intersections. Le modèle proposé prend en compte l'influence que la manœuvre d'un véhicule exerce sur la manœuvre des autres véhicules. Il incorpore aussi des informations sur l'influence de la géométrie et topologie de l'intersection sur le comportement d'un véhicule. L'approche proposée a été validée par des tests en environnement réel avec des véhicules communicants, ainsi qu'en simulation. Les résultats montrent que l'algorithme est capable de détecter les situations dangereuses et qu'il est compatible avec des applications sécuritaires temps-réel.

Acknowledgements

Over the course of this PhD I worked in three different places. In each of them I have met people who had a great influence on my professional and personal life, and today I feel both lucky and grateful for this opportunity.

I would like to start by thanking my two supervisors: Christian Laugier from Inria and Javier Ibañez-Guzmán from Renault. Thank you for your constant support, for your patience with me (I know I can be stubborn...), and for encouraging me when my self-confidence was low. I learned a lot from you both, not only about research but also about random things like feeding cats and fighting sharks. The guidance provided by the two of you was complementary in many respects, and as a result this PhD really was a collaboration between Inria and Renault. I benefited from the expertise and resources of both: research methods and Bayesian formalism at Inria and driver assistance systems and test vehicles at Renault. I realize how lucky I am to have had access to so much knowledge and so many tools during my PhD.

The first year of my PhD was spent at Renault learning about driver assistance systems and participating in the Safespot project. Thanks to Sylvain the long hours spent in cars running tests in the middle of winter were enjoyable. I am grateful to Olivier for passing along to me the precious knowledge he had accumulated during his PhD working with cars. Many thanks to the ADAS team for sharing their knowledge and for the good atmosphere at work: Sébastien, Charles, Javier, Christian S., Florent, Christophe, Christian C., Gaetan, Jean-François, Jean-Philippe, Johann, Stéphanie, and everybody else.

At the end of the Safespot project I went to the Stanford Artificial Intelligence Laboratory for an internship with Sebastian Thrun's Driving Group. This was a great experience, and I would like to thank the whole team for welcoming me and sharing their tools. I would also like to thank Javier, Christian, and François for making this happen.

The next year was spent at Inria, where I could not have hoped for better team mates. We shared

offices but also thoughts, fun, and lots of food. Thank you for proof-reading my papers, for helping me improve my presentations, for lightening my mood in difficult moments. Mathias, Alexandros, Amaury, Igor, Alessandro, Manu, Pierre, Dizan, Christian, Myriam, Ago, Anne, Arturo, Vivien, Omar, Panos, Nicolas, Qadeer, Chiara, Jorge, and everybody else, thank you for being a wonderful team. Special thanks to my office mates Chiara and Jorge, for bringing good mood (and good food) in the office again and again and again. To Pierre, Manu, Dizan, Mathias, and Alexandros, for many hours discussing Bayesian methods. To Myriam, for always finding solutions to our various administrative issues, and for smiling so much.

Going back to Renault for the implementation of my algorithms and the tests with the real cars, I found that I was no longer the only PhD student in the team. Laëtitia and Clément had joined the group, and the “PhD student team spirit” that ensued was a great help during that period. Also, a major part of this thesis would be missing if not for Jong-Hoon, Laëtitia, Clément, Nabil, Jérémy, and Jonathan. In total you ran 110 tests for me in La Brosse. Sometimes it rained, sometimes the sun shone, but everytime you were there for me and I am extremely grateful for that.

I would also like to thank my jury members for making the effort to read and review this thesis, for taking a long trip to attend the defense, and for giving me useful feedback.

Finally I would like to thank my family. My parents, for showing me in many different ways that they care for me, for supporting my life choices, and for always being there for me. My brothers, sister, brother-in-law, and sisters-in-law: I always feel close to you, even when we are far apart. Mike, thank you for showing me that I can become who I want to be. Everyday you were a source of motivation, and much more.

Contents

List of Figures	xv
List of Acronyms	xxi
Chapter 1 Introduction	1
1.1 Context: connected vehicles for road safety	2
1.1.1 Accidents at road intersections	2
1.1.2 Wireless vehicular communications	6
1.2 Problem formulation	17
1.2.1 Scope	18
1.2.2 Challenges and classic approaches	20
1.3 Contributions	22
1.4 Thesis outline	23
Chapter 2 Risk assessment at road intersections: state of the art	25
2.1 ‘Physical entities’ motion models	28
2.1.1 Evolution models	29
2.1.2 Trajectory prediction	30
2.1.3 Limitations	33
2.2 ‘Maneuvering entities’ motion models	33
2.2.1 Prototype trajectories	34
2.2.2 Maneuver intention estimation and maneuver execution	37

2.2.3	Limitations	40
2.3	'Interacting maneuvering entities' motion models	40
2.3.1	Models based on trajectory prototypes	40
2.3.2	Models based on Dynamic Bayesian Networks	41
2.3.3	Limitations	42
2.4	Collision detection and risk computation	42
2.4.1	Collision detection	42
2.4.2	Risk computation	44
2.5	Other approaches and conclusions	46
2.5.1	Other approaches	46
2.5.2	Conclusions	47
Chapter 3 Proposed motion model and risk estimation		49
3.1	Overview of the proposed approach	51
3.1.1	Background: representing vehicle motion with MSSMs	51
3.1.2	Proposed approach for general traffic situations	56
3.1.3	Application to road intersections	59
3.2	Variable definition	60
3.2.1	Intended maneuver I_t^n	60
3.2.2	Expected maneuver E_t^n	62
3.2.3	Measurements Z_t^n	63
3.2.4	Physical state Φ_t^n	63
3.2.5	Summary and notations	64
3.3	Joint distribution	65
3.4	Parametric forms	67
3.4.1	Expected longitudinal motion Es_t^n	67
3.4.2	Intended longitudinal motion Is_t^n	69
3.4.3	Intended lateral motion Ic_t^n	70
3.4.4	Pose P_t^n	71

3.4.5	Speed S_t^n	72
3.4.6	Turn signal state T_t^n	74
3.4.7	Measured pose Pm_t^n	77
3.4.8	Measured speed Sm_t^n	77
3.4.9	Measured turn signal state Tm_t^n	78
3.4.10	Summary	78
3.5	Bayesian risk estimation	78
3.6	Conclusions	79
3.6.1	Summary	79
3.6.2	Discussion	79
Chapter 4	Experiments and results	83
4.1	Functional and safety evaluation in simulation	85
4.1.1	Evaluation strategy	85
4.1.2	Functional performance evaluation	90
4.1.3	Safety performance evaluation	93
4.2	Robustness evaluation with field trials	96
4.2.1	Evaluation strategy	96
4.2.2	Results	103
4.3	The importance of context - Interactions between vehicles	106
4.3.1	Evaluation strategy	108
4.3.2	Qualitative evaluation	109
4.3.3	Statistical results	113
4.4	The importance of context - Intersection layout	115
4.4.1	Evaluation strategy	116
4.4.2	Qualitative evaluation	118
4.4.3	Statistical results	120
4.5	Conclusions	123
4.5.1	Summary	123
4.5.2	Discussion	124

Chapter 5 Conclusions	127
5.1 Summary	127
5.2 Conclusions	129
5.3 Perspectives	130
Bibliography	135
Appendix A DSRC standards	147
Appendix B The importance of context	151
B.1 Intersection layout	152
B.2 Interactions between vehicles	152
B.2.1 Intuition	153
B.2.2 Case study: non-dangerous scenario	154
B.2.3 Case study: dangerous scenario	157
B.3 Summary	158
Appendix C Digital maps	159
C.1 Definitions	159
C.2 High-level features	160
C.3 XML code	161
Appendix D Precision and recall analysis	167
D.1 Metrics	167
D.2 Results	168
Appendix E Accounting for physical limitations in the evolution model	171
Appendix F Inference for risk estimation	175
F.1 Variables	176
F.2 Initialization	176
F.3 Recursive risk computation	177

List of Figures

1.1	Distribution of road accidents at intersections and out of intersections [1].	3
1.2	Distribution of common accident scenarios in Germany and France [2].	5
1.3	Distribution of accident causes [1].	6
1.4	Impact of passive safety measures on the reduction of the number of accidents [3]. . .	7
1.5	Expected reduction in the number of accidents, and associated technologies [4].	8
1.6	Set of applications targeted by the DRIVE-C2X project [5].	10
1.7	V2X communications (in red) overcome the field-of-view limitations of on-board sensors (in green) at road intersections.	11
1.8	Set of applications targeted by the CICAS project [6].	12
1.9	Cellular-based wireless vehicular communications.	14
1.10	DSRC-based wireless vehicular communications.	14
1.11	Satellite-based wireless vehicular communications.	15
1.12	Comparison of the different wireless communication technologies [7].	15
1.13	Forecast deployment of DSRC [8].	17
1.14	Intelligent vehicle architecture, based on the JDL model [9]. The hatched area defines the focus of this thesis.	18
2.1	Examples of motion prediction with the different types of motion models.	28
2.2	Trajectory prediction with a Constant Velocity motion model (source: [10]).	31
2.3	Trajectory prediction with a Constant Velocity motion model and Gaussian noise simulation. Ellipses represent the uncertainty on the predicted positions (source: [11]). . .	32
2.4	Trajectory prediction (green area) with Monte Carlo simulation (source: [12]).	33

2.5	Clustered trajectories: each cluster corresponds to a typical motion pattern (source: [13]).	34
2.6	Comparative review of works on maneuver intention estimation.	38
2.7	Maneuver execution using four different approaches.	39
2.8	Trajectory prediction with an ‘Interacting maneuvering entities’ motion model. The model is able to predict that Agent 35 should yield to Agent 17 (source: [14]).	41
3.1	Graphical representation of a base MSSM modeling the motion of a vehicle n between time $t - 1$ and t	52
3.2	Graphical representation of a HMM-based motion model for a vehicle n	53
3.3	Graphical representation of a MSSM which accounts for the dependencies between the vehicles. Bold arrows represent multi-vehicle dependencies, i.e. the influences of the other vehicles on vehicle n	54
3.4	Graphical representation of a CHMM-based motion model.	55
3.5	Graphical representation of the proposed MSSM. Bold arrows represent multi-vehicle dependencies, i.e. the influences of the other vehicles on vehicle n	56
3.6	List of accident scenarios addressed at X-shaped intersections.	60
3.7	Illustrative example for the "course" concept. The courses originating from one road are displayed as blue arrows.	61
3.8	Detection of dangerous situations by comparing <i>intention</i> Is_t^n and <i>expectation</i> Es_t^n : example scenarios.	65
3.9	Computation of $P(Es_t^n \mathbf{Ic}_{t-1} \mathbf{P}_{t-1} \mathbf{S}_{t-1})$ in an example situation.	69
3.10	Example average and maximum speed profiles generated for a maneuver ($c_t^n, is_t^n = stop$) and for a maneuver ($c_t^n, is_t^n = go$).	73
3.11	Example calculation of $\mu_s(s_{t-1}^n, p_{t-1}^n, c_t^n, is_t^n)$ following Equation 3.24, for two different values of the previous speed s_{t-1}^n	74
3.12	Example results for <i>entrance()</i> , <i>exit()</i> , <i>angle()</i> , <i>most()</i> , <i>unique()</i> , and <i>side()</i>	75
3.13	Pseudocode for the computation of the probability of a turn signal given the previous pose p_{t-1}^n , the previous intended course c_{t-1}^n , and current intended course c_t^n . Illustrative examples are provided in Figure 3.14.	76

3.14	Examples of turn signal probabilities for different situations. The corresponding line number in the pseudocode is indicated above each image. The probability of the turn signal state being <i>left</i> , <i>none</i> , or <i>right</i> is displayed below each image.	77
4.1	Definition of the collision prediction horizon $T_{prediction}$	86
4.2	PreScan interface for simulating road traffic scenarios.	87
4.3	<i>Dangerous</i> scenarios tested in simulation.	88
4.4	<i>Non-dangerous</i> scenarios tested in simulation.	89
4.5	Distribution of the collision prediction horizon $T_{prediction}$	91
4.6	Percentage of detected collisions as a function of the time-to-collision.	91
4.7	Comparison of the performance for priority violations scenarios (A1, B1, C1, D) and stop violations scenarios (A2, B2, C2).	92
4.8	Comparison of the performance for the different types of maneuver executed by the Other Vehicle: merging right (Scenarios A1, A2), merging left (Scenarios B1, B2), crossing (Scenarios C1, C2), and left turn across path (Scenario D).	93
4.9	Percentage of avoided collisions depending on the type of violation and on the action taken when a danger is detected.	95
4.10	Images displayed by the HMI of the vehicle running the risk estimation algorithm. (a) Image displayed when the system predicts a collision with a vehicle coming from the left. (b) Image displayed when the system predicts no collision. (c) Image displayed when the system predicts a collision with a vehicle coming from the right.	97
4.11	Test intersection in Guyancourt: T-shaped give-way intersection. The view is obstructed by the vegetation, therefore the drivers cannot not see each other while approaching the intersection.	98
4.12	System architecture in the test vehicles.	99
4.13	Antenna placement on the test vehicles.	99
4.14	Percentage of received packets as a function of the relative distance.	100
4.15	<i>Dangerous</i> test scenarios for the field trials.	101
4.16	<i>Non-dangerous</i> test scenarios for the field trials.	101
4.17	Range of variation of D (distance to the intersection) and S (speed) for the Priority Vehicle when the Other Vehicle violates the priority in the two <i>dangerous</i> scenarios. . .	102

4.18	<i>Dangerous</i> situation detected during the field trials (Scenario 2). The left column displays the position, heading and turn signal state of the vehicles. When the system detects a <i>dangerous</i> situation, the vehicle causing the danger is displayed in red. The right column shows the view from inside the Priority Vehicle.	104
4.19	Sample scenarios recorded during the field trials. When the system detects a <i>dangerous</i> situation, the vehicle causing the danger is displayed in red.	105
4.20	The differences between the ‘Interacting vehicles’ motion model and the ‘Independent vehicles’ motion model are shown in the graphical representation and in the joint distribution. The elements shown in red do not exist in the ‘Independent vehicles’ model.	108
4.21	Estimated intended course (left graph) and risk (right graph) associated to the Other Vehicle in a sample instance of a <i>non-dangerous</i> scenario (Scenario 5). Screenshots of the situation at t_1 , t_2 , and t_3 are displayed above the graphs.	109
4.22	Estimated intended course (left graph) and risk (right graph) associated to the Other Vehicle in a sample instance of a <i>dangerous</i> scenario (Scenario 1). Screenshots of the situation at t_1 , t_2 , and t_3 are displayed above the graphs.	111
4.23	Comparative results between the ‘Independent vehicles’ motion model and the ‘Interacting vehicles’ motion model, for <i>dangerous</i> scenarios (top) and <i>non-dangerous</i> scenarios (bottom).	114
4.24	Test intersections in Stanford: four-way-stop intersections with multiple entrance and exit lanes in each road.	116
4.25	"Junior", the autonomous car used for data collection. The Velodyne laser rangefinder mounted on the roof top provides a 3D point cloud of the car surroundings.	117
4.26	Examples of <i>consistent</i> and <i>inconsistent</i> behavior scenarios.	118
4.27	Estimation of the intended course in a <i>consistent</i> behavior scenario and in an <i>inconsistent</i> behavior scenario. The size of the green disks located on the exit lanes is proportional to the probability that the driver intends to reach that exit lane.	119
4.28	Percentage of <i>correct</i> , <i>incorrect</i> , and <i>undecidable</i> predictions as a function of the distance to the exit of the intersection for <i>consistent</i> behavior scenarios and for <i>inconsistent</i> behavior scenarios.	121
A.1	Organizations involved in the development of standards for DSRC.	148
A.2	Bandwidth allocated for DSRC applications in Europe and in the U.S. [15, 16].	149

B.1	Example situation where it is necessary to take into account the layout of the intersection to interpret correctly the yellow vehicle's behavior.	152
B.2	Example situation where it is necessary to take into account the dependencies between the vehicles to interpret the black car's behavior.	153
B.3	Non-dangerous scenario: impact of the independence assumption on situation assessment and risk estimation. The thicker the arrows, the higher the probability for a maneuver.	154
B.4	Sensitivity of the risk assessment: the range of variation of <i>Collision_risk</i> in dangerous and non-dangerous situations depends on the value set for λ	156
B.5	Dangerous scenario: impact of the independence assumption on situation assessment and risk estimation. The thicker the arrows, the higher the probability for a maneuver.	157
C.1	Illustrative example of the <i>node</i> and <i>way</i> concepts in digital maps.	160
C.2	Example digital map, superimposed on the corresponding satellite image. This intersection is defined by 6 courses (represented by 6 <i>ways</i>). The corresponding XML code is provided in Section C.3.	161
D.1	Precision, recall, and smallest collision prediction horizon for different values of λ	169
E.1	Two different models for the evolution of the pose: (a) Combination of a Constant Velocity model and a projection model. (b) Path following model taking into account the physical limitations of the vehicle.	172

List of Acronyms

ADAS	Advanced Driver Assistance Systems
DBN	Dynamic Bayesian Network
DSRC	Dedicated Short-Range Communication
FOT	Field Operational Tests
GP	Gaussian Process
GPS	Global Positioning System
HMM	Hidden Markov Model
IMU	Inertial Measurement Unit
ITS	Intelligent Transportation Systems
KBS	Knowledge-Based System
KF	Kalman Filter
LCS	Longest Common Subsequence
MLP	Multi-Layer Perceptron
OSM	OpenStreetMap
OV	Other Vehicle
PV	Priority Vehicle
RNDF	Road Network Definition Format

RRT	Rapidly-exploring Random Tree
SKF	Switching Kalman Filter
SVM	Support Vector Machine
TTC	Time-To-Conflict
TTR	Time-To-React
TTS	Time-To-Stop
TTX	Time-To-X
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-X

Chapter 1

Introduction

Contents

1.1	Context: connected vehicles for road safety	2
1.1.1	Accidents at road intersections	2
1.1.2	Wireless vehicular communications	6
1.2	Problem formulation	17
1.2.1	Scope	18
1.2.2	Challenges and classic approaches	20
1.3	Contributions	22
1.4	Thesis outline	23

Information and communication technologies are an integral part of modern vehicles. There has been a rapid introduction of navigation maps and exteroceptive sensors, with the underlying objective of using driver assistance systems to reduce the number of road accidents. Despite major advances in road safety, accidents at road intersections remain a challenge.

Wireless vehicular communications open new opportunities for safety-related applications and Intelligent Transportation Systems (ITS). The sharing of information between vehicles and the infrastructure extends the perception horizon of a vehicle, beyond the field-of-view limitations of on-board sensors. The research contributions described in this thesis address the problem of situation assessment and risk estimation for road traffic situations, with a focus on safety applications at road intersections and Vehicle-to-Vehicle wireless communications.

The remainder of this chapter is organized as follows. Section 1.1 provides the overall context by describing the issues related to safety at road intersections and the fundamentals of wireless vehicular communications. The problem is formulated in Section 1.2, including a description of the challenges tackled by this thesis and a presentation of the classic approaches. Section 1.3 describes the contributions resulting from this research. Finally, Section 1.4 presents the thesis outline.

1.1 Context: connected vehicles for road safety

This section presents the motivations for using wireless communication between vehicles to improve safety at road intersections. First, accident statistics are used to identify the context in which accidents typically occur. Next, wireless vehicular communications are presented and their potential for reducing the number of accidents at road intersections is outlined.

1.1.1 Accidents at road intersections

Accident statistics have been gathered over the years by different governmental and non-governmental organizations. They are fundamental for the design of safety mechanisms and applications. In this thesis, we use this information to determine the typical conditions and causes of road intersection accidents. The sources used are: the World Health Organization [17] for world statistics, the European Road Safety Observatory [18], TRACE project [19], and PReVENT-INTERSAFE project [20] for statistics in Europe, and the U.S. Department of Transportation [21, 22, 23] for statistics in the U.S..

1.1.1.1 Background

Road safety is a major societal challenge. The World Health Organization (WHO) estimates that the total number of road accident fatalities is approximately 1.3 million every year, with a further 20 to 50 million injured. Road crashes are the ninth cause of death or disability in the world, and it is predicted that it will move to the fifth place by 2030.

In Europe 2 million people are injured in road accidents every year, and the number of fatalities is about 50,000. The cost is estimated at €180 billion a year, or 2.1% of Europe's Gross Domestic Product (GDP).

In the U.S. road accidents are the leading cause of death for people between the ages of 4 and 34. The economic impact of road accidents is estimated by the National Highway Traffic Safety Administration (NHTSA) at \$230.6 billion a year, or 1.6% of the U.S.'s Gross Domestic Product (GDP).

1.1.1.2 Frequency

Comparatively to their density on the road network, intersections are the location of many accidents. In Europe in 2004, accidents at intersections represented 43% of all road accidents, 21% of fatalities, and 34% of the seriously injured (see Figure 1.1).




Figure 1.1. Distribution of road accidents at intersections and out of intersections [1].

These numbers vary across the different countries: in the United Kingdom, more than one third of the overall road accident fatalities occur at intersections, whereas in Greece fatalities at intersections are a minority. In the U.S. intersection-related accidents represent more than 20% of road accident fatalities, and more than 50% of the combined fatal and injury crashes occur at intersections. In Japan intersection crashes figures are even more devastating: they represent 58% of all road traffic accidents and 30% of road traffic fatalities.

1.1.1.3 Context

This section presents statistics on the context in which accidents at road intersections occur.

Road geometry: Most accidents occur at X-shaped intersections (39%), followed by T-shaped and Y-shaped intersections (25%), roundabouts (5%), and railroad crossings (1%).

Urban/rural areas: Road intersection accidents occur mainly in rural areas (80%).

Traffic control: There is a high variation in the occurrence of accidents at signalized and unsignalized intersections across countries. In the UK 58% of accidents happen at give-way intersections, while this case represents less than 10% of intersection accidents in France. Accidents at intersections controlled by traffic lights vary between 20% in the UK and 60% in France. Accidents at intersections ruled by traffic lights tend to be less severe than the ones at stop or give-way intersections. In the U.S. 74% of intersection accidents occur at intersections with some type of traffic control device in place: 46% at intersections controlled by traffic lights, 16% at two-way-stop intersections, 6% at four-way-stop intersections, and 5% at intersections with some other type of control.

Environmental conditions: Most accidents at road intersections occur in daylight conditions (70%), with good visibility (85%), and on a dry road (75%).

Involved actors: In most cases accidents at road intersections involve two vehicles (75%). The case with one pedestrian and one vehicle represents 10% of accidents.

Driver's age group: Young (or inexperienced) drivers and the elderly are the most often involved in intersection-related accidents. For persons older than 80, half of fatal crashes occur at intersections.

Driving situation: The vehicle configuration with respect to the intersection layout was studied in detail in the PReVENT-INTERSAFE project [20]. Figure 1.2 summarizes the results. The scenario where the two vehicles involved in the accident crossed path (scenario **a**) is more frequent and more severe than any other scenario.

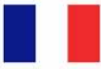


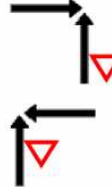



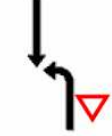




general accident type	description		% of injury accidents		% of injury accidents	% of severe injuries +fatalities	
a	turn into/straight crossing path straight crossing path, opponent coming from the left or right		34%		33,3%	28,5%	
b	turn into/straight crossing path left turn into path, opponent coming from the left		10,5%		13,5%	12,5%	
c	turn across path left turn across path, oncoming opponent (opposite direction)		10%		18,7%	16,6%	
d	turn into/straight crossing path right turn into path, opponent coming from the left		2%		4,1%	2,8%	
e	turn across path left turn across path, preceding opponent (same direction)		2%		2,2%	3,3%	
			58,5% of all intersection-related accidents = 15.636 accidents in France				71,8% of all intersection-related accidents = 106.260 accidents in Germany

Figure 1.2. Distribution of common accident scenarios in Germany and France [2].

1.1.1.4 Causes

The causes of accidents at road intersections have been investigated in the TRACE project [19]. Driver error was identified as being at the origin of 90% of accidents. From the statistics displayed in Figure 1.3, the major factors are:

- The “internal conditions”: non respect of the regulations, or misinterpretation of the situation.
- The “driver behavior”: typically perception failures. The driver took a quick look, or his attention was focused on something else than the other vehicle, or he did not look.
- The “driver state”: the driver’s capacities were impaired, e.g. by stress, alcohol, tiredness.




Figure 1.3. Distribution of accident causes [1].

1.1.2 Wireless vehicular communications

The concept of sharing information between vehicles goes back to the 1939 World Fair, where the Futurama exhibit by General Motor envisioned that communications would make road traffic safer and more efficient in the future. The sharing of information between vehicles (Vehicle-to-Vehicle communications, V2V) or between vehicles and infrastructure (Vehicle-to-Infrastructure communication, V2I) allows a vehicle to perceive its environment beyond the limits of the field-of-view of its on-board sensors. The numerous applications of Vehicle-to-X communication (V2X) for road safety, traffic management, and sustainable transportation explain the growing interest shown by governments and vehicle manufacturers.

In this section, the potential of V2X communications for road safety is addressed. Example V2X-based safety applications are listed, followed by the functional requirements for these applications and a comparison of the different wireless communication technologies which could be used for sharing information between vehicles.

1.1.2.1 The potential of V2X communications for road safety

A study about the trend in the number of road accidents between 1980 and 2000 indicates that *passive safety measures* have had a significant impact on reducing the number of road traffic fatalities, as shown in Figure 1.4.

Measure	Accident reduction
Seat belt use	15-20%
Alcohol countermeasures	15-20%
Specific measures for vulnerable road users	30-40%
Actions targeting the infrastructure	5-10%
Education / training / communication	7-18%

Figure 1.4. Impact of passive safety measures on the reduction of the number of accidents [3].

The introduction of *Advanced Driver Assistance Systems (ADAS)* in vehicles is expected to prevent an even larger number of accidents [24]. These electronic systems are designed to help drivers manage traffic situations safely. Examples of current ADAS are Adaptive Cruise Control, Adaptive Lighting, Night Vision Enhancement, Lane Departure Detection and Warning, Braking and Stability Control Assistance.

The potential of *V2X-based ADAS* should be even greater than standalone ADAS, as shown in Figure 1.5. The sharing of information leads to a better representation of the environment and to an extended perception horizon, since the system can “see” beyond the limitations of on-board sensors. It is estimated that V2X applications may potentially address up to 81% of crash scenarios [25].

The next paragraphs present a list of safety-oriented applications developed using V2X technologies, starting with general safety applications and then applications specifically designed for road intersections.



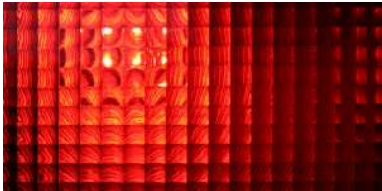




Figure 1.5. Expected reduction in the number of accidents, and associated technologies [4].

General safety applications: Numerous safety applications using V2X technologies have been demonstrated through different projects in Europe, the U.S., and Japan. The example safety applications presented below are part of the European FOT project DRIVE-C2X and are currently under test [5]. This set of applications focuses on driver warning, and does not consider actions on the vehicle's commands such as automatic braking:

Illustration	Description
 <p data-bbox="376 663 695 689">Approaching emergency vehicle</p>	<p data-bbox="815 495 1414 667">Wireless communications are used to broadcast messages about approaching emergency vehicles which claim the right-of-way. The information can be displayed on the head unit or another display device and may also be augmented by audio or haptic signals.</p>
 <p data-bbox="368 987 703 1014">Emergency electronic brake lights</p>	<p data-bbox="815 797 1414 1014">The goal is to avoid rear-end collisions which can occur if a vehicle driving ahead suddenly brakes, especially in dense driving situations or in situations with decreased visibility. The driver will be warned before he is able to realize that the vehicle ahead is braking hard, especially if he/she does not see the vehicle directly (vehicles in between).</p>
 <p data-bbox="432 1312 639 1339">Slow vehicle warning</p>	<p data-bbox="815 1178 1366 1283">This system warns the driver in case of a significant speed difference with the vehicle in front. The goal is to avoid or mitigate rear-end collisions.</p>
 <p data-bbox="440 1641 632 1668">Motorcycle warning</p>	<p data-bbox="815 1507 1390 1612">This system warns the driver if it detects a motorcycle nearby and the distance between the two vehicles is lower than a predefined safety margin.</p>

1.1. Context: connected vehicles for road safety

 <p>Road works warning</p>	<p>V2X devices mounted in areas where road works are taking place send messages to approaching vehicles, warning them of a potentially dangerous situation.</p>
 <p>In-vehicle signage</p>	<p>V2X devices mounted on traffic signs and in critical areas of the road network send messages to approaching vehicles, warning them of a potentially dangerous situation.</p>
 <p>Car breakdown warning</p>	<p>This function enhances the safety of vehicles by detecting an upcoming disabled vehicle or by warning other/following cars that the own car is about to stop. The information is relayed to vehicles driving towards the location where the incident occurred.</p>
 <p>Obstacle warning</p>	<p>The presence of an obstacle along a road could be detected by a vehicle's on-board sensors or entered manually by a driver. This information is broadcasted to vehicles nearby.</p>

Figure 1.6. Set of applications targeted by the DRIVE-C2X project [5].

Safety applications at road intersections: A major issue for safety applications at road intersections is the potential occlusion of part of the scene due to the geometry of the intersection, the presence of obstacles like trees, buildings, etc. Some of the vehicles can be detected by on-board exteroceptive sensors like cameras, radars or lasers, but others will be occluded or simply be beyond the field of view of the sensors. V2X communications do not suffer from this limitation, as illustrated in Figure 1.7.




Figure 1.7. V2X communications (in red) overcome the field-of-view limitations of on-board sensors (in green) at road intersections.

In the depicted situation, a forward-looking on-board sensor only discloses the presence of the vehicle in front (red car) while V2V communications detect the cars reaching and leaving the intersection on the sides. This extended perception horizon is crucial at intersections, where danger could come from the sides (cross-traffic) as well as from the front and back. One could argue that the same perception coverage (360°) could be obtained without vehicular communications by equipping vehicles with several cameras, lasers, or radars, however the cost would be prohibitive.

Another unique feature of vehicular communications is that they provide information which could never be sensed by on-board perception systems, such as the driver's actions on the vehicle controls (e.g. accelerator pedal, turn signals), or the destination entered in the navigation system. Access to such information is very useful to infer the intentions of the drivers, and therefore facilitates situation understanding.

As an example, the use cases addressed by the Cooperative Intersection Collision Avoidance Systems (CICAS) project in the U.S. are presented below [6]:

Illustration	Description
 <p data-bbox="197 669 636 696">Intersection assistant - Left turn across path</p>	<p data-bbox="695 497 1305 674">This function addresses safety problems due to a driver's poor judgment of gaps in oncoming traffic. It warns a driver when it is unsafe to make a left turn, taking into account oncoming vehicles as well as the presence of pedestrians or other road users. Gaps are assessed with infrastructure sensors.</p>
 <p data-bbox="233 1093 600 1120">Intersection assistant - Crossing path</p>	<p data-bbox="695 875 1286 1052">This function addresses safety problems due to a driver's poor judgment of lateral gaps in traffic. It enhances the driver's decision through information and warnings about the available gap. The information is conveyed via a dynamic sign. Gaps are assessed with infrastructure sensors.</p>
 <p data-bbox="248 1382 584 1408">Violation detection - Traffic signal</p>	<p data-bbox="695 1229 1302 1442">The objective is to assist drivers in avoiding crashes in the intersection by warning the driver of an impending violation of a traffic signal. The V2X-equipped intersection broadcasts the signal phase and timing information. Based on speed and distance to the stop location, the system detects whether or not the driver will violate the traffic signal.</p>
 <p data-bbox="264 1686 568 1713">Violation detection - Stop sign</p>	<p data-bbox="695 1552 1302 1729">The objective is to assist drivers in avoiding crashes in the intersection by warning the driver of an impending violation of a stop sign. Based on speed and distance to the stop location, the system detects whether or not the driver will violate the stop sign.</p>

Figure 1.8. Set of applications targeted by the CICAS project [6].

1.1.2.2 Functional requirements for V2X communications

The addition of V2X communications to the existing ADAS enables a number of new vehicle safety applications, as shown in the previous section. For these applications to be feasible in practice, the technology used for wireless communication should meet the requirements listed below.

Latency: Time is critical for safety applications, therefore they have strong latency requirements. While some of the applications described in Figure 1.6 for non-intersection scenarios could bear a few seconds' delay, the applications described in Figure 1.8 for road intersections are all time-critical and require a transmission delay lower than 100 ms [26].

Minimum frequency: The update rate (i.e. rate at which new information is sent/received) should be at least 10 Hz [26].

Communication range and packet losses: Nodes should be able to communicate when they are less than 300 m apart [26], with a packet error rate smaller than 10%.

Mobile nodes: V2X devices should be capable to operate in a rapidly varying environment. The communication needs to be maintained between moving nodes under difficult and fast changing propagation conditions.

Two-way communication: A V2X unit should be able to dialog with another unit or with multiple units.

Information: Precise timestamp and location information are crucial for safety applications, so that each vehicle can create a real-time virtual map of the environment and calculate whether there is a risk.

Privacy: The communication process should guarantee the anonymity of the users.

1.1.2.3 Technologies for V2X communications

The communication protocol used for V2X-based safety applications should meet the requirements stated above, and should also be common between the different brands and different countries. At this point, three technologies are considered:

4G cellular (or LTE) is an evolution of the GSM/UMTS standards for wireless communication using mobile phones. It provides a significantly higher capacity and performance compared to the previous cellular communication standards, with a theoretical peak downlink data rate of 300 Mbps and latency of 5 ms. In practice, the downlink data rate has been observed to be 10-30 Mbps, and the latency around 100 ms [27].




Figure 1.9. Cellular-based wireless vehicular communications.

5.9 GHz DSRC (Dedicated Short Range Communications) is a set of protocols and standards dedicated to short and medium-range wireless communications for ITS. A DSRC network consists of Road Side Units (RSUs) and On Board Units (OBUs) sharing information in the 5.9 GHz band. Typically, a RSU is static and located at one specific point of the infrastructure. OBUs are mounted on vehicles and are therefore mobile nodes. The communication between the nodes is done in an ad-hoc manner, i.e. it is peer-to-peer and does not rely on access points (contrary to cellular communications). More details about the 5.9 GHz DSRC standard are provided in Appendix A.




Figure 1.10. DSRC-based wireless vehicular communications.

Two-way satellite allows two distant users to both send and receive data via satellite, through what is named the Very Small Aperture Terminal (VSAT) service.




Figure 1.11. Satellite-based wireless vehicular communications.

The performance of these technologies are compared in Figure 1.12. The latency achieved by satellite communications is well above 100 ms, therefore this technology is only considered for applications which are not time-critical and can bear a delay of several seconds. Even then, the cost can be prohibitive. DSRC was designed to provide high data transfer rates and a low latency in small communication zones, therefore it matches all the requirements stated above. DSRC is also relatively inexpensive to deploy (no license fee). While the evolution from 2G to 3G and the deployment of 4G improved the performance of cellular communications in terms of transmission rates and delay, the link setup latency is still lower with DSRC. The major issues with non peer-to-peer communications are the availability of networks in rural areas, the time needed to initiate communication, and the radio interferences which can occur when mobile stations are within reach of more than one base station [27]. Finally, another advantage of DSRC is that it does not share the medium with other unbounded services. *The remainder of this section centers on DSRC communications.*

	4G cellular	5.9 GHz DSRC	Two-way satellite
Range	Cell tower coverage	100-1000 m	Thousands of km
Data rates	75-300 Mbps	3-27 Mbps	1-40 Mbps
Latency	<100 ms	200 μ s	>1000 ms
Cost	\$	None	\$\$\$

Figure 1.12. Comparison of the different wireless communication technologies [7].

1.1.2.4 Technical issues and deployment

The DSRC technology is undergoing Field Operational Tests (FOT) in Europe with the DRIVE-C2X project [28], in the U.S. with the Safety Pilot project [29], and in Japan with the ASV-5 project [30]. These large-scale projects aim at testing a variety of applications relying on DSRC and analyzing how the technology scales to a large number of vehicles. The results will help understand the limitations, user acceptance, and safety benefits of V2X applications.

Technical issues: Despite the progress made, several technical issues still need to be addressed before connected vehicles can be deployed. In particular, communication performance is severely impaired in non line-of-sight conditions due to inherent limitations of the radio channel [31, 32]. This is an issue for applications related to road intersection safety. Further, the ability of DSRC communications to scale in the presence of many vehicles still needs to be validated in real conditions [33, 34]. Other technical issues include data security and privacy [35]. Finally, further standardization work is required to ensure the interoperability between the different platforms.

Deployment: The deployment of DSRC is an economically challenging endeavor, as it faces the classic “chicken and egg” problem [36]. Why should vehicle manufacturers invest in the development and integration of DSRC devices in vehicles when there is no guarantee that there would be any infrastructure-based devices to communicate with? Similarly, why should road infrastructure stakeholders invest in the installation DSRC devices on the infrastructure when there is no guarantee that there would be any in-vehicle devices with which to communicate?

One solution to this problem is to enforce the deployment of DSRC through regulations issued by governments. In the U.S., the results obtained by the ongoing Field Operational Tests shall serve as a base for the government to decide on the future of this technology, including the potential role of the government in its deployment [29]. A summary of the U.S. Department of Transportation’s considerations about the deployment of DSRC is provided on their website [37].

The introduction of DSRC shall be gradual, and the rate of increase of the market penetration will depend on the deployment strategies. It shall take years before the majority of vehicles are equipped with DSRC devices, an issue to consider when deploying DSRC-based applications. This is illustrated by the graph in Figure 1.13. It is forecast that the initial applications (first 5 years) will be service-related, since these do not require a great penetration rate. Traffic management applications should be deployed next (5-8 years), followed by DSRC-based driving assistance functions (after 8 years).




Figure 1.13. Forecast deployment of DSRC [8].

1.2 Problem formulation

In the previous section, statistical data was used to justify our decision to address road intersection safety. Accidents at intersections represent 43% of the total number of road accidents, and contrary to the popular belief most of them occur during normal traffic conditions (i.e. daylight, dry roads) and in rural areas. Statistics show that X-shaped intersections, T-shaped intersections and Y-shaped intersections are the most dangerous, and that driver error is at the origin of the majority of accidents. The potential of wireless vehicular communications to reduce the number of accidents at road intersections was presented, along with a list of potential V2X-based safety applications. It was shown that through the sharing of information between vehicles (V2V) and with the infrastructure (V2I), the driver's situational awareness horizon can be extended. This is a fundamental advantage over on-board sensors when dealing with dangerous situations at road intersections.

It is predicted that V2V will have a larger impact than V2I on road safety: according to Najm et al. [25] V2V could address 79% of accidents, against 26% for V2I. This thesis only addresses road intersection safety using V2V communications. In this section we formulate the problem within the context of a typical 'Intelligent vehicle' architecture, to partition the problem and to define the scope of this thesis.

1.2.1 Scope

Figure 1.14 shows the architectural framework within which our research problem is defined.

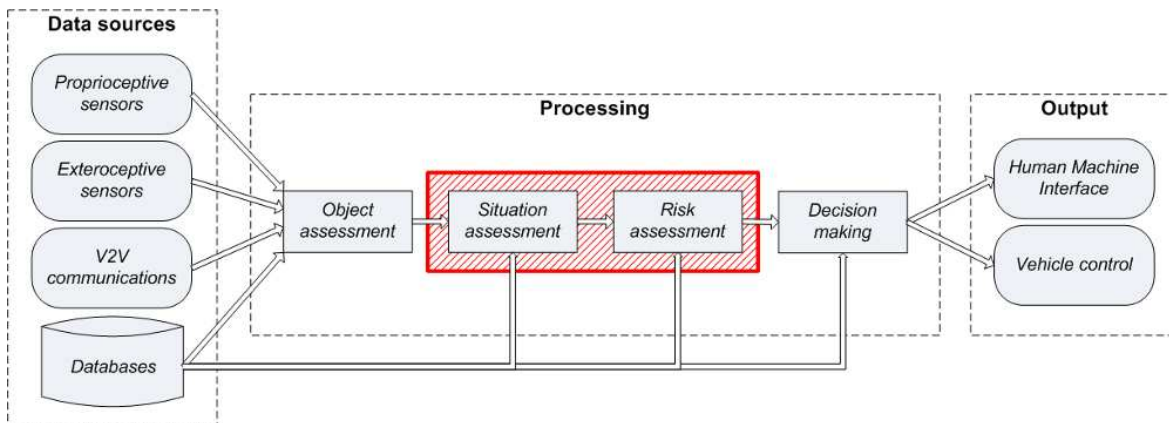


Figure 1.14. Intelligent vehicle architecture, based on the JDL model [9]. The hatched area defines the focus of this thesis.

The different components are defined as follows:

Data sources: Sources of information about the ego-vehicle, the surrounding environment, and the context.

- The proprioceptive sensors provide information about the state of the ego-vehicle such as its position, speed, or the state of the turn signal.
- The exteroceptive sensors return information about the environment. For example, a laser scanner gives the distance to the closest object in the directions the beams point at.
- V2V communications allow the sharing of information between vehicles.
- Databases contain information about the context. They include digital maps, which are available today as part of vehicle navigation (guidance) systems and store information about the road geometry, topology, signalization, etc. Other databases can store information on typical driver behavior (e.g. typical speed profiles, gap acceptance behaviors). Databases are fundamental for reasoning about situations and risk, as they help establish the relationships between the entities in the scene.

Processing: The information provided by the data sources is processed in order to extract useful characteristics of the current situation and make a decision about the necessity to intervene.

- Object assessment

This process combines the information provided by the data sources to detect the different entities in the scene and to estimate their attributes.

In the context of road safety, it consists in estimating the location and velocity of vehicles and pedestrians in the scene.

- Situation assessment

The goal of this process is to establish the relationships between the different entities, by analyzing the observed events.

In the context of road safety, it consists in inferring the intentions of drivers and pedestrians from their joint motion with respect to the road network.

- Risk assessment

The purpose of this process is to draw inferences about potential dangers.

In the context of road safety, the risk that the current situation will result in a collision has to be assessed.

- Decision making

It is the process of selecting the best actions to address the potential threats detected in the previous process.

In the context of road safety, it consists in taking the appropriate actions to avoid or mitigate accidents when a dangerous situation is detected.

Output: The output of the decision making process can be directed to the drivers through a Human-Machine Interface (HMI), or can take the form of a direct action on the vehicle.

- Human-Machine Interface

Its goal is to enhance the driver's situational awareness without distracting him from the driving task.

- Vehicle control

Examples of direct actions on the vehicle are autonomous braking, autonomous steering.

Our research and contributions focus on situation assessment and risk assessment. In this work exteroceptive sensors are not considered, therefore the detection of other entities is done through V2V communications only. Off-the-shelf DSRC modems are used, therefore issues concerning the communication protocols are not addressed in this research. Similarly, the choice of a strategy after a dangerous situation is detected (deciding what action should be taken) is not in the scope of this work.

1.2.2 Challenges and classic approaches

1.2.2.1 Challenges

This thesis is concerned with the formulation of algorithms which can assess the situation and estimate the risk at road intersections. The reasoning is based on the use of proprioceptive vehicle data, shared information between vehicles, and map data. Within this context there are two major sources of difficulty:

1. The uncertainties inherent to the input data

Due to inherent limitations of sensors, a degree of uncertainty is associated with the estimation of the ego-vehicle state (e.g. vehicle position). These estimates are shared between vehicles through V2V communication, but at times due to communication latencies or other factors these estimates will not be available. In addition digital maps might store inaccurate information in terms of road geometry or other parameters.

2. The complexity of traffic situations at road intersections

Road intersection situations are highly dynamic and involve complex interactions between vehicles. The motion of the vehicles is influenced not only by their physical properties, but also by high-level factors such as driver intention, the layout of the intersection, the presence of other vehicles, the traffic rules. These factors cannot be modeled in an exact manner; this is known as the model incompleteness problem [38]. The incompleteness of the model has to be transformed into uncertainty in the reasoning, so that the ambiguities in the interpretation can be addressed.

Therefore, in order to reliably assess situations and risks, it is necessary to take into account the uncertainties in the input data and in the reasoning. However this results in complex models whose

computational needs may not be compatible with the real-time constraints of vehicular safety applications. *Therefore, a tradeoff has to be found between the computational complexity of models and their ability to represent complex, uncertain, and highly dynamic situations.*

1.2.2.2 Classic approaches

The most standard approach for risk assessment in road traffic scenarios consists in predicting the likely future trajectories of the vehicles in the scene using vehicle motion models, and then checking if these trajectories intersect. Once the uncertainties associated to the input data and the future sequence of events are incorporated, there exists a large number of possible future trajectories. The whole process becomes computationally demanding and makes it difficult to comply with real-time constraints. In the literature, this issue is addressed by either ignoring the uncertainties or assuming independence between vehicles:

1. Ignoring the uncertainties

One solution is to assume that sensors provide perfect data and that the models used to assess the situation and risk are flawless. By excluding uncertainties, the computation of risk becomes simpler. However this results in a large sensitivity with respect to the input data, therefore the output of such models can be unstable and unreliable [39].

2. Assuming independence between vehicles

The other solution is to ignore the mutual influences between the motion of the vehicles. By this means the computational complexity is greatly decreased, since the reasoning about a vehicle's behavior is performed independently of the behavior of the other vehicles. However the independence assumption has some drawbacks. They are studied in Appendix B on toy examples, and can be summarized as follows:

- Situation assessment is affected by the independence assumption: some maneuvers become indistinguishable.
- Risk assessment is affected by the independence assumption: the value of risk is overestimated, which affects the sensitivity of the risk assessment.

1.3 Contributions

The problem addressed in this thesis is situation and risk assessment at road intersections. The previous section identified that the main difficulties arise from the complexity of traffic situations at road intersections and from the uncertainties inherent to the input data. In order to match the real-time constraints of vehicular safety applications, it is current practice to either ignore the uncertainties or assume independence between the vehicles in the scene. Neither of these solutions are satisfactory, since they have a significant impact on the ability of an algorithm to assess the current situation and its risk.

The approach proposed in this thesis takes into account both the inter-vehicle dependencies and the uncertainties, while keeping the computational complexity compatible with real-time execution. This is achieved through a novel formulation of risk combined with a context-aware motion model for vehicles negotiating an intersection. Our contributions are summarized below:

1. A novel formulation of risk: comparing *intention* and *expectation*

The classic approach to risk estimation consists in using a motion model to predict the future trajectories of the vehicles in the scene and to check whether or not they intersect. However, this process is computationally too expensive to be executed in real-time with “advanced” motion model. Instead, simpler models are used which either ignore uncertainties or assume independence between the vehicles in the scene.

As an alternative to this classic ‘Trajectory prediction + collision detection’ approach for risk estimation, we propose to detect dangerous situations by comparing what drivers intend to do with what they are expected to do. This is made possible by explicitly representing the traffic rules in the motion model.

A major advantage of this approach compared with the classic approach is that the computation of the future trajectories and their collision points is avoided, therefore the use of “advanced” motion model is no longer incompatible with real-time risk assessment. The proposed approach can in theory be applied to any type of traffic situations. In this thesis an implementation is proposed for the specific case of road intersections.

2. A context-aware motion model for vehicles at a road intersections which
 - (a) Accounts for the interactions between vehicles.

Instead of making the classic assumption that vehicles’ trajectories are independent, we model their mutual influences by introducing a prior knowledge that drivers tend to respect

traffic rules. The motion model therefore takes into account the priority rules and the presence of other vehicles to better interpret a vehicle's behavior.

The advantages are twofold. Firstly, we are able to better estimate the maneuver intention of the drivers, which means our situation assessment capabilities are improved. Secondly, risk is estimated with a higher sensitivity. We avoid risk overestimation while still being able to detect dangerous situations as early.

- (b) Accounts for the influences of the intersection layout on the behavior of the vehicles.

The geometrical and topological characteristics of the intersection are automatically extracted from the digital map and exploited in our motion model. The constraints exerted by the road network on the behavior of the vehicles are taken into account. We show that this allows us to identify situations where a vehicle's behavior is inconsistent, which to our knowledge has not been done before.

3. A framework designed for real traffic conditions

The framework for reasoning about situations and risk at road intersections has been designed to be compatible with real-world constraints. It is generic in the sense that theoretically it can be applied to any intersection layout (geometry, topology) and any number of vehicles. Uncertainties in the input data, which are inherent to real sensors, are taken into account. The algorithm has been tested in real-time on passenger vehicles involved in realistic hazardous situations at a real intersection.

1.4 Thesis outline

The remaining of this thesis is structured as follows:

Chapter 2 - Risk assessment at road intersections: state of the art presents related work on risk estimation at road intersections, with an emphasis on methods based on trajectory prediction. The motion models used for trajectory prediction are categorized into three families: 'Physical entities' motion models, 'Maneuvering entities' motion models, and 'Interacting maneuvering entities'. The advantages and drawbacks are described for each category, followed by a review of the different ways to compute a risk value based on the predicted trajectories.

Chapter 3 - Proposed motion model and risk estimation The proposed approach for modeling traffic situations is described in this chapter, as well as the proposed method for risk estimation. Firstly, the approach is described in the context of general traffic situations. Secondly, an implementation is proposed for road intersections. The description of the algorithm follows the four steps of the Bayesian Programming formalism: 1) Definition of the variables, 2) Definition of the joint distribution, 3) Definition of the parametric forms of the conditional probability terms, 4) Definition of the question(s) relevant for risk assessment.

Chapter 4 - Experiments and results This chapter is concerned with the evaluation of the proposed approach. For this purpose, a combination of simulations and field experiments were carried out. In order to evaluate our approach for risk assessment, both a ‘functional performance’ and a ‘safety performance’ evaluation are performed, i.e. both the system’s ability to detect dangerous situations and its ability to avoid accidents are analyzed. Next, the importance of context for situation and risk assessment at road intersections is studied.

Chapter 5 - Conclusions The final chapter summarizes the thesis and the achievements of this work. Perspectives for future developments are provided.

Chapter 2

Risk assessment at road intersections: state of the art

Contents

2.1	'Physical entities' motion models	28
2.1.1	Evolution models	29
2.1.2	Trajectory prediction	30
2.1.3	Limitations	33
2.2	'Maneuvering entities' motion models	33
2.2.1	Prototype trajectories	34
2.2.2	Maneuver intention estimation and maneuver execution	37
2.2.3	Limitations	40
2.3	'Interacting maneuvering entities' motion models	40
2.3.1	Models based on trajectory prototypes	40
2.3.2	Models based on Dynamic Bayesian Networks	41
2.3.3	Limitations	42
2.4	Collision detection and risk computation	42
2.4.1	Collision detection	42
2.4.2	Risk computation	44
2.5	Other approaches and conclusions	46

2.5.1 Other approaches	46
2.5.2 Conclusions	47

The literature of collision prediction for ITS has its roots mostly in the robotics domain, which can seem surprising since cars have existed for longer than robots¹. The reason is that artificial intelligence algorithms are an integral part of robots and therefore have been the focus of many works in the robotics domain, while the “intelligent” part of the driving task was left entirely to the driver until recently. Robotics and ITS share many research problems, as both are concerned with sensor-equipped mobile platforms which should localize themselves, perceive the environment, understand the situation, assess its risk, make decisions, etc. It is therefore natural that methods developed for robotic applications should be adopted by the ITS domain.

The adaptation of these algorithms to the prediction of road traffic accidents is not straightforward, mainly because most of them were designed for unconstrained environments while the road network is a highly constrained environment. It is important to account for this difference, especially in intersection areas where the complexity of the layouts and of the traffic rules make the progression of vehicles particularly constrained, interactive, and dynamic.

By far the most popular approach to collision risk estimation in robotics is the ‘Trajectory prediction + collision detection’ approach, and it is the prevailing approach among the ITS research community. It is composed of two steps:

1. Predict the potential future trajectories for all the moving entities in the scene. This is done with a motion model.
2. Detect collisions between pairs of trajectories, and derive a risk estimate based on the overall chance of collision.

The first three sections of this chapter focus on Step 1. Motion prediction has been the focus of many works in robotics, and presenting a review of existing approaches from a mathematical point of view would be redundant with previously published surveys. Instead we propose a classification at a conceptual level, based on the different ways to explain a vehicle’s motion. Three families of motion models have been identified and are presented in this chapter:

- ‘Physical entities’ motion models are the simplest models, they represent vehicles as dynamic entities governed by the laws of physics and are described in Section 2.1.
- ‘Maneuvering entities’ motion models are more advanced as they consider that the future motion of a vehicle is influenced by the maneuver that the driver intends to perform. These models are presented in Section 2.2.

¹The first modern car is generally considered to be Karl Benz’s Motorwagen from 1885. Shakey, the first mobile robot capable of reasoning and reacting to its environment, was build in 1970 at SRI.

- 'Interacting maneuvering entities' motion models take into account the inter-dependencies between vehicles' maneuvers. Such advanced representations are used only in a few works, which are reviewed in Section 2.3.

The three families of motion models are illustrated in Figure 2.1. In these examples, the 'Physical entities' motion model assumes a constant speed and orientation for the cars, the 'Maneuvering entities' motion model assumes that the black car goes straight and the blue car turns left, the 'Interacting maneuvering entities' motion model assumes that the black car goes straight, that the blue car turns left and that the joint motion of the cars is constrained by the traffic rules.

The fourth section addresses Step 2, i.e. the different approaches for detecting collisions between the predicted trajectories and for evaluating the risk of the situation.




Figure 2.1. Examples of motion prediction with the different types of motion models.

2.1 'Physical entities' motion models

'Physical entities' motion models represent vehicles as dynamic entities governed by the laws of physics. Future trajectories are computed using dynamic and kinematic models linking some control

inputs (e.g. steering, acceleration), car properties (e.g. weight), external conditions (e.g. friction coefficient of the road surface) to the evolution of the state of the vehicle (e.g. position, heading, speed).

Extensive work has been done on such ‘Physical entities’ motion models for vehicles, and they remain the most commonly used motion models for trajectory prediction and collision risk estimation in the context of road safety. Models are more or less complex depending on how fine the representation of the dynamics and kinematics of a vehicle is, how uncertainties are handled, whether or not the geometry of the road is taken into account, etc.

This section is divided into three parts. In the first part, the most standard evolution models are described. The second part provides a review of the different methods for predicting trajectories using these evolution models. Finally, the limitations of ‘Physical entities’ motion models are addressed.

2.1.1 Evolution models

2.1.1.1 Dynamic models

Dynamic models describe motion based on Lagrange’s equations, taking into account the different forces that affect the motion of a vehicle, such as the longitudinal and lateral tire forces, or the road banking angle [40].

Car-like vehicles are governed by complex physics (effect of driver actions on the engine, transmission, wheels etc.), therefore their derived dynamic models can get extremely large and involve many internal parameters of the vehicle. Such complex models are relevant for control-oriented applications, but for applications such as trajectory prediction simpler models are preferred. They are often based on a “bicycle” representation, which represents a car as a two-wheeled vehicle with front-wheel drive moving on a 2-D plane. Examples of such simple dynamic models are found in several works [41, 42, 43, 44, 45, 46].

2.1.1.2 Kinematic models

Kinematic models describe a vehicle’s motion based on the mathematical relationship between the parameters of the movement (e.g. position, velocity, acceleration), without considering the forces that affect the motion. The friction force is neglected, and it is assumed that the velocity at each wheel is in the direction of the wheel [40]. Kinematic models are far more popular than dynamic models for trajectory prediction because they are much simpler and usually sufficient for this type of

applications (i.e. applications which are not vehicle control). In addition the internal parameters of a vehicle needed by dynamic models are not observable by exteroceptive sensors, which rules out the use of dynamical models for ITS applications involving other vehicles than the ego-vehicle.

A survey of kinematic models for car-like vehicles was done by Schubert et al [47]. The simplest of these are the Constant Velocity (CV) and Constant Acceleration (CA) models, which both assume straight motion for vehicles [11, 48, 49, 50, 10]. The Constant Turn Rate and Velocity (CTRV) and Constant Turn Rate and Acceleration (CTRA) models take into account the variation around the z-axis by introducing the yaw angle and yaw rate variables in the vehicle state vector [51, 52, 53, 50, 54, 48]. The complexity remains low as the velocity and yaw rate are decoupled. By considering the steering angle instead of the yaw rate in the state variables, one obtains a “bicycle” representation, which takes into account the correlation between the velocity and the yaw rate. From this representation, the Constant Steering Angle and Velocity (CSAV) and the Constant Steering Angle and Acceleration (CSAA) can be derived.

2.1.2 Trajectory prediction

The evolution models described above can be used for trajectory prediction in various ways, the main difference being in the handling of uncertainties.

2.1.2.1 Single trajectory simulation

A straightforward manner to predict the future trajectory of a vehicle is to apply an evolution model (see Section 2.1.1) to the current state of a vehicle, assuming that the current state is perfectly known and that the evolution model is a perfect representation of the motion of the vehicle. This strategy can be used with dynamic models [41] or kinematic models [10, 49, 54]. This process is illustrated in Figure 2.2.

The advantage of this single forward simulation is its computational efficiency, which makes it suitable for applications with strong real-time constraints. However the predictions do not take into account the uncertainties on the current state nor the errors introduced by the evolution model, and as a result the predicted trajectories are not reliable for long term prediction. In addition, predicting future motion using a single trajectory fails to represent the variety of possible maneuvers at an intersection (e.g. turn left, right).




Figure 2.2. Trajectory prediction with a Constant Velocity motion model (source: [10]).

2.1.2.2 Gaussian noise simulation

Uncertainty on the current vehicle state and on its evolution can be modeled by a normal distribution [11, 53, 48, 50]. The popularity of the “Gaussian noise” representation of uncertainty is due to its use in the Kalman Filter (KF). Kalman filtering is a standard technique for recursively estimating a vehicle’s state from noisy sensor measurements. It is a special case of Bayesian filtering where the evolution model and the sensor model are linear², and uncertainty is represented using a unimodal normal distribution. In a first step (*prediction* step) the estimated state at time t is fed to the evolution model, resulting in a predicted state for time $t + 1$ which takes the form of a Gaussian distribution. In a second step (*update* step) the sensor measurements at time $t + 1$ are combined with the predicted state into an estimated state for time $t + 1$, which is also a Gaussian distribution. Looping on the *prediction* and *update* steps each time a new measurement is available is called tracking³.

By looping on the *prediction* step, one can obtain for each future timestep a mean and covariance matrix for the state of the vehicle, which can be transformed into a mean trajectory with associated uncertainty (normal distribution at each timestep) [11, 53, 52]. This process is illustrated in Figure 2.3.

²Extensions of the Kalman Filter exist which deal with non-linear models. A comparison of their performances for trajectory prediction is done in [55].

³More details about probabilistic filtering can be found in [56].




Figure 2.3. Trajectory prediction with a Constant Velocity motion model and Gaussian noise simulation. Ellipses represent the uncertainty on the predicted positions (source: [11]).

Compared to the “single trajectory simulation” approaches, these techniques have the advantage that they represent the uncertainty on the predicted trajectory. However they suffer from a similar limitation in the sense that modeling uncertainties using a unimodal normal distribution is insufficient to represent the different possible maneuvers at an intersection.

A solution to this problem is to represent uncertainty using mixtures of Gaussians. Switching Kalman Filters (SKF) [57] can be used for this purpose. They rely on a bank of Kalman Filters to represent the possible evolution models of a vehicle and switch between them [48, 58, 59]. An alternative to the SKF is to use heuristics to switch between the different kinematic models depending on the situation [50].

2.1.2.3 Monte Carlo simulation

In the general case, i.e. when no assumption is made on the linearity of the models or on the Gaussianity of the uncertainties, the analytical expression for the distribution on the predicted states is usually not known. Monte Carlo methods provide tools to approximate this distribution. The idea is to randomly sample from the input variables of the evolution model in order to generate potential future trajectories. In order to take into account the layout of the intersection, weights are applied to the generated trajectories to penalize the ones which do not respect the constraints of the road network.

The evolution models described in Section 2.1.1 can be used for Monte Carlo simulation by sampling on the inputs instead of considering them to be constant. Typical inputs to be sampled from are the

acceleration and steering angle or lateral deviation. In order to take into account the feasibility of a maneuver, one can either remove the generated trajectory samples with a higher lateral acceleration than what is physically allowed [12], or take into account the physical limitations of a vehicle in the evolution model so that the inputs are distributed in a more realistic manner and the post-processing step for removing unfeasible trajectories is not needed [45, 60]. Example predicted trajectories are displayed in Figure 2.4.

Monte Carlo simulation can be used to predict a vehicle's trajectory either from a perfectly known current state or from an uncertain current state estimated by a tracking algorithm.



Figure 2.4. Trajectory prediction (green area) with Monte Carlo simulation (source: [12]).

2.1.3 Limitations

Since they only rely on the low level properties of motion (dynamic and kinematic properties), 'Physical entities' motion models are limited to short-term motion prediction. Typically, they will be unable to anticipate any change in the motion of the car caused by the execution of a particular maneuver (e.g. slow down, turn at constant speed, then accelerate to make a turn at an intersection), or changes caused by external factors (e.g. slowing down because of a vehicle in front). This shortcoming is problematic especially at road intersections. Situations in these areas are highly dynamic, therefore it is necessary to be able to anticipate changes in order to predict trajectories in a reliable manner.

2.2 'Maneuvering entities' motion models

'Maneuvering entities' motion models represent vehicles as independent maneuvering entities, i.e. they assume that the motion of a vehicle on the road network corresponds to a series of maneuvers executed independently from the other vehicles. Here, a maneuver is defined as "a physical movement or series of moves requiring skill and care" [61]. The word "behavior" is sometimes

used in the literature for the same purpose [62, 63, 64, 65, 66], but for the sake of clarity the word “maneuver” will be used throughout this document.

Trajectory prediction with ‘Maneuvering entities’ motion models is based on the early recognition of the maneuvers that drivers intend to perform. If one can identify the maneuver intention of a driver, one can assume that the future motion of the vehicle will match that maneuver. Thanks to this a priori, trajectories derived from this scheme are more relevant and reliable in the long term than the ones derived from ‘Physical entities’ motion models.

‘Maneuvering entities’ motion models are either based on prototype trajectories or based on maneuver intention estimation. The first part of this section introduces prototype trajectories and how they can be used for trajectory prediction. Approached based on maneuver intention estimation are covered in the second part of this section. Finally, the limitations of ‘Maneuvering entities’ motion models are analyzed.

2.2.1 Prototype trajectories

The idea is that the trajectories of vehicles on the road network can be grouped into a finite set of clusters, each cluster corresponding to a typical motion pattern. Example clusters are displayed in Figure 2.5.




Figure 2.5. Clustered trajectories: each cluster corresponds to a typical motion pattern (source: [13]).

Motion patterns are represented using prototype trajectories which are learned from data during a

training phase. Subsequently prediction can be performed online given a partial trajectory by finding the most likely motion pattern(s) and using the prototype trajectories as a model for future motion.

2.2.1.1 Representation

Motion patterns are represented using prototype trajectories which are learned from sample (previously observed) trajectories.

Because the road network is a structured environment, it is generally assumed that the motion patterns can be identified in advance (they can for example be extracted from a digital map, by identifying all the possible maneuvers at a given location). In this case no clustering process is needed, i.e. each trajectory in the training dataset is already assigned to a cluster⁴. Starting from there, several possibilities exist for representing a motion pattern based on the sample trajectories.

One solution is to compute a unique prototype trajectory for each motion pattern, by agglomerating the previously observed trajectories. For example, a stochastic representation of a motion pattern can be derived by computing the mean and standard deviation of the sample trajectories [70]. Another way to account for the variations in the execution of a motion pattern is to have several prototypes for each class, e.g. a subset of the training samples [71].

More recently, several works showed that Gaussian Processes (GPs) are well-suited for representing motion patterns in the context of road traffic [72, 73, 63]. GPs can be seen as a generalization of Gaussian probability distributions. They model a process as a Gaussian distribution over a function. When applied in the context of vehicle trajectories, the assumption is that the trajectories in the learning dataset are sample functions from a Gaussian Process. Therefore the learning consists in fitting a Gaussian distribution over these functions. The main advantages of GPs are their robustness to noise in the observed trajectories (compared to the approaches presented above) and their ability to represent the variations in the execution of a motion pattern in a consistent and probabilistic manner. For example, a GP featuring the function $f(t) = (x, y)$, with t the time and (x, y) the 2D coordinates of the vehicle will be able to partially account for variations in the speed of execution of a maneuver [63]. An alternative is to use the function $f(x, y) = (\Delta x/\Delta t, \Delta y/\Delta t)$, which maps locations to velocities and therefore removes any time-dependent aspect from the model [72, 73]. The variability in the velocity due to varying traffic conditions is captured in the covariance function of the GP. However this capacity to represent the inherent variations of a pattern comes at a price; a naive implementation

⁴For readers interested in a state of the art of clustering, a comparative survey of the most popular methods was presented by Morris and Trivedi [67]. In particular, road traffic situations were investigated by Buzan et al. [68], Hu et al.[69], and Atev et al. [13].

of a GP has a complexity of $O(n^3)$, where n is the number of training input points.

2.2.1.2 Trajectory prediction

Starting from the partial trajectory executed by a vehicle so far, prediction can be performed by comparing it with the learned motion patterns, selecting the most likely one(s) and using the prototype trajectories as a model for future motion.

First, metrics need to be defined to measure the distance of a partial trajectory to a motion pattern. When motion patterns are represented by Gaussian Processes, the distance is computed as the probability that the partial trajectory corresponds to the GP, by integrating over the possible futures of the trajectory [72, 73, 63]. When motion patterns are represented by a finite set of prototype trajectories, the distance of a partial trajectory to a motion pattern is measured by its similarity with the prototype trajectories. A number of metrics have been defined to measure the similarity between two trajectories, including the average Euclidian distance between points of the trajectories [69], the modified Hausdorff [13], the Longest Common Subsequence (LCS) [68] and its translation-and-rotation invariant version the Quaternion-based Rotationally Invariant LCS [71].

The simplest solution to predict future motion after the distance to each motion pattern has been computed is to select the most likely motion pattern and to use it as a unique model [69]. Alternatively it is possible to consider a mixture of motion patterns: a probability distribution over the different motion patterns is computed, and then either the different motion models are combined (weighted) into one [72], or a set of potential future trajectories are generated with associated weights [73, 71].

2.2.1.3 Limitations

For a long time, the main limitation of prototype trajectories was their strictly deterministic representation of time. Indeed, when motion patterns are represented using a finite set of trajectories it is impossible to model the great variation in the execution of a motion pattern. In order to be able to recognize maneuvers involving a waiting period at a stop line for example, one has to resort to hard thresholds to identify waiting intervals and ignore them when computing the distance between two trajectories. Handling more subtle variations in velocity like the ones caused by heavy traffic is still an issue for such models.

To a certain extent, the introduction of Gaussian Processes solved this problem by allowing a time-independent representation of motion patterns [72, 73]. However GPs suffer from other limitations. In addition to their heavy computational burden, they lack the ability to take into account the physical

limitations of a vehicle and therefore may generate unrealistic sample trajectories. To our knowledge the only solution to this problem proposed in the literature is the one introduced by Aoude et al [73], which uses a Rapidly-exploring Random Tree (RRT) algorithm to randomly sample points toward dynamically feasible trajectories, using as inputs the current state of the vehicle and the sample trajectories generated by the GPs.

Another difficulty when using prototype trajectories is their adaptation to different intersection layouts. Because each motion model is trained for a specific intersection geometry and topology, they can be reused only at intersections with a similar layout. For example, motion models learned at a one-lane X-shaped intersection cannot be used on a two-lane X-shaped intersection because the motion patterns are not the same.

2.2.2 Maneuver intention estimation and maneuver execution

An alternative to trajectory prototypes is to first estimate the maneuver intention of the driver (e.g. waiting at the stop line, following another vehicle, executing a left turn) and then predict the successive physical states so that they correspond to a possible execution of the identified maneuver.

A major advantage over trajectory prototypes is that there is no need to match the partial trajectory with a previously observed trajectory. Instead, higher-level characteristics are extracted and used to recognize maneuvers, which makes the recognition process more flexible.

2.2.2.1 Maneuver intention estimation

Many cues can be used to estimate the maneuver intention of a driver, for example the physical state of the vehicle (position, speed, heading, acceleration, yaw rate, turn signal, etc.), information about the road network (geometry and topology of the intersection, speed limit, traffic rules, etc.), driver behavior (head movement, driving style, etc.). Maneuver intention estimation has been investigated by many works; this survey focuses on the ones concerned with intersection-related maneuvers.

Context and heuristics can be used to determine what maneuvers are likely to be performed in the near future in a deterministic manner [74]. For classifying maneuvers in more complex scenarios, discriminative learning algorithms are very popular, such as Multi-Layer Perceptrons (MLP) [66] or Support Vector Machines (SVM) [75]. An equally popular alternative is to break down each maneuver into a chain of consecutive events and to represent this sequence of events using a Hidden Markov Model (HMM). The transition probabilities between the different events can be learned from data, as well as the observation model (i.e. the relationship between the non-observable events and the

available observations). For a new sequence of observations, the maneuver intention is estimated by comparing the likelihood of the observations for each HMM [76, 62, 63].

A comparative review of works on maneuver intention estimation is provided in Figure 2.6.

	Estimated maneuvers	Estimation method	Features
Greene et al. [74]	Stop, go straight, left turn, right turn	Heuristics	Position, velocity, state of the nearby traffic light
Tamke et al. [77]	Go straight, left turn, right turn	Heuristics	Turn signal
Aoude et al. [75]	Safe, errant	SVM-BF	Distance to intersection, speed, longitudinal acceleration
Garcia-Ortiz et al. [66]	Stopped, brake, keep speed	MLP	Speed, gas and brake pedal information, distance to intersection, state of traffic light
Aoude et al. [62]	Compliant, violating	HMM	Distance to intersection, speed, longitudinal acceleration
Tay [63]	Go straight, overtake, left turn, right turn	Hierarchical HMM	Position, speed
Berndt et al. [76]	Go straight, left turn, right turn	HMM	Steering angle

Figure 2.6. Comparative review of works on maneuver intention estimation.

2.2.2.2 Maneuver execution


Trajectories are predicted so that they match the identified maneuver(s).

This can be done in a deterministic manner, by deriving the input controls (in this case the steering angle) corresponding to the recognized maneuver and then generating a single trajectory from a kinematic motion model [77]. One issue with this strictly deterministic approach is that it cannot take into account uncertainties on the current vehicle state, on the maneuver which is being performed or


on the execution of this maneuver.

In order to explore the space of the potential executions of a maneuver in a probabilistic manner, GPs or RRTs can be used (see Figure 2.7). A GP can be learned for each maneuver from training data and used in a generative manner to create sample trajectories for each maneuver [63, 78]. Alternatively, a RRT tree can be grown by sampling points in the input space of the vehicle's evolution model, applying a bias in the sampling according to the estimated maneuver intentions [75]. This method has the advantage that it always generates dynamically feasible trajectories for a maneuver.


As an alternative to trajectories, reachable states can be used to represent the future motion of vehicles. This representation can be stochastic [79] or geometric [74] (see Figure 2.7).




(a) RRT (source:[75]).



(b) GP (source:[78])



(c) Stochastic reachable states (source: [79])



(d) Segmented cones (source: [74])

Figure 2.7. Maneuver execution using four different approaches.

2.2.3 Limitations

In practice, the assumption that vehicles move independently from each other does not hold. Vehicles share the road with other vehicles, and the maneuvers performed by one vehicle will necessarily influence the maneuvers of the other vehicles. Inter-vehicle dependencies are particularly strong at road intersections, where priority rules force vehicles to take into account the maneuvers performed by the other vehicles. Disregarding these dependencies can lead to erroneous interpretations of the situations, and affects the evaluation of the risk. This problem is studied in more detail in Appendix B.

2.3 'Interacting maneuvering entities' motion models

'Interacting maneuvering entities' motion models represent vehicles as maneuvering entities which interact with each other, i.e. the motion of a vehicle is assumed to be influenced by the motion of the other vehicles in the scene. Taking into account the dependencies between the vehicles leads to a better interpretation of their moves compared with the 'Maneuvering entities' motion models described in the previous section. As a result, it contributes to a better understanding of the situation and a more reliable evaluation of the risk.

Despite this, there are few 'Interacting maneuvering entities' motion models in the literature. They are either based on prototype trajectories or based on Dynamic Bayesian Networks. Both solutions are presented in this section, followed by the limitations of 'Interacting maneuvering entities' motion models.

2.3.1 Models based on trajectory prototypes

For methods relying on trajectory prototypes, inter-vehicles influences cannot be taken into account during the learning phase because the resulting number of motion patterns would quickly become intractable.

However, it is possible to take into account the mutual influences during the matching phase by assuming that drivers have a strong tendency to avoid collisions when they can [80]. Pairs of trajectories which lead to an unavoidable collision are penalized in the matching process, and as a result safe trajectories are always considered to be more likely than hazardous ones.

This approach is an elegant workaround for taking into account inter-dependencies when using trajectory prototypes. However the issue of modeling other types of influences remains, since the influence

of one vehicle on the trajectory of another cannot be modeled directly.

2.3.2 Models based on Dynamic Bayesian Networks

All the other ‘Interacting maneuvering entities’ motion models we are aware of are based on Dynamic Bayesian Networks (DBN).

Pairwise dependencies between multiple moving entities can be modeled with Coupled HMMs (CHMMs) [81]. However, since the number of possible pairwise dependencies grows exponentially with the number of entities, the complexity is not manageable in the context of road intersection situations. A solution to simplify the model is to make CHMMs asymmetric by assuming that the surrounding traffic affects the vehicle of interest, but not vice versa [82].

The fact that influences between vehicles are regulated by traffic rules is exploited by Agamennoni et al. [14, 83, 84]. Multi-agent influences are decomposed as log-linear combinations of pairwise dependencies, with pairwise dependencies of the type “vehicles on a smaller road yield to vehicles on the main road”. This is illustrated in Figure 2.8.




Figure 2.8. Trajectory prediction with an ‘Interacting maneuvering entities’ motion model.

The model is able to predict that Agent 35 should yield to Agent 17 (source: [14]).

Finally, a general probabilistic framework for tracking vehicles and predicting their future motion was introduced by Gindele et al. [64]. Instead of modeling pairwise dependencies, the model accounts for mutual influences by using factored states. The causal dependencies between the vehicles are

modelled as a function of the local situational context, which reduces greatly the computational complexity. The approach was validated on simulated highway scenarios, but could in theory be applied to road intersections.

2.3.3 Limitations

The ‘Interacting maneuvering entities’ motion models are the most comprehensive models proposed so far in the literature. They allow longer-term predictions compared to ‘Physical entities’ motion models, and are more reliable than ‘Maneuvering entities’ motion models since they account for the dependencies between the vehicles. However, this exhaustiveness has some drawbacks: computing all the potential trajectories of the vehicles with these models is computationally expensive and not compatible with real-time risk assessment. For this reason, *none of the motion models presented in this section have been applied to collision risk estimation.*

2.4 Collision detection and risk computation

In the previous sections, motion models were introduced which can be used to predict the future motion of vehicles. Following this, algorithms are needed which can assess the risk of the situation based on these predictions.

A primary element of the problem is the detection of a collision between two individual trajectories. This “collision detection” function is at the base of the computation of risk. The different methods for this purpose are reviewed in the first part of this section.

The computation of a global risk value relies on the “collision detection” function, but is not straightforward. An important aspect of risk assessment is the definition of semantics for the term “risk”. Indeed, the variable of interest to characterize risk is not the same depending on the context of the final application (e.g. driver warning, autonomous driving). This matter is addressed in the second part of this section.

2.4.1 Collision detection

Given the future trajectories of two vehicles, many tools exist to extract information regarding the occurrence of a collision. While the most basic methods only provide basic information such as whether, where and when a collision will occur, the more advanced ones can compute its probability, type, severity, etc.

2.4.1.1 Binary collision detection

In the special case of linear ‘Physical entities’ motion models, the analytical solution for the state of the vehicles at a specific time can easily be derived by solving the linear differential equations of the motion model. It follows that the intersection point between two trajectories can be computed in an efficient manner [49, 10].

However, in the general case the motion equations are too complex for a closed-form solution to be derived. One solution is to approximate each trajectory by a piecewise-straight line trajectory [54]. A more common approach is to discretize the trajectories and to check iteratively for a collision at each discrete timestep. Following this reasoning, collisions can be detected in a simple manner by defining a threshold on the distance between two points (from two trajectories at the same timestep) [75]. In order to take into account the shape of the vehicles, this threshold can be replaced by a condition on the “overlap between the shapes of the two vehicles”. Although the exact shape is not always mentioned [60, 77], vehicles are often represented as polygons [41, 45, 63, 46, 12]. If information is available about the uncertainty on the state of the vehicles, and if this uncertainty is Gaussian, an ellipse can be used instead of a polygon by applying a threshold on the standard deviations [11, 53]. In order to simplify the calculation of the intersection area, ellipses can be approximated by a set of circles [11] or by a set of points [53].

2.4.1.2 Probabilistic collision detection

Few works carry out collision detection in a probabilistic manner.

For a normally distributed uncertainty on the current state, a solution based on stochastic linearization via the unscented transformation has been proposed [85].

In the case of stochastic reachable states, the collision probability can be computed on a discretized position space by calculating the probability that the center of both vehicles is in the same cell, for all the possible combinations of cells [60]. With the geometric version of reachable states, the collision probability can be measured as the percentage of overlap between the geometric shapes representing the future motion of vehicles [74].

2.4.1.3 Other collision indicators

By analyzing further the predicted trajectories and their intersecting points, it is possible to derive some indicators which give more information about the potential collision.

Popular indicators of the criticality of a potential collision are the velocity of the vehicles [46, 86], the amount of overlap between the shapes representing the vehicles [11], and the configuration of the collision [11].

The information provided by these indicators can be used to determine the best way to mitigate or avoid the potential collision.

2.4.2 Risk computation

The computation of risk relies on the detection of collisions between pairs of trajectories (approaches presented above), and is application-dependent. Indeed the variable of interest is not the same if the goal of the application is to assess the criticality of a situation to decide if an intervention is necessary on the ego-vehicle, or if the goal is to estimate the hazard level of different maneuvers to decide which one the ego-vehicle should execute. The first application is typical of ADAS, while the second one can be useful both for ADAS and for autonomous driving.

2.4.2.1 Binary risk computation

The computation of risk can be binary. In this case, risk is assigned the value 0 or 1 depending on whether there exists a collision-free maneuver that the driver can perform. Determining whether such a maneuver exists can be done in two ways.

The first one consists in computing escape maneuvers (i.e. how the vehicle should steer, brake or accelerate to avoid the collision) and check whether these maneuvers are feasible (with “feasible” meaning that the steering, braking or accelerating does not exceed the physical limitations of the vehicle) [41].

The second one is to consider the entire space of combined steering, braking and accelerating maneuvers, and to perform an optimized search for collision-free trajectories [46].

2.4.2.2 Probabilistic risk computation

Risk can be computed in a probabilistic manner, taking into account the uncertainty on the future motion of vehicles.

When the future motion of a vehicle is represented by a probability distribution on sample trajectories (which is typically the case with approaches relying on Monte Carlo simulations or Gaussian Processes), it is possible to compute risk as the “probability of a collision in the future” by integrating

over all the possible future trajectories and detecting collisions between each possible pair. This approach provides a lot of flexibility in the handling of uncertainties. For example, for a ‘Maneuvering entities’ motion model the calculation can either sum over both the maneuvers and their executions, or assume that the maneuvers are known and sum on the possible executions only [63]. Further, depending on the final application one can compute the risk of colliding with a specific vehicle or sum over all the vehicles and obtain a global collision risk [63]. If the collision detection is done in a probabilistic manner (as in Section 2.4.1), the uncertainty about the collision can also be taken into account in the calculation.

2.4.2.3 Other risk indicators

Several risk indicators exist which can complement the collision risk. They are generally based on a measure of the “Time-To-X” (or TTX) where X corresponds to a relevant event in the course toward the collision.

- **Time-To-Collision (TTC)**

The most standard indicator is the Time-To-Collision, which corresponds to the time remaining before the collision occurs⁵. It can be used as an indication of what action (if any) should be taken [52, 87, 11, 41, 45]. For example, when the TTC is still large it might be preferable to inform or warn the driver rather than to apply the brakes. For autonomous emergency braking applications, the TTC can be compared with the time required for the vehicle to come to a full stop in order to decide when to apply the brakes [88]. For driver warning applications, the driver reaction time needs to be added to the time to stop the vehicle [10]. The TTC can also be used as a tool to identify the least dangerous maneuver for an autonomous vehicle, by assuming that the risk of executing a specific trajectory for the autonomous vehicle is inversely proportional to the earliest TTC (the TTC is calculated for all the possible trajectories of the other vehicles in the scene) [75].

- **Time-To-React (TTR)**

A closely related indicator is the Time-To-React, which corresponds to the time available for the driver to act before the collision is inevitable. The idea is to simulate different driver actions (such as braking, accelerating, steering) and to identify the latest moment at which one of these maneuvers is able to avoid the collision [77, 49].

⁵The term Time-To-Collision is often used to refer specifically to collision detection methods which assume a constant velocity for the vehicles. Here the alternative definition of the Time-To-Collision is used, where it corresponds to the time remaining before the collision occurs.

2.5 Other approaches and conclusions

This chapter focused on the ‘Trajectory prediction + collision detection’ approach for the problem of estimating the risk of a situation at an intersection. This final section reviews the few methods which do not rely on trajectory prediction, and concludes on the status and remaining issues for risk assessment at road intersections.

2.5.1 Other approaches

2.5.1.1 Knowledge-Based Systems

An intuitive approach is to define a set of rules which detect danger based on the context and on the current observations of the state of the vehicles.

The rules can be simple heuristics on acceptable speeds when approaching an intersection [39], or can include more advanced concepts such as the semantics of the location, weather conditions or the level of fatigue of the driver [89]. Rule-based systems are a particular case of Knowledge-Based Systems (KBS), which aim at providing a complete framework for knowledge acquisition, knowledge representation, reasoning and decision making. For example, knowledge can be stored in the form of databases (containing detailed historical facts), expert knowledge bases (identified rules, norms, and patterns), case bases (a set of solved case studies), ontologies (formal description of domain-specific concepts) [90].

Because context is explicitly taken into account in KBS, the characteristics of the environment can easily be incorporated. However an established limitation of these systems is their inability to account for uncertainties (both on the data and on the model).

2.5.1.2 Data mining techniques

An alternative is use accident databases to learn typical collision patterns between two vehicles. This way potentially dangerous configurations can be identified when they occur again.

Data mining techniques have been used to map the relationship between vehicles states (input) and collision risk (output) directly [91, 92].

However obtaining the data to learn from remains an issue, since real data is not available and simulations will not be representative of real accident situations. Another limitation is that these techniques

learn collision patterns between pairs of vehicles without putting them in context with the other vehicles. Since the relationships between the different vehicles in the scene are not modeled, some dangerous situations will be detected very late (e.g. vehicle A following vehicle B and getting closer, preparing to overtake vehicle B, while a vehicle C is driving on the same road but in the opposite direction).

2.5.2 Conclusions

With the objective to improve road safety, the automotive industry is evolving toward more “intelligent” transportation systems. One of the major current challenges for vehicles is to be capable of detecting dangerous situations and react accordingly. In this respect the field of ITS benefited from the research conducted over the last 40 years in the mobile robotics domain, as they share a number of scientific problems (e.g. localization, perception, navigation). However, a notable particularity of the road network compared with other environments is that it is very structured, with specific rules governing the behavior of the traffic participants. This difference is not always acknowledged in approaches proposed in the literature.

The most common approach for risk estimation is based on the ‘Trajectory prediction + collision detection’ steps. First, a motion model is used to predict the likely future trajectories of the vehicles in the scene. Then all the combinations of future trajectories are inspected for collisions, and risk is computed based on the detected collisions. The motion model plays an important role: the model should be sophisticated enough that complex situations can be represented, but should also allow the computation of risk in real-time. ‘Physical entities’ motion models allow for an efficient computation of the risk, but are limited to short-term collision prediction. ‘Maneuvering entities’ motion models provide a more reliable estimation of long-term motion and risk by introducing a higher-level reasoning, but they ignore the dependencies between the vehicles in the scene. ‘Interacting maneuvering entities’ account for these dependencies, but their computational complexity make them incompatible with real-time risk assessment with the ‘Trajectory prediction + collision detection’ approach.

Overall the difficulty is that in order to estimate risk at an intersection it is necessary to reason at a high level about a set of interacting maneuvering entities, taking into account uncertainties as well as the context of the road network. This high-level reasoning is computationally expensive, and when it is followed by trajectory prediction and collision detection the whole process is no longer compatible with real-time applications. Therefore we argue that an important challenge for future risk assessment algorithms is to be able to integrate contextual information and to account for uncertainties, while

keeping the complexity manageable.

In the next chapter, a novel approach to risk estimation at road intersection is proposed which attempts to tackle this challenge. The risk estimation problem is formulated in such a way that it is no longer necessary to predict the future trajectories of the vehicles to compute the collision risk of a situation. By avoiding the time-consuming computation of the future trajectories and their collision points, we make the use of ‘Interacting maneuvering entities’ motion models compatible with real-time risk assessment.

Chapter 3

Proposed motion model and risk estimation

Contents

3.1	Overview of the proposed approach	51
3.1.1	Background: representing vehicle motion with MSSMs	51
3.1.2	Proposed approach for general traffic situations	56
3.1.3	Application to road intersections	59
3.2	Variable definition	60
3.2.1	Intended maneuver I_t^n	60
3.2.2	Expected maneuver E_t^n	62
3.2.3	Measurements Z_t^n	63
3.2.4	Physical state Φ_t^n	63
3.2.5	Summary and notations	64
3.3	Joint distribution	65
3.4	Parametric forms	67
3.4.1	Expected longitudinal motion Es_t^n	67
3.4.2	Intended longitudinal motion Is_t^n	69
3.4.3	Intended lateral motion Ic_t^n	70
3.4.4	Pose P_t^n	71

3.4.5	Speed S_t^n	72
3.4.6	Turn signal state T_t^n	74
3.4.7	Measured pose Pm_t^n	77
3.4.8	Measured speed Sm_t^n	77
3.4.9	Measured turn signal state Tm_t^n	78
3.4.10	Summary	78
3.5	Bayesian risk estimation	78
3.6	Conclusions	79
3.6.1	Summary	79
3.6.2	Discussion	79

The state of the art review conducted in the previous chapter led to the conclusion that an important challenge for risk assessment algorithms is to find a tradeoff between the amount of information which is taken into account in the motion model (constraints imposed by the road network, dependencies between vehicles' maneuvers, uncertainties) and the complexity of the computation.

This work proposes an alternative to the classic 'Trajectory prediction + collision detection' approach to risk estimation. The idea is to detect dangerous situations by comparing what drivers intend to do with what they are expected to do according to the traffic rules. The time-consuming computation of the future trajectories and their collision points is avoided, and the use of advanced motion model is no longer incompatible with real-time risk assessment. More precisely, we propose an 'Interacting maneuvering entities' motion model for vehicles in the form of a Dynamic Bayesian Network (DBN) where driver intentions and traffic rules are explicitly represented. The risk of a situation is estimated by performing Bayesian inference on the relevant variables given some observations.

This chapter describes the proposed motion model and risk estimation strategy. Section 3.1 introduces the principles behind the proposed approach, first in the general context of road traffic situations and then in the specific context of road intersections. The final four sections are dedicated to the detailed description of the DBN proposed for risk estimation at road intersections. The DBN is described in 4 steps, following the Bayesian Programming formalism [93]:

1. Define the variables (Section 3.2),
2. Define the joint distribution (Section 3.3),

3. Define the parametric forms of the conditional probability terms (Section 3.4),
4. Define the question(s) relevant for risk assessment (Section 3.5).

3.1 Overview of the proposed approach

In this work the motion of a vehicle is represented using a Markov State Space Model (MSSM). Markov State Space Models belong to the family of Dynamic Bayesian Networks. They describe continuous motion using a set of state variables, a set of measurement variables, and a set of conditional probability functions. The state variables are hidden (not observed) and are inferred from the successive measurements. The conditional probability functions specify the evolution of the state between two timesteps, and specify the dependencies between the measurement variables and the state variables. This section is organized in three parts. First, some background is provided about the use of MSSMs to represent the motion of vehicles on the road. Then the proposed MSSM is introduced in the context of general traffic situations. The application of this general framework to the specific case of road intersections is presented in the last part.

3.1.1 Background: representing vehicle motion with MSSMs

MSSMs have been used successfully in the past to represent vehicle motion [94, 64, 14, 63, 62]. They have the advantage of being very flexible: one motion pattern can be represented in many different ways with MSSMs. The complexity of the model depends on the choice of variables and on the form of the conditional probability functions. For this reason, the specification of the variables is a very important step in the design of MSSMs. The set of variables should be comprehensive enough that complex situations can be represented, but the relationships between them should be simple enough that inference on the hidden state variables can be performed in real-time.

There are two ways of representing maneuvering vehicles with MSSMs: either with a single MSSM [64, 14] or with a set of MSSMs [94, 63, 62].

3.1.1.1 With a single MSSM

In this case, the motion of a vehicle is typically modeled based on three layers of abstraction:

- Level 1: the highest level corresponds to the maneuver performed by the vehicle (e.g. overtake, turn right). The variables at this level are discrete and hidden (not observable).

We will call I the conjunction of these variables.

I_t^n therefore represents the maneuver being performed by vehicle n at time t . We call it I as “Intention”, since the maneuver performed by a vehicle reflects the *intended maneuver* of the driver.

- Level 2: this level corresponds to the physical state of the vehicle (e.g. position, speed). The variables at this level are hidden (not observable).

We will call Φ the conjunction of these variables.

Φ_t^n therefore represents the *physical state* of vehicle n at time t .

- Level 3: the lowest level corresponds to the measurements available (e.g. measurement of the vehicle’s position). The variables at this level are observable. They often correspond to a noisy version of a subset of the physical variables.

We will call Z the conjunction of these variables.

Z_t^n therefore represents the *measurements* of the state of vehicle n at time t .

The dependencies between these variables can take a variety of forms. A general form is shown in Figure 3.1. Well-known models such as Hidden Markov Models or Switching Markov Models belong to the family of MSSMs and are derived from this general model by making further independence assumptions.




Figure 3.1. Graphical representation of a base MSSM modeling the motion of a vehicle n between time $t - 1$ and t .

3.1.1.2 With a set of MSSMs

In this case, each possible maneuver is modeled by a dedicated MSSM. A maneuver is broken into a sequence of “primitive” maneuvers, with constraints on their order and duration. In order to represent this, a fourth level of abstraction is added between level 1 (maneuver intention) and level 2 (physical state):

- Level 1.5: this level corresponds to the primitive maneuver performed by the vehicle (e.g. slow down). The variables at this level are discrete and hidden (not observable).

We will call Π the conjunction of these variables.

Π_t^n therefore represents the *primitive maneuver* performed by vehicle n at time t .

Any type of MSSM can be used to represent each maneuver, but HMMs are by far the most popular. The graphical representation of a model relying on a set of HMMs to represent maneuvering vehicles is shown in Figure 3.2.




Figure 3.2. Graphical representation of a HMM-based motion model for a vehicle n .

3.1.1.3 Accounting for dependencies between vehicles

Vehicle do not move on the road independently from each other: their joint motion is regulated by traffic rules. Traffic rules are “compulsory for all participants in road traffic [...]. [They] establish

the meaning of the signals of traffic lights and traffic controllers, road signs, and road markings and describe the movements of participants in road traffic in the most typical conditions and situations. The rules regulate the positioning of vehicles on the roadway, the changing of position, passing, stopping, parking, and traversing intersections and railroad crossings.” [95].

The base MSSMs (described above) do not account for the dependencies between the vehicles, but solutions have been proposed in the literature.

In the case of a single MSSM, the solution proposed in the literature [64, 14] is to make the current intended maneuver of a driver dependent on the previous “situational context”, i.e. dependent on the previous intended maneuvers and physical states of all the vehicles in the scene. The corresponding graphical model is shown in Figure 3.3. The multi-vehicle dependencies implicitly encode the traffic rules, since the latter regulate the joint motion of vehicles on the road. For example if vehicle n is approaching a give-way intersection, and there are vehicles with right-of-way approaching the same intersection, we can assume based on traffic rules that there is a high probability that the driver of vehicle n intends to stop at the intersection and a low probability that he intends to proceed in the intersection. Another example is a vehicle n following a vehicle m on the highway: if vehicle m slows down, there is a high probability that vehicle n will either slow down or change lanes to overtake vehicle m , and there is a low probability that vehicle n will remain on the same lane with the same speed.




Figure 3.3. Graphical representation of a MSSM which accounts for the dependencies between the vehicles. Bold arrows represent multi-vehicle dependencies, i.e. the influences of the other vehicles on vehicle n .

In the case of a set of HMMs, an extension has been proposed in the form of Coupled HMMs (CHMMs) [81]. Each type of interaction between two vehicles is represented by two HMMs whose primitive maneuvers influence each other (see graphical model in Figure 3.4). For example, the interaction “vehicle n overtakes vehicle m ” will be represented with one CHMM where the primitive maneuver executed by vehicle m (e.g. “slow down”) will influence the primitive maneuver executed by vehicle n (e.g. “change lane left”) and vice-versa.




Figure 3.4. Graphical representation of a CHMM-based motion model.

3.1.2 Proposed approach for general traffic situations

Our risk assessment method is based on a single MSSM. Sets of HMMs were discarded for complexity reasons: as was mentioned in the previous chapter, the fact that a CHMM is needed for each possible interaction between two vehicles is not compatible with real-time applications. We propose a novel manner to take into account vehicle interactions in MSSMs. The resulting motion model makes it possible to estimate the risk of a situation without computing the future trajectories of the vehicles.

3.1.2.1 Proposed MSSM

The core idea is the following: instead of implicitly encoding the traffic rules in the multi-vehicle dependencies (as in Figure 3.3), we introduce an intermediate layer in the MSSM which explicitly represents what traffic rules expect from drivers in the current context. This is represented in Figure 3.5.




Figure 3.5. Graphical representation of the proposed MSSM. Bold arrows represent multi-vehicle dependencies, i.e. the influences of the other vehicles on vehicle n .

The *expected maneuver* E_t^n represents the behavior expected from a vehicle according to the traffic rules. It is a conjunction of variables analogous to the intended maneuver I_t^n : every variable in the conjunction I_t^n has an equivalent in the conjunction E_t^n . As can be seen in the graph, the expected maneuver E_t^n is derived from the previous situational context and has an influence on the intended maneuver I_t^n .

Our model of interactions between vehicles is similar to the solutions proposed in the literature (Figure 3.3) in the sense that the dependencies between the vehicles are modeled by making the current intended maneuver I_t^n dependent on the previous “situational context”. **The difference is that the dependency is not direct in our case: the expected maneuver E_t^n is inserted as an intermediate. The previous “situational context” influences what the driver is expected to do, which in turn influences what the driver intends to do.**

3.1.2.2 Proposed risk estimation strategy

Modeling explicitly what is expected from a vehicle at time t creates new possibilities for the estimation of the risk of a situation. **Instead of using the MSSM to predict the future trajectories of the vehicles, we use it to jointly infer what drivers currently intend to do (I_t) and what they are expected to do (E_t). The risk of a situation is computed based on the probability that intentions and expectations do not match, given the measurements:**

$$P([I_t^n \neq E_t^n] | Z_{0:t}) \quad (3.1)$$

This novel approach to risk assessment reflects the fact that most accidents are caused by driver error [1], and matches the intuitive notion that “dangerous” situations are situations where drivers are acting differently from what is expected of them.

Based on Eq. 3.1, a variety of safety-oriented applications can be derived.

Detection of hazardous vehicles: A “hazard probability” can be computed for every vehicle in the scene using Eq. 3.1. Subsequently, actions can be triggered depending on the value of the risk. An example ADAS application would be to warn all the drivers in the area when the risk is higher than a predefined threshold, i.e. if:

$$\exists n \in N : P([I_t^n \neq E_t^n] | Z_{0:t}) > \lambda \quad (3.2)$$

The warning message could be adapted to the level of risk, so that the driver is aware of the urgency of the situation.

If autonomous braking is considered, there are even more possibilities. Inference on the Dynamic Bayesian Network proposed in this work can provide relevant information about the situation, beyond

the risk value. The maneuvers at the origin of the danger can be identified, and since our model is generative it is then possible to predict the time and location of the accident by predicting the future trajectories of the vehicles involved in the accident. Depending on the risk value and on the time remaining before the collision, applications using our approach could trigger gradual responses to danger: e.g. warning the driver when the risk is moderate, slowing down the vehicle autonomously when the risk is high and there is not enough time to warn the driver, or applying emergency braking autonomously when a collision is imminent.

Risk of a specific maneuver: The risk of a specific maneuver I_t^n can be computed for a vehicle n :

$$P\left(\bigcup_{m=1}^N [I_t^m \neq E_t^m] | I_t^n Z_{0:t}\right) \quad (3.3)$$

This is an important feature for autonomous driving, but also for ADAS. One application is to find the best escape maneuver in a dangerous situation. Another one is to assist the driver for tasks such as changing lanes on the highway or negotiating an intersection by computing the risk of each possible maneuver and informing the driver about how safe each maneuver is.

Other applications: The proposed model can be used to estimate the intended maneuver of a driver, or to predict future trajectories. These are useful features for numerous applications which need to reason about traffic situations [96].

3.1.2.3 Note on trajectory prediction

In Section 3.1.2.2, we suggest using the proposed motion model to predict the future trajectories of vehicles. This can seem contradictory with the angle taken so far in this thesis about risk estimation: we stated repeatedly that one of the main advantages of our approach is that it does not require to predict the trajectories of the vehicles in the scene to compute the risk of a situation.

Both statements are actually compatible. The ‘Trajectory prediction + collision detection’ approach described in the state-of-the-art chapter relies on trajectory prediction to compute the risk. One issue is that predicting all the potential trajectories of all the vehicles for all their potential maneuvers and then checking for a collision between all the possible pairs of trajectories is very time-consuming and not always compatible with real-time risk estimation. When we suggest predicting trajectories in Section 3.1.2.2, the context is very different. The risk of the situation has already been estimated (by

comparing the *intention* I_t^n with the *expectation* E_t^n), along with the vehicle(s) causing the danger and their intended maneuver(s). The purpose of trajectory prediction here is to obtain some additional information about the collision such as the time remaining before it occurs. In this context trajectory prediction is performed only for specific vehicles performing a specific maneuver, the computational requirements are therefore very low.

We believe that trajectory prediction is complementary to our method. The difference with state-of-the-art approaches is that trajectory prediction is optional in our case: it is performed after risk estimation if the application requests additional information about the collision (see example in Section 3.1.2.2).

3.1.3 Application to road intersections

The model presented above could in theory be applied to any traffic situations, by defining the variables I_t^n , E_t^n , Z_t^n , and Φ_t^n accordingly.

Here the focus is on risk estimation at unsignalized road intersections, i.e. intersections ruled by anything but traffic lights (stop, give-way, priority to the right). More specifically we address accidents caused by traffic sign violations. All intersection layouts (geometry, topology) are considered. As an example, the list of accident scenarios addressed in this work are depicted in Figure 3.6 for X-shaped intersections.

The case of signalized road intersections (i.e. controlled by traffic lights) is not addressed, although the model described below can very easily be extended to traffic light violation cases. Neither rear-end collisions or driver error in the lateral execution of the maneuver (such as a hazardous lane change) are addressed in this work, but extensions of the motion model to handle these cases are proposed in the conclusion of this chapter.

The proposed model for risk estimation at road intersections is described in the remaining sections. The description follows the Bayesian Programming formalism: first the variables are defined, then the proposed joint distribution, the parametric forms, and finally the calculation of risk.




Figure 3.6. List of accident scenarios addressed at X-shaped intersections.

3.2 Variable definition

This section proposes definitions for the intended maneuver I_t^n , the expected maneuver E_t^n , the physical state Φ_t^n , and the measurements Z_t^n , in the context of road intersections.

3.2.1 Intended maneuver I_t^n

As was mentioned in the previous section, driver error in the lateral direction is not addressed in this work. The focus is on errors in the longitudinal execution of the maneuver. Since we want to be able to reason on the lateral motion and on the longitudinal motion separately, we make a distinction between the lateral and longitudinal components of a maneuver in our specification of the intended maneuver I_t^n .

3.2.1.1 Lateral component

For the lateral component, we exploit the fact that intersections are highly structured areas where the lateral motion of vehicles is constrained by the geometry and the topology of the intersection.

It is assumed that a digital map of the road network is available. From this digital map we extract a set of “courses”, where a course is defined for each authorized maneuver at the intersection as the typical path that is followed by a vehicle when executing that particular maneuver. The concept of a course is illustrated in Figure 3.7. For details about the extraction of courses from digital maps, see Appendix C.




Figure 3.7. Illustrative example for the "course" concept. The courses originating from one road are displayed as blue arrows.

The variable representing the lateral component of a maneuver is defined as:

- $Ic_t^n \in \{c_i\}_{i=1:N_C}$: the driver's **intended lateral motion**, which will also be called the driver's **intended course** in the context of road intersections. It corresponds to the course followed by vehicle n at time t . $\{c_i\}_{i=1:N_C}$ is the set of possible courses at the intersection, extracted from the digital map.

3.2.1.2 Longitudinal component

For the longitudinal component, we exploit the fact that intersections are highly structured areas where the longitudinal motion of vehicles is constrained by the geometry and the topology of the

intersection as well as by the traffic rules.

For a vehicle n at time t we define two possible intentions with respect to the motion in the longitudinal direction: *go* and *stop*. The variable representing the longitudinal component of a maneuver is defined as:

- $Is_t^n \in \{go, stop\}$: the driver's **intended longitudinal motion**, which will also be called the driver's **intention to stop** in the context of road intersections. It corresponds to the driver's intention regarding the longitudinal execution of the maneuver.

$Is_t^n = go$ means that the driver intends to adapt its speed to the layout of the intersection only. In other words, the driver intends to negotiate the intersection as if there were no constraints from the traffic rules (stop, give-way). This is typically the case for vehicles which have priority: drivers will adapt their speed to the topology and the geometry of the intersection (slowing down to make a turn) but will not slow down to stop or yield to another vehicle.

$Is_t^n = stop$ means that the driver intends to adapt its speed to the layout of the intersection (similarly to $Is_t^n = go$), but will also adapt his speed so that he can stop at the intersection. Typically, this behavior will be adopted by vehicles approaching a stop intersection with the intention to respect the stop, and by vehicles which do not have the right-of-way at an intersection and intend to yield to another vehicle.

3.2.2 Expected maneuver E_t^n

In the general framework presented in Section 3.1.2, each variable in the conjunction I_t^n has an equivalent in the conjunction E_t^n . The purpose is twofold: to model the influences of the surrounding vehicles on the maneuver performed by a vehicle, and to compute the risk based on the probability that the expected maneuver and the intended maneuver do not match.

If this principle is applied to our problem, the expected maneuver E_t^n should contain two variables: the “expected lateral motion” (analogous to Ic_t^n) and the “expected longitudinal motion” (analogous to Is_t^n). However in our case it is not necessary to include an “expected lateral motion” variable, for two reasons. The first reason is that in the context of road intersections the dependencies between the vehicles mostly concern the longitudinal motion: whether a driver will stop or not at the intersection is influenced by the presence of other vehicles, but the course followed by a driver is not. Therefore it is reasonable to assume independence between the lateral motion of vehicles. The second reason is that, as mentioned earlier, this work addresses risks in the longitudinal direction only. Therefore for a vehicle n at time t we define the following variable:

- $Es_t^n \in \{go, stop\}$: the **expected longitudinal motion**, which will also be called the **expectation to stop** in the context of road intersections. It corresponds to the expected longitudinal motion of the vehicle according to the traffic rules.

The definition of $Es_t^n = go$ and $Es_t^n = stop$ are analogous to the definitions provided for the $Is_t^n = go$ and $Is_t^n = stop$ in the previous section; the only difference is that it corresponds to what the driver should do (according to the traffic rules) instead of what he intends to do:

$Es_t^n = go$ means that the driver should adapt his speed to the layout of the intersection only.

$Es_t^n = stop$ means that the driver should adapt his speed to the layout of the intersection, and should also adapt his speed so that he can stop at the intersection.

3.2.3 Measurements Z_t^n

In this work, the following measurements of the state of a vehicle $n \in N$ at time t are available:

- $Pm_t^n = (Xm_t^n Ym_t^n \theta m_t^n) \in \mathbb{R}^3$: the **measured pose**, i.e. the position and orientation of the vehicle in the Universal Transverse Mercator (UTM) coordinate system.
- $Sm_t^n \in \mathbb{R}$: the **measured speed** of the vehicle.
- $Tm_t^n \in \{left, right, none\}$: the **measured turn signal state** of the vehicle.

The pose and speed measurements are obtained through a combination of GPS (Global Positioning System) and IMU (Inertial Measurement Unit) measurements, and the turn signal state is read on the CAN-bus of each vehicle. These measurements are shared by each vehicle via V2V communication, and are therefore accessible by all the vehicles in the scene¹. However, the model presented here can be applied independently of the type of sensors. For example, if the vehicles do not communicate, the information about the other vehicles could be obtained via embedded exteroceptive sensors such as cameras, lasers, radars, or a combination of those.

3.2.4 Physical state Φ_t^n

Based on the available measurements (see previous section), the following variables are selected to represent the physical state of a vehicle $n \in N$ at time t :

¹More details about the measurements and about the sharing of information between vehicles in our experiments are given in the results chapter (Chapter 4).

- $P_t^n = (X_t^n Y_t^n \theta_t^n) \in \mathbb{R}^3$: the **true pose** of the vehicle.
- $S_t^n \in \mathbb{R}$: the **true speed** of the vehicle.
- $T_t^n \in \{left, right, none\}$: the **true turn signal state** of the vehicle.

3.2.5 Summary and notations

The variables of our DBN were defined by adapting the general model proposed in Section 3.1.2 to the context of road intersections. The following variables were defined for a vehicle $n \in N$ at time t :

- For the intended maneuver: $I_t^n = (Ic_t^n Is_t^n)$

Ic_t^n represents the driver's intended lateral motion, i.e. the path that the vehicle will follow to negotiate the intersection. Is_t^n represents the driver's intended longitudinal motion, i.e. whether or not he intends to stop at the intersection.

- For the expected maneuver: $E_t^n = Es_t^n$

Es_t^n represents the expected longitudinal motion, i.e. whether or not the driver is expected to stop at the intersection according to the traffic rules.

- For the measurements: $Z_t^n = (Pm_t^n Sm_t^n Tm_t^n)$

The variables correspond to the measurements available in this work.

- For the physical state: $\Phi_t^n = (P_t^n S_t^n T_t^n)$

The variables represent the true states behind the measurements.

The variables Z_t^n and Φ_t^n were defined according to the measurements available, and the variables I_t^n and E_t^n were defined with the objective to detect dangerous situations at a road intersection by comparing driver *intention* and driver *expectation*, as illustrated in Figure 3.8 with some example scenarios.




Figure 3.8. Detection of dangerous situations by comparing *intention* Is_t^n and *expectation* Es_t^n : example scenarios.

In the next sections we carry on with the specification of the DBN, still following the Bayesian Programming formalism. For more clarity in the equations, in the remaining of this thesis a bold symbol will be used to represent the conjunction of variables for all the vehicles in the scene. For example, for a variable \mathbf{X} :

$$\mathbf{X} \triangleq (X^1 \dots X^N) \quad (3.4)$$

with X^n the variable associated with vehicle n .

3.3 Joint distribution

For the general model proposed in Figure 3.5, the joint distribution over all the vehicles is as follows:

$$\begin{aligned}
 P(\mathbf{E}_{0:T} \mathbf{I}_{0:T} \Phi_{0:T} \mathbf{Z}_{0:T}) &= P(\mathbf{E}_0 \mathbf{I}_0 \Phi_0 \mathbf{Z}_0) \\
 &\times \prod_{t=1}^T \times \prod_{n=1}^N [P(E_t^n | \mathbf{I}_{t-1} \Phi_{t-1}) \times P(I_t^n | \Phi_{t-1}^n I_{t-1}^n E_t^n) \\
 &\times P(\Phi_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \times P(Z_t^n | \Phi_t^n)] \quad (3.5)
 \end{aligned}$$

In this section we adapt this general model to road intersection situations, with the variables defined in the previous section. We start by making the following classic independence assumptions:

For the intended maneuver: The current intended lateral motion and intended longitudinal motion are conditionally independent given $(\Phi_{t-1}^n I_{t-1}^n E_t^n)$. Therefore the following simplification is obtained:

$$P(I_t^n | \Phi_{t-1}^n I_{t-1}^n E_t^n) = P(Ic_t^n | \Phi_{t-1}^n I_{t-1}^n E_t^n) \times P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n E_t^n) \quad (3.6)$$

For the physical state: The current pose, speed and turn signal state are conditionally independent given $(\Phi_{t-1}^n I_{t-1}^n I_t^n)$. Therefore the following simplification is obtained:

$$P(\Phi_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) = P(P_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \times P(S_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \times P(T_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \quad (3.7)$$

For the measurements: A classic sensor model is used, i.e. the measurements are conditionally independent given the physical quantities they are associated with. Therefore the following simplification is obtained:

$$P(Z_t^n | \Phi_t^n) = P(Pm_t^n | P_t^n) \times P(Sm_t^n | S_t^n) \times P(Tm_t^n | T_t^n) \quad (3.8)$$

After applying these independence assumptions, and taking into account that $E_t^n = Es_t^n$, the joint distribution (Equation 3.5) becomes:

$$\begin{aligned} P(\mathbf{E}_{0:T} \mathbf{I}_{0:T} \Phi_{0:T} \mathbf{Z}_{0:T}) &= P(\mathbf{E}_0 \mathbf{I}_0 \Phi_0 \mathbf{Z}_0) \\ &\times \prod_{t=1}^T \times \prod_{n=1}^N [P(Es_t^n | \mathbf{I}_{t-1} \Phi_{t-1}) \\ &\times P(Ic_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) \times P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) \\ &\times P(P_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \times P(S_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \times P(T_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) \\ &\times P(Pm_t^n | P_t^n) \times P(Sm_t^n | S_t^n) \times P(Tm_t^n | T_t^n)] \end{aligned} \quad (3.9)$$

3.4 Parametric forms

In this section the parametric forms of the conditional probability terms in Equation 3.9 are described, along with the hypotheses they build on.

3.4.1 Expected longitudinal motion Es_t^n

The expected longitudinal motion of a vehicle is derived from the previous intended course, pose and speed of all the vehicles in the scene:

$$P(Es_t^n | \mathbf{I}_{t-1} \Phi_{t-1}) = P(Es_t^n | \mathbf{I}_{t-1} \mathbf{P}_{t-1} \mathbf{S}_{t-1}) \quad (3.10)$$

What is expected of vehicles on the road is regulated by traffic rules, but a lot is left to the judgment of the driver. If we take as an example give-way intersections in France, the rules specify that the driver which does not have the right-of-way must “yield to vehicles driving on the other road(s) and make sure there is no danger before entering the intersection”². There exists no formula to calculate whether it is legal or not for a driver to enter an intersection at time t in a specific context. Instead, our expectation model is based on typical driver behavior, i.e. on a statistical analysis of what drivers consider to be acceptable. *The necessity for a vehicle to stop given the context is derived using probabilistic gap acceptance models* [97, 98]:

$$P([Es_t^n = stop] | [\mathbf{I}_{t-1} = \mathbf{c}_{t-1}] [\mathbf{P}_{t-1} = \mathbf{p}_{t-1}] [\mathbf{S}_{t-1} = \mathbf{s}_{t-1}]) = f(\mathbf{c}_{t-1}, \mathbf{p}_{t-1}, \mathbf{s}_{t-1}) \quad (3.11)$$

with \mathbf{c}_{t-1} the conjunction of courses for the N vehicles in the scene (i.e. $\mathbf{c}_{t-1} = (c_{t-1}^0 \dots c_{t-1}^N)$), \mathbf{p}_{t-1} the conjunction of positions for the N vehicles in the scene (i.e. $\mathbf{p}_{t-1} = (p_{t-1}^0 \dots p_{t-1}^N)$), \mathbf{s}_{t-1} the conjunction of speeds for the N vehicles in the scene (i.e. $\mathbf{s}_{t-1} = (s_{t-1}^0 \dots s_{t-1}^N)$), and f a function which computes the probability that the gap available for vehicle n to execute its maneuver is sufficient given the previous situational context ($\mathbf{c}_{t-1}, \mathbf{p}_{t-1}, \mathbf{s}_{t-1}$). For a vehicle n heading towards a give-way intersection, the calculation is detailed below and illustrated in Figure 3.9 on an example:

1. Project forward (or backward) the position of vehicle n until the time t^n when it reaches the intersection, using the vehicle’s previous state ($c_{t-1}^n, p_{t-1}^n, s_{t-1}^n$) and a constant speed model.

²Original version in French (Code de la route - Article R415-7): "A certaines intersections indiquées par une signalisation dite cédez le passage, tout conducteur doit céder le passage aux véhicules circulant sur l'autre ou les autres routes et ne s'y engager qu'après s'être assuré qu'il peut le faire sans danger."

2. Let V_{ROW} be the set of vehicles whose maneuvers have the right-of-way w.r.t. the maneuver of vehicle n . For each vehicle $m \in V_{ROW}$ project forward (or backward) the position of vehicle m until the time t^m when it reaches the intersection, using the vehicle's previous state $(c_{t-1}^m, p_{t-1}^m, s_{t-1}^m)$ and a constant speed model.
3. Find the vehicle $k \in V_{ROW}$ which is the most likely to cause vehicle n to stop, by finding the smallest positive time gap available for vehicle n to execute its maneuver:

$$\begin{cases} k &= \arg \min_{m \in V_{ROW}} (t^m - t^n), \text{ for } t^m - t^n \geq 0 \\ g_{min} &= t^k - t^n \end{cases} \quad (3.12)$$

4. The necessity for vehicle n to stop at the intersection is calculated as the probability that the gap g_{min} is not sufficient, using a probabilistic gap acceptance model:

- For merging cases, the following gap acceptance model is used [97]:

$$f(\mathbf{c}_{t-1}, \mathbf{p}_{t-1}, \mathbf{s}_{t-1}) = 1.0 - \frac{1}{1 + e^{-\lambda(\ln(g_{min}) + (1-\alpha)\ln(s_{t-1}^k) - \ln(\gamma))}} \quad (3.13)$$

with g_{min} the gap computed in the previous step, s_{t-1}^k the speed of vehicle k at time $t - 1$, and $(\lambda, \alpha, \gamma)$ the parameters of the logistic function set to the average values $\lambda = 3.611$, $\alpha = 0.602$, $\gamma = 19.347$ [97].

- For other cases, the following gap acceptance model is used (adapted from [98]):

$$f(\mathbf{c}_{t-1}, \mathbf{p}_{t-1}, \mathbf{s}_{t-1}) = 1.0 - \frac{1.05}{1 + \left(\frac{g_{min}}{\delta}\right)^\sigma} \quad (3.14)$$

with g_{min} the gap computed in the previous step and (δ, σ) the parameters of the logistic function set to the values $\delta = 6.1$, $\sigma = -4.0$.

For a vehicle approaching a stop intersection the calculation is similar, except the probability that the vehicle is expected to stop is set to 1 until it reaches the intersection (i.e. $P([Es_t^n = stop] | \mathbf{I}_{t-1} \mathbf{P}_{t-1} \mathbf{S}_{t-1}) = 1$), and after that point the last two steps of the calculation above are used (i.e. steps 3 and 4).

This context-aware reasoning about the necessity for a vehicle to stop at the intersection will allow us to detect vehicles running stop signs, or vehicles entering an intersection when they should have waited for another vehicle to pass (as illustrated earlier in Figure 3.8). A similar calculation can be done for intersections controlled by traffic lights, but this is not the focus of this work.

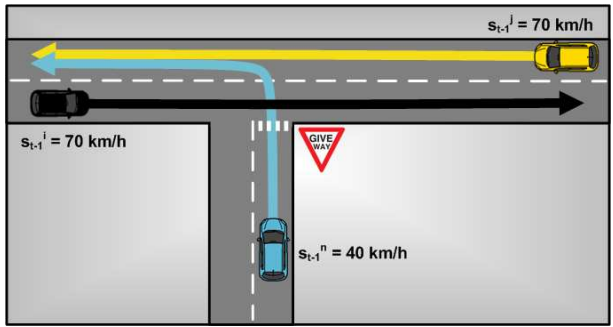
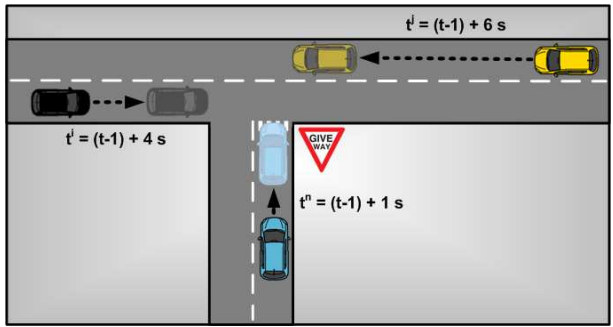
<p>Previous situational context: pose, speed, and intended course of vehicle n (blue car), vehicle i (black car), and vehicle j (yellow car).</p>	
<p>Steps 1 and 2</p>	
<p>Step 3</p>	<p>$k = i$ $g_{min} = 3 s$</p>
<p>Step 4</p>	$P([Es_t^n = stop] [Ic_{t-1} = c_{t-1}] [P_{t-1} = p_{t-1}] [S_{t-1} = s_{t-1}])$ $= 1.0 - \frac{1}{1 + e^{-\lambda(\ln(g_{min}) + (1-\alpha)\ln(s_{t-1}^k) - \ln(\gamma))}}$ $= 0.92$

Figure 3.9. Computation of $P(Es_t^n | Ic_{t-1} P_{t-1} S_{t-1})$ in an example situation.

3.4.2 Intended longitudinal motion Is_t^n

A number of different strategies could be applied for the intended longitudinal motion model, to model drivers' habits in terms of compliance with traffic rules. In this work, the model is based on the comparison between the previous intention Is_{t-1}^n and the current expectation Es_t^n :

$$P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) = P(Is_t^n | Is_{t-1}^n Es_t^n) \quad (3.15)$$

If the driver's intention at time $t - 1$ coincides with what is currently expected of him, it is assumed that the driver will comply with a probability P_{comply} . Otherwise a uniform prior (0.5) is assumed:

$I s_{t-1}^n$	$E s_t^n$	$P([I s_t^n = go] I s_{t-1}^n E s_t^n)$	$P([I s_t^n = stop] I s_{t-1}^n E s_t^n)$
<i>go</i>	<i>go</i>	P_{comply}	$1.0 - P_{comply}$
<i>go</i>	<i>stop</i>	0.5	0.5
<i>stop</i>	<i>go</i>	0.5	0.5
<i>stop</i>	<i>stop</i>	$1.0 - P_{comply}$	P_{comply}

When setting the value of P_{comply} , it is important to understand that different values correspond to different psychological models of drivers:

- $P_{comply} \in [0.0 \ 0.5[$ corresponds to the assumption that most of the time drivers disobey the rules.
- $P_{comply} = 0.5$ means that there is no prior assumption about whether drivers follow the rules or not.
- $P_{comply} \in]0.5 \ 1.0]$ corresponds to the assumption that most of the time drivers comply with the rules.

The probability P_{comply} is set to $P_{comply} = 0.9$ to reflect our assumption that **chances are high that the driver will comply**, but ideally it should be learned from data.

3.4.3 Intended lateral motion $I c_t^n$

In the general case of a vehicle driving from point A to point B on the road network, the lateral motion will change with time as one maneuver is performed after another. In this thesis the focus is on road intersections, and the possible lateral motions were defined as a set of possible courses. Courses cover the entire maneuver at the intersection (approaching phase + execution inside the intersection + exit phase), and there is no reason to believe that a driver will change his mind about the course he wants to follow. For this reason, it is assumed that the probability of a course at time t is dependent on the previous intended course only (Cf. Equation 3.16) and that drivers keep the same intention between two timesteps with probability P_{same} , all the other courses being equally probable (Cf. Equation 3.17).

$$P(I s_t^n | \Phi_{t-1}^n I_{t-1}^n E s_t^n) = P(I c_t^n | I c_{t-1}^n) \quad (3.16)$$

$$P([Ic_t^n = c_t^n] | [Ic_{t-1}^n = c_{t-1}^n]) = \begin{cases} P_{same} & \text{if } c_t^n = c_{t-1}^n \\ \frac{1.0 - P_{same}}{N_C - 1} & \text{otherwise} \end{cases} \quad (3.17)$$

The value of P_{same} was set manually to $P_{same} = 0.9$ to indicate that **drivers rarely change their intended course**, but should ideally be learned from data.

3.4.4 Pose P_t^n

It is assumed that **a vehicle performing a maneuver will follow the course corresponding to that maneuver**. The evolution of the pose of a vehicle is computed from its previous pose, previous speed, and current intended course:

$$P(P_t^n | \Phi_{t-1}^n I_t^n) = P(P_t^n | P_{t-1}^n S_{t-1}^n I_t^n) \quad (3.18)$$

The likelihood of a pose is defined as a trivariate normal distribution with no correlation between x , y and θ :

$$P(P_t^n | [P_{t-1}^n = p_{t-1}^n] [S_{t-1}^n = s_{t-1}^n] [Ic_{t-1}^n = c_{t-1}^n]) = \mathcal{N}(\mu_{xy\theta}(p_{t-1}^n, s_{t-1}^n, c_{t-1}^n), \sigma_{xy\theta}) \quad (3.19)$$

where $\mu_{xy\theta}(p_{t-1}^n, s_{t-1}^n, c_{t-1}^n)$ is a function which computes the mean pose $(\mu_x, \mu_y, \mu_\theta)$, and $\sigma_{xy\theta} = (\sigma_x, \sigma_y, \sigma_\theta)$ is the standard deviation. These terms are described in the paragraphs below.

Mean: The mean pose $(\mu_x, \mu_y, \mu_\theta)$ is computed by combining two models. The first model is a classic Constant Velocity (CV) model: a pose is computed from p_{t-1}^n and s_{t-1}^n by assuming that the speed and orientation of the vehicle is constant between $t - 1$ and t . The second model uses the pose computed by the Constant Velocity model and projects it orthogonally on the course c_{t-1}^n . The poses computed by the two models are averaged to produce a final predicted pose. Mixing these two models makes a compromise between the current physical state of the vehicle and the “ideal” pose that the vehicle would have if following the course perfectly.

The drawback of computing a pose as described above is that the physical limitations of the vehicle are not taken into account: this calculation predicts that the vehicle will “jump” in the direction of the course c_{t-1}^n . It can seem like a more elaborate path following model would be more representative of the true behavior of a vehicle. For example, a Proportional-Integral-Derivative (PID) controller could be applied on a “bicycle” model (see Section 2.1) to take into account the physical limitations of the

vehicle when it tries to follow the course. In practice in our experiments, taking into account the physical limitations of the vehicle systematically led to worse results compared with the simple model proposed above. The reasons for that are analyzed in Appendix E.

Standard deviation: The standard deviation values were manually set to $\sigma_x = \sigma_y = 0.2$ m and $\sigma_\theta = 0.1$ rad, but should ideally be learned from data.

3.4.5 Speed S_t^n

It is assumed that *drivers adapt their speed to their intentions and to the geometry of the road*. The evolution of the speed of a vehicle is computed from its previous speed, previous pose, and current intended maneuver:

$$P(S_t^n | \Phi_{t-1}^n I_t^n) = P(S_t^n | S_{t-1}^n P_{t-1}^n I c_t^n I s_t^n) \quad (3.20)$$

The distribution on S_t^n is assumed normal and defined as:

$$\begin{aligned} P(S_t^n | [S_{t-1}^n = s_{t-1}^n] [P_{t-1}^n = p_{t-1}^n] [I c_t^n = c_t^n] [I s_t^n = i s_t^n]) \\ = \mathcal{N}(\mu_s(s_{t-1}^n, p_{t-1}^n, c_t^n, i s_t^n), \sigma_s(s_{t-1}^n, p_{t-1}^n, c_t^n, i s_t^n)) \end{aligned} \quad (3.21)$$

where $\mu_s(s_{t-1}^n, p_{t-1}^n, c_t^n, i s_t^n)$ is a function which computes the mean speed and $\sigma_s(s_{t-1}^n, p_{t-1}^n, c_t^n, i s_t^n)$ is a function which computes the standard deviation. The paragraph below describes the speed profiles which were used as a model for the evolution of the speed, and is followed by the description of the calculation of the mean speed and the standard deviation.

Generic speed profiles: A number of statistical analyses of the behavior of drivers approaching an intersection can be found in the literature, e.g. [99]. From these it is possible to derive *generic speed profiles for vehicles negotiating an intersection*. For each possible pair $(c_t^n, i s_t^n)$, it is possible to define typical speed profiles of the type $f : d \rightarrow s$ where d is the distance to the intersection and s is the speed of the vehicle. We define:

$$\begin{cases} s_A : d \rightarrow s \\ s_M : d \rightarrow s \end{cases} \quad (3.22)$$

with s_A the average speed profile and s_M the maximum speed profile (i.e. the highest speed at which it is possible to execute the maneuver). The constraints imposed by the geometry of the road are taken into account: the values in the speed profiles are bounded by the highest speed at which it is physically possible for a vehicle to drive on course c_t^n . More specifically, the upper bound s_{max} for course c_t^n is calculated as:

$$s_{max} = \sqrt{\mu \times g \times r} \quad (3.23)$$

with μ the coefficient of friction (set to a conservative value of 0.65), $g = 9.81 \text{ m/s}^2$ is the gravity, and r the radius of curvature of course c_t^n . In Figure 3.10 the general aspect of the speed profiles is shown.




Figure 3.10. Example average and maximum speed profiles generated for a maneuver ($c_t^n, is_t^n = stop$) and for a maneuver ($c_t^n, is_t^n = go$).

Mean: The evolution of the speed of a vehicle is predicted using the following model (illustrated in Figure 3.11):

$$\mu_s(s_{t-1}^n, p_{t-1}^n, c_t^n, is_t^n) = s_A(d_t) - \frac{s_A(d_t) - s_M(d_t)}{s_A(d_{t-1}) - s_M(d_{t-1})} \times (s_A(d_{t-1}) - s_{t-1}^n) \quad (3.24)$$

with d_{t-1} the previous distance to the intersection calculated by projecting the pose p_{t-1}^n on course c_t^n , and d_t the current distance to the intersection calculated with a constant velocity model as $d_t = d_{t-1} + s_{t-1}^n \times \Delta t$.

This calculation accounts for some variations in the driving styles. Indeed, for one combination (c_t^n, is_t^n) , the predicted speed will depend on the previous speed and therefore is different for a “sporty” driver and for a “relaxed” driver, while still following the pattern defined for the pair (c_t^n, is_t^n) . This is illustrated in Figure 3.11.




Figure 3.11. Example calculation of $\mu_s(s_{t-1}^n, p_{t-1}^n, c_t^n, is_t^n)$ following Equation 3.24, for two different values of the previous speed s_{t-1}^n .

Standard deviation: The standard deviation is set dynamically based on the difference between the average and maximum speed at the current distance: $\sigma_s(s_{t-1}^n, p_{t-1}^n, c_t^n, is_t^n) = f(s_A(d_t) - s_M(d_t))$.

3.4.6 Turn signal state T_t^n

The presence of a turn signal is assumed to be dependent on the previous pose of the vehicle, and on both the previous and the current course followed by the driver, so as to consider that **turn signals can be used to indicate a turn or to indicate an intention to change lanes:**

$$P(T_t^n | \Phi_{t-1}^n I_{t-1}^n I_t^n) = P(T_t^n | P_{t-1}^n I_{t-1}^n I_t^n) \quad (3.25)$$

The computation is based on an analysis of the geometrical and topological characteristics of the intersection. In particular, we are interested in extracting the relationships between the different entrance

lanes (leading to the intersection) and exit lanes (leaving from the intersection) of the maneuvers. To this end, the following functions are defined (and illustrated in Figure 3.12):

- $entrance(c)$ returns the entrance lane of course c ,
- $exit(c)$ returns the exit lane of course c ,
- $angle(en_i, ex_j)$ returns the angle between entrance lane en_i and exit lane ex_j ,
- $most(en_i, ex_j)$ is true if exit lane ex_j is the leftmost or rightmost exit lane that can be reached from entrance lane en_i ,
- $unique(en_i)$ is true if only one exit lane can be reached from entrance lane en_i ,
- $side(en_i, en_j)$ returns 0 if entrance lane en_i and entrance lane en_j do not belong to the same road, 1 if en_j is located to the right of en_i , and 2 if en_j is located to the left of en_i .




Figure 3.12. Example results for $entrance()$, $exit()$, $angle()$, $most()$, $unique()$, and $side()$.

The pseudocode for the computation of the probability in Equation 3.25 is displayed in Figure 3.13, detailed in the paragraphs below, and illustrated in Figure 3.14.

```

1   $en_{t-1} = entrance(c_{t-1}^n)$ 
2   $ex_{t-1} = exit(c_{t-1}^n)$ 
3   $en_t = entrance(c_t^n)$ 
4   $ex_t = exit(c_t^n)$ 
5  IF  $en_{t-1} \neq en_t$ 
6      IF  $side(en_{t-1}, en_t) == 0$ 
7           $P(T_t^n | [P_{t-1}^n = p_{t-1}^n][Ic_{t-1}^n = c_{t-1}^n][Ic_t^n = c_t^n]) = f(p_{t-1}^n, angle(ex_{t-1}, ex_t), \gamma)$ 
8      ELSE
9           $P(T_t^n | [P_{t-1}^n = p_{t-1}^n][Ic_{t-1}^n = c_{t-1}^n][Ic_t^n = c_t^n]) = f(p_{t-1}^n, side(en_{t-1}, en_t), \beta)$ 
10 ELSE
11      $P(T_t^n | [P_{t-1}^n = p_{t-1}^n][Ic_{t-1}^n = c_{t-1}^n][Ic_t^n = c_t^n])$ 
12      $= f(p_{t-1}^n, angle(en_t, ex_t), most(en_t, ex_t), unique(en_t), \alpha)$ 

```

Figure 3.13. Pseudocode for the computation of the probability of a turn signal given the previous pose p_{t-1}^n , the previous intended course c_{t-1}^n , and current intended course c_t^n . Illustrative examples are provided in Figure 3.14.

The first case (line 6 to line 9) handles the situations where the turn signal indicates an intention to change lanes, since the entrance lane of the previous and current intended courses are different. When the current intended course is not legal (line 7, en_{t-1} and en_t do not belong to the same road, therefore c_t^n is not a course that is authorized for a vehicle approaching from en_{t-1}), the probability of the turn signal is a function of the angle between the exit of the previous intended course c_{t-1}^n and the exit of the current intended course c_t^n . When the current intended course is legal (line 9, en_{t-1} and en_t belong to the same road), the probability of the turn signal is a function of the relative position of en_{t-1} and en_t (left or right lane change).

The second case (line 11 to line 12) handles the situations where the turn signal is used to indicate an intention to make a turn. In this situation, the probability of a turn signal is computed as a function of the angle between the entrance lane and exit lane of the course. The function considers that the chance that a driver uses the turn signals is higher if the exit lane is on the extreme left or extreme right, and lower if the entrance lane gives access to a unique exit lane.

γ , ζ and ξ are multivariate constant variables which represent the probability that a driver puts on a turn signal in various scenarios (e.g. γ is the probability that a driver puts on a turn signal while taking a forbidden path). They were set manually, but should ideally be learned from data.




Figure 3.14. Examples of turn signal probabilities for different situations. The corresponding line number in the pseudocode is indicated above each image. The probability of the turn signal state being *left*, *none*, or *right* is displayed below each image.

3.4.7 Measured pose Pm_t^n

The uncertainties on the measurements of the pose are modeled using a **classic sensor model**. A trivariate normal distribution is assumed, centered on the true state and with no correlation between x , y and θ :

$$P(Pm_t^n | [P_t^n = p_t^n]) = \mathcal{N}(p_t^n, \sigma_{xy\theta}) \quad (3.26)$$

with $\sigma_{xy\theta} = (\sigma_x, \sigma_y, \sigma_\theta)$ the standard deviation. It was manually set to $\sigma_x = \sigma_y = 2$ m and $\sigma_\theta = \pi/6$ rad, according to the specifications of our sensors (GPS + IMU).

3.4.8 Measured speed Sm_t^n

The speed measurements given by our sensors (GPS + IMU) are accurate enough that the uncertainties can be ignored (accuracy < 1 m/s). Therefore a **Dirac function** is used for the speed sensor model:

$$P([Sm_t^n = sm_t^n] | [S_t^n = s_t^n]) = \delta(s_t^n - sm_t^n) \quad (3.27)$$

3.4.9 Measured turn signal state Tm_t^n

The turn signal state measurements given by the CAN-bus of a vehicle are reliable enough that the uncertainties can be ignored. Therefore a **Kronecker delta function** is used for the turn signal state sensor model:

$$P([Tm_t^n = tm_t^n] | [T_t^n = t_t^n]) = \delta_{t_t^n tm_t^n} \quad (3.28)$$

3.4.10 Summary

Throughout this section a number of independence assumptions were made. As a result the joint distribution in Equation 3.9 becomes:

$$\begin{aligned} P(\mathbf{E}_{0:T} \mathbf{I}_{0:T} \Phi_{0:T} \mathbf{Z}_{0:T}) &= P(\mathbf{E}_0 \mathbf{I}_0 \Phi_0 \mathbf{Z}_0) \\ &\times \prod_{t=1}^T \times \prod_{n=1}^N [P(Es_t^n | \mathbf{I}_{t-1} \mathbf{P}_{t-1} \mathbf{S}_{t-1}) \\ &\times P(Ic_t^n | Ic_{t-1}^n) \times P(Is_t^n | Is_{t-1}^n Es_t^n) \\ &\times P(P_t^n | P_{t-1}^n S_{t-1}^n Ic_t^n) \times P(S_t^n | S_{t-1}^n P_{t-1}^n Ic_t^n Is_t^n) \times P(T_t^n | P_{t-1}^n Ic_{t-1}^n Ic_t^n) \\ &\times P(Pm_t^n | P_t^n) \times P(Sm_t^n | S_t^n) \times P(Tm_t^n | T_t^n)] \end{aligned} \quad (3.29)$$

3.5 Bayesian risk estimation

We propose to detect dangerous situations by comparing what drivers intend to do with what they are expected to do. When applied to the specific case of road intersections (defined in Section 3.1.3), this principle leads to computing the risk based on the probability that a driver does not intend to stop while he is expected to:

$$P([Is_t^n = go][Es_t^n = stop] | \mathbf{P}_{\mathbf{m}_{0:t}} \mathbf{S}_{\mathbf{m}_{0:t}} \mathbf{T}_{\mathbf{m}_{0:t}}) \quad (3.30)$$

Exact inference on such a non-linear non-Gaussian model is not tractable. A number of solutions exist for approximate inference, the most popular being variational methods [100] and Markov Chain Monte Carlo (MCMC) methods [101]. Here, inference is performed using the classic **Sequential**

Importance Sampling (SIR) filter (also known as bootstrap filter or condensation algorithm) [102]. The inference equations are provided in Appendix F.

The bootstrap filter approach was implemented and tested for scenarios involving two vehicles (see the results in Chapter 4). A good compromise between computation time and quality of the estimation for two vehicles was achieved for a number of particles $N_{particles} = 400$.

3.6 Conclusions

3.6.1 Summary

This chapter introduced the proposed approach for risk estimation. Our solution is based on a Markov State Space Model which represents the joint motion of vehicles on the road. As opposed to the MSSMs proposed in the literature, ours represents explicitly both what drivers intend to do and what they are expected to do. This model allows to detect dangerous situations by comparing *intentions* and *expectations*, without having to compute the future trajectories of the vehicles. As mentioned in the previous chapter, to our knowledge there exists no ‘Interacting maneuvering entities’ motion model in the literature which is compatible with real-time constraints. The main source of computational complexity of state-of-the-art approaches is the prediction of future trajectories and their collision points. The MSSM proposed in this thesis offers a solution to this problem and therefore makes ‘Interacting maneuvering entities’ motion model compatible with real-time safety applications.

The proposed approach can in theory be applied to any type of traffic situation. An implementation for unsignalized road intersections was proposed in this chapter. The algorithm was described in four steps, following the Bayesian Programming formalism. The model takes into account the mutual influences between the maneuvers performed by the vehicles, and can be used with an arbitrary number of vehicles. It models the influence of the topology and the geometry of the intersection on the behavior of a vehicle, and therefore can automatically adapt to any intersection layout.

3.6.2 Discussion

In this section we address the possible application of the proposed approach to non-intersection traffic situations. As an example, we consider highway driving. The model defined above for road intersection situations is used as a base, and we investigate the extensions needed to apply it in the context of highways. We identified the following needs:

1. In the case of road intersections, we decomposed the possible maneuvers into their lateral and longitudinal component (see Section 3.2). We included both components in the *intention*, but only the longitudinal component in the *expectation*. This was done for two reasons:
 - In the specific context of road intersections, it is reasonable to assume that the presence of other vehicles does not influence the lateral motion of a vehicle.
 - The focus of this work was on accidents caused by traffic sign violations; driver errors in the lateral direction were not addressed.

For highway traffic situations, the lateral component has to be included in the *expectation*. The independence assumption for the lateral motion is no longer valid, since lane change maneuvers are highly dependent on the relative positions and speeds of the surrounding vehicles. Hazardous lane changes constitute a major risk in highway traffic situations, therefore driver errors in the lateral direction should be addressed if the model is to be used in this context.

2. In the case of road intersections, we characterized the longitudinal motion of a vehicle by the driver's intention to stop only (see Section 3.2.1). This is sufficient if the focus is on traffic sign violation scenarios, but should be extended if rear-end collisions are to be addressed (at intersections or on a highway). A solution would be to define an additional possible value for the intended longitudinal motion to represent the necessity for a vehicle to adapt its speed to the vehicle in front.

Following these recommendations, the intended maneuver $I_t^n = (Ic_t^n Is_t^n)$ and the expected maneuver $E_t^n = (Ec_t^n Es_t^n)$ could be defined as follows:

- $Ic_t^n \in \{c_i\}_{i=1:N_C}$: the driver's ***intended lateral motion***, corresponding to the course followed by vehicle n at time t . $\{c_i\}_{i=1:N_C}$ the set of possible courses is extracted from the digital map, with a course defined as the centerline of a lane on a highway and branching out into several courses near an intersection.
- $Is_t^n \in \{go, follow, stop\}$: the driver's ***intended longitudinal motion***, corresponding to the driver's intention regarding the longitudinal execution of the maneuver.

$Is_t^n = go$ means that the driver intends to adapt its speed to the geometry of the road only. This behavior is sometimes called "free flow" in the context of highways.

$Is_t^n = follow$ means that the driver intends to adapt its speed to the geometry of the road (similarly to $Is_t^n = go$), but will also adapt his speed so that he does not run into the vehicle

in front. The model for the evolution of the speed (Section 3.4.5) should take into account the relative speed and distance to the vehicle in front. For example, the Intelligent Driver Model [103] could be implemented.

$Is_t^n = stop$ means that the driver intends to adapt its speed to the layout of the intersection (similarly to $Is_t^n = go$), but will also adapt his speed so that he can stop at the next intersection.

- $Ec_t^n \in \{c_i\}_{i=1:N_C}$: the **expected lateral motion**. The definition is analogous to the definition of the intended lateral motion Ic_t^n ; the only difference is that it corresponds to what the driver should do (according to the traffic rules) instead of what he intends to do.
- $Es_t^n \in \{go, follow, stop\}$: the **expected longitudinal motion**. The definition is analogous to the definition of the intended longitudinal motion Is_t^n ; the only difference is that it corresponds to what the driver should do (according to the traffic rules) instead of what he intends to do.

This extended model could be used to evaluate the risk of situations in the context of highways or road intersections. By comparing driver *intention* $I_t^n = (Ic_t^n Is_t^n)$ and driver *expectation* $E_t^n = (Ec_t^n Es_t^n)$ it is possible to predict dangerous situations such as a vehicle changing lanes when it should not, a vehicle failing to adapt its speed to the vehicle in front, or a vehicle violating the rules at a road intersection.

Chapter 4

Experiments and results

Contents

4.1	Functional and safety evaluation in simulation	85
4.1.1	Evaluation strategy	85
4.1.2	Functional performance evaluation	90
4.1.3	Safety performance evaluation	93
4.2	Robustness evaluation with field trials	96
4.2.1	Evaluation strategy	96
4.2.2	Results	103
4.3	The importance of context - Interactions between vehicles	106
4.3.1	Evaluation strategy	108
4.3.2	Qualitative evaluation	109
4.3.3	Statistical results	113
4.4	The importance of context - Intersection layout	115
4.4.1	Evaluation strategy	116
4.4.2	Qualitative evaluation	118
4.4.3	Statistical results	120
4.5	Conclusions	123
4.5.1	Summary	123
4.5.2	Discussion	124

In the previous chapter, a Dynamic Bayesian Network (DBN) was proposed to model the motion of vehicles at road intersections. By performing inference on the relevant variables, it is possible to extract high-level information about the current situation (e.g. intended maneuvers of the drivers) and about the risk of the situation. Unlike the standard approaches in the literature, our approach does not rely on trajectory prediction to estimate the risk. Instead, risk is computed based on the comparison between what drivers intend to do and what they are expected to do.

The first two sections of this chapter are concerned with the evaluation of our risk estimation approach. Evaluating the performance of risk assessment algorithms is not straightforward: the ground truth of the risk of a situation is not available, therefore the evaluation cannot be conducted on the output of the algorithm directly. Instead it is generally conducted at the “application” level, by applying a threshold on the risk value to separate *dangerous* and *non-dangerous* situations. The algorithm is then evaluated based on the rate of false alarms, the rate of missed detections, and the collision prediction horizon. More specifically in this chapter, *dangerous* and *non-dangerous* situations are classified using the “hazard probability” criterion introduced in Section 3.1.2.2. The idea is to compute a “hazard probability” for every vehicle in the scene by comparing the driver’s *intention to stop* Is_t^n with the *expectation to stop* Es_t^n . Subsequently a situation is classified as *dangerous* if the hazard probability is higher than a threshold for at least one vehicle, i.e. iff:

$$\exists n \in N : P([Is_t^n = go][Es_t^n = stop] | \mathbf{Pm}_{0:t} \mathbf{Sm}_{0:t} \mathbf{Tm}_{0:t}) > \lambda \quad (4.1)$$

The threshold λ is set with the double objective to detect *dangerous* situations as early as possible and to avoid false alarms in *non-dangerous* situations.

Another difficulty for testing risk assessment algorithms is that *dangerous* situations are difficult to reproduce with real vehicles. Simulation can be used instead, but there are some limitations as to how realistic the simulated accident scenarios can be. The strategy adopted in this work was to conduct experiments both in simulation and with real vehicles. In Section 4.1, trajectories are simulated for vehicles involved collisions at a two-way stop intersection. Statistics are drawn about how early before the collision our algorithm is able to detect the danger. A preliminary study about the impact of different accident avoidance strategies is also presented. Section 4.2 describes the experimental setup for the field trials that were carried out. Two passenger vehicles were equipped with wireless communication modems and used to test the online execution the algorithm, for realistic scenarios at a T-shaped give-way intersection.

The aim of the final two sections is to evaluate the importance of contextual information in motion models. In Section 4.3 we compare two versions of our model to evaluate the impact of the classic assumption that vehicles move independently on the road. In Section 4.4 we show that taking into account the layout of the intersection allows us to recognize complex situations where a vehicle's behavior is inconsistent.

4.1 Functional and safety evaluation in simulation

Simulation can generate a large number of trajectories of vehicles involved in crashes at road intersections. This kind of data is very difficult to obtain with real vehicles: reproducing accidents is dangerous and costly, and collecting a database of real-life accidents would take a tremendous amount of time since accidents are rare events.

4.1.1 Evaluation strategy

A distinction should be made between functional performance and safety performance [86]. The former concerns the ability of the system to perform the functions it was designed for. The latter relates to the consequences in terms of safety after some action was taken based on the output of the system. More specifically, in our case:

- The **functional performance** is the the ability of the system to separate *dangerous* situations and *non-dangerous* situations. It is evaluated based on three criteria:
 1. The rate of false alarms, i.e. *non-dangerous* situations classified as *dangerous* by the algorithm,
 2. The rate of missed detections, i.e. *dangerous* situations classified as *non-dangerous* by the algorithm,
 3. The collision prediction horizon, i.e. how early before a collision the algorithm is able to classify the situation as *dangerous*. The collision prediction horizon $T_{prediction}$ is defined as:

$$T_{prediction} = t_{collision} - t_{detection} \quad (4.2)$$

with $t_{collision}$ the time at which the collision happens and $t_{detection}$ the earliest time when Equation 4.1 is true. This is illustrated in Figure 4.1.




Figure 4.1. Definition of the collision prediction horizon $T_{prediction}$.

The results of the functional performance evaluation are presented in Section 4.1.2.

- The *safety performance* can be measured as the impact of the system on the reduction or mitigation of accidents.

As was mentioned in Section 1.2.1 the focus of this thesis is situation assessment and risk estimation, and accident avoidance / mitigation strategies are out of the scope of this work. However, a preliminary study was conducted to assess the safety potential of four different strategies. The safety impact of each strategy was evaluated by looking at the percentage of accidents which could be avoided in the dataset if that specific strategy was triggered on a vehicle after a *dangerous* situation is detected by the algorithm.

The results of this study are presented in Section 4.1.3.

4.1.1.1 Data generation

The PreScan simulator [104] provides tools to create a custom road layout and to specify trajectories for vehicles driving on that road. The user can then synchronize the trajectories of two vehicles so that they intersect at a specific location. In order to generate realistic data, the simulated road network was replicated from a real intersection, as shown in Figure 4.2. The type of intersection, an X-shaped intersection with two stops, was chosen because it is a very common layout worldwide.

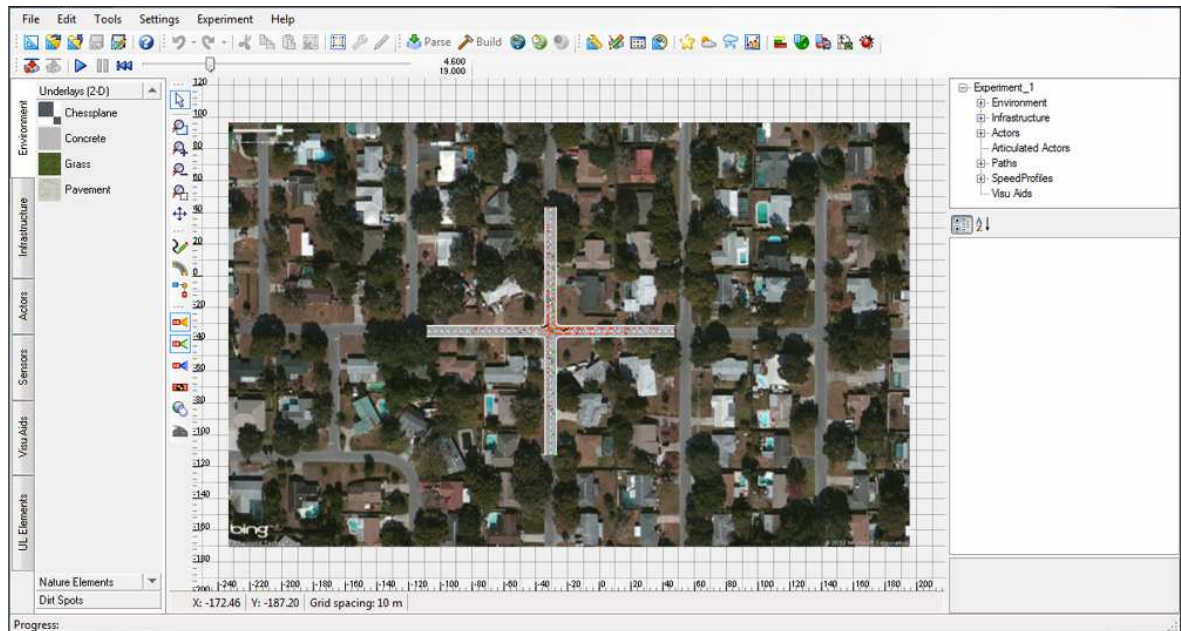


Figure 4.2. PreScan interface for simulating road traffic scenarios.

Ideal perception capabilities were assumed for the tests in simulation. The wireless communication link between the vehicles was assumed to never suffer packet losses and to provide instantaneous data transmission. The position measurements were also assumed to be perfect, i.e. the true position was used directly as a measurement without adding any noise.

The turn signal state was assumed to be equal to *none* for both vehicles all the time. The goal was to test the scenarios where the intended course of the drivers is not given away by the presence of a turn signal.

4.1.1.2 Test scenarios

Both *dangerous* and *non-dangerous* scenarios were simulated, all of them involving a Priority Vehicle (PV) driving on the main road and an Other Vehicle (OV) performing various maneuvers. The scenarios are illustrated in Figure 4.3 and in Figure 4.4, and described below.




Figure 4.3. Dangerous scenarios tested in simulation.




Figure 4.4. *Non-dangerous* scenarios tested in simulation.

The *dangerous* scenarios were defined in order to cover 70% of all accident scenarios at road intersections in Europe [105]. Four maneuvers were considered for the Other Vehicle: merging right (Scenarios A1, A2), merging left (Scenarios B1, B2), crossing (Scenarios C1, C2), and left turn across path (Scenario D). For each of these maneuvers except for the last one, there is a priority violation scenario and a stop violation scenario. A priority violation scenario is a situation where the Other Vehicle stops at the stop line but then proceeds into the intersection when it should have yielded to the Priority Vehicle, causing an accident. Such accidents are typically caused by the driver of the Other Vehicle failing to notice the presence of the Priority Vehicle, or misjudging the speed and distance of the Priority Vehicle [1]. A stop violation scenario is a situation where the Other Vehicle does not stop at the stop sign, causing an accident. Such accidents are typically caused by the driver of the Other Vehicle failing to notice the presence of the stop sign, or choosing to ignore it [1]. There is no stop violation scenario for the left turn across path maneuver since the regulation does not require the Other Vehicle to stop when executing this maneuver. A total of 240 instances of these scenarios were simulated, by varying the speed profiles of the Other Vehicle and the synchronization between the trajectories of the two vehicles.

For the *non-dangerous* scenarios, the same number of trajectories were generated for the same configurations as for the *dangerous* scenarios, this time without violating the traffic rules and with a safety distance of 3 seconds: the vehicles were always at least 3 seconds apart when passing the point where their paths intersected (Scenarios E, F, G, H).

4.1.2 Functional performance evaluation

As mentioned in Section 4.1.1 the functional performance is evaluated based on three criteria: the rate of false alarms, the rate of missed detections, and the collision prediction horizon. The threshold separating *dangerous* and *non-dangerous* situations was set to $\lambda = 0.3$, which was found to be optimal for these criteria (see precision and recall analysis in Appendix D). This value was used for all the tests presented in this thesis.

4.1.2.1 False alarms and missed detections

There were no false alarms, i.e. no *non-dangerous* situations classified as *dangerous* by the algorithm.

There were no missed detections, i.e. no *dangerous* situations classified as *non-dangerous* by the algorithm.

4.1.2.2 Collision prediction horizon

The distribution of the collision prediction horizon is displayed in Figure 4.5 for the 240 *dangerous* tests. The corresponding percentage of detected collisions as a function of the time remaining before the collision is shown in Figure 4.6.




Figure 4.5. Distribution of the collision prediction horizon $T_{prediction}$.




Figure 4.6. Percentage of detected collisions as a function of the time-to-collision.

The graphs show that every collision in the dataset was predicted at least 0.6 s before it occurred, and that a majority of collisions are predicted between 2 s and 3 s ahead of time. In 80% of the cases the algorithm is able to predict collisions at least 2 s before they occur. Based on these results, the potential benefits on safety are discussed in Section 4.1.3.

The next two paragraphs investigate further the functional performance by looking at the impact of the type of violation and the impact of the type of maneuver.

Impact of the type of violation: In order to analyze the influence of the type of violation on the functional performance of our algorithm, results for priority violation scenarios (Scenarios A1, B1, C1, D) are compared with results for stop violation scenarios (Scenarios A2, B2, C2). Figure 4.7 displays the percentage of detected collisions as a function of the time-to-collision for both types of violations.

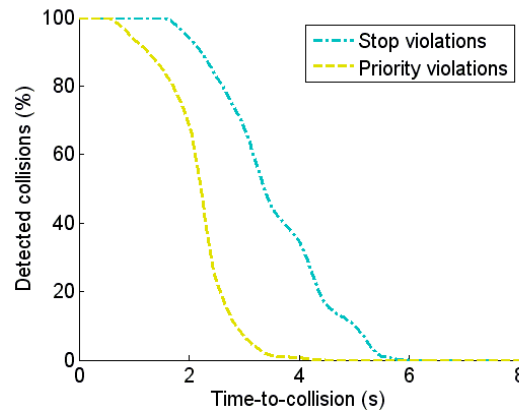


Figure 4.7. Comparison of the performance for priority violations scenarios (A1, B1, C1, D) and stop violations scenarios (A2, B2, C2).

On average, collisions which are caused by a stop violation are detected 1 s earlier than the ones caused by a priority violation. All of them are detected more than 1.5 s before the crash, against 0.6 s for priority violations. This is explained by the fact the Other Vehicle's intention to violate the stop is given away by the evolution of the vehicle's speed while it is approaching the intersection, while priority violations can be detected only as the Other Vehicle accelerates to enter the intersection.

Impact of the type of maneuver: In order to analyze the influence of the type of maneuver executed by the Other Vehicle on the functional performance of our algorithm, results for the four *dangerous* maneuvers present in our dataset are compared: merging right (Scenarios A1, A2), merging left (Scenarios B1, B2), crossing (Scenarios C1, C2), and left turn across path (Scenario D). Figure 4.8 displays the percentage of detected collisions as a function of the time-to-collision for the four types of maneuver.




Figure 4.8. Comparison of the performance for the different types of maneuver executed by the Other Vehicle: merging right (Scenarios A1, A2), merging left (Scenarios B1, B2), crossing (Scenarios C1, C2), and left turn across path (Scenario D).

The graph shows that 100% of the collisions are predicted at least 1.5 s before they occur for crossing and merging maneuvers. This means that all the collisions which were detected less than 1.5 s in advance in the previous graphs (Figure 4.5 and Figure 4.7) correspond to left turn across path maneuvers. These results highlight the fact that collisions caused by a priority violation during a left turn across path maneuver are more difficult to predict than other collisions. Indeed, contrary to crossing and merging maneuvers there is no marking on the road to indicate where a vehicle making a left turn across path should stop to yield to a vehicle with right-of-way. As a result there is a lot of variation in the way vehicles execute a left turn across path maneuver, and this makes it more challenging to estimate the intentions of drivers. The high variation in driver behavior during left turn across path maneuvers has been pointed out in the past by other works [98].

4.1.3 Safety performance evaluation

As was mentioned in the evaluation strategy (Section 4.1.1), safety performance evaluation is out of the scope of this work. The purpose of the preliminary study conducted below is to give an indication of the safety potential of different strategies when they are triggered after our algorithm detects a *dangerous* situation. For this purpose, four basic strategies were selected:

1. Autonomous emergency braking on the Priority Vehicle
2. Autonomous emergency braking on the Other Vehicle

3. Warning the driver of the Priority Vehicle

4. Warning the driver of the Other Vehicle

In this study the efficiency of a strategy is measured as its ability to avoid the accident. Accident mitigation (i.e. the reduction of the strength of the collision) would be another relevant criterion, but it is not addressed in this work.

In order to assess whether or not an upcoming accident can be avoided by triggering a specific strategy at time t , we compute the time needed by the vehicle to reach a full stop; this is the Time-To-Stop (TTS). If the TTS is smaller than the time remaining before the collision, the collision is considered to be avoidable.

When autonomous emergency braking is applied, the TTS is computed as follows [105]:

$$TTS = \frac{s_t}{\delta} + T_{machine} \quad (4.3)$$

with s_t the speed of the vehicle at time t , $\delta = 7 \text{ m/s}^2$ the deceleration, and $T_{machine} = 0.4 \text{ s}$ the average braking system response time [105].

When the strategy is to warn the driver instead of applying emergency braking, the response time of the driver has to be taken into account in the computation of the TTS [105]:

$$TTS = \frac{s_t}{\delta} + T_{machine} + T_{driver} \quad (4.4)$$

with $T_{driver} = 1.4 \text{ s}$ the average driver brake response time [105].

The results in terms of percentage of avoided accidents for the different strategies are displayed in Figure 4.9 and commented below.

- Auto-brake is always more efficient than driver warning, since the brakes are applied immediately after the danger is detected.
- In priority violation scenarios, actions on the Other Vehicle are always more efficient than actions on the Priority Vehicle. The reason is that the speed of the Other Vehicle is much lower than the speed of the Priority Vehicle at the instant when the violation is detected (the Other Vehicle was stopped and is accelerating to enter the intersection). A lower speed leads to a smaller TTS, which is why accidents are more easily avoided with actions triggered on the Other Vehicle.

- Actions on the Priority Vehicle have a significantly different impact in stop violation instances and in priority violation instances, while the difference is small for the Other Vehicle. Once again this can be explained by the speed difference.

The Priority Vehicle drives at the same speed in stop violation instances and priority violation instances, therefore the TTS is constant. However stop violations are detected earlier (see Figure 4.7), which leaves more time for the vehicle to stop compared with priority violation instances. As a result, on average 66.1% of accidents are avoidable by triggering an action on the Priority Vehicle in stop violation scenarios, but only 17.4% in priority violation scenarios.

The Other Vehicle drives at a higher speed in stop violations instances, compared with priority violation instances. The TTS is therefore higher in stop violation scenarios, and one could expect a lower percentage of avoided collisions in stop violation instances. However stop violations are detected earlier than priority violations (see Figure 4.7), which compensates for the high speed of the vehicle. As a consequence, the results are very similar for stop violations and priority violations: on average 83.9% of accidents are avoidable by triggering an action on the Other Vehicle in stop violation scenarios, and 80.6% in priority violation scenarios.

- The outermost numbers are obtained for priority violation instances: while 99.3% of the collisions in the dataset could be avoided by applying emergency braking on the Other Vehicle, only 0.7% could be avoided by warning the driver of the Priority Vehicle.

These results show the necessity to take into account the configuration of the accident when selecting an avoidance strategy.

	Auto-brake on OV	Warn driver of OV	Auto-brake on PV	Warn driver of PV
% of avoided accidents (stop violations)	93.1%	74.7%	90.8%	41.4%
% of avoided accidents (priority violations)	99.3%	61.8%	34.0%	0.7%

Figure 4.9. Percentage of avoided collisions depending on the type of violation and on the action taken when a danger is detected.

4.2 Robustness evaluation with field trials

In the previous section, simulated trajectories were used to evaluate the ability of our algorithm to detect *dangerous* situations. An analysis of the collision prediction horizon allowed us to characterize the efficiency of the algorithm in different situations, as well as the potential of different strategies to avoid accidents after the algorithm detected a *dangerous* situation. Simulation has the advantage that it can generate a large number of colliding trajectories, while this type of data is difficult to obtain with real vehicles. However our simulated data is not fully representative of real-life situations, mainly for two reasons:

1. There is little variation in the generated trajectories. All of them are very smooth, while real trajectories are rarely so. Besides, in practice there exists different driving styles, and drivers do not all negotiate an intersection in the same way (e.g. variations in the deceleration and acceleration profile).
2. The perception capabilities are assumed to be perfect. The wireless communication link between the vehicles was not simulated, which means the evaluation conducted in the previous section assumes perfect communication (no packet loss, instantaneous data transmission). The position measurements were also assumed to be perfect, i.e. they correspond to the real position of the vehicles. In real situations the positioning is not perfect and communication errors result in missing measurements for our algorithm, which can make the risk estimation process challenging.

For these reasons, it is important to validate the algorithm not only in simulation but also with real data. The objective of the evaluation conducted in this section is to validate the online functioning of the algorithm as well as its robustness to imperfect or missing input data.

4.2.1 Evaluation strategy

The field trials were conducted at a real road intersection with two passenger cars sharing information via wireless communication links. Our risk estimation algorithm was run online in one of the vehicles. Both *dangerous* and *non-dangerous* situations were created in order to evaluate the ability of our algorithm to analyze situations and estimate the risk in real-time.

Whenever the algorithm detected a *dangerous* situation (see definition in the introduction, Equation 4.1), an auditory and a visual warning were triggered. The goal of the auditory warning was to

catch the attention of the driver, while the visual warning provided the driver with some information about the nature of the detected danger (see Figure 4.10).




Figure 4.10. Images displayed by the HMI of the vehicle running the risk estimation algorithm. (a) Image displayed when the system predicts a collision with a vehicle coming from the left. (b) Image displayed when the system predicts no collision. (c) Image displayed when the system predicts a collision with a vehicle coming from the right.

The driver of the vehicle running the risk estimation algorithm was instructed to brake as soon as the warning appeared, while the driver of the other vehicle was instructed to not brake and continue the maneuver. Naturally, the driver of the car running the risk assessment algorithm was instructed to brake before a crash became unavoidable even if no warning was triggered. Moreover, in order to avoid confusions in the drivers' roles, each driver who took part in the trials was assigned one vehicle for the entire series of tests: a person either drove the car running the algorithm and had to brake in case of danger, or drove the other car and never had to perform emergency braking.

As with the tests in simulation, the performance of the algorithm was evaluated by looking at these three criteria:

1. The rate of false alarms, i.e. *non-dangerous* situations where the algorithm triggered a warning,
2. The rate of missed detections, i.e. *dangerous* situations where the algorithm did not trigger a warning,
3. The collision prediction horizon, i.e. how early before a collision the algorithm triggers a warning. Since we did not create real collisions, the actual collision prediction horizon could not be computed. Instead we looked at whether or not the algorithm was able to trigger a warning early enough to avoid accidents in *dangerous* situations.

4.2.1.1 Experimental setup

The tests were carried out at a T-shaped give-way intersection in a rural area near Guyancourt (France). As can be seen in Figure 4.11 it is a blind intersection, i.e. the vehicles approaching the intersection on the main road and on the secondary road cannot see each other until they reach the intersection.




Figure 4.11. Test intersection in Guyancourt: T-shaped give-way intersection. The view is obstructed by the vegetation, therefore the drivers cannot see each other while approaching the intersection.

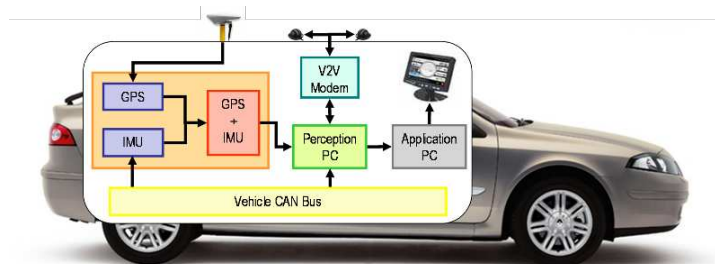
A Renault Espace and a Renault Laguna passenger vehicles were equipped with Cohda MK2 DSRC modems and shared their pose, speed and turn signal state information at a rate of 10 Hz. In each vehicle the pose and speed information was obtained via an IXSEA LandINS GPS + IMU unit with a precision of $\sigma = 2$ m for the position. The CAN-bus provided the turn signal state information. The position, orientation, speed, turn signal state, and DSRC messages were recorded in both vehicles during the experiments.

The system architecture in the test vehicles is illustrated in Figure 4.12. The Perception PC is responsible for reading information from the different sensors, reading and writing the messages received and sent by the V2V modem, and forwarding the relevant information to the Application PC. The Perception PC is equipped with an Intel Core 2 Duo 1.6 GHz processor, and 2 GB memory. The Application PC runs the risk estimation algorithm, and is present in the Laguna test vehicle only. It is equipped with

an Intel Core 2 Duo 2.26 GHz processor, and 3.48 GB memory. No attempt was made to optimize the code; in its current state the algorithm is implemented in C++ and runs at 10 Hz. The interactions between the different components of this architecture, as well as the logging of data, are handled by Intempora's RTMaps software [106] which runs on both the Perception PC and the Application PC.



(a) Espace test vehicle



(b) Laguna test vehicle

Figure 4.12. System architecture in the test vehicles.

Antenna placement: In order to optimize the communication performance, two antennas were used on each vehicle and placed on the rooftop at the front and at the back of the vehicle (see photo in Figure 4.13). The number of antennas and their placement has a large impact on the communication performance [107].



Figure 4.13. Antenna placement on the test vehicles.

DSRC performance tests: Some communication tests were conducted before the risk estimation tests, in order to evaluate the communication performance of our modems at the test site. Indeed, the efficiency of DSRC varies greatly with the environmental context, and no performance results have ever been reported in the context of a blind intersection where the obstacles to the communication are not buildings but a dense vegetation.

During the DSRC performance tests, one vehicle drove on the main road at 50 km/h and the other one on the secondary road at 30 km/h, both vehicles reaching the intersection at approximately the same time. The percentage of received packets as a function of the relative distance between the vehicles (calculated with the Haversine formula) is plotted in Figure 4.14.




Figure 4.14. Percentage of received packets as a function of the relative distance.

The results show a 100% packet reception when the relative distance between the vehicles is 50 m or less. This would not be acceptable for safety applications at high speed [108], but it is sufficient for our risk assessment tests since we limited the speed to 60 km/h for safety reasons.

It is important to keep in mind that our rural blind intersection scenario is one of the most challenging environments for DSRC. The signal is reflected on the vegetation and scattered in all directions (back to the sender car or towards the open field nearby), making communication difficult. Even the ‘multi-path estimation’ algorithm used by the Cohda modems to counter the multi-path effect offers no benefit in this situation.

4.2.1.2 Test scenarios

The purpose of the field trials was to evaluate the ability of our algorithm to correctly interpret a situation and estimate the risk in real-time for typical situations at an intersection. Both *dangerous* and *non-dangerous* scenarios were tested, all of them involving a Priority Vehicle (PV) driving on the main road and an Other Vehicle (OV) performing various maneuvers. The scenarios are illustrated in Figure 4.15 and Figure 4.16, and described below.




Figure 4.15. Dangerous test scenarios for the field trials.




Figure 4.16. Non-dangerous test scenarios for the field trials.

For the *dangerous* scenarios, two maneuvers were considered for the Other Vehicle: merging left (Scenario 1) and left turn across path (Scenario 2).

For the *non-dangerous* scenarios, tests were run for the same configurations as for the *dangerous* scenarios, but this time without violating the traffic rules (Scenarios 3, 4). An additional scenario was tested, where the Other Vehicle makes a right turn while the Priority Vehicle is approaching the intersection from the right (Scenario 5).

In total 90 *dangerous* and 20 *non-dangerous* tests were run. We alternated between 6 different drivers both for the Priority Vehicle and the Other Vehicle. In order to generate some diversity in the scenario instances, the drivers of the Other Vehicle were not given precise instructions about the execution of the maneuvers; they were told to execute the maneuver in a manner that they felt was *dangerous* or *non-dangerous* depending on the scenario. As a result we obtained some significant variations in the instances of the scenarios. As an example, the range of variation of the position and speed for the Priority Vehicle when the Other Vehicle violates the priority in the *dangerous* scenarios are displayed in Figure 4.17. In addition to the variations in positions and speeds, we obtained variations in the behavior of the Other Vehicle when approaching the intersection: in some instances the driver stopped at the intersection before executing his maneuver, and in some other instances proceeded without stopping.

Out of the 90 *dangerous* trials, 60 were performed with the warning system running on the Priority Vehicle, and 30 with the warning system running on the Other Vehicle.




Figure 4.17. Range of variation of D (distance to the intersection) and S (speed) for the Priority Vehicle when the Other Vehicle violates the priority in the two *dangerous* scenarios.

4.2.2 Results

The first part of this section reports the observations made on-site during the field trials. In the second part, screenshots of offline replays are used to give an idea of the system's behavior on specific examples.

4.2.2.1 On-site observations

The first tests with *dangerous* and *non-dangerous* situations were performed without using turn signals, but we found that this was too risky in *dangerous* scenarios (Scenarios 1, 2). The system was always able to trigger a warning before the accident became unavoidable, but too late to run our tests in safe conditions. This observation backs up the results obtained in simulation, which showed that driver warning is not an efficient strategy in priority violation scenarios. Of course the scenarios tested in simulation and in the field trials are different, but the priority violation scenarios are similar enough that we can draw parallels: Scenario 1 is of the same nature as Scenarios A1 and B1 (priority violation during a merging maneuver), and Scenario 2 is of the same nature as Scenario D (priority violation during a left turn across path maneuver).

For safety reasons, the subsequent *dangerous* scenarios tests were always performed with the Other Vehicle indicating its maneuver with a turn signal. Thanks to the presence of the turn signal, there is less uncertainty on the driver's intended course. Warnings are triggered earlier, which makes the field trials safer. For every of the 90 test runs, the system was able to issue a warning early enough that the driver avoided the collision by braking.

Turn signals were not used in *non-dangerous* scenarios (Scenarios 3, 4, 5). In the 20 test instances there was no false alarm, i.e. no warning was ever triggered by the system.

4.2.2.2 Replay of sample examples

Figure 4.18 shows a sample *dangerous* left turn across path scenario during which we recorded the view from inside the Priority Vehicle. The danger is detected as soon as the Other Vehicle starts to execute the left turn. The driver of the Priority Vehicle receives a warning indicating that a *dangerous* situation was detected and that the danger comes from a vehicle on the left side.




Figure 4.18. *Dangerous* situation detected during the field trials (Scenario 2). The left column displays the position, heading and turn signal state of the vehicles. When the system detects a *dangerous* situation, the vehicle causing the danger is displayed in red. The right column shows the view from inside the Priority Vehicle.

The scenario in Figure 4.19 (a) is a sample *non-dangerous* scenario. There is no false alarm, even when the Other Vehicle proceeds into the intersection and the Priority Vehicle is approaching at high speed ($t = 3.1$ s).




Figure 4.19. Sample scenarios recorded during the field trials. When the system detects a *dangerous* situation, the vehicle causing the danger is displayed in red.

The scenario in Figure 4.19 (b) is a sample *dangerous* scenario with challenging errors in the localization of the Other Vehicle. The position measurements are shifted several meters in the North direction,

and as a result it appears as if the car is still behind the give-way line when in reality it has already entered the intersection. Nevertheless the system is able to detect the danger, thanks to the fact that the algorithm handles measurement uncertainties.

The scenario in Figure 4.19 (c) is a sample *dangerous* scenario with wireless communication issues. The communication between the two cars is lost during 1 s after the first message was received (at $t = 0.1$ s). Between $t = 0.2$ and $t = 1.2$ s, the Priority Vehicle does not receive any new measurements to indicate the intentions of the Other Vehicle. When the second message is received, the Other Vehicle is already in the middle of the intersection ($t = 1.3$ s). The system is able to detect immediately that this is a *dangerous* situation.

Since we did not go so far as to create real collisions, the statistical analysis which was performed on the simulated data cannot be performed on the real data. However, the field trials proved that our approach can operate with success in real-life situations where passenger vehicles share data via a V2V communication link. The experiments also showed the robustness of the algorithm. The tests involved 6 different drivers, who were not given specific instructions about the execution of the scenarios. The algorithm was able to detect *dangerous* situations despite bad localization and communication errors, and did not trigger any false alarms.

4.3 The importance of context - Interactions between vehicles

The two previous sections evaluated the ability of the proposed risk estimation algorithm to assess the risk of a situation. This section and the next one focus on evaluating the importance of taking into account the context in motion models for situation assessment and risk estimation.

We claimed throughout this thesis that it is necessary to account for the mutual influences between the vehicles' maneuvers to correctly interpret situations and risks. This claim is supported by the analytical study conducted on example scenarios in Appendix B. The conclusions of this study is that accounting for inter-dependencies between the vehicles:

- Should in theory lead to a better estimation of drivers' maneuver intentions
- Should in theory not delay or prevent the detection of *dangerous* situations
- Should in theory lead to a better sensitivity of the risk assessment

The purpose of this section is to evaluate these claims with real data. To this end, we will compare the performances of two version of our motion model: the 'Interacting vehicles' motion model and the

‘Independent vehicles’ motion model.

The ‘Interacting vehicles’ motion model corresponds to the model presented in the previous chapter and evaluated in the previous sections in simulation and in field trials. This model takes into account the dependencies between the vehicles by making the intended longitudinal motion (Is_t^n) dependent on the expected longitudinal motion (Es_t^n):

$$P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) = P(Is_t^n | Is_{t-1}^n Es_t^n) \quad (4.5)$$

Details were given in the previous chapter in Section 3.4.2.

The ‘Independent vehicles’ motion model ignores the dependencies between the vehicles. The assumption that vehicles move independently on the road is implemented by removing the link between Es_t^n and Is_t^n from the ‘Interacting vehicles’ motion model. The assumption that “drivers tend to comply with traffic rules” made in the ‘Interacting vehicles’ motion model is replaced by a more conservative assumption independent of the traffic rules.

Two solutions were explored. The first one assumes that “no prior knowledge is available about a driver’s intention to stop”. Therefore the intended longitudinal motion for a vehicle n at time t becomes a uniform distribution:

$$P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) = P(Is_t^n) = 0.5 \quad (4.6)$$

The second one assumes continuity in the driver’s intended longitudinal motion:

$$P(Is_t^n | \Phi_{t-1}^n I_{t-1}^n Es_t^n) = P(Is_t^n | Is_{t-1}^n) = \begin{cases} P_{same} & \text{if } Is_t^n = Is_{t-1}^n \\ 1.0 - P_{same} & \text{otherwise} \end{cases} \quad (4.7)$$

with P_{same} the probability that the driver’s intention stays the same between two timesteps.

Both solutions were implemented and tested, and it was found that they lead to very similar results when compared with the ‘Interacting vehicles’ motion model. In this section, only the results obtained with the “no prior knowledge” assumption (Equation 4.6) will be described, but the conclusions drawn also apply to the model which assumes “continuity in the intention” (Equation 4.7).

The differences between the ‘Interacting vehicles’ motion model and the ‘Independent vehicles’ motion model are summarized on the graphical model and joint distribution in Figure 4.20.




Figure 4.20. The differences between the ‘Interacting vehicles’ motion model and the ‘Independent vehicles’ motion model are shown in the graphical representation and in the joint distribution. The elements shown in red do not exist in the ‘Independent vehicles’ model.

4.3.1 Evaluation strategy

The comparative evaluation was done using the data recorded during the field trials. As described in Section 4.2, the field trials were conducted with two vehicles involved in typical *dangerous* and *non-dangerous* situations at an intersection. The logs recorded during the field trials were replayed offline to compare the ‘Independent vehicles’ motion model and the ‘Interacting vehicles’ motion model.

During the field trials, the Other Vehicle always indicated its intended course using the turn signal for *dangerous* scenarios. This was done for safety reasons, as explained in Section 4.2.2. The turn signals were removed from the replays so that we could compare the ability of each model to infer the intended course of drivers without the strong clue given by the turn signal.

The performances of the two models were measured by looking at their ability to:

- Correctly estimate the intended course of the drivers: This is measured by $P([Ic_t^n = c_{GT}^n])$, with c_{GT}^n the ground truth for the maneuver of vehicle n (the ground truth course is defined by the scenario).
- Correctly assess the situation in general: This will be reflected by N_{eff} , the effective sample size of the particle filter.

- Correctly assess the risk: In the case of a *dangerous* situation, this means detecting that the situation is *dangerous* as early as possible. Therefore the detection time $t_{detection}$ should be as early as possible. In the case of a *non-dangerous* situation the risk should not be overestimated: this feature can be evaluated by monitoring $P([Is_t^n = go][Es_t^n = stop])$.

4.3.2 Qualitative evaluation

In this section, the performance differences between the two models are analyzed on sample instances of a *non-dangerous* scenario and of a *dangerous* scenario.

Qualitative evaluation on a *non-dangerous* scenario (Figure 4.21)




Figure 4.21. Estimated intended course (left graph) and risk (right graph) associated to the Other Vehicle in a sample instance of a *non-dangerous* scenario (Scenario 5). Screenshots of the situation at t_1 , t_2 , and t_3 are displayed above the graphs.

In this scenario the Other Vehicle makes a right turn at the intersection while the Priority Vehicle

is approaching the intersection from the right. The Other Vehicle can execute its maneuver without considering the Priority Vehicle since the two vehicles always stay on different lanes. The graphs can be broken into three phases:

- Before time t_1 , both vehicles are driving toward the intersection. The Other Vehicle is slowing down, but this does not give any indication about the driver's intended course since all the potential courses require slowing down. For this reason both motion models consider that a left turn and a right turn are equally likely. The risk is considered low by the two models, since the Other Vehicle seems to comply with the traffic rules.
- Between time t_1 and time t_2 , the Other Vehicle is proceeding into the intersection. At the beginning, the heading of the car does not give any clue about the intended course of the driver. A left turn would clearly be a violation of the priority rules, since there is not enough time to merge before the Priority Vehicle reaches the intersection. At this stage, the 'Independent vehicles' motion model has no way to tell whether the driver of the Other Vehicle intends to turn left or right. For this reason, the risk value reaches 0.5. The 'Interacting vehicles' motion model favors the right turn as the most likely intention of the Other Vehicle, and as a result the risk remains low.
- Between time t_2 and time t_3 , the Other Vehicle executes its maneuver inside the intersection. The clues provided by its position and heading leave less and less doubt about the intended course. Both motion models are able to estimate correctly the intended course of the driver. Making a right turn while the Priority Vehicle is approaching is not against the rules, therefore the risk remains low.

In this *non-dangerous* scenario, the difference between the two motion models is visible in the time interval $t_1 - t_2$. The 'Interacting vehicles' motion model is able to predict the intended course of the Other Vehicle earlier, and to keep the risk value low. The 'Independent vehicles' motion model produces a peak in the risk value. As a consequence, the threshold λ which is used to separate *dangerous* situations from *non-dangerous* situations (see Equation 4.1) has to be set to a different value for the two motion models. As was mentioned earlier, it was set to $\lambda = 0.3$ for the 'Interacting vehicles' motion model. In the case of the 'Independent vehicles' motion model, λ has to be set to a value higher than 0.5 in order to avoid systematic false alarms in situations like the one in Figure 4.21. A statistical analysis of our results led us to set $\lambda = 0.65$ for the 'Independent vehicles' motion model. It corresponds to the lowest value which does not trigger any false alarm on the dataset. In *dangerous* situations, the value of the risk varies in the range $[0.3 \ 1.0]$ if the 'Interacting vehicles' motion model

is used, and in [0.65 1.0] if the ‘Independent vehicles’ motion model is used. Therefore the sensitivity of the risk estimation [109] is significantly better with the motion model which accounts for the dependencies between the vehicles.

Qualitative evaluation on a *dangerous* scenario (Figure 4.22)




Figure 4.22. Estimated intended course (left graph) and risk (right graph) associated to the Other Vehicle in a sample instance of a *dangerous* scenario (Scenario 1). Screenshots of the situation at t_1 , t_2 , and t_3 are displayed above the graphs.

In this scenario the Other Vehicle makes a left turn at the intersection while the Priority Vehicle is approaching the intersection from the right. This is a priority violation case, and without an emergency braking the situation would have resulted in a crash. The graphs can be broken into four phases:

- Before time t_1 , the situation is the same as in the *non-dangerous* scenario presented above.
- Between time t_1 and time t_2 , the Priority Vehicle is approaching the intersection and the Other Vehicle remains stationary at the give-way line. Over this interval, the probability that the Other

Vehicle will turn left increases with the ‘Interacting vehicles’ motion model and stays stable with the ‘Independent vehicles’ motion model. The ‘Interacting vehicles’ motion model is able to interpret the stillness of the Other Vehicle as an indication that the driver’s intention is to turn left. A human observing the scene would reach the same conclusion, since there is no reason for the Other Vehicle to wait at the give-way line if its intention is to turn right. The ‘Independent vehicles’ motion model does not have this ability, since it assumes independence between the vehicles.

- At time t_2 , the Other Vehicle accelerates to enter the intersection. At first, the heading of the car does not give any clue about the intended course of the driver. The probability for the maneuver intention remains stable for the ‘Independent vehicles’ motion model, while the probability that the vehicle is turning left decreases with the the ‘Interacting vehicles’ motion model. Indeed, turning right is compatible with the priority rules while turning left is not, therefore it is the preferred interpretation. Here again, a human would reach the same conclusion as the ‘Interacting vehicles’ motion model. The value of the risk reaches 0.5 with the ‘Independent vehicles’ motion model and remains low with the ‘Interacting vehicles’ motion model, similarly to the *non-dangerous* scenario studied in the paragraph above. After one second the heading of the Other Vehicles reflects the fact that the driver intends to make a left turn. Both motion models are able to catch this clue, and the value of risk rises fast. Using the values of λ defined above, we can see that the two models are able to predict the collision at exactly the same time $t_{detection} = t_3$.
- After time t_3 , the probability that the Other Vehicle intends to turn left keeps on rising steadily as more clues about the driver’s intended course (position, heading) become available. The same observation can be made for the risk, which becomes maximal a few hundred milliseconds later.

In this *dangerous* scenario, the two motion models initially have a different interpretation of the behavior of the Other Vehicle but they detect the danger at the same time.

The objective of this section was to analyze qualitatively the behavior of the ‘Independent vehicles’ motion model and the ‘Interacting vehicles’ motion model in sample scenarios. The results are in accordance with the claim made at the beginning of the section. In the *non-dangerous* scenario, taking into account the dependencies between the vehicle improves the estimation of a driver’s intended course. The assumption that drivers generally respect traffic rules does not prevent or delay the detection of the *dangerous* situation, and improves the sensitivity of the risk estimation.

In the next section, statistical results are derived using the whole field trials data.

4.3.3 Statistical results

We computed statistical results over the 90 *dangerous* and 20 *non-dangerous* test instances from the field trials in order to verify if the differences observed on specific instances in the previous section can be generalized. The comparison between the ‘Independent vehicles’ motion model and the ‘Interacting vehicles’ motion model is carried out by subtracting the results obtained by the two models. We define:

$$\left\{ \begin{array}{l} \Delta P([\mathbf{Ic}_t = \mathbf{c}_{GT}]) \\ \Delta N_{eff} \\ \Delta t_{detection} \\ \Delta P([\mathbf{Is}_t = \mathbf{go}][\mathbf{Es}_t = \mathbf{stop}]) \end{array} \right. \begin{array}{l} = P([\mathbf{Ic}_t = \mathbf{c}_{GT}]^{Int} - P([\mathbf{Ic}_t = \mathbf{c}_{GT}]^{Ind} \\ = (N_{eff}^{Int} - N_{eff}^{Ind})/N_{particles} \\ = t_{detection}^{Int} - t_{detection}^{Ind} \\ = P([\mathbf{Is}_t = \mathbf{go}][\mathbf{Es}_t = \mathbf{stop}]^{Int} - P([\mathbf{Is}_t = \mathbf{go}][\mathbf{Es}_t = \mathbf{stop}]^{Ind} \end{array}$$

where the superscript indicates the motion model used for the calculation (*Int* for ‘Interacting vehicles’ and *Ind* for ‘Independent vehicles’). The number of efficient particles is divided by the total number of particles $N_{particles}$ so that a relative value between 0 and 1 is obtained. The results are displayed in Figure 4.23 and commented below.

Estimation of the intended course: In *dangerous* situations, the two approaches perform similarly on average. When there is no danger, the two approaches estimate the intended course equally well 53% of the time. However the ‘Interacting vehicles’ motion model leads to a better estimation on average and performs better 42% of the time.

From the case-based reasoning conducted in Appendix B.2 and from the qualitative evaluation performed in the previous section, we know that the situations where the ‘Interacting vehicles’ motion model performs better are either 1) when a vehicle is waiting at the entrance of an intersection for a priority vehicle to pass or 2) when a vehicle is proceeding into an intersection while a vehicle with priority is approaching the intersection.

The ability to estimate the intended course of a driver is crucial for intelligent vehicles. In this thesis, the estimated intended course is used to compute the risk of a situation, but there are other applications. An autonomous vehicle negotiating an intersection needs to adapt its behavior to the maneuvers performed by the other vehicles, for example it has to decide whether or not to proceed into the intersection based on the priority rules and on the intended maneuvers of the other vehicles.




Figure 4.23. Comparative results between the ‘Independent vehicles’ motion model and the ‘Interacting vehicles’ motion model, for *dangerous* scenarios (top) and *non-dangerous* scenarios (bottom).

Neff: Accounting for the dependencies between the vehicles has a positive and significant impact on the number of efficient particles both when the situation is *dangerous* and when it is not. In both cases the performance is increased 57% of the time and mostly equivalent the rest of the time.

These results validate our modeling of how vehicles influence each other. Indeed, the value of N_{eff} reflects the quality of the tracking, which itself is dependent on how well the motion model matches the actual motion of the vehicles.

Risk estimation: The plot of $\Delta P([I_{s_t} = \text{go}][E_{s_t} = \text{stop}])$ shows that the risk is systematically lower when estimated using the ‘Interacting vehicles’ motion model in *non-dangerous* situations. The reason behind it is the assumption that drivers tend to respect traffic rules. A direct consequence is that the threshold λ has to be set to a different value for the two models, and that the sensitivity of the risk assessment is better with the ‘Interacting vehicles’ motion model. This was explained in more detail in Appendix B.2 and in Section 4.3.2.

In *dangerous* situations, one could expect that the ‘Interacting vehicles’ motion model would be less conservative than the ‘Independent vehicles’ motion model because of the assumption that drivers tend to respect traffic rules. However in 42% of the cases the warning is issued at the exact same time by the two approaches. On average the ‘Interacting vehicles’ motion model even detects *dangerous* situations slightly earlier. The sample *dangerous* situation studied in the previous section illustrates why: in *dangerous* situations the risk starts rising earlier with the ‘Independent vehicles’ motion model, but since the threshold λ is different for the two approaches the final danger detection time is similar.

4.4 The importance of context - Intersection layout

The work presented in this section was conducted during a 3-month internship in the Stanford Artificial Intelligence Laboratory. Like the previous section, this section focuses on evaluating the importance of taking into account context in motion models for situation assessment and risk estimation. The previous section addressed the importance of accounting for the dependencies between the vehicles in the motion model. This section focuses on the benefits of taking into account the geometry and topology of the intersection when estimating a driver’s intended course in the presence of turn signals. The motion model proposed in this thesis models the influence of the layout of the intersection on the turn signal state (see Section 3.4.6), and we will show using real data that this modeling allows us to handle complex situations where the vehicle’s behavior is inconsistent. Since our risk estimation approach relies on the estimated intended course of the drivers, taking into account the layout of the

intersection also has consequences on the quality of the risk assessment. However we did not evaluate this impact; this section focuses on the impact on the estimation of the intended course I_c .

4.4.1 Evaluation strategy

Data was collected in Stanford (U.S.) at the two road intersections shown in Figure 4.24. This type of layout is very standard in the U.S.: four-way-stop intersections with multiple entrance and exit lanes in each road. Some lanes are reserved for specific maneuvers, as indicated by the arrows painted on the road surface. Our goal was to collect real trajectories of vehicles negotiating relatively complex intersections where the layout of the intersection provides clues about the maneuver intention of drivers. We tested the ability of our algorithm to estimate the intended course of drivers in two types of situations: situations where the driver puts a turn signal which is consistent with the executed maneuver, and situations where it is inconsistent. This way we were able to investigate the ability of the algorithm to balance the clues provided by the turn signal state and the clues provided by the layout of the intersection in challenging situations.

The tests were performed using an early version of the DBN. The early model is based on the same principles as the model presented in the previous chapter, except the kinematic behavior of the vehicles is not modeled: only the relationships between the pose of a vehicle, its turn signal state, and the intended course of the driver are represented. Implementation details can be found in [110, 111]. Working with this simplified model made it easier to analyze the results and to identify the benefits of taking into account the intersection layout in the model.



Figure 4.24. Test intersections in Stanford: four-way-stop intersections with multiple entrance and exit lanes in each road.

4.4.1.1 Data collection

Data was collected using the autonomous car “Junior”, a replica of the robotic car which took second place at the DARPA Urban Challenge in 2007 [112]. The vehicle is a Volkswagen Passat wagon equipped with a variety of sensors including a Velodyne 64-beam laser rangefinder, four cameras, six radars, two one-beam laser scanners, and an inertial GPS system. The autonomous navigation capabilities of the car arise from the software modules implemented on it, e.g. mapping and localization, object detection, trajectory planning and control [113].

In this work, Junior was used as a data collection platform. We parked the car by the side of the intersection of interest (so that to not disrupt traffic) and recorded Velodyne data while normal traffic was proceeding (see Figure 4.25). The trajectories of 42 vehicles were manually annotated offline using Stanford’s Velodyne data labeling tool.



Figure 4.25. "Junior", the autonomous car used for data collection. The Velodyne laser rangefinder mounted on the roof top provides a 3D point cloud of the car surroundings.

4.4.1.2 Test scenarios

The aim was to test our algorithm on two types of scenarios:

- *Consistent* behavior scenarios:

A vehicle is approaching, then negotiating the intersection. The turn signal state of the vehicle is consistent with the executed maneuver (see examples in Figure 4.26).

- *Inconsistent* behavior scenarios:

A vehicle is approaching, then negotiating the intersection. The turn signal state of the vehicle is not consistent with the executed maneuver (see examples in Figure 4.26).




Figure 4.26. Examples of *consistent* and *inconsistent* behavior scenarios.

Two scenario instances were generated for each of the 42 recorded trajectory: one with a *consistent* turn signal and one with an *inconsistent* turn signal. For example, a left turn trajectory was labeled with a left turn signal (or no turn signal) to generate a *consistent* behavior scenario instance, and with a right turn signal to generate an *inconsistent* behavior scenario instance. The labels were automatically generated for each trajectory in order to incorporate some randomness, and the turn signal state was assumed to stay the same for the whole trajectory.

4.4.2 Qualitative evaluation

The objective is to evaluate the algorithm's ability to make reliable interpretations of a vehicle's behavior in *consistent* and *inconsistent* scenarios. Figure 4.27 shows the output of the algorithm for a left turn maneuver scenario. In this example the allowed maneuvers for a vehicle approaching the intersection from the south are: turn left / go straight from the left lane and turn right / go straight from the right lane.




Figure 4.27. Estimation of the intended course in a *consistent* behavior scenario and in an *inconsistent* behavior scenario. The size of the green disks located on the exit lanes is proportional to the probability that the driver intends to reach that exit lane.

The results can be interpreted as follows:

Consistent behavior scenario: The vehicle reaches the intersection through the left lane with the left turn signal on (*a1*, *a2*). The algorithm interprets this as a high probability that the driver's intention is to turn left, and a much smaller probability that the driver's intention is to go straight. The probability for a right turn is negligible. The algorithm becomes increasingly confident that the driver's intention is to turn left as the vehicle makes its way through the intersection, thanks to the clues given by the vehicle's position and heading (*a3*, *a4*).

Inconsistent behavior scenario: The vehicle reaches the intersection through the left lane with the right turn signal on ($b1$, $b2$). When the vehicle is still far away from the intersection ($b1$), this behavior is interpreted as an intention to change lanes. There is also a chance that the vehicle is actually on the right lane (i.e. positioning error), which is why the probability of turning right is higher than the probability of going straight. When the vehicle gets closer to the intersection ($b2$) the algorithm has no preferred interpretation because this behavior makes little sense. A human would come to the same conclusion and consider all maneuvers to be equally likely. As the vehicle executes the maneuver inside the intersection it becomes more and more obvious that the vehicle will turn left since the other maneuvers are no longer feasible, therefore the uncertainty of the algorithm's estimation decreases ($b3$, $b4$).

The algorithm was tested qualitatively on the 42 trajectories, and this led to the conclusion that it is able to reason about the intended course of drivers in a reliable manner even in complex scenarios where the behavior of the vehicle is inconsistent. In the next section, a statistical analysis of these results is performed.

4.4.3 Statistical results

Conventionally the performance of an algorithm which estimates a driver's intended course is evaluated by looking at how early and/or frequently it is able to make correct predictions. A prediction is considered correct if the true course is found to be much more probable than the others, incorrect otherwise. This approach is not appropriate in our case because we evaluate our algorithm in situations where it is not reasonable to expect such correct predictions. For example in situation $b2$ of Figure 4.27 it is not possible for the algorithm (or for a human) to estimate the driver's intended course with high certainty from the clues provided by the map and the vehicle's behavior. Our evaluation metrics should take this into account, and reflect that a large uncertainty on the output is to be expected in this situation.

We define $c_A \in \{c_i\}_{i=1:N_C}$ the most probable intended course according to the algorithm at time t , $c_B \in \{c_i\}_{i=1:N_C}$ the second most probable intended course according to the algorithm at time t , and $c_{GT} \in \{c_i\}_{i=1:N_C}$ the ground truth course (i.e. the course which is eventually followed by the driver). We consider three possible labels for a prediction at time t :

- *Correct:* $P(c_A) - P(c_B) > 0.2$ and $c_A = c_{GT}$.
- *Incorrect:* $P(c_A) - P(c_B) > 0.2$ and $c_A \neq c_{GT}$.

- *Undecidable*: $P(c_A) - P(c_B) \leq 0.2$. Intuitively, *undecidable* predictions correspond to cases where the algorithm is not able to interpret a vehicle's behavior.

The statistical evaluation was run on the 42 trajectories for both types of scenarios (*consistent* / *inconsistent*). The graphs in Figure 4.28 show the obtained results and are commented below.




Figure 4.28. Percentage of *correct*, *incorrect*, and *undecidable* predictions as a function of the distance to the exit of the intersection for *consistent* behavior scenarios and for *inconsistent* behavior scenarios.

Consistent behavior scenarios: There are no *incorrect* predictions, which is not surprising since the behavior of the vehicles is never misleading in *consistent* behavior scenarios. The occasional *undecidable* predictions originate from situations where a vehicle is temporarily driving in between two lanes and the clue given by the turn signal (or the absence of turn signal) is not discriminating enough to identify which exit the driver is aiming at.

Inconsistent behavior scenarios: In these scenarios the turn signal information is misleading, which explains why the intended course is identified correctly much later than in the *consistent* scenarios. The graph can be broken into three phases:

- Phase 1 corresponds to the early approach phase, when the vehicle is still more than 10 m away from the entrance of the intersection. On average 55% of the predictions are *incorrect* and 45% are *undecidable*. The *incorrect* predictions correspond situations where there exists a consistent explanation for the presence of the turn signal. An example is the situation *b1* of Figure 4.27: the vehicle is still far away from the intersection, so the right turn signal is interpreted as an intention to change lanes. A human would make the same interpretation, but it turns out to be an *incorrect* prediction. Another example is a vehicle on the left lane with the left turn signal on that eventually goes straight. The *undecidable* predictions correspond to cases where the algorithm cannot find a consistent explanation for the vehicle's behavior. As a result the uncertainty on the intended course is high.
- Phase 2 corresponds to the end of the approach phase, when the vehicle is less than 10 m away from the entrance of the intersection. The trend in this phase is for the percentage of *undecidable* predictions to rise quickly while the percentage of *incorrect* predictions drops considerably. This inversion (*incorrect* predictions being replaced by *undecidable* predictions) is explained by the fact that as vehicles get closer to the intersection it is less and less likely that the drivers intend to change lanes. This is illustrated by situations *b1* and *b2* in Figure 4.27: the algorithm's prediction is *incorrect* in *b1* (Phase 1), and *undecidable* in *b2* (Phase 2) because there no longer exists a consistent explanation for the vehicle's behavior.
- Phase 3 corresponds to the execution of the maneuver inside the intersection. The position and heading of the vehicles provide new clues which compete with the misleading turn signal information. As the vehicles get closer to an exit lane, a larger number of maneuvers become unfeasible and therefore the algorithm is able to identify the driver's intention despite the misleading turn signal information. We observe a steady rise of the percentage of *correct* predictions until it reaches 100% approximately 10 m away from the exit of the intersection.

By incorporating information about the intersection layout, our model is able to 1) correctly estimate a driver's intended course in cases where there are enough clues to identify it and 2) keep the uncertainty on the prediction high when clues are lacking or inconsistent. The ability to assess that the situation is uncertain is crucial, since the level of confidence in the output can be used for making decisions at an applicative level.

4.5 Conclusions

4.5.1 Summary

The performance of the proposed risk assessment algorithm was evaluated based on the rate of false alarms, the rate of missed detections, and the collision prediction horizon. Tests were conducted both in simulation and in real conditions with passenger cars.

Our simulations assumed ideal perception and communication, and considered typical accident scenarios at a two-way-stop X-shaped intersection. An analysis of the collision prediction horizon led to the following main conclusions:

1. In our dataset, the proposed algorithm did not trigger any false alarm and was able to predict collisions at least 0.6 s before they occurred.
2. All the “late” detections (less than 1.5 s before the collision) correspond to left turn across path maneuvers.
3. Accidents caused by stop violations are detected on average 1 s earlier than the ones caused by priority violations.
4. There is a great variation in the efficiency of the different accident avoidance strategies which could be triggered after a *dangerous* situation is detected. The strategy should be selected depending on the situation (stop violation or priority violation).

Field trials were conducted using two vehicles equipped with off-the-shelf V2V wireless communication modems. Six different drivers took part in the experiments to recreate realistic *dangerous* and *non-dangerous* situations. The risk estimation algorithm proposed in the previous chapter was run online in one of the vehicles, and triggered a warning for the driver when it detected a *dangerous* situation. No false alarms were triggered in *non-dangerous* situations. In *dangerous* situations the warning was always triggered early enough that accidents were avoided by performing an emergency braking. The field trials proved that our approach can operate with success in real-life situations and trigger warnings in real time. They also showed the robustness of the algorithm, since the experiments were carried out with several drivers, a positioning system with a precision of $\sigma = 2$ m, and challenging wireless communication conditions.

The second half of this chapter focused on evaluating the importance of context in motion models. In Section 4.3 we showed that accounting for the dependencies between the vehicles’ maneuvers improves the estimation of the drivers’ intended course and improves risk assessment, while the collision

detection time is not affected. Section 4.4 evaluated the ability of our algorithm to estimate the intended course of drivers in challenging situations where the driver's behavior is inconsistent. The results showed that our model is able to balance the information provided by the turn signal state and by the layout of the intersection and to adapt the uncertainty of its output depending on the concordance of the different clues.

4.5.2 Discussion

The evaluation conducted in this thesis focused on two aspects:

1. The functional performance, safety performance, and robustness of the proposed risk assessment algorithm.
2. The benefits obtained by taking into account context in motion models.

Some other aspects were not investigated, such as the quantitative impact of communication limitations on the algorithm's performance. Such a study could be conducted in simulation, using a network simulator (e.g. ns-2 [114] or ns-3 [115]) to generate realistic communication errors. Another relevant study would be the impact of the driving style: our model was designed to adapt to different driving styles to a certain extent (see Section 3.4.5) and our algorithm was tested with several drivers, but this aspect was not investigated in detail. Extreme situations like drivers committing violations at very low or very high speeds would be interesting to test.

Further, the algorithm's performance could be evaluated in more complex scenarios. The evaluations presented in this chapter involved a maximum of two vehicles, but in theory the algorithm can be used with an arbitrary number of vehicles. Similarly, four different intersection layouts were tested in this work but the algorithm can in theory be applied to any intersection geometry and topology. In the future, experiments should be performed with more vehicles and more intersections to prove the generality of our approach.

No performance comparisons with state-of-the-art algorithms were presented in this work. In order to be compared, risk assessment algorithms must be tested on the same dataset. To the best of our knowledge there exists no publicly available dataset on which risk assessment algorithms have been tested and the results published. The non-availability of a dataset with real accidents can be explained by the difficulty to collect data, but simulation could be used as a mean to generate collision data which could be used by the ITS community to compare risk estimation algorithms. In the future we

would like to investigate this matter, as well as ways to standardize the performance evaluation of risk assessment algorithms.

Chapter 5

Conclusions

Contents

5.1 Summary	127
5.2 Conclusions	129
5.3 Perspectives	130

5.1 Summary

The topic addressed by this thesis is situation assessment and risk estimation for road traffic situations, with a focus on road intersections and connected vehicle safety applications.

The context of this research was described in **Chapter 1**. Using statistical studies on road accidents, we showed that intersections are the most dangerous areas of the road network and that wireless vehicular communications have the potential to improve safety. In this work, Vehicle-to-Vehicle communications (V2V) were used as a tool to extend the perception horizon of vehicles. We identified two main challenges for situation assessment and risk estimation at road intersections. Firstly, traffic situations at road intersections are very difficult to model. They are highly dynamic and involve complex interactions between the vehicles: drivers adapt their maneuvers according to the presence of other vehicles, the layout of the intersection, and the traffic rules. Secondly, there is uncertainty both in the available measurements and in the interpretation that should be made of a situation. Therefore, the challenge is to propose a comprehensive reasoning framework which reflects the complexity of

road intersection situations while keeping the computational complexity compatible with real-time constraints.

State-of-the art approaches to the risk estimation problem were reviewed in **Chapter 2**. The classic approach is based on vehicle motion models. Likely future trajectories are computed for every vehicle in the scene, then risk is calculated as the likelihood that these trajectories will collide. The main drawback is that this process is computationally expensive, since there are many possible trajectories for each vehicle in the scene. In order to keep the complexity manageable, it is common practice to simplify the problem either by ignoring uncertainties (in the data and / or in the motion model) or by assuming independence between the motion of the vehicles. The first solution leads to overconfident predictions, while the second one limits the ability to interpret of a vehicle's behavior.

A novel approach to risk assessment was proposed in **Chapter 3**. It consists of a general framework for reasoning about traffic situations and risk which explicitly models the intentions of drivers and what is expected of them according to the current context and traffic rules. The proposed approach is based on the comparison between *intentions* and *expectations*. This intuitive formulation of risk is very flexible in terms of applications and has the advantage that it does not require predicting the future trajectories of the vehicles. Inference on *intentions* and *expectations* is performed using a Dynamic Bayesian Network, taking into account uncertainties in the input data and uncertainties in the reasoning. This general framework was implemented for the specific case of road intersections. In the proposed motion model, the influence of the geometry and topology of the road intersection is taken into account, as well as the dependencies between the motion of the vehicles. What is expected of a driver at an intersection is derived from the estimated context (maneuvers performed by drivers and state of their vehicles), using the priority rules at the intersection and probabilistic gap acceptance models.

The proposed model was evaluated in **Chapter 4** using simulation and trials in real situations. The performance of the risk estimation was measured at an applicative level, by defining a threshold on the risk to separate *dangerous* and *non-dangerous* situations. Simulation was used to analyze statistically the collision prediction horizon for typical collision scenarios. The results showed the ability of the algorithm to detect dangerous situations before a collision occurs, and showed that some scenarios are more challenging than others (e.g. priority violations during a left turn across path maneuver are detected later). A preliminary study was conducted to assess the potential of different strategies for avoiding collisions once a dangerous situation has been detected. We found that strategies perform very differently depending on the configuration of the accident. An experimental setup was built to test the algorithm in real conditions with several drivers. Two passenger vehicles were equipped with

wireless communication modems to enable the sharing of information about their current states. The algorithm ran in real-time on one of the vehicles and warned the driver when dangerous situations arose. Despite localization errors and communication losses, the warnings were always triggered early enough that the upcoming collision could be avoided with an emergency braking. The importance of taking into account the layout of the intersection and the dependencies between the vehicles was evaluated using real data replayed offline. Two versions of our algorithm were compared in order to evaluate the consequences of the standard assumption of independence between the vehicles. The results confirmed our claim that accounting for the dependencies leads to a better estimation of the drivers' intentions and to a better sensitivity of risk assessment. We also showed that our algorithm can estimate the intended course of drivers in challenging situations where the driver's behavior is inconsistent, thanks to the information extracted from the digital map.

5.2 Conclusions

After conducting this research, our two main conclusions are the following:

- ***Context should and can be taken into account in algorithms addressing road safety.***

This statement can be broken into two parts, which are commented below:

- Context should be taken into account in algorithms addressing road safety

The advantages of taking into account the layout of the intersection and the dependencies between vehicles were highlighted using toy examples, and were evaluated quantitatively using real data. The results showed that incorporating information about the current local context allows us to interpret complex situations such as a vehicle yielding to another vehicle or vehicles exhibiting inconsistent behaviors.

- Context can be taken into account in algorithms addressing road safety

We showed that taking into account context is not incompatible with real-time constraints. To the best of our knowledge, no risk estimation algorithm has been proposed in the past which accounts for the mutual influences between the vehicles and has been shown to run in real-time. The proposed approach reaches this goal by computing the risk without the need to predict the future trajectories of vehicles.

- ***Risk is not only about trajectories intersecting.***

Throughout this thesis, risk estimation was presented as a tool to predict collisions. This is the most common interpretation of risk, however we believe that predicting collisions is only part of the road safety problem.

While classic approaches estimate the risk of a situation by predicting the future trajectories of vehicles and looking for collisions, our approach suggests that risk is about drivers behaving differently from what is expected. This formulation extends the concept of risk beyond collisions. For example, a vehicle proceeding in an intersection when only a short gap is available will not necessarily result in a collision, but most drivers will consider it to be dangerous since they expected the vehicle to wait for a longer gap. If the classic definition of risk is used, a warning in this situation will be considered as a false alarm.

This thesis focused on the prediction of collisions, but we think that our formulation of risk opens new opportunities for safety-related ADAS applications.

5.3 Perspectives

Future work will include additional performance evaluation, as mentioned in the previous chapter. Further, the contributions presented in this thesis can be extended in the following directions:

- The first one addresses ***context-aware motion models for vehicles.***

Motion models are at the core of most risk assessment algorithms, however they often fail to account for the context in which the vehicles evolve. The motion model proposed in this thesis takes a step toward context-awareness by taking into account the influence of the traffic rules, the interactions with the other vehicles, and the layout of the road network on the motion of vehicles. Additional information could be incorporated into the motion model to improve the ability to infer driver intentions.

- Drivers' actions such as pressure on pedals, steering, gear shifting are strong indicators of their intentions. Monitoring these actions would allow us to anticipate the motion of a vehicle, and therefore to estimate the intentions of drivers earlier. For example, a driver preparing to enter an intersection will turn the steering wheel before the vehicle's orientation reveals his intention.

- Further, the parameters of the motion model could be made context-dependent and driver-dependent. For example the windshield wipers running indicate that it may be raining, and in this case the speed profiles could be adapted to reflect that the stopping distance of vehicles will be longer than in dry weather. Another example is light sensors, which can identify day and night context.

The speed profiles and gap acceptance models used in Chapter 3 correspond to average driver behavior, and would be more representative of a driver's actual behavior if they were to be learned for each driver.

- Another source of information which could be considered is Vehicle-to-Infrastructure (V2I) communications. One of the purpose of roadside wireless communication units is to provide vehicles with local contextual information such as additional information about the layout, the state of the traffic lights, temporary modifications of the traffic rules, etc. [116]. Incorporating this information into the reasoning would improve situation understanding.
 - Finally, sensor fusion should be considered in order to have a comprehensive understanding of the situation. The work presented here builds a representation of the environment using proprioceptive sensors, maps, and V2V communication, but does not integrate information about potential non-communicating objects in the scene. Pedestrians, non-communicating vehicles, obstacles on the road, etc. can be detected using exteroceptive sensors (e.g. cameras, lasers, radars). This information and its associated uncertainty can be combined with a vehicle's own representation of the scene to make it more reliable (through the redundancy of information) and more complete (through additional information about non-communicating entities).
- The second one addresses *the application of the proposed risk estimation approach to other scenarios*.

In Section 3.1.2 we proposed that dangerous road traffic situations shall be detected based on the comparison between what drivers intend to do and what is expected of them. The remaining of the thesis focused on unsignalized road intersections, and we showed that the proposed approach is able to detect dangerous situations caused by the non-respect of traffic rules. We believe that the concept of comparing *intention* and *expectation* is relevant for a number of other applications.

- The proposed motion model could be extended to other types of intersections, such as roundabouts or intersections controlled by traffic lights. For each new type of intersection,

a new “expected longitudinal motion” model would have to be defined to match the traffic rules. For roundabouts, that model would take the form of a probabilistic gap acceptance model, similarly to the intersections addressed in this thesis. For intersections controlled by traffic lights, the probability that a vehicle should stop would be a function of the time needed by the vehicle to reach the intersection and of the dynamic state of the traffic light. The design of these new models could be based on previous work in the literature such as [117] for roundabouts and [118, 119] for traffic lights.

- The idea of comparing *intention* and *expectation* can also be used in traffic situations which are not related to intersections. Examples were given in Section 3.6.2 for highway traffic situations. Rear-end collisions and hazardous lane changes on a highway could be predicted by monitoring the driver’s intention to change lanes or to follow the vehicle in front. If a driver gets too close to the vehicle in front and does not seem to intend to change lanes or to adapt its speed, there is a risk of a rear-end collision. If a driver intends to change lanes while another vehicle is approaching, there is a risk of sideways collision.
- The third addresses ***the design of safety applications for Advanced Driver Assistance Systems (ADAS) and Autonomous Driving***.

Our framework provides tools for reasoning about situations and risk, but risk estimation is only part of the solution to make roads safer. Following the detection of dangerous situations, decisions have to be made to enhance the safety of drivers. This is not an easy task, as it raises a number of questions in terms of user acceptance, liability, and ethics.

- The approach formulated in this thesis is relevant to both ADAS and Autonomous Driving applications. The motion model can be used to estimate the risk of a situation and to infer relevant information about potential collisions such as the vehicles involved, the configuration (e.g. frontal or side collisions), the imminence (time-to-collision), and the severity (depending on the speed of the vehicles). These indicators are valuable for safety applications, as they help to make decisions about which action(s) to take to avoid or mitigate the upcoming collision.

For example they can be used by an ADAS to decide if the driver should be warned, when, and how. An ADAS could also use these indicators to decide that a direct action on the brakes is necessary and take control over (part of) the vehicle’s commands, as it is the case in Automated Braking Systems.

In the case of Autonomous Driving applications, the vehicle can adapt its maneuver to

reduce the risk, for example by slowing down or by taking a less risky course.

- If the driver is left in control of the vehicle, the difficulty is to convey information from the computer to the human. The computer is in possession of a large amount of information regarding the specifics of the danger. It is also capable of computing the risk of potential avoidance maneuvers. However the driver can process only a limited amount of information, and time is critical. The system can either warn the driver (providing information about the situation to the driver, who then decides what to do) or send the driver an injunction (e.g. “apply emergency braking now”).

In the first case (warning) the information has to be cut down and simplified before it is displayed (no probabilities, simple graphics as in Figure 4.10).

In the second case (injunction) the system itself decides of the appropriate action, which could save some time. However there are some liability issues, since the final situation (collision avoided or not) will be partially imputable to the system.

- If the vehicle is partially or totally computer-controlled, similar liability issues exist.

Ethics is a major challenge, as the computer would have to make decisions related to life and death of humans in critical situations. This issue is linked to the problem of driver acceptance: part of the reluctance that some drivers have to leave the control of their car to a computer is caused by the lack of trust in the machine’s capabilities.

On a technical level, the advantages of letting the computer control part of the vehicle in specific cases are numerous. As was mentioned in the previous paragraph, the computer has access to a lot of information which could be used to handle dangerous situations in an efficient manner. But beyond emergency situations, we believe that much of the power of computer control in vehicles resides in preventive measures. In particular, it could address situations where the risk is rising but is not high enough to warn the driver or apply emergency braking. In this case, the computer could make the vehicle slow down. Even small changes in speed have a significant impact in the severity of an accident, and they are likely to be easily accepted by drivers.

Bibliography

- [1] TRACE project, “Accident causation and pre-accidental driving situations - In-depth accident causation analysis,” Deliverable D2.2, 2008, available at <http://www.trace-project.org/publication/archives/trace-wp2-d2-2-v2.pdf>. [Accessed 12 Aug. 2012].
- [2] PREVENT project - INTERSAFE subproject, “Project evaluation and effectiveness of the intersection safety system,” Final report (Deliverable D40.75), 2007, available at http://prevent.ertico.webhouse.net/en/prevent_subprojects/intersection_safety/intersafe. [Accessed 12 Aug. 2012].
- [3] SUNflower project, “SUNflower: a comparative study of the development of road safety in Sweden, the United Kingdom and the Netherlands,” Final report, 2002, available at <http://www.swov.nl/rapport/sunflower/sunflower.pdf>. [Accessed 12 Aug. 2012].
- [4] J. Ibañez-Guzmán, C. Laugier, J.-D. Yoder, and S. Thrun, *Handbook of intelligent vehicles*. Springer, 2012, ch. Autonomous driving: context and state-of-the-art, pp. 1271–1310.
- [5] DRIVE-C2X project, “Use cases,” available at <http://drive-c2x.eu/use-cases>. [Accessed 12 Aug. 2012].
- [6] CICAS project, <http://www.its.dot.gov/cicas/index.htm>. [Accessed 12 Aug. 2012].
- [7] Visteon Corporation, California PATH, and Industrial Technology Research Institute, “Enabling accelerated installation of aftermarket on-board equipment for connected vehicles - Task 1: State of the industry report,” 2011, available at http://cts.virginia.edu/PFS_AFTERO2_Task1.pdf. [Accessed 12 Aug. 2012].
- [8] G. Segarra, “Road co-operative systems - Societal and business values,” in *Proc. International Conference on Intelligent Transport Systems Telecommunications*, 2009, pp. 610–615.

- [9] D. Hall and J. Llinas, "An introduction to multisensor data fusion," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 6–23, 1997.
- [10] R. Miller and Q. Huang, "An adaptive peer-to-peer collision warning system," in *Proc. IEEE Vehicular Technology Conference*, 2002, pp. 317–321.
- [11] S. Ammoun and F. Nashashibi, "Real time trajectory prediction for collision risk estimation between vehicles," in *Proc. IEEE Intelligent Computer Communication and Processing*, 2009, pp. 417–422.
- [12] A. Broadhurst, S. Baker, and T. Kanade, "Monte Carlo road safety reasoning," in *Proc. IEEE Intelligent Vehicles Symposium*, 2005, pp. 319–324.
- [13] S. Atev, G. Miller, and N. Papanikolopoulos, "Clustering of vehicle trajectories," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 647–657, 2010.
- [14] G. Agamennoni, J. I. Nieto, and E. M. Nebot, "A Bayesian approach for driving behavior inference," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 595–600.
- [15] European Telecommunication Standards Institute, "News release," 2008, available at http://www.etsi.org/WebSite/NewsandEvents/2008_09_Harmonizedstandards_ITS.aspx. [Accessed 12 Aug. 2012].
- [16] Federal Communications Commission, "News release," 1999, available at http://transition.fcc.gov/Bureaus/Engineering_Technology/News_Releases/1999/nret9006.html. [Accessed 12 Aug. 2012].
- [17] World Health Organization, "Global status report on road safety," 2009, available at http://www.who.int/violence_injury_prevention/road_safety_status/2009/en. [Accessed 12 Aug. 2012].
- [18] European Road Safety Observatory, "Traffic safety basic facts 2010 - Junctions," 2010, available at http://ec.europa.eu/transport/road_safety/pdf/statistics/dacota/bfs2010_dacota-trl-1-3-junctions.pdf. [Accessed 12 Aug. 2012].
- [19] TRACE project, <http://www.trace-project.org>. [Accessed 12 Aug. 2012].
- [20] PREVENT project - INTERSAFE subproject, http://prevent.ertico.webhouse.net/en/prevent_subprojects/intersection_safety/intersafe. [Accessed 12 Aug. 2012].

BIBLIOGRAPHY

- [21] U.S. Department of Transportation - Federal Highway Administration, "Intersection safety," available at <http://safety.fhwa.dot.gov/intersection>. [Accessed 12 Aug. 2012].
- [22] U.S. Department of Transportation - National Highway Traffic Safety Administration, "Crash factors in intersection-related crashes: an on-scene perspective," Report DOT-HS-811-366, 2010, available at www-nrd.nhtsa.dot.gov/Pubs/811366.pdf. [Accessed 12 Aug. 2012].
- [23] —, "The economic impact of motor vehicle crashes," Report DOT HS-809-446, 2002, available at <http://www.nhtsa.gov/Driving+Safety/The+Economic+Impact+of+Motor+Vehicle+Crashes+2000>. [Accessed 12 Aug. 2012].
- [24] TRACE project, "Review of crash effectiveness of Intelligent Transport Systems," Deliverable D4.1.1 - D6.2, 2007, available at <http://www.trace-project.org/publication/archives/trace-wp4-wp6-d4-1-1-d6-2.pdf>. [Accessed 12 Aug. 2012].
- [25] U.S. Department of Transportation - National Highway Traffic Safety Administration, "Frequency of target crashes for IntelliDrive safety systems," Report DOT-HS-811-381, 2010, available at <http://www.nhtsa.gov/DOT/NHTSA/NVS/Crashes/2010/811381.pdf>. [Accessed 12 Aug. 2012].
- [26] G. Karagiannis, O. Altintas, E. Ekici, G. Heijenk, B. Jarupan, K. Lin, and T. Weil, "Vehicular networking: a survey and tutorial on requirements, architectures, challenges, standards and solutions," *IEEE Communications Surveys Tutorials*, vol. 13, no. 4, pp. 584–616, 2011.
- [27] H. Onishi and F. Mlinarsky, "Wireless technology assessment for automotive applications," in *Proc. ITS World Congress*, 2012.
- [28] DRIVE-C2X project, <http://www.drive-c2x.eu/project>. [Accessed 12 Aug. 2012].
- [29] Safety Pilot project, http://www.its.dot.gov/research/safety_pilot_overview.htm. [Accessed 12 Aug. 2012].
- [30] T. Inagaki, "Technological and legal considerations for the design of interactions between human driver and advanced driver assistance systems," in *Proc. NeTWork Workshop: Control and Accountability in Highly Automated Systems*, 2011.
- [31] R. Meireles, M. Boban, P. Steenkiste, O. Tonguz, and J. Barros, "Experimental study on the impact of vehicular obstructions in VANETs," in *Proc. IEEE Vehicular Networking Conference*, 2010, pp. 338–345.

-
- [32] F. Martelli, M. E. Renda, and P. Santi, "Measuring IEEE 802.11p performance for active safety applications in cooperative vehicular systems," in *Proc. IEEE Vehicular Networking Conference*, 2011, pp. 1–5.
- [33] S. Eichler, "Performance evaluation of the IEEE 802.11p WAVE communication standard," in *Proc. IEEE Vehicular Networking Conference*, 2007, pp. 2199–2203.
- [34] L. Le, A. Festag, R. Baldessari, and W. Zhang, "Vehicular wireless short-range communication for improving intersection safety," *IEEE Communications Magazine*, vol. 47, no. 11, pp. 104–110, November 2009.
- [35] Q. Wu, J. Domingo-Ferrer, and U. Gonzalez-Nicolas, "Balanced trustworthiness, safety, and privacy in vehicle-to-vehicle communications," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 2, pp. 559–573, 2010.
- [36] ITS America, "The connected vehicle - Next generation ITS," 2012, available at <http://www.itsa.org/industryforums/connectedvehicle>. [Accessed 12 Aug. 2012].
- [37] U.S. Department of Transportation - Research and Innovative Technology Administration, "Principles for a connected vehicle environment," 2012, available at http://www.its.dot.gov/connected_vehicle/principles_connectedvehicle_environment.htm. [Accessed 12 Aug. 2012].
- [38] F. Colas, J. Diard, and P. Bessière, "Common Bayesian models for common cognitive issues," *Acta biotheoretica*, vol. 58, no. 2-3, pp. 191–216, 2010.
- [39] J. Ibañez-Guzmán, S. Lefèvre, A. Mokkaedem, and S. Rodhain, "Vehicle to vehicle communications applied to road intersection safety, field results," in *Proc. IEEE International Conference on Intelligent Transportation Systems*, 2010, pp. 192–197.
- [40] R. Rajamani, *Vehicle dynamics and control*. Birkhauser, 2006.
- [41] M. Brännström, E. Coelingh, and J. Sjöberg, "Model-based threat assessment for avoiding arbitrary vehicle collisions," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 658–669, 2010.
- [42] C.-F. Lin, A. G. Ulsoy, and D. J. LeBlanc, "Vehicle dynamics and external disturbance estimation for vehicle path prediction," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 3, pp. 508–518, 2000.

- [43] J. Huang and H.-S. Tan, "Vehicle future trajectory prediction with a DGPS/INS-based positioning system," in *Proc. American Control Conference*, 2006, pp. 5831–5836.
- [44] R. Pepy, A. Lambert, and H. Mounier, "Reducing navigation errors by planning with realistic vehicle model," in *Proc. IEEE Intelligent Vehicles Symposium*, 2006, pp. 300–307.
- [45] A. Eidehall and L. Petersson, "Statistical threat assessment for general road scenes using Monte Carlo sampling," *IEEE Transactions on Intelligent Transportation Systems*, vol. 9, no. 1, pp. 137–147, 2008.
- [46] N. Kaempchen, B. Schiele, and K. Dietmayer, "Situation assessment of an autonomous emergency brake for arbitrary vehicle-to-vehicle collision scenarios," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 4, pp. 678–687, 2009.
- [47] R. Schubert, E. Richter, and G. Wanielik, "Comparison and evaluation of advanced motion models for vehicle tracking," in *Proc. International Conference on Information Fusion*, 2008, pp. 1–6.
- [48] N. Kaempchen, K. Weiss, M. Schaefer, and K. C. J. Dietmayer, "IMM object tracking for high dynamic driving maneuvers," in *Proc. IEEE Intelligent Vehicles Symposium*, 2004, pp. 825–830.
- [49] J. Hillenbrand, A. M. Spieker, and K. Kroschel, "A multilevel collision mitigation approach: situation assessment, decision making, and performance tradeoffs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, pp. 528–540, 2006.
- [50] A. Polychronopoulos, M. Tsogas, A. Amditis, and L. Andreone, "Sensor fusion for predicting vehicles' path for collision avoidance systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 3, pp. 549–562, 2007.
- [51] A. Barth and U. Franke, "Where will the oncoming vehicle be the next second?" in *Proc. IEEE Intelligent Vehicles Symposium*, 2008, pp. 1068–1073.
- [52] H.-S. Tan and J. Huang, "DGPS-based vehicle-to-vehicle cooperative collision warning: engineering feasibility viewpoints," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, pp. 415–428, 2006.
- [53] T. Batz, K. Watson, and J. Beyerer, "Recognition of dangerous situations within a cooperative group of vehicles," in *Proc. IEEE Intelligent Vehicles Symposium*, 2009, pp. 907–912.

- [54] P. Lytrivis, G. Thomaidis, and A. Amditis, "Cooperative path prediction in vehicular environments," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2008, pp. 803–808.
- [55] B. Mourllion, D. Gruyer, A. Lambert, and S. Glaser, "Kalman filters predictive steps comparison for vehicle localization," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 565–571.
- [56] S. Thrun, W. B. Fox, and Dieter, *Probabilistic robotics*. MIT Press, 2005.
- [57] K. P. Murphy, "Dynamic Bayesian networks: representation, inference and learning," Ph.D. dissertation, University of California, Berkeley, USA, 2002.
- [58] H. Veeraraghavan, N. Papanikolopoulos, and P. Schrater, "Deterministic sampling-based switching Kalman filtering for vehicle tracking," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2006, pp. 1340–1345.
- [59] H. Dyckmanns, R. Matthaei, M. Maurer, B. Lichte, J. Effertz, and D. Stuker, "Object tracking in urban intersections based on active use of a priori knowledge: active interacting multi model filter," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 625–630.
- [60] M. Althoff and A. Mergel, "Comparison of Markov chain abstraction and Monte Carlo simulation for the safety assessment of autonomous cars," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1237–1247, 2011.
- [61] Definition from Wordreference,
available at <http://www.wordreference.com/definition/maneuver>. [Accessed 12 Aug. 2012].
- [62] G. S. Aoude, V. R. Desaraju, L. H. Stephens, and J. P. How, "Behavior classification algorithms at intersections and validation using naturalistic data," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 601–606.
- [63] C. Tay, "Analysis of dynamic scenes: application to driving assistance," Ph.D. dissertation, Institut National Polytechnique de Grenoble, France, 2009.
- [64] T. Gindele, S. Brechtel, and R. Dillmann, "A probabilistic model for estimating driver behaviors and vehicle trajectories in traffic environments," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2010, pp. 1625–1631.
- [65] I. Dagli and D. Reichardt, "Motivation-based approach to behavior prediction," in *Proc. IEEE Intelligent Vehicles Symposium*, 2002, pp. 227–233.

- [66] M. Garcia-Ortiz, J. Fritsch, F. Kummert, and A. Gepperth, "Behavior prediction at multiple time-scales in inner-city scenarios," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 1068–1073.
- [67] B. Morris and M. Trivedi, "Learning trajectory patterns by clustering: experimental studies and comparative evaluation," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 312–319.
- [68] D. Buzan, S. Sclaroff, and G. Kollios, "Extraction and clustering of motion trajectories in video," in *Proc. International Conference on Pattern Recognition*, vol. 2, 2004, pp. 521–524.
- [69] W. Hu, X. Xiao, Z. Fu, D. Xie, T. Tan, and S. Maybank, "A system for learning statistical motion patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 9, pp. 1450–1464, 2006.
- [70] D. Vasquez and T. Fraichard, "Motion prediction for moving objects: a statistical approach," in *Proc. IEEE International Conference on Robotics and Automation*, vol. 4, 2004, pp. 3931–3936.
- [71] C. Hermes, C. Wohler, K. Schenk, and F. Kummert, "Long-term vehicle motion prediction," in *Proc. IEEE Intelligent Vehicles Symposium*, 2009, pp. 652–657.
- [72] J. M. Joseph, F. Doshi-Velez, and N. Roy, "A Bayesian nonparametric approach to modeling mobility patterns," in *Proc. AAAI Conference on Artificial Intelligence*, 2010.
- [73] G. Aoude, J. Joseph, N. Roy, and J. How, "Mobile agent trajectory prediction using Bayesian nonparametric reachability trees," in *Proc. of AIAA Infotech@Aerospace*, 2011, pp. 1587–1593.
- [74] D. Greene, J. Liu, J. Reich, Y. Hirokawa, A. Shinagawa, H. Ito, and T. Mikami, "An efficient computational architecture for a collision early-warning system for vehicles, pedestrians, and bicyclists," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 4, pp. 1–12, 2010.
- [75] G. S. Aoude, B. D. Luders, K. K. H. Lee, D. S. Levine, and J. P. How, "Threat assessment design for driver assistance system at intersections," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2010, pp. 25–30.
- [76] H. Berndt, J. Emmert, and K. Dietmayer, "Continuous driver intention recognition with hidden Markov models," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2008, pp. 1189–1194.

- [77] A. Tamke, T. Dang, and G. Breuel, "A flexible method for criticality assessment in driver assistance systems," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 697–702.
- [78] C. Laugier, I. Paromtchik, M. Perrollaz, M. Yong, J. Yoder, C. Tay, K. Mekhnacha, and A. Negre, "Probabilistic analysis of dynamic scenes and collision risks assessment to improve driving safety," *IEEE Intelligent Transportation Systems Magazine*, vol. 3, no. 4, pp. 4–19, 2011.
- [79] M. Althoff, O. Stursberg, and M. Buss, "Model-based probabilistic collision detection in autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 2, pp. 299–310, 2009.
- [80] E. Kaefer, C. Hermes, C. Woehler, H. Ritter, and F. Kummert, "Recognition of situation classes at road intersections," in *Proc. IEEE International Conference on Robotics and Automation*, 2010, pp. 3960–3965.
- [81] M. Brand, N. Oliver, and A. Pentland, "Coupled hidden Markov models for complex action recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1997, pp. 994–999.
- [82] N. Oliver and A. P. Pentland, "Graphical models for driver behavior recognition in a SmartCar," in *Proc. IEEE Intelligent Vehicles Symposium*, 2000, pp. 7–12.
- [83] G. Agamennoni, "Towards probabilistic reasoning in multi-vehicle scenarios for intelligent transportation systems," Ph.D. dissertation, The University of Sydney, Australia, 2012.
- [84] G. Agamennoni, J. I. Nieto, and E. M. Nebot, "Estimation of multivehicle dynamics by considering contextual information," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 855–870, 2012.
- [85] A. Berthelot, A. Tamke, T. Dang, and G. Breuel, "Handling uncertainties in criticality assessment," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 571–576.
- [86] C.-Y. Chan, "Defining safety performance measures of driver assistance systems for intersection left-turn conflicts," in *Proc. IEEE Intelligent Vehicles Symposium*, 2006, pp. 25–30.
- [87] R. Labayrade, C. Royere, and D. Aubert, "A collision mitigation system using laser scanner and stereovision fusion and its assessment," in *Proc. Intelligent Vehicles Symposium*, 2005, pp. 441–446.

- [88] Y. Liu, O. Ozguner, and E. Ekici, "Performance evaluation of intersection warning system using a vehicle traffic and wireless simulator," in *Proc. IEEE Intelligent Vehicles Symposium*, 2005, pp. 171–176.
- [89] S. Worrall, D. Orchansky, F. Masson, and E. Nebot, "Improving vehicle safety using context based detection of risk," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2010, pp. 379–385.
- [90] A.-C. Boury-Brisset and N. Tourigny, "Knowledge capitalisation through case bases and knowledge engineering for road safety analysis," *Knowledge-Based Systems*, vol. 13, no. 5, pp. 297–305, 2000.
- [91] A. Chinaea and M. Parent, "Risk assessment algorithms based on recursive neural networks," in *Proc. International Joint Conference on Neural Networks*, 2007, pp. 1434–1440.
- [92] F. D. Salim, S. W. Loke, A. Rakotonirainy, B. Srinivasan, and S. Krishnaswamy, "Collision pattern modeling and real-time collision detection at road intersections," in *Proc. IEEE Intelligent Transportation Systems Conference*, 2007, pp. 161–166.
- [93] O. Lebeltel, P. Bessiere, J. Diard, and E. Mazer, "Bayesian robot programming," *Autonomous Robots*, vol. 16, no. 1, pp. 49–79, 2004.
- [94] D. Meyer-Delius, C. Plagemann, and W. Burgard, "Probabilistic situation recognition for vehicular traffic scenarios," in *Proc. IEEE International Conference on Robotics and Automation*, 2009, pp. 4161–4166.
- [95] Definition from The Free Dictionary, available at <http://encyclopedia2.thefreedictionary.com/Road+Traffic+Rules>. [Accessed 12 Aug. 2012].
- [96] L. Fletcher, S. Teller, E. Olson, D. Moore, Y. Kuwata, J. How, J. Leonard, I. Miller, M. Campbell, D. Huttenlocher, A. Nathan, and F. Kline, "The MIT-Cornell collision and why it happened," *Journal of Field Robotics*, vol. 25, pp. 775–807, 2008.
- [97] A. Spek, P. Wieringa, and W. Janssen, "Intersection approach speed and accident probability," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 9, no. 2, pp. 155–171, 2006.

- [98] D. R. Ragland, A. Sofia, S. E. Shladover, and J. A. Misener, "Gap acceptance for vehicles turning left across on-coming traffic: implications for intersection decision support design," in *Proc. Transportation Research Board Annual Meeting*, 2006.
- [99] H. Berndt, S. Wender, and K. Dietmayer, "Driver braking behavior during intersection approaches and implications for warning strategies for driver assistant systems," in *Proc. IEEE Intelligent Vehicles Symposium*, 2007, pp. 245–251.
- [100] M. Beal, "Variational algorithms for approximate Bayesian inference," Ph.D. dissertation, University College London, England, 2003.
- [101] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [102] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [103] M. Liebner, M. Baumann, F. Klanner, and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model," in *Proc. IEEE Intelligent Vehicles Symposium*, 2012, pp. 1162–1167.
- [104] TASS - PreScan, <http://www.tass-safe.com/prescan>. [Accessed 12 Aug. 2012].
- [105] PREVENT project - INTERSAFE subproject, "Requirements for intersection safety applications," Deliverable D40.4, 2005, available at http://prevent.ertico.webhouse.net/en/prevent_subprojects/intersection_safety/intersafe. [Accessed 12 Aug. 2012].
- [106] RTMaps, <http://www.intempora.com>. [Accessed 12 Aug. 2012].
- [107] S. Kaul, K. Ramachandran, P. Shankar, S. Oh, M. Gruteser, I. Seskar, and T. Nadeem, "Effect of antenna placement and diversity on vehicular network communications," in *Proc. IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, 2007, pp. 112–121.
- [108] Cohda Wireless, "DSRC field trials," Whitepaper, 2009, available at <http://www.cohdawireless.com/technology/whitepapers.html>. [Accessed 12 Aug. 2012].

- [109] A. Saltelli, "Sensitivity analysis for importance assessment," *Risk Analysis*, vol. 22, no. 3, pp. 579–590, 2002.
- [110] S. Lefèvre, J. Ibañez-Guzmán, and C. Laugier, "Context-based prediction of vehicle destination at road intersections," in *Proc. IEEE Symposium on Computational Intelligence in Vehicles and Transportation Systems*, 2011, pp. 67–72.
- [111] S. Lefèvre, C. Laugier, and J. Ibañez-Guzmán, "Exploiting map information for driver intention estimation at road intersections," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 583–588.
- [112] 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," *Journal of Field Robotics*, vol. 25, no. 9, pp. 569–597, 2008.
- [113] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Kolter, D. Langer, O. Pink, V. Pratt, M. Sokolsky, G. Stanek, D. Stavens, A. Teichman, M. Werling, and S. Thrun, "Towards fully autonomous driving: systems and algorithms," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 163–168.
- [114] ns-2 network simulator, http://nslam.isi.edu/nslam/index.php/User_Information. [Accessed 12 Aug. 2012].
- [115] ns-3 network simulator, <http://www.nslam.org>. [Accessed 12 Aug. 2012].
- [116] M. Maile, F. Ahmed-Zaid, S. Bai, L. Caminiti, P. Mudalige, M. Peredo, and Z. Popovic, "Objective testing of a cooperative intersection collision avoidance system for traffic signal and stop sign violation," in *Proc. IEEE Intelligent Vehicles Symposium*, 2011, pp. 130–137.
- [117] M. Muffert, T. Milbich, D. Pfeiffer, and U. Franke, "May I enter the roundabout? A time-to-contact computation based on stereo-vision," in *Proc. IEEE Intelligent Vehicles Symposium*, 2012, pp. 565–570.
- [118] C.-Y. Chan, "Characterization of driving behaviors based on field observation of intersection left-turn across-path scenarios," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 3, pp. 322–331, 2006.

- [119] N. Elmitiny, X. Yan, E. Radwan, C. Russo, and D. Nashar, "Classification analysis of driver's stop/go decision and red-light running violation," *Accident Analysis & Prevention*, vol. 42, no. 1, pp. 101–111, 2010.
- [120] Status of DSRC standards, available at <http://www.standards.its.dot.gov/news.asp#jun12>.
- [121] ETSI TC ITS WG1, "Vehicular Communications; Basic Set of Applications," Document no. EN 302 637-2, 2012, available at <http://www.etsi.org/WebSite/Technologies/CooperativeITS.aspx>. [Accessed 12 Aug. 2012].
- [122] SAE International, "Dedicated Short Range Communications (DSRC) message set dictionary," available at <http://www.sae.org/standardsdev/dsrc>. [Accessed 12 Aug. 2012].
- [123] DARPA Urban Challenge, "Route Network Definition File (RNDF) and Mission Data File (MDF) formats," available at http://archive.darpa.mil/grandchallenge/docs/RNDF_MDF_Formats_031407.pdf. [Accessed 12 Aug. 2012].
- [124] OpenStreetMap project, <http://www.openstreetmap.org>. [Accessed 12 Aug. 2012].
- [125] Navteq, <http://www.navteq.com>. [Accessed 12 Aug. 2012].
- [126] TomTom, http://www.tomtom.com/en_gb/maps. [Accessed 12 Aug. 2012].
- [127] G. Agamennoni, J. I. Nieto, and E. M. Nebot, "Robust and accurate road map inference," in *Proc. IEEE International Conference on Robotics and Automation*, 2010, pp. 3946–3953.
- [128] —, "Robust inference of principal road paths for intelligent transportation systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 1, pp. 298–308, 2011.

Appendix A

DSRC standards

5.9 GHz DSRC is a set of protocols and standards dedicated to short and medium-range wireless communications for ITS. Standardization work started in 1999 when the U.S. Federal Communication Commission allocated part of the 5.9 GHz band to wireless vehicular communications. Subsequently the ASTM and the IEEE jointly developed the IEEE 802.11p standard, also known as the Wireless Access to Vehicular Environment (WAVE) standard. This standard is based on the IEEE 802.11a (WiFi) standard, but was designed to operate with mobile nodes.

Standards for the other layers of the OSI model are being developed by a number of standards organizations, which are listed in Figure A.1. Some standards have already been finalized while others are in the development or approval process. A list of these standards and their development status is available on the U.S. Department of Transportation's website [120].




Figure A.1. Organizations involved in the development of standards for DSRC.

Communication channels: DSCR channels are divided into two categories: a control channel and several service channels. The control channel is reserved for broadcasting periodic messages and for coordinating the communications which take place on the service channels. The service channels are used for occasional exchanges of information. The frequencies allocated for DSRC applications in Europe and in the U.S. are shown in Figure A.2, along with the envisioned channel assignments.




Figure A.2. Bandwidth allocated for DSRC applications in Europe and in the U.S. [15, 16].

Messages specifications: Two types of messages are under standardization.

- Periodic messages: they are short messages which are periodically broadcasted in the control channel by all ITS stations located in an ad-hoc network, at a frequency varying between 2 Hz and 10 Hz depending on the time criticality of locally managed applications. These messages contain key information such as the station type, 3D position, heading, velocity, current dynamic and static states, capabilities.
 - In Europe, the standardized periodic message is the Cooperative Awareness Message (CAM) [121].
 - The equivalent message in the U.S. is the Basic Safety Message (BSM) [122].
- Event-driven messages: they are sent only when specific safety-related events are detected, e.g. an immobilized vehicle on the road, bad weather conditions, road / traffic problems, drivers with a dangerous behavior. The messages are broadcasted at a frequency varying between 2 Hz and 10 Hz until the event that caused the signal to be triggered disappears. Messages can be relayed by other ITS stations and that way be propagated to some predefined geographical areas.
 - In Europe, the standardized event-driven message is the Decentralized Notification Message (DENM) [121].
 - In the U.S., the equivalent messages have different names depending on the event type (ex: Emergency Vehicle Alert (EVA), Road Side Alert (RSA)) [122].

Appendix B

The importance of context

The motion of vehicles on the road is strongly constrained by the local context, i.e. by the shape of the road layout and by the presence of other vehicles. The latter is particularly true at road intersections, where priority rules force vehicles to take into account the maneuvers performed by the other vehicles. Therefore, a motion model representing vehicles as entities evolving independently from each other in an unconstrained environment will be valid for short-term motion only. ***For long-term motion prediction it is necessary to take into account the local context of the road network, i.e. to take into account the constraints imposed by the layout of the road network and by the traffic rules on the collective motion of vehicles in the scene.*** This statement is justified in the sections below, using example real-life scenarios. Section B.1 analyzes the necessity to take into account the layout of the intersection. Section B.2 analyzes the necessity to take into account the interactions between vehicles.

B.1 Intersection layout

For a vehicle approaching a road intersection, it is common practice to interpret the presence of a turn signal as an intention to make a turn [77]. In reality turn signals cannot be interpreted that directly, as illustrated in Figure B.1.



Figure B.1. Example situation where it is necessary to take into account the layout of the intersection to interpret correctly the yellow vehicle's behavior.

In this image, the situation is the following: the yellow vehicle is approaching the intersection with the left turn signal on. Without any information about the intersection layout, it will be (incorrectly) assumed that the driver intends to turn left. Taking into account the geometrical and topological characteristics of the intersection allows us to understand that the actual intention of the driver is to go straight. Indeed, left turns are allowed only from the leftmost lane but this lane is too far away to be reached by the yellow car, and the only two reachable lanes located on the left of the car are reserved for vehicles going straight.

This example highlights the importance of taking into account the geometrical and topological characteristics of the road intersection when interpreting turn signals and estimating driver intention.

B.2 Interactions between vehicles

As mentioned in the state of the art (Chapter 2), most of the existing motion models assume independence between the vehicles. This simplification is generally motivated by the complexity reduction that it provides. However, ignoring the mutual influences between vehicles in the motion model leads to limitations for situation assessment and risk assessment. These consequences are analyzed below,

based on simple traffic scenarios at typical road intersections.

B.2.1 Intuition

When driving, one has to analyze other traffic participants' behavior and infer their intentions in order to adapt one's behavior to the current situation and to the traffic rules. Humans performing this task will assume that the other drivers generally respect the traffic rules. In other words, *when the motion of a vehicle in the scene is compatible with several maneuvers, the maneuvers which match the traffic rules will be considered by a human to be more likely than the other ones*. This assumption leads to a better understanding of situations.

This statement is illustrated in Figure B.2.



Figure B.2. Example situation where it is necessary to take into account the dependencies between the vehicles to interpret the black car's behavior.

The situation involves two cars at a two-way stop intersection: the white car is approaching on the main road and driving at constant speed, while the black car is waiting at the stop line. In order to interpret correctly the black car's behavior, it is necessary to take into account both the layout of the intersection and the dependencies between the two cars:

- Algorithms which do not use layout information generally assume that slowing down when approaching an intersection is an indication that the driver intends to make a turn. An algorithm which takes into account the layout will be able to understand that in this case the black car is slowing down because of the stop, and that this behavior should not be interpreted as an intention to make a turn.

- When the black car remains stationary after reaching the stop line, an algorithm which takes into account the presence of the white car will be able to interpret this behavior as an indication that the driver's intention is to go straight or to turn left. Indeed, these two maneuvers involve waiting for the white car to pass while there is no reason to wait before making a right turn. In comparison, an algorithm which assumes independence between the two vehicles will not be able to infer anything about the driver's maneuver intention.

In the next two sections, a more detailed study of the impact of the independence assumption on situation assessment and risk assessment is carried out, using example scenarios at a T-shaped give-way intersection.

B.2.2 Case study: non-dangerous scenario

An analysis is conducted on an example non-dangerous scenario and illustrated in Figure B.3.




Figure B.3. Non-dangerous scenario: impact of the independence assumption on situation assessment and risk estimation. The thicker the arrows, the higher the probability for a maneuver.

A vehicle is proceeding in a T-shaped intersection from the secondary road while another vehicle is approaching at high speed from the right, without decelerating. A human observing the scene would conclude that the blue car is executing a right turn, and that the black car is driving at high speed because it has priority and intends to go straight at the intersection. This is actually a very common scenario; in this case a warning would be considered to be a false alarm. Taking into account the dependencies between the vehicles or not leads to very different outputs:

- Figure B.3, top row: the maneuvers of the vehicles are estimated independently from each other, based only on each vehicle's behavior and using knowledge about the layout of the intersection. Maneuvers 1 and 2 are therefore equally probable for the blue car, and maneuver 4 is much more probable than maneuver 3 for the black car. Let us assume that the following probabilities are obtained:

	Blue vehicle	Black vehicle
Maneuver 1	0.5	-
Maneuver 2	0.5	-
Maneuver 3	-	0.1
Maneuver 4	-	0.9

For the sake of simplicity in the calculation of the collision risk, let us assume that the pairs of maneuvers (1,3) and (1,4) result in a collision with 100% certainty. The collision risk can then be computed as:

$$Collision_risk = 0.5 \times 0.1 + 0.5 \times 0.9 = 0.5 \quad (B.1)$$

- Figure B.3, bottom row: this time, the dependencies between the vehicles are taken into account in the maneuver intention estimation process. This additional information does not change the interpretation of the black car's behavior but allows us to identify that the blue car is probably performing maneuver 2. Let us assume that the following probabilities are obtained:

	Blue vehicle	Black vehicle
Maneuver 1	0.1	-
Maneuver 2	0.9	-
Maneuver 3	-	0.1
Maneuver 4	-	0.9

Still assuming that the pairs of maneuvers (1,3) and (1,4) result in a collision with 100% certainty, the collision risk can then be computed as:

$$Collision_risk = 0.1 \times 0.1 + 0.1 \times 0.9 = 0.1 \quad (B.2)$$

For the same non-dangerous situation, the first approach estimates that the collision risk is 0.5 and the second one that it is 0.1. Therefore the sensitivity of the risk assessment [109] is different between the two approaches: in non-dangerous situations the calculated risk reaches comparatively higher values with the independence assumption, therefore the threshold λ used by applications to separate non-dangerous situations from dangerous situations has to be set to a higher value. This has significant consequences on safety applications. For example, let us consider an application which can control the brakes of the car and slows the car down when it detects danger ahead, i.e. when $Collision_risk \in [\lambda \ 1.0]$. A small λ results in a large range of variation of the $Collision_risk$ value in dangerous situations (see Figure B.4). It is then possible to adjust the action on the brakes (from soft to hard braking) depending on the value of $Collision_risk$. Comparatively, a larger λ results in a smaller range of variation of the $Collision_risk$ value in dangerous situations and will limit the choices of action to “no action” and “emergency braking”.

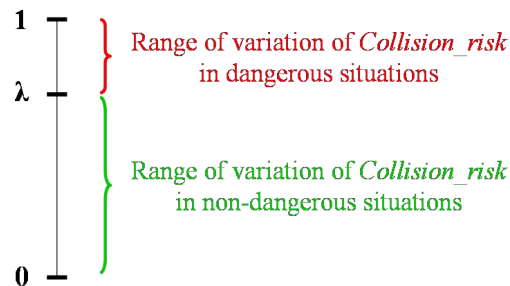


Figure B.4. Sensitivity of the risk assessment: the range of variation of $Collision_risk$ in dangerous and non-dangerous situations depends on the value set for λ .

It is also interesting to note that the difference between the two approaches is even more significant in X-shaped intersections. If we consider the same non-dangerous situation as in Figure B.3 but with a X-shaped intersection instead of a T-shaped intersection, the risk value computed by the first method will be 0.66, and still 0.1 with the second method.

B.2.3 Case study: dangerous scenario

One can wonder about potential drawbacks of the assumption that “drivers generally respect traffic rules”. Won’t such an assumption prevent or delay the detection of dangerous situations? Here we focus on a simple example scenario and show that a system that makes this assumption is still able to detect dangerous situations where a driver does not comply with the traffic rules.

The scenario depicted in Figure B.5 is similar to the previous one except for the presence of a third car.




Figure B.5. Dangerous scenario: impact of the independence assumption on situation assessment and risk estimation. The thicker the arrows, the higher the probability for a maneuver.

The yellow car is approaching the intersection from the main road at constant speed, in the opposite direction compared to the black car. This time, the traffic rules do not provide any indication on the maneuver performed by the blue car, since neither executing maneuver 1 nor executing maneuver 2 is in accordance with the traffic rules. It will be concluded that both maneuvers are equally probable, whether independence is assumed between the vehicles or not. The estimated collision risk will be high in both cases.

One can also wonder what would happen in the scenario of Figure B.3 if the blue car happened to turn left in the end. In this case, shortly after the car enters the intersection, the clues given by its position and orientation will counterbalance the traffic rules. The laws of physics (turning limitations of the car) will leave no doubt on the driver’s intention, and the estimated collision risk will be high.

B.3 Summary

The goal of this appendix was to state the importance of contextual information in motion models. Context refers here to the layout of the intersection and to the presence of other vehicles. The necessity to take into account the geometrical and topological characteristics of the intersection was demonstrated on a real-life example for the interpretation of turn signals. The advantages of accounting for the mutual influences between vehicles' maneuvers was illustrated by a real-life example, and then analyzed in more detail on toy examples. By comparing a model which assumes independence between vehicles with a model which assumes that drivers generally respect traffic rules, we reached the conclusion that accounting for the inter-dependencies between the vehicles:

1. Should in theory lead to a better estimation of drivers' maneuver intentions
2. Should in theory not delay or prevent the detection of dangerous situations
3. Should in theory lead to a better sensitivity of the risk assessment

These claims are evaluated in the results chapter (Chapter 4) using data collected at real intersections.

Appendix C

Digital maps

In this work it is assumed that a digital map of the road network is available from which we can extract high-level features of road intersections such as the set of authorized maneuvers and the traffic rules that apply. In our early experiments we used the Road Network Definition Format (RNDF) format [123], while our more recent work uses the OpenStreetMap (OSM) [124] format. Some alternatives exist, for example Navteq [125] and TomTom [126] commercial maps. It has also been shown that such maps can be inferred from position data [127, 128].

This appendix focuses on the RNDF and OSM formats since they are the ones used in this work. In the first section some definitions are provided about the representation of the road network in digital maps. In the second section, the procedure for extracting high-level features of the intersection from a digital map is described. The last section provides the xml code for one of our test intersection as an example.

C.1 Definitions

Both RNDF and OSM formats are based on a very standard representation of the road network consisting of *nodes*, *ways*, *relations*, and *attributes*:

- A *node* is a single geospatial point defined by its XYZ-coordinates in the Universal Transverse Mercator (UTM) coordinate system.
- A *way* is an ordered set of nodes.
- The logical or geographic relationships between *nodes* and/or *ways* are defined by *relations*.
- *Nodes*, *ways* and *relations* can have *attributes*, which describe the characteristics of an element.

This representation is illustrated in Figure C.1.




Figure C.1. Illustrative example of the *node* and *way* concepts in digital maps.

The RNDF format provides strict guidelines as to the generation of the maps (e.g. nodes are to be placed at the center of lanes, there is a predefined list of attributes which must be set for every node and way), while the OSM format leaves a lot of freedom in the representation of the road network.

C.2 High-level features

A set of high-level features can be extracted from digital maps:

- Authorized maneuvers

A maneuver at an intersection is defined as the process of entering the intersection through one lane and exiting the intersection through another lane. A maneuver is authorized iff it corresponds to a legal move. The information about the authorized maneuvers can be encoded in different manners in digital maps. For example in the RNDF format a *way* is defined for each authorized maneuver, while in the OSM format *ways* are sometimes shared by multiple maneuvers and the information about authorized moves is encoded in the *attributes* of the *ways* (e.g. “number of lanes”, “right turn only from right lane” *attributes*).

- Courses

For each authorized maneuver at an intersection, a course is defined as the typical path for executing that particular maneuver. Depending on how precise the digital map is, *ways* can be used directly as typical paths or they will have to be interpolated to make them smoother (e.g. using splines).

- Priority rules

A priority rule defines which maneuver has priority with respect to another maneuver. A priority rule can either be explicitly encoded through a *relation* between two *ways* or be implicitly encoded through *attributes* of *nodes* (e.g. “give way”, “stop” attributes).

This representation is illustrated in Figure C.2. The corresponding xml code for the digital map is given in Section C.3.




Figure C.2. Example digital map, superimposed on the corresponding satellite image. This intersection is defined by 6 courses (represented by 6 ways). The corresponding XML code is provided in Section C.3.

C.3 XML code

```
<?xml version='1.0' encoding='UTF-8'?>
<osm version='0.6' generator='JOSM'>
  <node id='-82' timestamp='2012-04-09T23:17:21Z' visible='true' lat='
    48.72659483361505' lon='2.0010682466535825' />
  <node id='-80' action='modify' timestamp='2012-04-09T23:17:21Z' visible='true'
    lat='48.726599870381015' lon='2.0011008755352964' />
  <node id='-78' action='modify' timestamp='2012-04-09T23:17:21Z' visible='true'
    lat='48.726599898457174' lon='2.0011008118373046' />
  <node id='-72' action='modify' timestamp='2011-07-12T09:21:40Z' visible='true'
    lat='48.72657426510799' lon='2.0011052987503395' />
```



```
<tag k='name' v='31' />
</node>
<node id='-70' action='modify' timestamp='2011-07-12T09:21:40Z' visible='true'
  lat='48.726581573167174' lon='2.0011220022874583'>
  <tag k='name' v='2' />
</node>
<node id='-68' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72678097971454' lon='2.001184386557687'>
  <tag k='name' v='41' />
</node>
<node id='-66' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.7266095123778' lon='2.0011055362211323'>
  <tag k='name' v='51' />
</node>
<node id='-64' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.7265597738326' lon='2.001126172491377'>
  <tag k='name' v='23' />
</node>
<node id='-60' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.726595190421726' lon='2.0010863034655095'>
  <tag k='name' v='45' />
</node>
<node id='-56' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72659903650132' lon='2.0010386126670845'>
  <tag k='exit' v='3' />
  <tag k='name' v='4' />
</node>
<node id='-52' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.7270559598962' lon='2.001310841399108'>
  <tag k='name' v='40' />
</node>
<node id='-50' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72704294841827' lon='1.9989236160232802'>
  <tag k='name' v='6' />
</node>
```

```
<node id='-48' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72591792978391' lon='2.00429861403403'>
  <tag k='name' v='0' />
</node>
<node id='-46' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72589496670608' lon='2.0042798065757887'>
  <tag k='name' v='24' />
</node>
<node id='-44' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'
  lat='48.72702102734542' lon='1.998907013430011'>
  <tag k='name' v='20' />
</node>
<node id='-40' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.726586401284486' lon='2.001100560347422'>
  <tag k='name' v='3' />
</node>
<node id='-34' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72660858523698' lon='2.001104321734558'>
  <tag k='name' v='44' />
</node>
<node id='-32' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.726569430972624' lon='2.001179982778936'>
  <tag k='name' v='1' />
</node>
<node id='-30' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72659903517921' lon='2.001152566459314'>
  <tag k='name' v='11' />
</node>
<node id='-28' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72661341520267' lon='2.0011490250825235'>
  <tag k='exit' v='2' />
  <tag k='name' v='12' />
</node>
<node id='-26' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72704658161592' lon='2.0013440908419757'>
  <tag k='name' v='14' />
```

```
</node>
<node id='-18' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.7265728137508' lon='2.001086262802265'>
  <tag k='name' v='30' />
</node>
<node id='-16' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72661341884052' lon='2.0011486219077823'>
  <tag k='exit' v='2' />
  <tag k='name' v='34' />
</node>
<node id='-10' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'
  lat='48.72658521809297' lon='2.0011601593408055'>
  <tag k='name' v='10' />
</node>
<node id='-6' action='modify' timestamp='2011-10-05T08:13:13Z' visible='true'
  lat='48.72658062288928' lon='2.0010276651917884'>
  <tag k='entrance' v='2' />
  <tag k='name' v='25' />
</node>
<way id='-94' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'>
  <nd ref='-52' />
  <nd ref='-68' />
  <nd ref='-66' />
  <nd ref='-78' />
  <nd ref='-40' />
  <nd ref='-72' />
  <nd ref='-64' />
  <nd ref='-46' />
  <tag k='name' v='15' />
  <tag k='speed_limit' v='50' />
</way>
<way id='-92' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'>
  <nd ref='-52' />
  <nd ref='-68' />
  <nd ref='-34' />
  <nd ref='-80' />
```

```
<nd ref='-60' />
<nd ref='-82' />
<nd ref='-56' />
<nd ref='-50' />
<tag k='name' v='14' />
<tag k='speed_limit' v='50' />
</way>
<way id='-90' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'>
  <nd ref='-48' />
  <nd ref='-32' />
  <nd ref='-70' />
  <nd ref='-40' />
  <nd ref='-56' />
  <nd ref='-50' />
  <tag k='name' v='10' />
  <tag k='speed_limit' v='70' />
</way>
<way id='-88' action='modify' timestamp='2011-07-12T06:47:18Z' visible='true'>
  <nd ref='-44' />
  <nd ref='-6' />
  <nd ref='-64' />
  <nd ref='-46' />
  <tag k='name' v='12' />
  <tag k='speed_limit' v='70' />
</way>
<way id='-86' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'>
  <nd ref='-48' />
  <nd ref='-32' />
  <nd ref='-10' />
  <nd ref='-30' />
  <nd ref='-28' />
  <nd ref='-26' />
  <tag k='name' v='11' />
  <tag k='speed_limit' v='70' />
</way>
<way id='-84' action='modify' timestamp='2011-07-12T08:32:15Z' visible='true'>
```

```
<nd ref='-44' />
<nd ref='-6' />
<nd ref='-18' />
<nd ref='-72' />
<nd ref='-70' />
<nd ref='-16' />
<nd ref='-26' />
<tag k='name' v='13' />
<tag k='speed_limit' v='70' />
</way>
<relation id='-102' action='modify' timestamp='2011-10-05T08:13:14Z' visible='
  true '>
  <member type='way' ref='-90' role='1' />
  <member type='way' ref='-84' role='0' />
  <member type='way' ref='-94' role='0' />
  <member type='way' ref='-92' role='0' />
</relation>
<relation id='-100' timestamp='2011-10-05T08:13:14Z' visible='true '>
  <member type='way' ref='-86' role='1' />
  <member type='way' ref='-84' role='0' />
</relation>
<relation id='-98' timestamp='2011-10-05T08:13:14Z' visible='true '>
  <member type='way' ref='-88' role='1' />
  <member type='way' ref='-94' role='0' />
</relation>
<relation id='-96' timestamp='2011-10-05T08:13:14Z' visible='true '>
  <member type='way' ref='-84' role='1' />
  <member type='way' ref='-94' role='0' />
</relation>
</osm>
```

Appendix D

Precision and recall analysis

For evaluation purposes in Chapter 4, a threshold λ is applied on the risk value to separate *dangerous* and *non-dangerous* situations. In order to set the value of λ , we used the simulated dataset presented in Section 4.1 and conducted an analysis on the precision, recall, and smallest collision prediction horizon for different values of λ . The selected value for λ was subsequently used for all the tests presented in Chapter 4.

D.1 Metrics

Precision: For one value of λ , the precision is defined as follows.

$$Precision = \frac{NC}{NC + NF} \quad (D.1)$$

with NC the number of correct detections and NF the number of false alarms.

A correct detection corresponds to a *dangerous* scenario instance where the risk exceeded λ at some point before the collision occurred.

A false alarm corresponds to a *non-dangerous* scenario instance where the risk exceeded λ at some point.

Recall: For one value of λ , the recall is defined as follows.

$$Recall = \frac{NC}{NC + MD} \quad (D.2)$$

with NC the number of correct detections and MD the number of missed detections.

A correct detection corresponds to a *dangerous* scenario instance where the risk exceeded λ at some point before the collision occurred.

A missed detection corresponds to a *dangerous* scenario instance where the risk never exceeded λ .

Smallest collision prediction horizon: For one value of λ , we find the smallest $T_{prediction}$ (see definition in Equation 4.2) among all the *dangerous* scenario instances.

D.2 Results

The precision, recall, and smallest collision prediction horizon were computed for values of λ between 0.0 and 1.0. The results are shown in Figure D.1. Optimal values for the precision and recall are obtained for $\lambda = 0.3$ and $\lambda = 0.35$. The threshold was set to $\lambda = 0.3$ since the collision prediction horizon is larger.

λ	Precision	Recall	Smallest $T_{prediction}$ (s)	
0.05	0.7	1.0	3.4	
0.1	0.7	1.0	1.6	
0.15	0.73	1.0	1.0	
0.2	0.79	1.0	0.8	
0.25	0.89	1.0	0.6	
0.3	1.0	1.0	0.6	← Selected value
0.35	1.0	1.0	0.4	
0.4	1.0	0.996	0.0	
0.45	1.0	0.996	0.0	
0.5	1.0	0.987	0.0	
0.55	1.0	0.983	0.0	
0.6	1.0	0.970	0.0	
0.65	1.0	0.965	0.0	
0.7	1.0	0.957	0.0	
0.75	1.0	0.957	0.0	
0.8	1.0	0.948	0.0	
0.85	1.0	0.939	0.0	
0.9	1.0	0.918	0.0	
0.95	1.0	0.900	0.0	

Figure D.1. Precision, recall, and smallest collision prediction horizon for different values of λ .

Appendix E

Accounting for physical limitations in the evolution model

In Section 3.4.4 we discuss about two possible models for the evolution of the pose of a vehicle.

Model 1 is a combination of two simple models. The first one predicts the pose at time t by applying a Constant Velocity (CV) model to the pose at time $t - 1$. The second one projects the pose obtained with the CV model on the course that the driver intends to follow. The poses computed by the two models are averaged to produce a final predicted pose.

Model 2 is a more advanced path following model which takes into account the physical limitations of the vehicle. The predicted pose at time t is one which is physically reachable by the vehicle. For example, this path-following behavior could be implemented using a Proportional-Integral-Derivative (PID) controller.

In our experiments, Model 1 led to better estimations of the intentions of drivers. The objective of this appendix is to explain why, using the example situation depicted in Figure E.1.




Figure E.1. Two different models for the evolution of the pose: (a) Combination of a Constant Velocity model and a projection model. (b) Path following model taking into account the physical limitations of the vehicle.

Analysis of the situation in Figure E.1: For $c_{t-1}^n = c_1$, the predicted pose at time t is very similar with the two models : both prediction are located very close to course c_1 .

For $c_{t-1}^n = c_2$, the predictions are very different. With Model 1, the predicted pose consists in a “jump” towards course c_2 . With Model 2, the predicted pose consists in a sharp steering.

As a result, the driver’s intention will be better estimated with Model 1. The true intent of the driver is obviously to make a left turn at the intersection, so the measurement of the pose of the vehicle at time t will be close to course c_1 . With Model 1, the blue prediction is far enough from the measurement that the algorithm will be able to conclude that the intended course of the driver is c_1 . This is not the case with Model 2: course c_1 is of course more likely, but the blue prediction is still close to the measurement and the algorithm will not be able to discard course c_2 as a potential intended course.

Hints for future work: We showed on a toy example why taking into account the physical limitations of the vehicle for the evolution of the pose is not suitable with the motion model proposed in

this thesis. In this paragraph, a solution is suggested to incorporate the physical limitations without the drawback mentioned above.

The idea is the following: in situations like the one in Figure E.1, it should be possible to discard course c_2 in another part of the motion model. If unlikely courses are discarded in another part of the motion model, it is no longer a problem to use a model like Model 2 for the evolution of the pose.

In order to discard unlikely courses, we could modify the joint distribution proposed in Equation 3.29 to make the current intended course Ic_t^n dependent of the previous pose P_{t-1}^n . Instead of the “continuity” model proposed for the intended course in Section 3.4.3, a model of the form $P(Ic_t^n | Ic_{t-1}^n P_{t-1}^n)$ could be defined where courses which are far from the previous pose are very unlikely to be the current intended course.

Appendix F

Inference for risk estimation

Inference in Equation 3.30 is performed using a bootstrap filter [102]. This recursive filtering approach is based on the approximation of the probability density function (pdf) by a set of weighted samples called particles. The set of $N_{particles}$ particles at time t is denoted:

$$\{H_{i,t}, w_{i,t}\}_{i=1:N_{particles}} \quad (\text{F.1})$$

with $H_{i,t}$ the state of particle i at time t and $w_{i,t}$ the weight of particle i at time t .

The set of particles has to be initialized when the first measurements are received. Then, each time some new measurements Z_t become available, the set of particles is adapted following 3 steps:

1. In the prediction step, the state of each particle at time t is predicted from its state at time $t - 1$, based on a proposal function π . In the bootstrap filter, the transition prior is used as a proposal function:

$$\pi(H_t|H_{0:t-1}Z_{0:t}) = P(H_t|H_{t-1}Z_{t-1}) \quad (\text{F.2})$$

2. In the update step, the measurements at time t are used to modify the weight of each particle so that it reflects the likelihood of the predicted state:

$$w_{i,t} \propto w_{i,t-1} \times P(Z_t|Z_{t-1}H_{i,t-1:t}) \quad (\text{F.3})$$

3. In the resampling step, a new set of particles is generated based on their weight in order to avoid their degeneracy and focus on particles with more likely states.

In the next sections, the bootstrap filter is applied to our inference problem.

Firstly, in order to optimize the efficiency of the inference process, the joint distribution is simplified. Since $P(Sm_t^n|S_t^n)$ is a Dirac function (see Section 3.4.8), it is possible to replace S_t^n by Sm_t^n and to remove $P(Sm_t^n|S_t^n)$. Similarly, T_t^n can be replaced by Tm_t^n and $P(Tm_t^n|T_t^n)$ can be removed. The joint distribution proposed in Equation 3.29 becomes:

$$\begin{aligned}
 P(\mathbf{E}_{0:T}\mathbf{I}_{0:T}\mathbf{\Phi}_{0:T}\mathbf{Z}_{0:T}) &= P(\mathbf{E}_0\mathbf{I}_0\mathbf{\Phi}_0\mathbf{Z}_0) \\
 &\times \prod_{t=1}^T \times \prod_{n=1}^N [P(Es_t^n|\mathbf{I}_{t-1}\mathbf{P}_{t-1}\mathbf{S}_{t-1}) \\
 &\times P(Ic_t|Ic_{t-1}) \times P(Is_t^n|Is_{t-1}^nEs_t^n) \times P(P_t^n|P_{t-1}^nSm_{t-1}^nIc_t) \\
 &\times P(Pm_t^n|P_t^n) \times P(Sm_t^n|Sm_{t-1}^nP_{t-1}^nIc_tIs_t^n) \times P(Tm_t^n|P_{t-1}^nIc_{t-1}Ic_t)]
 \end{aligned} \tag{F.4}$$

F.1 Variables

In our problem the hidden variable H_t and the observed variable Z_t are defined as:

$$\begin{cases} H_t &= (\mathbf{E}_t\mathbf{I}_t\mathbf{I}_t\mathbf{P}_t) \\ Z_t &= (\mathbf{P}_t\mathbf{S}_t\mathbf{T}_t) \end{cases} \tag{F.5}$$

F.2 Initialization

The initial values of the particles $H_0 = (\mathbf{e}_{s_0}, \mathbf{c}_0, \mathbf{i}_{s_0}, \mathbf{p}_0)$ are computed based on the first set of observations $Z_0 = (\mathbf{p}_{m_0}, \mathbf{s}_{m_0}, \mathbf{t}_{m_0})$.

First, for each vehicle $n \in N$ a distribution on the intended course $P(Ic_0^n)$ is computed based on the ‘‘pose measurement model’’ defined in Section 3.4.7. The projection of the measured pose pm_0^n on the intended course c_0^n is used as an input:

$$P([Ic_0^n = c_0^n]) \propto P([Pm_0^n = pm_0^n][P_t^n = proj(pm_0^n, c_0^n)]) \tag{F.6}$$

with $proj(p, c)$ the orthogonal projection of the pose p on course c .

Then, for each particle i the initial values $(\mathbf{e}_{s_{i,0}}, \mathbf{c}_{i,0}, \mathbf{i}_{s_{i,0}}, \mathbf{p}_{i,0})$ are computed as follows:

1. For each vehicle $n \in N$, draw a value for the intended course based on the distribution computed above:

$$c_{i,0}^n \sim P(Ic_0^n) \quad (\text{F.7})$$

2. For each vehicle $n \in N$, draw a value for the pose using the “pose model” defined in Section 3.4.4. Since the value of the previous pose is not known, the projected pose is used instead and a null speed is assumed:

$$p_{i,0}^n \sim P(P_t^n | [P_{t-1}^n = \text{proj}(pm_0^n, c_{i,0}^n)] [S_{t-1}^n = 0.0] [Ic_t^n = c_{i,0}^n]) \quad (\text{F.8})$$

3. For each vehicle $n \in N$, draw a value for the expectation to stop using the “expected longitudinal motion model” defined in Section 3.4.1. The values computed in the previous steps are used as an input:

$$es_{i,0}^n \sim P(Es_t^n | [\mathbf{Ic}_{t-1} = \mathbf{c}_{i,0}] [\mathbf{P}_{t-1} = \mathbf{p}_{i,0}] [\mathbf{S}_{t-1} = \mathbf{sm}_0]) \quad (\text{F.9})$$

4. For each vehicle $n \in N$, draw a value for the intention to stop using the “intended longitudinal motion model” defined in Section 3.4.2. Since the value of the previous intention to stop is not known, a random value is generated. For the expectation to stop, the value computed in the previous step is used.

$$is_{i,0}^n \sim P(Is_t^n | [Is_{t-1}^n = \text{random}] [Es_t^n = es_{i,0}^n]) \quad (\text{F.10})$$

F.3 Recursive risk computation

1. Prediction step

The general equation for the prediction step was given above (Equation F.2) and is recalled below in the context of our problem:

$$\pi(H_t | H_{0:t-1} Z_{0:t}) = P(\mathbf{Es}_t \mathbf{Ic}_t \mathbf{Is}_t \mathbf{P}_t | \mathbf{Es}_{t-1} \mathbf{Ic}_{t-1} \mathbf{Is}_{t-1} \mathbf{P}_{t-1} \mathbf{Pm}_{t-1} \mathbf{Sm}_{t-1} \mathbf{Tm}_{t-1}) \quad (\text{F.11})$$

In our case (see joint distribution in Equation F.4), this equation simplifies as:

$$\begin{aligned} \pi(H_t|H_{0:t-1}Z_{0:t}) &= \prod_{n=0}^N [P(Es_t^n|\mathbf{Ic}_{t-1}\mathbf{P}_{t-1}\mathbf{Sm}_{t-1}) \times P(Ic_t^n|Ic_{t-1}^n) \\ &\quad \times P(Is_t^n|Is_{t-1}^nEs_{i,t}^n) \times P(P_t^n|P_{t-1}^nSm_{t-1}^nIc_t^n)] \end{aligned} \quad (\text{F.12})$$

2. Update step

The general equation for the update step was given above (Equation F.3) and is recalled below in the context of our problem:

$$w_{i,t} \propto w_{i,t-1} \times P(\mathbf{Pm}_t\mathbf{Sm}_t\mathbf{Tm}_t|\mathbf{Pm}_{t-1}\mathbf{Sm}_{t-1}\mathbf{Tm}_{t-1}\mathbf{Es}_{t-1:t}\mathbf{Ic}_{t-1:t}\mathbf{Is}_{t-1:t}\mathbf{P}_{t-1:t}) \quad (\text{F.13})$$

In our case (see joint distribution in Equation F.4), this equation simplifies as:

$$w_{i,t} \propto w_{i,t-1} \times \prod_{n=1}^N P(Pm_t^n|P_t^n) \times P(Sm_t^n|Sm_{t-1}^nP_{t-1}^nIc_t^nIs_t^n) \times P(Tm_t^n|P_{t-1}^nIc_{t-1}^nIc_t^n) \quad (\text{F.14})$$

3. Resampling

A systematic resampling approach was adopted.

Finally, the risk as defined in Equation 3.30 can be approximated by summing the weights of the particles which verify the condition ($[Is_{i,t}^n = go]$, $[Es_{i,t}^n = stop]$):

$$P([Is_t^n = go][Es_t^n = stop]|\mathbf{Pm}_{0:t}\mathbf{Sm}_{0:t}\mathbf{Tm}_{0:t}) = \sum_{i: ([Is_{i,t}^n = go], [Es_{i,t}^n = stop])} w_{i,t} \quad (\text{F.15})$$