Collision Risk Assessment to Improve Driving Safety

Christian Laugier*, Igor E. Paromtchik*, Christopher Tay†, Kamel Mekhnacha†,
Gabriel Othmezouri‡, Hiromichi Yanagihara‡
*INRIA Grenoble Rhône-Alpes, 38334 Saint Ismier, France
†ProBayes, 38334 Saint Ismier, France
‡Toyota Motor Europe, B-1930 Zaventem, Belgium

***Abstract*—Robust analysis of dynamic scenes in urban traffic environments is needed to estimate and predict collision risk level during vehicle driving. The risk estimation relies on monitoring of the traffic environment of the vehicle by means of on-board lidars and a stereo camera. The collision risks are considered as stochastic variables. Hidden Markov Model and Gaussian process are used to estimate and predict collision risks and the likely behaviors of multiple dynamic agents in road scenes. The proposed approach to risk estimation is tested in a virtual environment with human-driven vehicles and during a highway driving. The obtained results have proven the feasibility of our approach to assist the driver in avoiding potentially dangerous situations.**

***Index Terms*—Collision risk, urban road, driver assistance, Hidden Markov Model, Gaussian process**

I. Introduction

A. Problem statement

The urban traffic environment with multiple participants contains risks of potential collision and damage. The vehicle safety technologies (e.g. seat belts, airbags, safety glass, energy-absorbing frames) mitigate the effects of accidents. The advanced technologies will be capable of monitoring the environment to estimate and predict collision risks during vehicle driving, in order to help reduce the likelihood of accidents occuring. The risk management by traffic participants is an efficient way to improve traffic safety toward *zero-collision* driving. The key problem is to correctly interpret the traffic scene by means of processing information from a variety of sensors [1].

A collision risk level can be predicted for a few seconds ahead to warn the driver about unnoticed potential risks. The estimated risk of collision can also be used to select a trajectory that minimizes the risks for an autonomous vehicle. The estimation of collision risk relies on the sensor information about the surrounding environment and the driver's behavior. The obtained risk values must be interpreted by the dedicated application. The following set of processed sensor inputs is assumed to be available.

Road geometry. The road width and the road curvature are obtained by processing raw information from camera images and lidars or, alternatively, from a Geographic Information System (GIS) with a pre-built map and a localization device such as Global Positioning System (GPS).

Object tracking. The detection and tracking of moving objects is accomplished by the dedicated algorithms, i.e. positions and velocities of the objects are available.

Auxiliary sensors. Information about the signal light status of other vehicles is an indicator of the motion intention. Additional "virtual" sensors are capable of detecting distances between the vehicles and the lane borders, which might indicate an intention to perform a lane change.

We use a term "ego-vehicle" to distinguish our vehicle from other vehicles. The ego-vehicle is assumed to be equipped with the appropriate sensors for obtaining a set of inputs mentioned above. The estimated risk is a numerical value which expresses quantitatively the collision risk of the ego-vehicle with another vehicle during the next few seconds.

Estimating the future collision risk requires the models which describe the vehicle motion in the sensor visibility range of the ego-vehicle. This model must be capable of reasonably predicting the future states of the vehicle in terms of the probability. We present a fully probabilistic model of the vehicle's motion evolution for obtaining and inferring beliefs on the future states of vehicles in an urban traffic environment. Consequently, the estimated risk of collision is obtained from the models in terms of the probability in a theoretically consistent manner.

B. Related work

Current commercially available crash warning systems aim at preventing front, rear, or side collisions. Such systems are usually equipped with radar based sensors on the front, rear or sides to measure the velocity and distance to obstacles. The algorithms for determining the risk of collision are based on variants of time-to-collision (TTC) [2], giving the time remaining before one vehicle collides with another one, assuming the both vehicles are maintaining their linear velocities.

Some systems are capable of directly controlling the brakes and possibly the steering to perform the necessary corrective actions. Systems based on TTC use the observations made at a reasonably high frequency in order to adapt to a dynamic environment. Current commercial systems work reasonably well on highways or straight sections of the city roads. However, the linearity assumption does not hold on curved roads, as shown in figure 1, where the risk level tends to be overestimated.

Several research projects overcome this problem by taking into account the structure of the environment, especially at intersections where the rate of accidents is higher. These projects aim at providing the collision warning systems, which use wireless communication either between the vehicles or between the vehicle and a road infrastructure, such as traffic

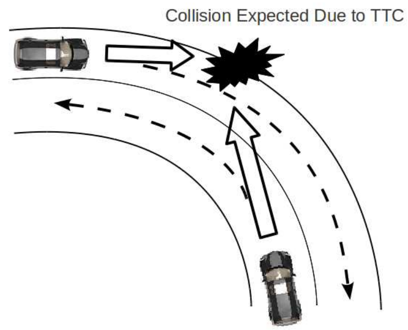

Figure 1: An example of triggering a false alarm about collision because of the invalid linearity assumption on curved roads in the TTC based systems.

lights [3], [4], [5], [6], [7]. These systems are equipped with a pair of detectors (radars or laser scanners) at the left and right front corners of the vehicle in order to detect cross-traffic vehicles at intersections. The obtained speeds of the objects and the respective TTC are then evaluated to determine the collision risk. Although the environmental structures are taken into consideration, yet the collision risk calculation assumes the straight motion. The time horizon of risk prediction is short, and the crucial environmental information and sensor data are not fully employed.

C. Outline of the approach

The knowledge about an object being at a certain location at a specific time does not provide sufficient information to assess its impact on the safety of the ego-vehicle. In addition, environmental constraints should be taken into account, especially on urban roads. We propose a framework for understanding behaviors of other vehicles and present our approach in the next section.

The relevant sensors include stereo vision, lidars, an inertial measurement unit (IMU) combined with the GPS, and odometry. The local environment is represented by a grid. The fusion of sensor data is accomplished by means of the Bayesian Occupancy Filter (BOF) [8], [9], that provides to assign probabilities of *cell occupancy* and *cell velocity* for each cell in the grid. The collision risks are considered as stochastic variables. Hidden Markov Model (HMM) and Gaussian process (GP) are used to estimate and predict collision risks and the likely behaviors of multiple dynamic agents in road scenes [10], [11].

Consider vehicle A and ego-vehicle B traveling in the same direction on the adjacent lanes, as shown in figure 2. The collision risk must be estimated for vehicle B. From the driver's viewpoint, the road structure is implicitly described by such maneuvers as: move straight, turn right/left, or change a lane. These maneuvers are referred to as behaviors, and a set of the possible behaviors is predefined. However, some behaviors are unavailable at all instances, e.g. it might be impossible to turn left at an intersection because of the road geometry.

The lane following for a given behavior is represented by means of a GP, i.e. a probability distribution over the possible future realizations of the paths with the mean corresponding to the exact following of the lane middle. The GP samples for such behaviors as "lane change" and "moving straight" are depicted in figure 2, where the dotted lines represent the paths sampled from the GP. For a lane turning on a curved road, the GP is adapted to reflect the road geometry. The set of GPs for each feasible behavior, in combination with the probability of vehicle A executing a certain behavior, gives a probabilistic model of the future evolution of vehicle A in the scene.

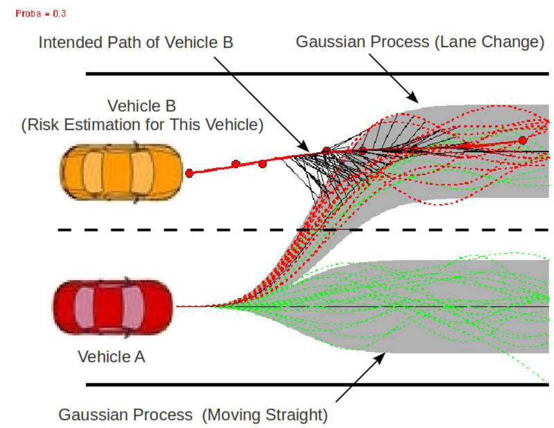

Figure 2: Collision risk estimation for vehicle B relies on predicting the path of vehicle A by sampling from the GP for two possible behaviors in this example: "moving straight" or "lane change". The collision risk is obtained as a weighted sum of paths leading to collision.

Similar to the TTC approach, the evaluation of collision risk is performed for vehicle B against vehicle A. In contrast to the TTC's linearity assumption about the future paths for the vehicles, we evaluate the collision risk of the intended path of vehicle B against all possible paths to be taken by vehicle A. The weights are assigned according to the probabilistic model of the behaviors' evolution of vehicle A.

II. Collision Risk Estimation

An overall architecture of the risk estimation module is shown in figure 3. It comprises three sub-modules, such as: driving behavior recognition, driving behavior realization, and collision risk estimation [11], [12].

Driving behavior recognition. The behavior recognition aims at estimating the probability distribution of feasible behaviors, e.g. $P(turn\_left)$ represents the probability of turning left by the vehicle. The behaviors provide an implicit high-level representation of a road structure, which contains semantics. The probability distribution over behaviors is obtained by HMM. Our current model includes the following four behaviors: move straight, turn left, turn right, and overtake.

Driving behavior realization. The collision risk evaluation requires the road geometry. Driving behavior realization takes the form of GP, i.e. a probabilistic representation of a possible evolution of the vehicle motion for a given behavior [10].

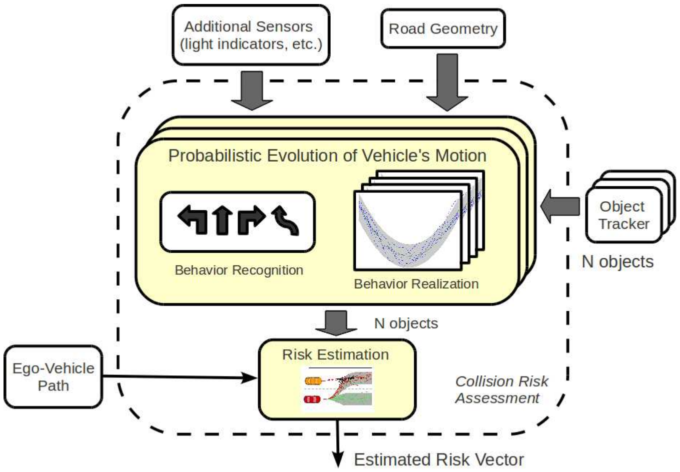

Figure 3: Architecture of the risk assessment module.

The adaptation of GP according to the behavior is based on the geometric transformation known as the Least Squares Conformal Map (LSCM) [13].

Collision risk estimation. A complete probabilistic model of the possible future motion is given by the probability distribution over behaviors from the driving behavior *recognition* and the driving behavior *realization*. The collision risk is calculated from this model. Intuitively, the result can be explained as "collision risk for a few seconds ahead". However, the precise mathematical definition of risk depends on the meaning and interpretation of risk [11].

Behavior recognition and modeling

The behavior recognition aims at assigning a label and a probability measure to sequential data, i.e. observations from the sensors. Examples of sensor values are: distance to lane borders, signaling light status, or a proximity to an intersection. The objective is to obtain the probability values over behaviors, i.e. the behaviors are hidden variables.

The behavior modeling contains two hierarchical layers. The upper layer is a single HMM, where its hidden states represent high-level behaviors, such as: move straight, turn left, turn right, and overtake. For each hidden state or each behavior in the upper layer HMM, there is a corresponding HMM in the lower layer to represent the sequence of the finer state transitions of a single behavior, as depicted in figure 4.

Let us define the following hidden state semantics in the lower layer HMMs for each of the following behaviors of the higher layer HMM:

- *Move straight (1 hidden state):* move forward.
- *Turn left or turn right (3 hidden states):* Decelerate before a turn, execute a turn, and resume a cruise speed.
- *Overtake (4 hidden states):* lane change, accelerate (while overtaking a vehicle), lane change to return to the original lane, resume a cruise speed.

In order to infer the behaviors, we maintain a probability distribution over the behaviors represented by the hidden states of the upper layer HMM. The observations of vehicles (i.e. sensor data) interact with the HMM in the lower layer, and the information is then propagated to the upper layer.

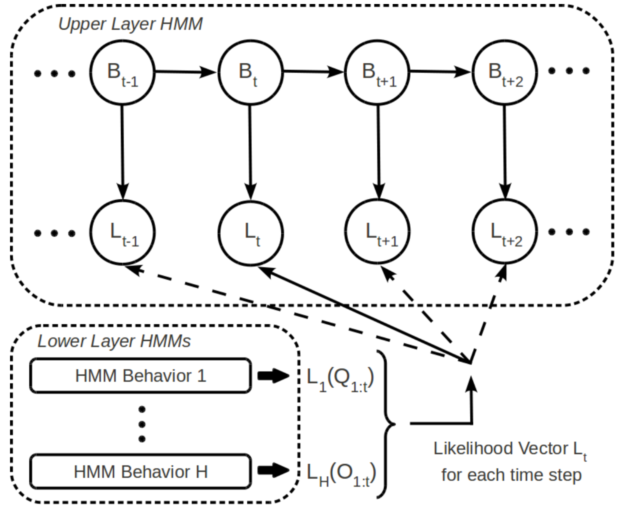

Figure 4: Layered HMM, where each lower layer HMM's observation likelihood is the upper layer HMM's observation.

Driving behavior realization

A behavior is an abstract representation of the vehicle motion. For a given behavior, a probability distribution over the physical realization of the vehicle motion is indispensable for risk estimation. The GP allows us to obtain this probability distribution by assuming that usual driving is represented by the GP, i.e. lane following without drifting too far off to the lane sides. On a straight road, this is a *canonical* GP with the mean corresponding to the lane middle.

To deal with the variations of lane curvature or such behaviors as "turn left" or "turn right", we propose an adaptation procedure, where the canonical GP serves as a basis and it is deformed according to the road geometry. The deformation method is based on LSCM. Its advantage is a compact and flexible representation of the road geometry. The canonical GP can be calculated once and, then, can be reused for different situations, thus, resulting in a better computational efficiency. An example is shown in figure 5 for a curved road.

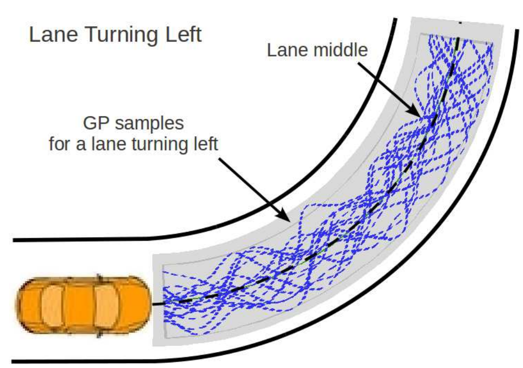

Figure 5: Deformed GP model example for a lane turning left.

Collision risk estimation

The layered HMM approach assigns a probability distribution over behaviors at each time instance, and a GP gives the probability distribution over its physical realization for each behavior. Because the behavioral semantics are propagated from the layered HMM down to the physical level, it is now possible to assign semantics to risk values.

One should note that the definition of risk can take a variety of forms, which is largely dependent on how the risk output is going to be used. A risk scalar value might be sufficient for a crash warning system, or an application might require the risk values against each vehicle in the traffic scene.

The risk calculation is performed by first sampling of the paths from the GP. The fraction of samples in collision gives the risk of collision, which corresponds to the behavior represented by the GP. A general risk value is obtained by marginalizing over behaviors based on the probability distribution over behaviors obtained from the layered HMM. It is possible to calculate risk of taking a certain path, a certain behavior, or a general risk value of a certain vehicle against another vehicle. A systematic framework for evaluation of different types of risk can be found in [11].

III. Experiments

A. Driving simulation

The simulation of crash situations in a virtual environment is used instead of dealing with them in real experiments. The virtual environment is a 3D geometric model of a road network with vehicles, where each vehicle is driven by a human driver. The simulator was developed by Toyota Motor Europe (TME). Each human driver controls his or her virtual vehicle by means of a steering wheel, the acceleration and brake pedals. Recording a scenario with multiple vehicles, which are driven concurrently, requires a large number of human drivers. An alternative is to generate the scenario iteratively, with one human-driven vehicle at a time and "adding" human drivers iteratively, with a replay of the previously recorded human-driven vehicles. The resulting virtual environment allows us to simulate crash situations safely.

The layered HMM evaluates the behavior of every vehicle in the scene for different time horizons, except the ego-vehicle. The training data are obtained by collecting sequences for a series of human-driven cases, where each driver uses the steering wheel as an interface to the virtual environment of the simulator. The driving sequences are then annotated manually by means of an annotation tool of ProBayes. Subsequently, the annotated data are used to train the layered HMM.

The TME simulator provides a 3D road view for the driver and a 2D view of the road network, as shown in figure 6. The collision risk is calculated for a yellow vehicle, while other vehicles are shown by red rectangles. The relevant area of the scene is inside a large yellow circle. The right-hand traffic rule is assumed. The trail behind the yellow vehicle in 2D view indicates the risk levels estimated previously. At each instant, the probabilities of the possible behaviors of the nearest neighbor (red vehicle) are estimated by the layered HMM and are displayed by the vertical white bars. The speed of the yellow vehicle is shown in 3D view, where the right-side vertical bar shows the risk encoding by color from "low" (green) to "high" (red). The left-side vertical bar in 3D view indicates the current risk value for the yellow vehicle.

Figure 6: Virtual environment of the TME simulator.

The speed warning in the case of a potential danger of frontal collision is available in most commercial systems. Additionally to this functionality, our algorithm evaluates risk at intersections, where the linearity assumption about the vehicle motion would result in underestimated values of collision risk. The combination of the behavior estimation by the layered HMM and the use of semantics (e.g. turn right or move straight) at the geometric level allows us to obtain the appropriate risk values.

The training data for the layered HMM were collected with ten human drivers who were asked to show different driving behaviors. The collected data is split uniformly distributed into the training data and the test data (30% of total data examples). The behavior recognition is trained on the training data and is evaluated against the test data.

Figure 7 summarizes the recognition performance of the layered HMM. The results are presented as a confusion matrix, where the columns correspond to the true class and the rows correspond to the estimated class. The diagonal values of the confusion matrix give the correctly predicted class, while non-diagonal values show the percentage of mislabeling for each class. The highest recognition rate is for "move straight" behavior (91.9%) as well as "turn right" or "turn left" ones (82.5% and 81.1%, respectively). The "overtake" behavior has a relatively low recognition rate of 61.6%. Intuitively, this lower rate can be explained by a composite structure of the overtaking maneuver because it consists of such behaviors as: accelerating, lane changing, returning to the original lane, and resuming a cruise speed. Consequently, it also takes longer than a three-second period (current prediction horizon) to complete an overtaking maneuver.

The approach to risk assessment is illustrated by figure 8, where the probability of collision is estimated for a period of three seconds ahead of each collision for ten different traffic scenarios. The rapid increase in the probability of collision and its certainty are observed when the collision instant approaches.

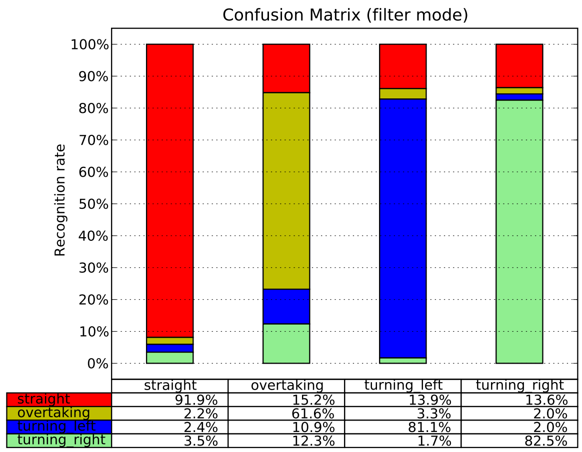

Figure 7: Performance summary of the behaviors recognition with layered HMM.

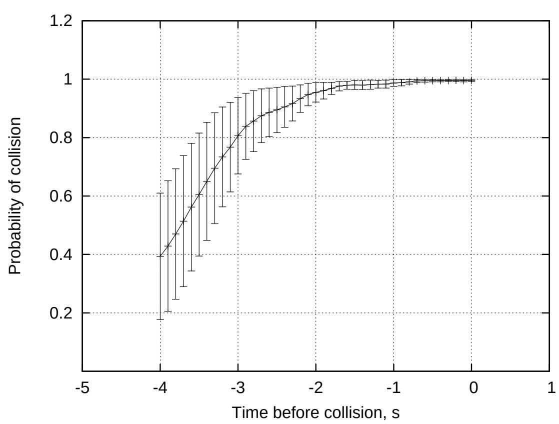

Figure 8: Example of collision risk assessment for ten human-driven scenarios and a three-second prediction horizon.

B. *Behavior estimation on a highway sequence*

The first phase is to gather experimental data when driving on a highway to estimate behaviors of other vehicles. The experiments have been conducted jointly by the TME and ProBayes. The data acquisition was performed for four scenarios on a highway, with each scenario lasting for ten minutes approximately and the sensor data (stereo camera images, vehicle odometry, and GPS information) being recorded. The behaviors to be estimated are: move straight, a lane change to the left, and a lane change to the right. An example of the behavior estimation on a highway is shown in figure 9.

The detection of vehicles is performed by clustering of the disparity points obtained from the stereo camera mounted behind the windshield. The clustering is performed in the image areas, which are indicated by the image based detection using support vector machines (SVMs). The positions of vehicles are tracked on the road plane by means of the BOF [8], [9].

The observation variables for behavior recognition include

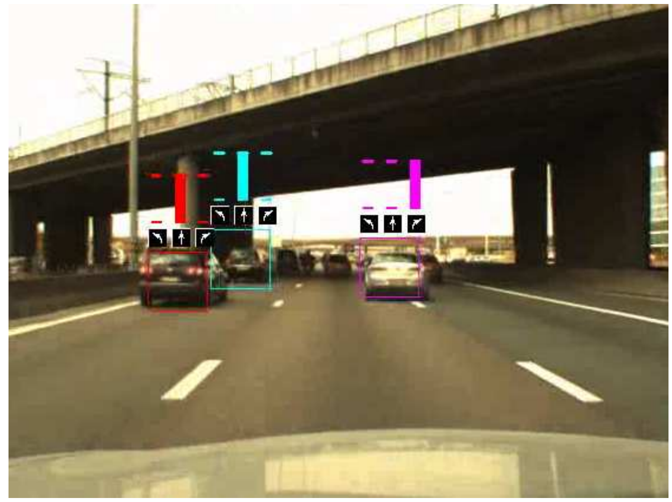

Figure 9: Example of a highway scenario where a vehicle on the middle lane performs a lane change to the right.

the vehicle's speed, the distances to the lane borders, and the information about the presence of other vehicles on the adjacent lanes. In order to obtain the observation variables in a global reference frame, a particle filter is used for localizing the vehicle on the highway map obtained from the Geographic Information System (GIS). The particle filter allows us to estimate the position and direction of the vehicle at each time instant and to employ the observations from stereo-vision (lanes detection), GPS and vehicle odometry. A similar approach is used for the training phase, when the acquired data are divided into the training and evaluation sets annotated manually to indicate the current behavior for each time instance of the data acquired.

An example of the behavior estimation on a highway is shown in Fig. 9. The positions of the tracked vehicles are projected onto the image plane and are represented by the rectangles. The probability distribution of the estimated behaviors is shown by the height of the color bars above the vehicles, e.g. the "lane change to the right" behavior of the vehicle on the middle lane and the "move straight" behavior of the two vehicles on the left lane are evaluated correctly. These results illustrate the validity of the proposed approach for behavior estimation. The different probability decomposition of the observation variables, the selection of the observation variables and the reactivity of the behavior estimation are topics of our ongoing work to generalize the approach.

The results on behavior estimation are currently preliminary. The current work involves experimenting with different probability decomposition of observation variables and observation variable selection. Furthermore, for purposes of risk evaluation, we will also be able to evaluate the reactivity of behavior estimation. The ongoing work will allow us to generalize the results and evaluate them quantitatively.

IV. Conclusion

Collision risk estimation and prediction will be mandatory for future vehicles. A fraction of a second of the driver's reaction time can help save human lives. Our data processing approach, sensor models and software modules allow us to monitor the urban traffic environment. The analysis and interpretation of traffic scenes rely on evaluation of driving behaviors as stochastic variables to estimate and predict collision risks for a short period ahead. Our initial experiments on behavior estimation during vehicle driving allowed us to verify the validity of the approach. Our future work will deal with its integration and evaluation on a Lexus vehicle equipped with sensors and shown in figure 10.

Figure 10: Experimental platform on a Lexus LS600h equipped with two IBEO Lux lidars, a TYZX stereo camera, and an Xsens MTi-G inertial sensor/GPS.

References

[1] I. E. Paromtchik, C. Laugier, M. Perrollaz, A. Nègre, M. Yong, and C. Tay, "The ArosDyn project: Robust analysis of dynamic scenes," in *Int. Conf. on Control, Automation, Robotics, and Vision*, (Singapore), December 2010.

[2] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437–459, 1976.

[3] J. Pierowicz, E. Jocoy, M. Lloyd, A. Bittner, and B. Pirson, "Intersection collision avoidance using its countermeasures," Tech. Rep. DOT HS 809 171, NHTSA, U.S. DOT, 2000.

[4] "Vehicle-based countermeasures for signal and stop sign violation," Tech. Rep. DOT HS 809 716, NHTSA, U.S. DOT, 2004.

[5] K. Fuerstenberg and J. Chen, "New European approach for intersection safety - results of the EC project INTERSAFE," in *Proc. International Forum on Advanced Microsystems for Automotive Application*, 2007.

[6] S. Lefèvre, J. Ibanez-Guzmán, and C. Laugier, "Context-based prediction of vehicle destination at road intersections," in *Proc. of the IEEE Symp. on Computational Intelligence in Vehicles and Transportation Systems*, (France), 2011.

[7] S. Lefèvre, J. Ibanez-Guzmán, and C. Laugier, "Exploiting map information for driver intention estimation at road intersections," in *Proc. of the IEEE Intelligent Vehicles Symp.*, (Germany), 2011.

[8] C. Coué, C. Pradalier, C. Laugier, T. Fraichard, and P. Bessière, "Bayesian occupancy filtering for multitarget tracking: An automotive application," *Int. J. Robotics Research*, no. 1, 2006.

[9] M. K. Tay, K. Mekhnacha, C. Chen, M. Yguel, and C. Laugier, "An efficient formulation of the Bayesian occupation filter for target tracking in dynamic environments," *Int. J. Autonomous Vehicles*, vol. 6, no. 1-2, pp. 155–171, 2008.

[10] C. Tay and C. Laugier, "Modelling smooth paths using Gaussian processes," in *Proc. of the Int. Conf. on Field and Service Robotics*, 2007.

[11] C. Tay, *Analysis of Dynamics Scenes: Application to Driving Assistance*. PhD Thesis, INRIA, 2009.

[12] C. Laugier *et al.*, "Vehicle or traffic control method and system," *Patent application no. 09169060.2-1264*, August 2009.

[13] B. Lévy, S. Petitjean, N. Ray, and J. Maillot, "Least squares conformal maps for automatic texture atlas generation," in *ACM SIGGRAPH conference proceedings*, 2002.