Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Risk estimation for driving support and behavior planning in intelligent vehicles

Julian Eggert

2018

Preprint:

This is an accepted article published in Automatisierungstechnik Sonderheft "The car becomes a robot". The final authenticated version is available online at: https://doi.org/10.1515/auto-2017-0132

Übersicht

Julian Eggert*

Risk estimation for driving support and behavior planning in intelligent vehicles

Risikoprädiktion für die Fahrerunterstützung und Verhaltensplanung in intelligenten Fahrzeugen

https://doi.org/10.1515/auto-2017-0132
Received December 7, 2017; accepted January 18, 2018

Abstract: Vehicles will be equipped with sensors and functions for highly automated driving in the foreseeable future. A big topic of research on the way to this goal is how to convey to these vehicles an understanding of the driving situations that is comparable to that of humans. For safe driving, this requires predicting how a scene will evolve and anticipating how dangerous it will potentially be. Risk estimation is a central ingredient in this process. In this paper, we describe how risk modeling frameworks help in managing the complexity of the driving task. We approach risk from the perspective of rare probabilistic events in environments where predictions might be inherently uncertain, and explain how this leads to a survival-based formulation which allows to model different types of risks encountered in driving situations within a single unified concept. In addition, we show how the framework can be used for driving behavior evaluation and risk-avoiding trajectory planning.

Keywords: Risk Modelling, Risk Maps, Foresighted/Predictive Driving, Highly Automated Driving, Survival Probability, Behavior and Trajectory Planning

Zusammenfassung: In naher Zukunft werden Fahrzeuge mit einer Vielzahl von Sensoren und Systemen für das hochautomatisierte Fahren ausgerüstet sein. Eine wichtige Forschungsfrage ist, wie diese Systeme ein Verständnis der Fahrsituationen erlangen können, welches mit dem von Menschen vergleichbar ist. Sicheres automatisiertes Fahren erfordert dafür eine verlässliche Risikoabschätzung, wobei prädiziert werden muss, wie sich eine Verkehrsszene entwickeln wird, und was das für das eigene Verhalten bedeutet. In diesem Artikel skizzieren wir Modelle und Systeme für eine Risikoabschätzung, die auf einem Ansatz von spärlichen probabilistischen Ereignissen und der Berechnung einer sogenannten „Überlebenswahrscheinlichkeit“ basieren. Der Ansatz eignet sich für eine einheitliche Erfassung von verschiedenen Arten von Risiken, und ermöglicht Anwendungen in der risikovermeidenden Fahrerunterstützung und der Trajektorienplanung bei intelligenten Fahrzeugen.

Schlagwörter: Risikomodellierung, Risikokarten, Vorausschauendes Fahren, Hochautomatisiertes Fahren, Überlebenswahrscheinlichkeit, Verhaltens- und Trajektorienplanung

1 Introduction

Advances in information technology, in form of technical components, computing capacity as well as better algorithms for data analysis, provide a broad basis for a radical change in the way we drive. Systems that for a limited time can help control either braking, acceleration, or steering, in form of, e. g., Adaptive Cruise Control (ACC), Lane-Keeping Assistance (LKA), Lane Departure Warning (LDW), Frontal Collision Warning (FCW) or Automatic Emergency Braking (AEB), complemented by sensor suites with enhanced environment analysis capabilities are now common in modern vehicles. More intelligent Advanced Driver Assistance Systems (ADAS) enable the vehicle to conduct some parts of the driving task under human supervision. However, it is the advent of highly automated and autonomous driving (AD) that promises to relieve the driver from most of the burdens and perils of the driving task.

For a better synchronization of the development efforts, regulatory bodies like the SAE (Society of Automotive Engineers) have introduced a taxonomy of automatization levels ranging from 0 to 5, with 0 meaning that the human driver does everything of the driving task, whereas at level 5, the automated system can perform all driving tasks under the same conditions as a human driver would [3].

***Corresponding author: Julian Eggert,** Honda Research Institute (HRI) Europe, Carl-Legien-Str. 33, 63073 Offenbach/Main, Germany, e-mail: julian.eggert@honda-ri.de

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

SAE level	0	1	2	3	4	5
Degree of automation	None	Assisted	Partial	Conditional	High (some modes)	Full
Examples	LDW FCW	ACC	Parking assist	Traffic jam chauffeur	Parking garage	Robot taxi
Monitoring	D	D	D	V	V	V
Execution	D	D+V	V	V	V	V
Fallback	D	D	D	D	V	V

Figure 1: SAE driving automatization levels. D: Driver, V: Vehicle.

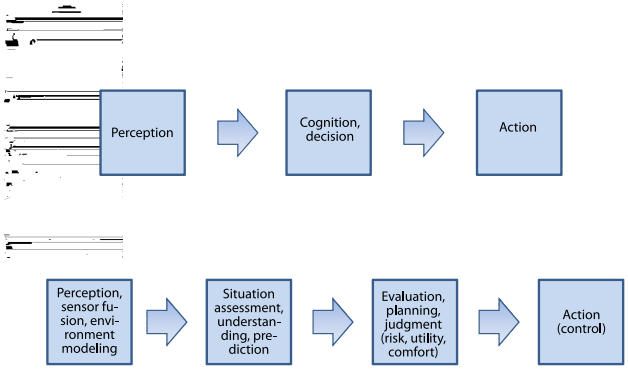

Figure 2: Architecture scheme for automatized driving.

See Fig. 1 for a short explanation. Starting already from SAE 2, but especially for SAE levels 4 and 5, the system also has to take decisions in critical situations. Vehicles can already get a detailed impression of their environment using modern sensors based on radar, lidar, infrared, cameras and ultrasonic. Environmental analysis capabilities like detection, classification and tracking of road users nearby can be tackled by modern algorithms from the artificial intelligence and machine learning domains and continue to be constantly improved, supported by an exponentially increasing amount of available data.

Nevertheless, understanding and interpreting the driving situation in terms of consequences for the vehicle and its passengers remains a challenging task. On the one side, this is because there is a huge (basically infinite) variability and complexity of the driving situation itself. On the other, it is very hard for the systems to anticipate what might happen, since the future evolution of a situation depends on numerous unknown factors like the intent of the other traffic participants, and since possible interactions between involved entities lead to high nonlinearities.

In future, similar to human drivers, systems for highly automated driving will have to process vast amounts of complex information about the environment to arrive at a human-like understanding of possible behaviors of the other traffic participants, as well as the implications of their behavior for the future evolution of the scene, and in particular, of the own state. In addition, this has to occur under consideration of the own action and interaction capabilities. Then, a decision has to be taken and executed on how to drive best with the flow.

From the requirements, this largely resembles architecture schemes as proposed for general robotics systems, which roughly follow a perception-cognition-action/sense-think-act scheme, as shown in top of Fig. 2. Currently, the focus of development is increasingly shifting to the "cognition" block, concentrating here first on the development of systems for understanding the driving scene (e. g., that the presence of children on a walkway might require special attention due to their unforseeable behavior, or that at occluded entrances to a road we have to slow down because of unexpectedly appearing cars) resp. predicting its future development, as shown in Fig. 2 bottom, and then on the evaluation/judgement of the relevance of the different possible outcomes of the scene evolution. Finally, with help of the judgment, a decision is taken and an appropriate action is selected and executed. We will come back to this scheme in section 3, where we explain a framework for behavior planning of intelligent vehicles in more detail.

One central ingredient in the evaluation/judgment block is risk. Safety and safe behavior support in all situations is claimed as the most important contribution and number one priority [30] of highly automated driving functions, with the potential of saving thousands of lives, reducing injuries, and saving follow-up costs. Other benefits like more comfortable and productive individual mobility time as well as more efficient overall traffic systems also play a role, but overall, safety should be increased in terms of less frequent and less severe accidents and accident consequences, since for around 93 % of crashes human behavior has been identified as the primary critical reason [28, 29]. Based on risk, different behavior options can be evaluated and compared with each other to arrive at the necessary behavior decisions. A proper risk estimation allows to detect safety-critical situations, to quantify their criticality, to communicate them to the driver or the system and to support behaviors which avoid or mitigate them.

Risk estimation serves to anticipate a critical driving situation before this situation actually occurs, such as a possible upcoming collision accident. If a risk indication occurs with sufficient temporal anticipation, it will then be used to warn, to act preventively, or to mitigate the occurrence of the situation as such (so that it may be avoided) or to alleviate the consequences of the situation, in case that it becomes unavoidable.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

Risk is a direct consequence of the driving context, embracing both the driver, its behavior and the state of its ego-vehicle, as well as the surrounding traffic participants and environment structures that have a relevant influence on the future possible developments of a driving scene. This can be seen directly by considering a car driven with low constant velocity along an empty road, resulting in nearly no risk at all. If the same driver with the same velocity now approaches an intersection with other cars crossing the first driver's trajectory, we suddenly encounter a situation with considerable risk. This means that, for risk assessment, we have to consider the hypothetical future spatiotemporal relationship between the ego entity and the surrounding world. Where you drive is of course important – it makes a difference if you practice driving skills on an empty parking lot as compared to a road passing a school – however, the entire set of dynamic variables from the ego-car as it is situated within the environment is crucial for risk estimation and prediction.

One of the big challenges of risk assessment is that the critical events which lead to accidents are extremely sparse (e. g., in Germany on the order of 0.1–5 accidents per 1 million driven km, [8]), and that the amount of factual counterexamples (e. g., where a very similar situation setting sometimes lead to accidents, and sometimes not) is huge. In addition, the variability as well as the number of variables that influence the outcome of a situation can also be very large and grows combinatorially with the complexity of the environment. This makes it hard to apply data-driven approaches, such as statistical and machine learning techniques, directly. A further challenge is that the involved processes which lead to a critical situation are inherently probabilistic, due to non-linearities (e. g., the deformations that take place during a physical crash process, or because of feedback and interaction loops), sensor inaccuracies (e. g., leading to false positive or false negative detections), unknown environment parameters (e. g., caused by occlusions, limited sensor range, or missing map data) as well as unobservable facts, like the driver's state of mind.

2 Modeling risk

Risk assessment inherently involves the consideration of the future: A situation is risky, if there is a hypothetical, dangerous event that might happen within a future time span. As such, it involves a prediction of the future evolution of a situation, and in addition, it is probabilistic, that is, it deals not with certainties, but with probabilities that a situation might evolve in a way which leads to quantifiable consequences.



Figure 3: Risk estimation: Prediction and survival, critical event probability, and event consequences.

In common-sense terms, risk is usually rather vaguely defined, dealing with a potential future gain or loss due to some choice of action. E.g., it is seen as the uncertain interaction with a potential, unpredictable, and uncontrollable outcome, with the risk itself being a "consequence of action taken despite this uncertainty" [40].

In more technical terms, especially in the context of risk management for technical and economical processes, risk is often defined as the probability that a gain or loss might occur, and this again is expressed more precisely as the consequences of a future event (in case of negative consequences, this term is sometimes called the "severity" from now on), associated with the probability that an event (identified as critical) will occur [17]. In short,

Event-Related Risk R_E = Probability of a Critical Event × Consequences of the Event

This is indicated schematically in Fig. 3. Starting from a time t, risk first involves a prediction of the future evolution of the scene, comprising the future states (and uncertainties) of all involved traffic entities. At any future point in time with predicted states, the probability of a critical event involving, e. g., the ego-car will depend on the probability that the ego-car has not been engaged in any event earlier (the so-called "survival probability") and a state-dependent event rate. Once a critical event (like an accident) happens, we use the predicted states to estimate the consequences of the event.

In driving-related research, however, the two terms are most often treated separately: On one side, the severity of critical events is usually handled in investigations which deal with the real modeling or the evaluation of the consequences of, e. g., accidents. On the other side, investigations about risk neglect the severity component or treat it as constant, concentrating on the event probability, like it is the case for collision risk of 2 entities by evaluation of the overlaps of their position probabilities. This might be justifiable in case that the conditions are sufficiently homogeneous for the severity to be neglected, as, e. g., in the case of similar cars involved in collisions with constant and comparable velocities. Considering, however, a risk-

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

based decision on adjusting gap sizes with respect to a standard car vs. a truck, or a slow passing car vs. a fast passing car, it becomes intuitively clear that risk prediction always involves the incorporation of both event probability and severity.

2.1 Event probability and severity

While being intuitively understandable, the risk definition from section 2 is not yet sufficiently precise to use it for models of quantitative risk prediction. Two issues arise for the involved terms. First, the event probability is conditioned on a particular event, but this event is only one of a large number of different events that can potentially happen. And second, it is necessary to consider possible consequences of an event in units that makes different types of risk comparable with each other, which can be accomplished by, e. g., linking potential damage for material and involved persons to their monetary value based on follow-up costs or insurance statistics.

The first term of definition from section 2 needs to be extended to encompass all possible **critical events** E. This has several dimensions, concretely:

1. All possible different types of risks that might arise during the time course of the scene evolution. Typical risk types are vehicle-to-X collision, loss of control (e. g., in curves) or rule violation.
2. All possible relevant entities that might get involved in any of the critical events. E. g., for collision risks, this implies the pairwise combination of traffic participants.
3. For a particular type of risk and a particular selection of involved traffic participants, different evolutions of the scene lead to different events in different situations.
4. For a particular situation, a countless number of critical events can happen at different predicted states such as positions, velocities, etc. Points 3 and 4 basically comprise the variability in the temporal evolution in the scene, with the first one of more discrete and the second one of more continuous character. In section 3, this distinction is used to separate the situation classification and scene prediction steps in our risk framework.

The risk R of the future time course of a driving scene is then given by combining the risk contributions from of all possible events, so that $R = \sum_E R_E$.

Regarding the critical events E, in a further step we need to incorporate the time more explicitly, since we are referring to scene evolution and hypothetical events that will occur in the future. Since predicted events may occur at any future time $t + s$ starting at t, we get

$$r(t+s|t)\,\Delta s = \sum_E p_E(t+s|t)\,\Delta s\; X_E(t+s)\ . \tag{1}$$

Here, $r(t + s|t)$ is the time-continuous risk rate that any type of critical event from the possible critical events E occurs around the future time (t+s) if we start predicting at the current timestep t. Similarly, p_E is the history-sensitive probability density for an event to occur at $t + s$ (number of events per time unit), given that we start our prediction at t, and where $p_E(t + s|t)\,\Delta s$ must be between $[0, 1]$. The consequences $X_E(t + s)$ might be positive or negative, for safety-critical events the unfavorable consequences are sometimes also called the **severity** of the event.

Looking at this last definition, we may start to understand risk as a weighted average of the consequences X_E. In fact, we regard $p(t + s|t) := \sum_E p_E(t + s|t)$ as a general probability density of *any* event to happen in the future time. Since events are assumed to be disruptive, after an event kicks in, the involved entity will not continue to move according to the original prediction, so that it is out of the game (which is understandable in case of a accident). This means that events are exclusive, and only one *or* another can happen, so that the total probability is then the or-ed combination of the probabilities of the single events. For $p(t+s|t)$ then to be a real probability density, it has to be normalized, which in this case means integration over the future time so that $\int_0^\infty \sum_E p_E(t + s|t)\,ds = 1$.

This normalization constraint has wide-reaching modeling consequences. Expressively, it means that along the (infinite) future time course, (a) there is no survival of the regarded entity on the long run, so that some event will *always* kick in and (b) event probabilities do not accumulate infinitely. Point (b) leads to a higher concentration of the event probability in the more closer time horizon for more dangerous scenarios. The normalization condition forces the consideration of the *history* of possible events for each predicted scene evolution, providing a discount on more distant future event probabilities $p_E(t + s|t)$ after high risk state combinations. This can be understood by regarding the *survival probability*

$$S(t+s|t) := 1 - \int_0^s \sum_E p_E(t+s'|t)\,ds'\ , \tag{2}$$

and considering that the history-dependent event probability $p_E(t + s|t)$ is a combination of the probability that a vehicle "survives" until $t + s$ and then gets engaged into

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

a critical event. As one would expect, the survival probability decreases monotonously (and faster for larger p_E's) from 1 to 0 for increasing future time horizon s ($S(t|t) = 1$ and $S(\infty|t) \to 0$).

With the event probability normalization, we arrive at new means of calculating and interpreting risk for driving situations. If we integrate the time-continuous risk equation for $r(t + s)$ over the (infinite) prediction time, for each moment in time t we get a single scalar risk value

$$R(t) := \int_0^\infty r(t + s|t)\, ds \tag{3}$$

so that the total risk becomes

R(t) = < X > = Total accumulated expected future consequences from critical events

for a given, predicted scene evolution (e. g., in terms of assumed trajectories and state evolutions of each involved traffic participant). Due to the survival property, this value will automatically give a higher weight to risks which are very proximal (in terms of prediction time s) as compared to risks which are further ahead, considering the history of critical encounters.

The detailed risk modeling now occurs via the history-sensitive probability density for critical events $p_E(t + s|t)$, and the consequence resp. severity function $X_E(t + s)$. In [10], it was shown that the future event probability $p_E(t + s|t)$ can be used to capture prediction uncertainties by consistently modeling it as an inhomogeneous Poisson process with renewal which depends on the history of event-relevant parameters (like the time-course of the distance between two cars for collision risks, or the curvature-dependent maximally allowed velocity for curve risks). The severity function can be modeled by the underlying physical event using the predicted parameters that are assumed at the time of the event.

The parameterization of $p_E(t + s|t)$ and $X_E(t + s)$ for simulations and real-world applications is an ongoing research effort based on two approaches: On one side, databases with reconstructions of individual critical situations are taken to optimize the behavior of risk-avoiding simulated vehicles. On the other side, for selected situations, macroscopic accident statistics can be derived from the risk models and compared to available data. This has been done e. g. in [11], where a fit between the curves provided a rough estimate of the probability density $p_E(t + s|t)$ for accident events in multilane roadways. Similarly, accident severity statistics from insurance companies, national institutions and clinical records provide a basis for the parameters of the severity functions $X_E(t + s)$.

2.2 Risk evaluation and risk maps

Traditionally, for collision accidents, conventional risk measures from the TTX type, like e. g. Time-To-Collision (TTC) [36] or the closely related Time-To-React (TTR) (see e. g. [18]) serve as scalar risk quantifiers. Here, X is any relevant event related to the accident, and the TTX quantifies the assumed time that remains until the critical event occurs. It is assumed that a larger TTX implies a larger time remaining to avoid an accident, with the consequence of a lower event probability for larger TTX. Drawback of TTX measures are that (i) they can only address special situations (2 cars, longitudinal scenarios, slower preceding car), (ii) non-collision cases lead to zero risk and (iii) that it is missing severity considerations. Several extensions have been introduced, such as Time-Exposed-TTC (TET) or Time-Integrated-TTC (TIT) [27], non-longitudinal extensions [38], or Kamm's circle type positional uncertainty [37]. Typically, TTX type risk measures rely on a constant velocity assumption for the prediction of trajectories, but this condition can be and is relaxed in the TTX extensions.

Numerous other approaches for collision risk prediction exist, e. g., based on analyzing geometrically the predicted trajectories and estimating their intersection points. Among these are occupancy grid methods (checking if the two involved traffic participants will occupy the same spatial grid cell in some future moment) or probabilistic collision prediction methods (relying, e. g., on Monte-Carlo simulations or based on Gaussian probability distributions for the vehicle positions). One example of a modern extension of TTC is given by [33, 34], where a Time-To-Critical-Collision-Probability (TTCCP) is determined as the time when a collision probability exceeds for the first time a critical threshold.

Current risk indicator approaches are, however, limited in the type of risks, usually concentrating, e. g., on collision risks, in their specialization for a narrow set of situations, and in the number of events resp. road users (usually only one risk type and only other road user) they can handle. In addition, the estimated risk is dominated by the first critical event that can be possibly encountered, so that a sequence (history) of events is neglected.

The survival risk concept is compatible with the previous approaches, providing a basis for a generic and probabilistic treatment of risk and risk-based behavior evaluation in driving situations including uncertainties. E. g.,

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

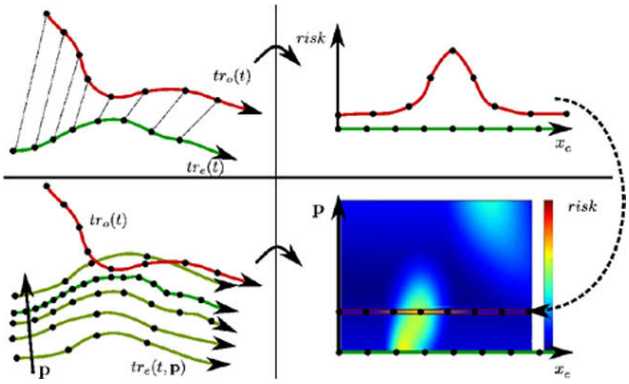

Figure 4: Trajectory prediction and dynamic risk maps, collision risk example (see [5, 6] for more details). TL: The time course of the ego and other trajectories is compared. TR: The future risk function $r(t + s|t)$ (here plotted over driven distance) is predicted. BL, BR: The risk functions for different ego trajectory options are calculated and integrated into a *risk map*. Ego-car longitudinal distance x_e, trajectory tr_e, trajectory parameters (e. g., velocity) **p**, other vehicle trajectory tr_o.

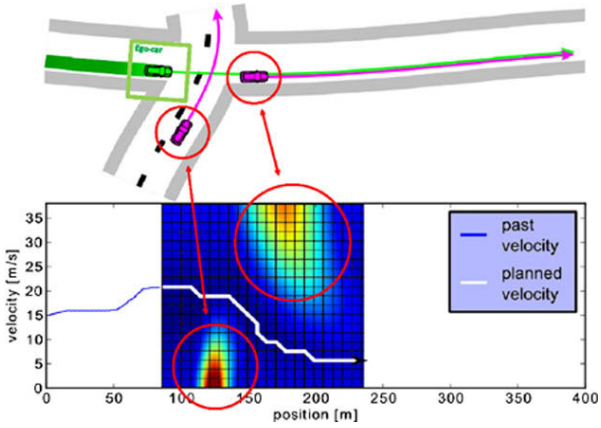

Figure 5: A dynamic *risk map* for risk prediction and behavior planning. Top: A situation with an ego car (green) and two other vehicles conveying collision risks (longitudinal and crossing). Bottom: A plot of the risk function $r(t + s|t)$ over future position vs. ego-car velocity. The two colored peaks indicate the risks caused by the two other vehicles. The ego-car has to accelerate to avoid the crossing vehicle by passing first and then decelerate before reaching the front vehicle.

the typical increase of event probability with decreasing time TTX remaining until a critical event happens can be derived from the properties of the survival probability. In addition, the survival risk concept exhibits several desirable properties such that it can handle multiple events and multiple risk types in a single risk quantifier, that it considers the history and sequential effects of risk consistently (potential risks that are predicted to occur before others are considered as more important), and that it includes behavioral uncertainty in the sense that risks, which are still further away in the future are considered as less likely to kick-in, because there is more time for the driver to react in a way that could lead to an avoidance of the critical event. In addition, it incorporates severity as an integral part of the risk estimation from the start. In [14], the presented survival-based probabilistic risk approach was compared to extensions of TTC and Gaussian probability based risk measures for the case of collision risks in recorded accident trajectories, showing beneficial properties in terms of earlier risk detection in case of real collision events and less false positives in case of near-miss and uncritical conditions.

Each hypothetical future scene evolution can now be analyzed in terms of its potential risk $R(t)$. Given several behavior alternatives for the involved traffic participants, e. g., in form of trajectory alternatives, $R(t)$ can then be used on each alternative scene evolution based on predicted risk. In detail, the $r(t + s|t)$ curve can be used to analyze the future time course of the risk involved in the development of a situation, as shown in Fig. 4. Varying an acting entities' possible behaviors, we get different curves $r(t+s|t)$ which can be used to provide a landscape of potential risk as a function of future time vs. behavioral parameters. These predictive "risk maps" indicate how critical (i. e., how probable and severe) a certain behavior will be in the future, and under which conditions (where, when). Spots of high risks can then be avoided.

Figure 5 shows an example of a risk map for a crossing situation involving the ego-car (green) and two other vehicles (red). Each other vehicle gives rise to a risk peak in the risk map; the crossing vehicle at low ego-car velocities and the preceding vehicle at high ego-car velocities. The white curve in the risk map shows a planned velocity profile which avoids both risk peaks by first maintaining speed high to avoid the crossing vehicle and then decelerating to avoid a crash into the preceding vehicle. Figure 6 shows a simulated intersection scenario with corresponding risk-based behavior planning. The ego car approaches an intersection and has to adjust to the risks conveyed by other vehicles numbered (1)–(4) as well as to the risk (5) conveyed by increased speed at high curvature. Red areas in the risk maps indicate the different risk peaks which have to be avoided. A way to find appropriate risk-avoiding behaviors from the risk map is, e. g., by using stochastic optimization processes, in this case, a variant of Rapidly-Exploring Random Trees (RRT), see [25, 19, 20] for more details. Different red lines represent possible behaviors for which the integral risk $R(t)$ is calculated until a trajectory of low risk is found.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

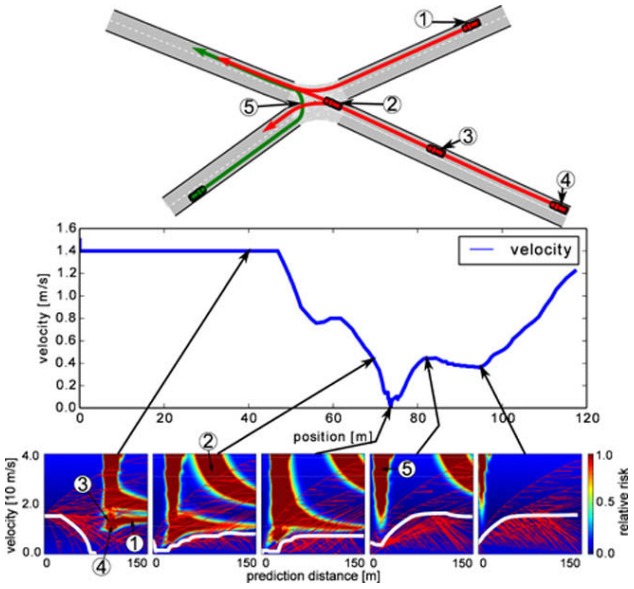

Figure 6: Risk based behavior planning example. Top: Complex intersection scenario, with the ego-car in green and other users (1)–(4) and curve (5). Bottom: Temporal snapshots of the risk map comprising curvature risk (5) and the collision risk from cars (1)–(4) (red strokes on the risk map show a stochastic search method used to find the risk-avoiding behavior, white line indicates the proposed trajectory). Middle: Resulting ego-car behavior, accelerating and decelerating to time for the right gap. Adapted from [6].

3 Driving behavior support based on risk assessment

3.1 A risk assessment framework

Different to current Advanced Driving Assistance Systems (ADAS) which have been developed and validated for relatively specific functional domains, systems designed for AD as well as future ADAS will need to cover more complex situations such as mixed longitudinal and lateral behavior support while overtaking on highway in the presence of multiple other traffic participants or a multilane crossing with combined frontal and intersecting urban traffic. In the SAE's international definition of automated driving systems [3], level 4 and 5 systems even need to address the fully dynamic driving task including the ability to bring the vehicle to a minimal risk condition like a safe stop, without the expectation that a human driver has to take over [39]. To achieve this, future driving systems need to have a deep contextual situation understanding comparable to that of humans, which foresees the consequences of the available behavioral options to decide on the best situation support in terms of expected risk and utility.

3.2 Framework components

The framework for risk-based driving support consists of several components which roughly follow the perception-decision-action scheme from section 1. In particular, a driving system with assistance and autonomous functions that is scalable to general-purpose situations comprises a series of modules which can be roughly systematized as follows, see e. g. [13] and [39]:

- **Scene observation: Where am I and what is around me?** At this step, the road itself as well as objects on the road are detected, classified and their states (e. g., speed, heading, acceleration) are estimated, map information is extracted and the ego-car is located relative to the map. This helps to identify different road users (e. g., pedestrians, cyclists, vehicles of different kinds), road elements (e. g., traffic signs, traffic signals, road markings), driving-related road structures and driving paths. The results may be integrated in an environment representation scheme which allows to construct a consistent real-time view of the world accessible by the other framework components. This representation then constitutes the basis for a semantic understanding of the situation the vehicle is in, including the type of objects and their meaning for us, their likely behaviors, and the consequences for our behavior. As an example, the detection of a preceding car that is driving on an adjacent lane with blinker on might imply with high probability that we are in a situation where a lane change will be initiated.
- **Situation classification: What does the current driving scene mean to me?** A situation is defined as a stereotypical pattern of how a driving scene evolves, such as the pattern of a car entering from right in front of the ego-car into an intersection. A classification of the current scenario into different situation hypotheses is done from the ego-vehicle perspective, based on context information and the interaction between the traffic participants. Because of the prediction uncertainty, and depending on the behavioral alternatives of the traffic participants, each driving scene might give rise to a combinatorial number of possible situations, with each situation representing a different way how the scene might evolve. Situation classification can go along with a holistic estimation of the involved risks and benefits (as gained from empirical statistics), and its main purpose is the reduction of situation hypotheses. The result is a subset of situation hypotheses applicable to the current driving scene which are considered relevant for the ego-car behavior, together with the likelihoods of each hy-

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

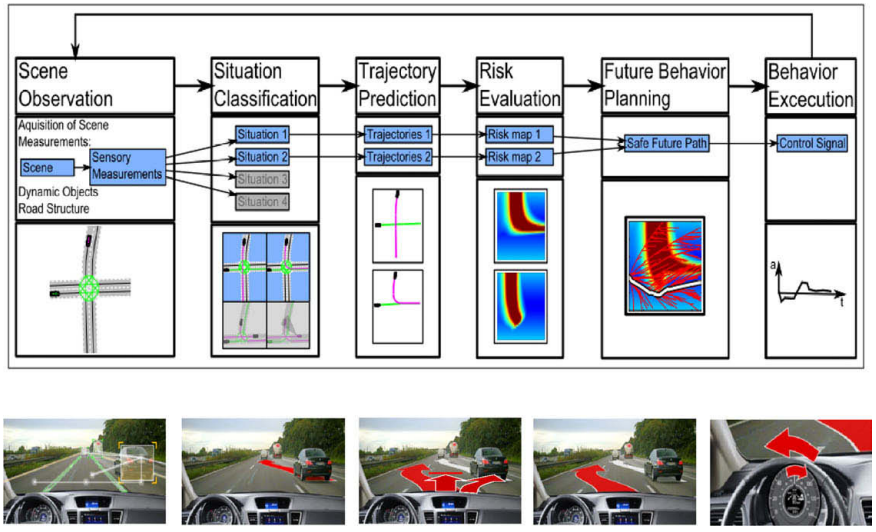

Figure 7: Risk estimation and behavior planning framework. Top: A schematic diagram of the involved components, following [13]. Bottom: Explanation of the functions from the drivers perspective.

pothesis to occur. See e. g. [22, 23, 9] and several others.

- **Scene prediction: What will happen next?** The future development of each of the selected situation hypotheses resp. the states of the involved expected traffic participants is extrapolated. This implies modeling how different road users are likely to behave, including their intent. Different levels of prediction accuracy can be explored, starting with purely kinematic extrapolation models up to models which take constraints on road structures and motion paths as well as interactions between the road users into account, see [26] for a review. In this step, it is, e. g., modeled that vehicles will mainly move along the road and that pedestrians may behave in a more erratic, direction-changing way than other road users. For each situation, different spatial paths with different spatiotemporal trajectories are modeled to incorporate behavior variability.
- **Risk evaluation: Assessment of behavioral risks and benefits** For each situation hypothesis and for a variety of own future behaviors (eventually in form of trajectories), the expected risks and benefits are estimated. The risk evaluation incorporates all the information from the scene prediction, playing through different behavior alternatives, and estimating the probability and severity of critical events of different types, such as collision and loss of control. Different risk and safety measures can be applied here, including the one presented in the previous sections.
- **Future behavior planning: What should I do?** For each situation hypothesis, from the variety of ego-behaviors, we get a risk landscape (*risk map*) that can be used to identify future points of high criticality. Retrieving the most appropriate behavior for our vehicle then implies finding the trajectory that minimizes accumulated risk by avoiding the future points of high criticality while keeping benefits high. Benefits in this case are expressed in terms of travel gain and comfort maximization. In this step, the results from the different applicable situations are integrated and a resulting target behavior is calculated.
- **Behavior execution: How should I do?** In this last step, the target behavior is used either for evaluation of/comparison with the current driving behavior for ADAS support and warning functions, or used as a control signal to execute the corresponding actuations in case of autonomous driving. See for example [2] for a review.

In figure 7, the architecture components of the risk estimation framework are shown together with an internal view from an operating car. Central elements of the risk estimation framework like situation understanding and behavior prediction already appear in modern ADAS functions, like in the intelligent Adaptive Cruise Control (i-ACC) [32, 21], which allows the system to foresee a lane change of an adjacent car even before a lateral movement of the lane-changing car occurs, leading to a more anticipatory and therefore smoother reaction.

In general, the target of such risk estimation frameworks is a longer prediction time horizon and a better concentration of resources on safety-relevant behaviors. Longer prediction time horizons enable more strategical

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

planning of trajectories (like delayed lane change or overtaking maneuvers), however at the cost of having to deal with the increase of prediction and behavioral uncertainties.

The framework is amenable for the incorporation of inherent risks by missing, unreachable or unknown information. In a heterogenous driving environment composed of cars with and without AD capabilities, risk estimation that is able to deal with incomplete information becomes necessary. Examples for such risks are given, e. g., when the own car is overtaking with limited view to the front, or when it has to operate with sensor occlusions, as present in inner-city intersections.

Figure 8 shows an example of such a risk-based behavior planning with explicit consideration of sensory limitations. In inner city crossing traffic, limitation of the field of view by buildings is a major source of uncertainty: We simply don't know if a car will appear from an occluded area of an intersection. Nevertheless, the "virtual" cars that might appear at the occluded regions can be included into the prediction component for risk estimation. In Fig. 8, top, a satellite view of an intersection with the ego car (green), and its viewing range is indicated (yellow lines). Below that, simulation views and risk maps are shown for 3 consecutive timesteps. At timestep 1, the predicted risk by the hypothetical red car causes the ego-car to slow down (decreasing yellow curve). At timestep 2, when the insight into the intersection has increased sufficiently, the ego-car foresees that it can also pass "in front of" the hypothetical risk (rising yellow curve). At timestep 3, it accelerates back to cruising speed because the insight into the intersection is such that it can be estimated that a hypothetically appearing car will not get in time to the intersection any more to pose a threat to the ego-car. The purple curve is a velocity profile measured from human driving behavior, showing that humans also break preemptively and accelerate again when the risk decreases with increased insight.

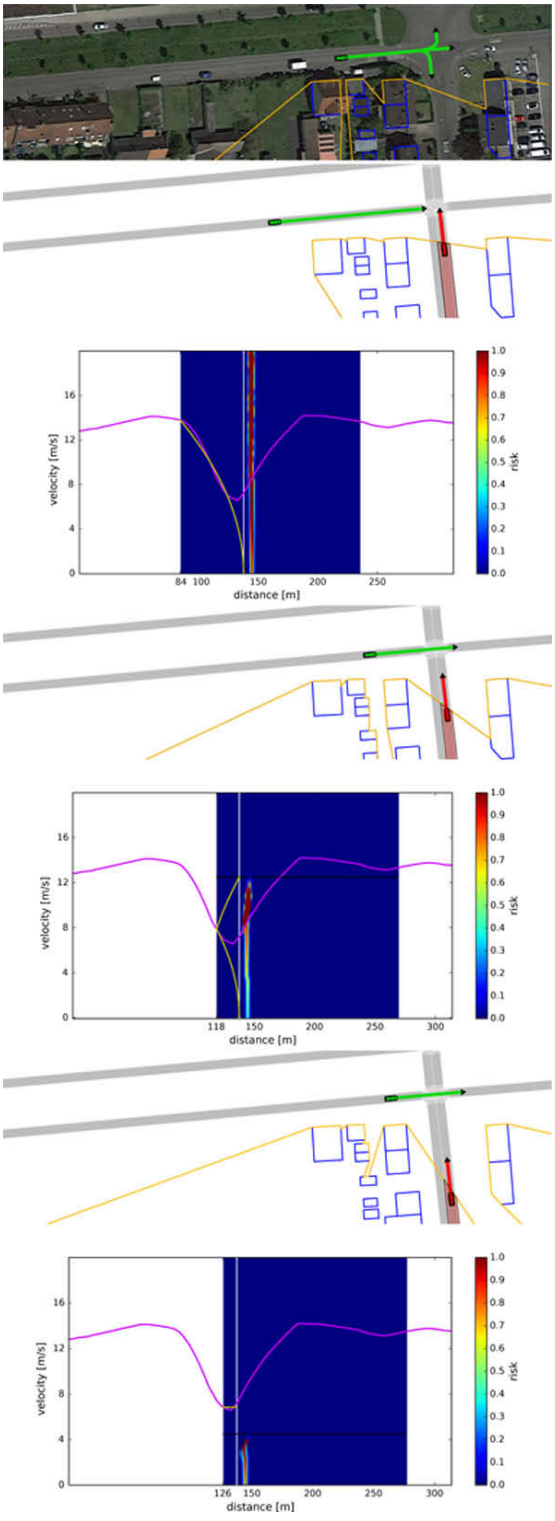

Figure 8: Risk estimation using incomplete information and "virtual" cars as risk sources at occluded intersections (adapted from [7]). See text for the explanation.

4 Discussion and conclusions

Since safety is a major concern in ADAS and AD, behavior planning based on predicted risk under uncertainty is a cornerstone for future intelligent driving systems. We presented current research that is making progress in this field, and a general framework which leads the direction towards systems that start to get an understanding of the behavioral consequences of driving scenes, based on risk. The target is to account and generalize for the various types of encountered risks with a unified approach and within a single system.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

4.1 Interactions with other traffic participants and motion planning

Interactions between the ego-car and other traffic participants occur at two levels. On a first level, different interaction patterns are clustered into different situations at the situation classification step from section 3.2. That is, we assume a discrete set of behaviors and behavior combinations and accordingly, situation classes, and for each situation class a prototypical trajectory of each traffic participant. E. g., at an intersection with 2 vehicles, we could use 3 situation classes which exhibit behavioral consequences for the ego-car: First, the other vehicle accelerates and passes the intersection before the ego-vehicle does; second, the other vehicle brakes to let the ego-car pass, and third, the other vehicle is unaware and continues driving as if the ego-car was not present. As we see, the first two classes involve an reaction of the other vehicle whereas the third class does not. For work on interaction-aware situation classification see, e. g., [23, 4, 24]. In [23], a framework for an interaction-aware situation classification system is introduced and in [4], it has been shown how the presence or lack of interaction at intersections can be detected several seconds in advance in recorded accident cases.

For each situation class identified as relevant for the ego-car and applicable to the current scene, prototypical trajectories of all involved traffic participants are then predicted, including the ego-vehicle. This is the second level where interactions between the ego-car and other traffic participants are incorporated. For this purpose, the ego-trajectory is kept fixed and simplified, interaction-aware prediction models are used for the other vehicles, such as modified versions of the Intelligent Driver Model (IDM) [35] or the Foresighted Driver Model (FDM) [12]. The resulting prototypical trajectories then constitute, together with the variation of the ego-vehicle trajectory, the basis for the calculation of the time-continuous risk rate $r(t + s|t)$, the total risk $R(t)$ from section 2.1 and the risk maps from section 2.2.

4.2 Relation to motion planning techniques

The increased research efforts during the last years have led to a substantial progress in the aspect of trajectory planning, see e. g. [31] and [16] for more comprehensive reviews on this topic. The focus of this contribution is not on the algorithmic details of the trajectory optimization, but on the consistent and generalizable incorporation of risk measures and risk-related uncertainties into the behavior evaluation.

Similar to several other approaches, if used for trajectory planning decision hierarchies, the presented framework shares the coarse subdivision into a behavioral layer, a motion planning layer and a local control layer (see for example [31]). In Fig. 7, the behavioral layer corresponds to the situation classification component, whereas the motion planning layer is based on the trajectory prediction, risk evaluation and future behavior planning components. From the trajectory finding methodology, as "sampling-based" planners we use variants of the RRT family of algorithms, which are applicable for fast and high-dimensional online path planning. This allows to consider non-holonomic constraints, but has the disadvantage of delivering jerky and not curvature-continuous trajectories, so that in the local control layer further smoothening becomes necessary. However, if the predicted trajectories are represented in a parametric way using, e. g., splines or polynomials, direct optimization techniques can also be used in combination with the risk evaluation.

Future key research aspects of motion planning techniques are the holistic treatment of all involved risks, the consideration of perception limitations, as well as the consideration of uncertainty in measurements and in the prediction of future behavior. The survival-based risk estimation provides some entry points on dealing with these aspects (risk treatment and prediction uncertainty), but others (like the perception limitations) require efforts in several steps of the behavior support framework from Fig. 7.

4.3 Balancing risk against utility

We have so far concentrated mainly on behavior evaluation and behavior support based on risk prediction. However, beyond risk, a series of additional considerations have to be taken for behavior planning. On one side, we have the drivability constraints from the vehicle: The solutions have to be sufficiently close to physically drivable trajectories so that they can be supported or executed with suitable control methods. This is usually implemented by incorporating explicit ego-vehicle and ego-driver models. In the presented framework, this occurs at the trajectory prediction level and during the stochastic search of trajectory solutions in the risk maps. On the other side, the solutions have to provide maximal driving return in terms of utility such as, e. g., travel gain, comfort, naturalness, as well as adherence to rules, regulations and "normal" behavior.

Considering only the target of risk minimization from section 2, we have so far not detailed the utility aspect.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

Different to the risk estimation, where disruptive events lead to a deviation from the anticipated trajectory, the utility is a continuous additional cost factor which accumulates over time. Without utility, many recommended solutions would consist in a full stop, because of the low risk involved, which certainly is not a desired behavior. It is important to understand that under the uncertainties involved in driver and environment prediction, we always have to consider a *risk vs. utility tradeoff*, since there is no motion that is completely free of risk. If the risk estimation model is to conservative, the utility effects vanish, see for example [18]. The remaining question is the extent to which a driver (or a legal framework) is willing to trade-in an increase in utility by an acceptable, low amount of risk.

The utility of a predicted trajectory accumulates during the time for which the ego-vehicle is still assumed to be on track, i. e., not affected by disruptive, critical events of any kind. In section 2.1, this was expressed by the survival probability eq. 2. The utility of a predicted trajectory is therefore calculated as a sum of all utility contributions $u_T(t+s)$ weighted by the survival probability $S(t+s|t)$:

$$U(t) := \int_0^\infty S(t+s|t) \sum_T u_T(t+s)\, ds \tag{4}$$

Each term u_T accounts for a specific utility type with the dimensions of a utility rate (i. e., utility per time interval). As an example, since one of the most important considerations is whether the trajectory takes us forward towards our goal, a "transportation utility" u_0 favors larger traveled distances, e. g. given by $u_0(t+s) \sim v(t+s)$ (with the longitudinal "moving forward" velocity $v(t+s)$ of the predicted ego-vehicle trajectory). In similar way, we can favor smooth, jerk-free driving, obeying traffic rules and driving on the right side, while penalizing high accelerations, deviations from a desired cruising speed and violation of regulations.

We then arrive at a behavior evaluation based on risk vs. utility costs

$$R(t) + U(t)\ , \tag{5}$$

where during normal, risk-free driving the utility terms dominate but before risky conditions arise the risk terms kick in rapidly. To make the costs comparable, we can base them on a common scale, requiring a monetarization of all the involved hypothetical consequences X_E and the utility rates u_T. This is currently a topic of interest especially for the legal regulation and insurance sectors. Knowing (and recording) the potential overall driving behavior including the driver's propensity to take risks is an important step towards regulated autonomous driving, compare e. g. the recently issued guidelines of the ethics commission for autonomous driving at SAE levels 4 and 5 in Germany [15] and the derived legal consequences [1].

Once we have terms for the total costs, behavior planning mainly amounts to finding the trajectories which optimize them. Nevertheless, the issue of how to weight the different contributions by the severity-dependent cost terms remains open, and is rather an ethical and legal debate. From the ethics guidelines [15], things to consider are that material damage should always be preferred over personal damage, that there should not be a selection or discrimination of any kind of involved participants, that it is not allowed to sacrifice other participants and that there should not be an unfair distribution of responsibilities. In terms of risk cost modeling, this has direct consequences in the sense that human injury should have much higher costs than any material damage, that the costs of other involved traffic participants also have to be considered, and that the involvement of bystanders should be avoided. Addressing these issues explicitly in the risk modeling stage would then contribute to the transparency of the decisions underlying future ADAS and AD systems.

Acknowledgment: This paper is an overview of ongoing efforts to develop an all-situation risk estimation theory for autonomous driving at HRI, and would not have been possible without the intensive support and contributions of students and coworkers. The author is especially and greatly indebted to Dr. Florian Damerow, Stefan Klingelschmitt, and Tim Puphal, who contributed with ideas, discussions and material.

References

1. BMVI. Maßnahmenplan der Bundesregierung zum Bericht der Ethik-Kommission Automatisiertes und Vernetztes Fahren. Technical report, Die Bundesregierung, 2017. www.bmvi.de/SharedDocs/DE/Publikationen/DG/massnahmenplan-zum-bericht-der-ethikkommission-avf.pdf?__blob=publicationFile.
2. Ashwin Carvalho, Stéphanie Lefévre, Georg Schildbach, Jason Kong, and Francesco Borrelli. Automated driving: The role of forecasts and uncertainty – a control perspective. *European Journal of Control*, 24(Supplement C):14–32, 2015.
3. On-Road Automated Driving (Orad) Committee. Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems. Technical report, Society of Automotive Engineers (SAE), 2014. standards.sae.org/j3016_201401/.
4. Florian Damerow and Julian Eggert. Spatio-temporal trajectory similarity and its application to situation classification and the evaluation of lack of interaction. In *International Conference*

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

on Intelligent Transportation Systems (ITSC), pages 2512–2519. IEEE, 2016.
5. Florian Damerow and Julian Eggert. Predictive risk maps. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 703–710. IEEE, 2014.
6. Florian Damerow and Julian Eggert. Balancing risk against utility: Behavior planning using predictive risk maps. In *Intelligent Vehicles Symposium (IV)*, pages 857–864. IEEE, 2015.
7. Florian Damerow, Tim Puphal, and Julian Eggert. Risk-based driver assistance for approaching intersections of limited visibility. In *International Conference on Vehicular Electronics and Safety (ICVES)*, pages 178–184. IEEE, 2017.
8. International Traffic Safety Data and Paris Analysis Group. Road safety annual report 2015. Technical report, OECD/ITF (International Transport Forum), 2015. Retrieved 2016-01-29, data from 2013.
9. Thao Dang, Gabi Breuel, Andreas Tamke, Andreas Wedel, Wolfgang Rosenstiel, Dietmar Kasper, and Galia Weidl. Object-oriented bayesian networks for detection of lane change maneuvers. *IEEE Intelligent Transportation Systems Magazine*, 4(3):19–31, 2012.
10. Julian Eggert. Predictive risk estimation for intelligent ADAS functions. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 711–718. IEEE, 2014.
11. Julian Eggert. Solomon curve 2020: Relating microscopic risk models with accident statistics. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 2293–2300. IEEE, 2016.
12. Julian Eggert, Florian Damerow, and Stefan Klingelschmitt. The foresighted driver model. *2015 IEEE Intelligent Vehicles Symposium (IV)*, pages 322–329, 2015.
13. Julian Eggert, Stefan Klingelschmitt, and Florian Damerow. The foresighted driver: Future ADAS based on generalized predictive risk estimation. In *FAST-zero Symposium, Gothenburg, Sweden*, 2015.
14. Julian Eggert and Tim Puphal. Continuous risk measures for adas and ad. In *FAST-zero Symposium, Nara, Japan*, 2017.
15. Ethik-Kommission. Automatisiertes und Vernetztes Fahren. Technical report, Bundesministerium für Verkehr und digitale Infrastruktur, 2017. www.bmvi.de/bericht-ethikkommision.
16. D. González, J. Pérez, V. Milanés, and F. Nashashibi. A review of motion planning techniques for automated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 17(4):1135–1145, April 2016.
17. ISO31000. Risk management – principles and guidelines. *International Organization for Standardization, Geneva, Switzerland*, 2009.
18. Kristian Kroschel, Jörg Hillenbrand, and Andreas M. Spieker. A multilevel collision mitigation approach – its situation assessment, decision making, and performance tradeoffs. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):528–540, 2006.
19. Sertac Karaman and Emilio Frazzoli. Optimal kinodynamic motion planning using incremental sampling-based methods. In *Proceedings of the 49th IEEE Conference on Decision and Control, CDC 2010, Atlanta, Georgia, USA, December 15–17, 2010*, pages 7681–7687, 2010.
20. Sertac Karaman, Matthew R. Walter, Alejandro Perez, Emilio Frazzoli, and Seth J. Teller. Anytime motion planning using the RRT. In *IEEE International Conference on Robotics and Automation, ICRA 2011, Shanghai, China, 9–13 May 2011*, pages 1478–1483, 2011.
21. Marcus Kleinehagenbrock, Morimichi Nishigaki, Robert Kastner, Jens Schmüdderich, Sven Rebhan, Thomas Weisswange, Hiroyuki Kamiya, Shunsuke Kusuhara, Naoki Mori, and Shinnosuke Ishida. Introduction of Intelligent Adaptive Cruise Control (i-ACC) – a predictive safety system. In *FAST-zero Symposium, Gothenburg, Sweden*, 2015.
22. Stefan Klingelschmitt, Florian Damerow, and Julian Eggert. Managing the complexity of inner-city scenes: An efficient situation hypotheses selection scheme. In *Intelligent Vehicles Symposium (IV)*, pages 1232–1239, 2015.
23. Stefan Klingelschmitt, Florian Damerow, Volker Willert, and Julian Eggert. Probabilistic situation assessment framework for multiple, interacting traffic participants in generic traffic scenes. In *Intelligent Vehicles Symposium (IV)*, pages 1141–1148, 2016.
24. Stefan Klingelschmitt, Volker Willert, and Julian Eggert. Probabilistic, discriminative maneuver estimation in generic traffic scenes using pairwise probability coupling. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 1269–1276. IEEE, 2016.
25. Steven M. Lavalle. Rapidly-exploring random trees: A new tool for path planning. Technical report, Computer Science Dept., Iowa State University, 1998.
26. Stéphanie Lefèvre, Dizan Vasquez, and Christian Laugier. A survey on motion prediction and risk assessment for intelligent vehicles. *ROBOMECH Journal*, 1(1):1, 2014.
27. Michiel M. Minderhoud and Piet H. L. Bovy. Extended time-to-collision measures for road traffic safety assessment. In *Accid. Anal. Prev.*, volume 33, pages 89–97, 2014.
28. US Department of Transportation. National motor vehicle crash causation survey. Technical report, NHTSA, 2008. crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811059.
29. US Department of Transportation. Critical reasons for crashes investigated in the national motor vehicle crash causation survey. Technical report, NHTSA, 2015. crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115.
30. US Department of Transportation. Automated driving systems (ADS): A vision for safety 2.0. Technical report, NHTSA, 2017. www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf.
31. Brian Paden, Michal Cap, Sze Zheng Yong, Dmitry Yershov, and Emilio Frazzoli. A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles. *Intelligent Vehicles Journal*, 1:33–55, 2016.
32. Jens Schmüdderich, Sven Rebhan, Thomas Weisswange, Marcus Kleinehagenbrock, Robert Kastner, Morimichi Nishigaki, Shunsuke Kusuhara, Hiroyuki Kamiya, Naoki Mori, and Shinnosuke Ishida. A novel approach to driver behavior prediction using scene context and physical evidence for intelligent Adaptive Cruise Control (i-ACC). In *FAST-zero Symposium, Gothenburg, Sweden*, 2015.
33. Matthias Schreier, Volker Willert, and Jürgen Adamy. Bayesian, maneuver-based, long-term trajectory prediction and criticality assessment for driver assistance systems. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 703–710. IEEE, 2014.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM

34. Matthias Schreier, Volker Willert, and Jürgen Adamy. An integrated approach to maneuver-based trajectory prediction and criticality assessment in arbitrary road environments. *IEEE Transactions on Intelligent Transportation Systems*, 17:2751–2766, 2016.
35. Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Phys. Rev. E*, 62:1805–1824, 2000.
36. A. van der Horst. Time-to-Collision as a cue for decision making in braking, In *Vision in Vehicles III*, pages 19–26. Elsevier, 1991.
37. Walther Wachenfeld, Philipp Junietz, Raphael Wenzel, and Hermann Winner. The Worst-Time-To-Collision Metric for Situation Identification. In *Intelligent Vehicles Symposium (IV)*, 2016.
38. James Ward, Gabriel Agamennoni, Stewart Worrall, and Eduardo Nebot. Vehicle Collision Probability Calculation for General Traffic Scenarios Under Uncertainty. In *Intelligent Vehicles Symposium (IV)*, pages 986–992, 2014.
39. Waymo. Waymo safety report: On the road to fully self-driving. Technical report, Waymo, 2017. storage.googleapis.com/sdc-prod/v1/safety-report/waymo-safety-report-2017.pdf.
40. Wikipedia. Risk – Wikipedia, the free encyclopedia, 2014.

Bionotes

Julian Eggert
Honda Research Institute (HRI) Europe, Carl-Legien-Str. 33,
63073 Offenbach/Main, Germany
julian.eggert@honda-ri.de

Dr. rer. nat. Julian Eggert is Chief Scientist at the Honda Research Institute Europe. His main research is on environment perception, situation analysis and knowledge-based prediction and reasoning for artificial cognitive systems.

Authenticated | julian.eggert@honda-ri.de author's copy
Download Date | 2/13/18 3:53 PM