Collision Avoidance: a Literature Review on Threat-Assessment Techniques

John Dahl, Gabriel Rodrigues de Campos, Claes Olsson, and Jonas Fredriksson

***Abstract*—For the last few decades, a lot of attention has been given to intelligent vehicle systems, and in particular to automated safety and collision avoidance solutions. In this paper, we present a literature review and analysis of threat-assessment methods used for collision avoidance. We will cover algorithms that are based on single-behavior threat metrics, optimization methods, formal methods, probabilistic frameworks and data driven approaches, i.e., machine learning. The different theoretical algorithms are finally discussed in terms of computational complexity, robustness and most suited applications.**

***Index Terms*—Threat-assessment algorithms, decision-making methods, intelligent vehicles, active-safety systems,**

I. Introduction

IN order to help drivers deal with complex traffic situations, several types of automated safety systems have been developed during the last decades. These Advanced Driver Assistance Systems (ADAS) have revolutionized the automotive industry and better solutions are continuously being introduced in new passenger cars. Leveraging arrays of sensors that detect nearby objects and obstacles, these features can keep the vehicle within the lane, maintain a safe distance, and even avoid collisions in several critical situations. They are also included in the safety rating procedures of different national New Car Assessment Programmes (NCAP) in a effort to encourage automotive manufacturers to build even safer cars and, ultimately, reduce road fatalities worldwide.

In complex traffic scenarios, situation assessment is paramount for an effective automotive safety system. An illustration for human-driving and ADAS process is given in Fig. 1. One can see that: i) the *human-driving* process includes *Perception*, *Driver Intention* and *Driver Action* submodules; ii) the *ADAS* process comprises *Sensing & Estimation*, *Threat-assessment/Decision-making* and *Actuation* functionalities. Naturally, ADAS processes (also referred to in some literature as semi-autonomous systems) are design to be an abstraction of the human-driving process, and great progresses have been made to enlarge the range and complexity of the scenarios handled today. Nevertheless, a major theoretical challenge in the presence of human drivers remains how to precisely distinguish a safe (though perhaps aggressive) driving behaviour from an unsafe one.

In this paper, we will focus our attention solely on Threat-Assessment (TA) and Decision-Making (DM) aspects for semi-autonomous systems. While some literature review papers on motion prediction algorithms exist, see [1] and [2] for example, there is not, to the best of our knowledge, any major survey paper on threat-assessment. We will also include in our discussion pertinent TA-approaches designed for fully Automated Driving (AD), that unlike ADAS systems are aimed at guaranteeing perpetual safety in absence of a skilled driver and in all type of circumstances. We propose here a wide scope article on threat-assessment and decision-making approaches, covering research areas such as safety verification of dynamical systems, scheduling-based approaches, set-based algorithms and machine learning-based methods. Moreover, we will leverage in the discussions from an industrial perspective and experience on the design of critically safe systems. Note that vehicle control or path planning-related aspects are excluded of our discussion, even if they present fundamental theoretical problems worth of discussion and a deeper study. The reader is nevertheless referred to [1]–[3] for a recent overview of motion planning and control techniques for automated vehicles.

The remaining of the paper is organized as follows. Section II presents a literature review on TA-methods, while Section III summarizes the TA-methods from an application and actuation point of view. Section IV discusses the different methods in terms of computational complexity, robustness and potential for future developments.

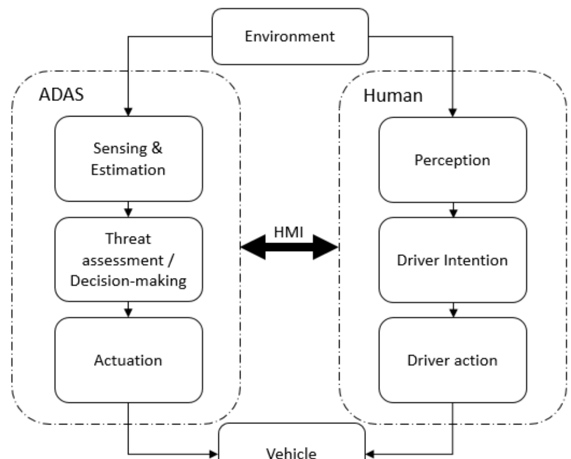

Fig. 1: Driving task: an engineering-designed ADAS process versus a human-driving process. HMI stands for Human-Machine Interface, and can take different forms (e.g., sound feedback, haptic feedback and displays).

J. Dahl, G. Rodrigues de Campos and C. Olsson are with Zenuity, Gothenburg, Sweden. e-mail: john.dahl@zenuity.com, gabriel.campos@zenuity.com, claes.olsson@zenuity.com.

J. Dahl and J. Fredriksson are with the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden. e-mail: jonas.fredriksson@chalmers.se.

II. Threat-assessment methods

In the following we provide a literature overview on existing TA-methods. Methods found in the literature can be divided, on a higher level, into physical model-based methods and data driven-based methods. The main difference among the two classes is that in the first one, threat-assessment and decision-making are made based on learning from physical insights and models, while data driven methods rely on data-based learning, i.e., black-box modeling. Furthermore, physical-based methods can be further divided into several categories, and in this paper we divided such methods into the following categories: single-behavior threat metrics (SBTM), optimization-based methods, formal methods and probabilistic approaches. While the classification of the different works mentioned in this paper is not always straightforward or unambiguous, we have tried to objectively classify them based on the inherent decision-making process. If the threat-assessment is made assuming a single future behavior of the different traffic participants, or if in some cases a closed-form expression can be found, it will fall into the category of single-behavior threat metrics. If the decision-making is a result from an optimization problem it will fall into the optimization-based methods category. If it is possible to find formal guarantees regarding the presence of a threat, these methods will fall into the category of formal methods (including set-based approaches). If the decision-making is made on a probabilistic setup, it will fall into the probabilistic methods category. Finally, if the decision-making is data-driven, such works will be lumped together into one separate category. It is worth noting that sometimes the borderline between different categories is subtle, and therefore overlaps between the chosen categories are unavoidable. Hence, some papers can naturally fit into several categories.

A. Single-Behavior Threat Metrics (SBTMs)

In presence of perfect measurements and a comprehensive knowledge of the intention of each participant (e.g., its destination), it can be possible to retrieve an exact deterministic motion prediction using model-based state propagation. Naturally, this is rarely the case in reality, and simplifications are usually introduced in the problem formulation. For example, it is reasonable to assume that vehicles will continue to follow the current path during a certain time window, i.e., only a single future behavior is considered. Another simplification can be to assume that the measurement data is noise-free. Hence, by using reasonable simplifications one can reformulate the TA problem as being driven by threat metrics based on single future behaviors of the different traffic participants. This type of Threat Metrics (TM) can be defined in any domain such as, e.g., time, distance, or acceleration, and can be used to identify hazardous situations.

1) Time domain TMs: A vastly used TM is Time-to-Collision (TTC), which represents the time until a collision between two objects occurs. In the literature, TTC values are normally used as a threshold in the decision-making process for enabling a warning or an automatic intervention, see [4]–[6]. While there exist several approaches to derive TTC, the most frequently used one computes it as the distance to the closest object divided by the relative velocity. A more general definition of TTC can be found in [7], where TTC is calculated with respect to the relative distance and relative velocity, given a constant relative acceleration. Using the intuition behind TTC, [8] used Time-to-Lane Crossing (TLC) for lane departure warning systems.

Several similar metrics have also been proposed by different authors. For instance, the Inverse TTC (ITTC). ITTC is naturally defined as the inverse of TTC and therefore increases with the risk of collision, see e.g., [9]. For example, the inverse ITC concept was used in [10] for the derivation of a forward collision warning system, based on an experimental data set on maneuvers where drivers execute last-second braking and steering maneuvers.

As an alternative to TTC, Inter-Vehicle-Time (IVT) was also introduced in [11]. It is defined as the time it takes for the host vehicle, given the current host's velocity, to travel the distance equal to the relative distance with respect to the obstacle ahead. Thus, a short relative distance and a high host velocity yields a low IVT time. An important difference between the IVT and TTC is that IVT is calculated based on the host's velocity while TTC is calculated based on the relative velocity. Hence, as highlighted in [11], IVT is a good complement in situations where TTC metric falls short in identifying the threat. For example, in cases where two vehicles are traveling at the same velocity but close to each other, TTC will attain large values even though the situation is hazardous. Another relevant TM is the Time-to-Manoeuvre (TTM) (also referred to as Time-to-X), which is the time until an automated manoeuvre must be initiated to avoid a collision, e.g., Time-to-Brake (TTB), Time-to-Steer (TTS) and Time-to-Kickdown (TTK). Moreover, Time-to-React (TTR) defines the last point in time at which an evasive trajectory still exists. For instance, [12] have used a set-based approach to determine TTR, see Section II-C for a more detailed description. Another approach is presented in [13], where a modified binary search algorithm computes approximated values for different TTMs to derive the TTR for multiple object scenarios.

2) Acceleration domain TMs: Some common acceleration-based metrics include for example the Brake Threat Number (BTN) and the Steering Threat Number (STN), defined as the ratio between the longitudinal (lateral) acceleration required to avoid a collision, and the maximum longitudinal (lateral) acceleration, respectively. Hence, if the required deceleration reaches the maximum achievable threshold, then the ratio is equal to one, a ratio higher than one corresponds to an unavoidable collision, i.e., collision mitigation solutions are the only option.

The authors of [7] and [14] computed the required acceleration to avoid a collision based on constant curvature escape paths while [15] assumed manoeuvres with constant jerk (third time-derivative of the distance). Using a recursive search-tree approach, [16] focused on collision avoidance in a multi-vehicle setup, relying on the assumption that the ego vehicle's optimal path (i.e., with the least lateral acceleration) is either going straight forward or it will always tangent at least one of the objects. A real-time implementation of this algorithm was later presented and validated in [17], [18].

Working with both BTN and STN metrics, [19] presented an estimation of worst-case performance for collision avoidance systems. In particular, the authors derived worst-case decision timing and identified scenarios that are guaranteed not to exhibit unnecessary interventions. The authors considered both steering and braking maneuvers, assuming measurement errors, non-ideal state prediction, as well as sensor and actuator delays. More recently, [20] proposed the notion of predictive steering threat number (PSTN), using a bicycle model with a time-varying lateral acceleration to represent the dynamics of the vehicle, what the authors claimed to be a more realistic modelling when compared to commonly used models in STN-related works. The PSTN has been used to design an adaptive Collision Warning Algorithm (CWA) based on an adaptive, speed-dependent warning threshold.

3) Distance domain TMs: A commonly used distance-based metric is the Minimal Safe Distance (MSD), defined as the minimum distance to be kept between the host and the obstacle, see [11]. This metric is aimed for situations where spatial margins are important, e.g., when queuing at a traffic light or in a traffic jam. A review on different distance domain metrics is given in [21].

4) Multi-domain TMs: In complex scenarios, however, a single TM may not be enough to fully characterize a situation, and therefore multiple TMs may be needed. For example, let a scenario consist of two vehicles traveling next to each other with approximately the same velocity. If only TTC is used, the threat level will be low (a high TTC) even if the inter-vehicle distance decreases to a minimum. Hence, the author of [11] proposed to use a combination of TMs (e.g., TTC, IVT and MSD) to better reflect the real threat level.

B. Optimization methods

Dynamic optimization has become a standard tool for decision-making in a wide range of areas, able to provide fuel-efficient rocket controllers or operational-efficient control strategies for complex chemical production facilities. Indeed, many practical problems can often be expressed as dynamic optimization problems, and a popular optimization framework widely used in the literature is Model Predictive Control (MPC). It consists of an optimization problem minimizing an objective cost function subject to constraints such as a dynamical model and boundaries on the states and/or inputs. The optimization is performed over a finite-time horizon and the optimal solution yields the best control law for the given cost function. While optimization-based approaches may often require high computation power, it is possible in some cases to leverage the structure of the problem to retrieve an explicit control law. This can allow off-line pre-computation of the explicit feedback policies, reducing the on-line computation in receding horizon control setup to a function evaluation, therefore avoiding the on-line solution of complex optimization programs. This is of particular interest for safety-critical and time-critical constrained applications such as automotive threat-assessment algorithms. A detailed introduction to MPC can be found in, e.g., [22] and [23].

In [24] and [25], an MPC-based method was introduced to assess the threat level for forward collision avoidance, i.e., where another vehicle is ahead of the host vehicle. This approach relies on the calculation of the minimum front wheel angle necessary for collision avoidance using an optimal control framework. The solution is also assumed to be able to violate some of the constraints, as the problem's hard constraints have been relaxed to guarantee the feasibility of the MPC problem. In order to cope with more complex situations, the lateral acceleration metric is combined with a different metric defining how much constraints are violated, such that an intervention is only triggered if the level of violation of the constraints or the lateral acceleration exceeds a given threshold. In [26], the authors combined a look-ahead driver model and a vehicle model, and used an MPC controller to compute the minimum steering action necessary to avoid an obstacle. In order to cope with modelling uncertainties on the driver's behaviour, such uncertainties are considered as probabilistic constraints in this work. By computing an upperbound on the uncertainties' deviation, the constraints on the MPC problem can be refined according, which yields a robust control scheme.

The authors in [27] proposed a grid occupancy and optimization based method, that can be used to find the optimal braking point for collision avoidance at low speeds, e.g., for parking applications. For conflict resolution at traffic intersections, [28] leveraged the structure of the problem in order to derive a hierarchical decomposition of the original optimization problem where the central coordination problem is separated from the local optimal control problems on each vehicle, which the authors claim to significantly reduce demands on computational capabilities and information exchange. More precisely, assuming that vehicles are following a predefined path through the intersection, the combinatorial part of the problem (i.e., defining the vehicle crossing order and collision free time slots that are feasible under the vehicle dynamics and physical constraints) is separated from the problem of finding the appropriate control inputs for a given crossing order. Partly based on this decomposition, [29] focused later on the properties of the underlying coordination problem. The authors formulated a finite-time optimal control problem and proposed a primal decomposition, showing that the problem can be efficiently tackled using a standard sequential quadratic programming such that most computations can be performed in a distributed manner by the vehicles. More recently, [30] tried to tackle some of the communication aspects of a distributed Sequential Quadratic Programming (SQP) framework for solving the optimization problem. In order to reduce the communication time and burden, the authors proposed an asynchronous algorithm where the sensitivity of the optimal control problem are only updated for a subset of all agents at each step, and show how one can decide on this subset of agents so to optimize the contraction properties of the algorithm.

C. Formal methods

In complex situations, a great challenge in the design of threat-assessment/decision-making systems is to determine whether a situation can evolve to a threatful situation, especially when subject to complicated requirements composed of safety, task sequentiality and restrictiveness arguments, for example. Hence, a lot of attention has been given in the literature to formal methods as a way to handle these aspects. The idea behind formal methods is that performing appropriate mathematical analysis can contribute to the reliability and robustness of a system, and allow the derivation of correct-by-design control systems. Such methods have been greatly used in computer science, specifically in software engineering and hardware engineering. More recently, researchers from the control theory communities have also been increasingly interested in combining control-theoretic tools for complex physical systems with formal methods for accommodating complex specifications, see [31] and [32]. Hence, the literature in this field is vast and covers different approaches, from temporal logics [33] and robust control techniques [34], to set-based and supervisory control approaches [35] and [36]. Within the scope of this article, several formal approaches for autonomous and semi-autonomous vehicles exist in the literature, and usually fall into one of two sub-categories: logic-based and set-based approaches.

1) Logic-based approaches: Logic-based approaches focus on formalizing a requirement by translating it into logical sentences, and are aimed at verifying the design of a system instead of specifying complex requirements. In [37], the authors proposed a distributed control system for collision-free highway driving, where each vehicle is controlled by adaptive cruise control. The authors also propose a formal safety proof for the proposed control approach, by leveraging hybrid systems models and concepts (e.g, quantified hybrid programs (QHP) and quantified differential dynamic logic). An interesting aspect of the proposed safety verification approach is that the safety proof structure is strictly modular, which reduces the proof to modular stages that can be verified without the details in lower levels of abstraction. In [38] and [39], the authors used Multi-Lane Spatial Logic (MLSL) for proving traffic safety on multi-lane roadways and double-sided traffic country roads, respectively. A very interesting idea behind this set of works is that the authors separate the purely spatial reasoning from the car dynamics for safety analysis, such that the control layer is given by spatial properties formalized in MLSL. Inspired by Interval Temporal Logic, Duration and Shape Calculus, the MLSL approach presented in [38] can be seen as a two-dimensional extension of interval temporal logic. This yields that safety amounts to proving that, under certain assumptions, the occupancy of certain parts of the road are always disjoint. The abstract traffic model for multi-lane motorways given in [38] was later refined by taking traffic directions into account in order to extend this approach to country roads with two-way traffic in [39]. A formal design and verification methodology for cooperative driver assistance systems have also been proposed in [40], that formally verifies timed probabilistic requirements on the successful completion of a safe (i.e., collision-free) driving task. Finally, [41] focused on early specification phases. The authors used Higher Order Logic (HOL), a purely logical environment, to formalize traffic rules, and showed that it is possible to formally check the compliance of traffic rules by autonomous vehicles.

2) Set-based approaches: In set-based approaches, on the other hand, requirements are often formalized by specifying a set of acceptable/unacceptable behaviours or system configurations. For example, [42] proposed a hybrid architecture for collision avoidance at intersections that guarantees safety. This approach consists of an interaction between a centralized component, i.e., a scheduler that assigns a time slot to vehicles, and distributed vehicles that need to determine if they can go through the intersection in the assigned time slot without violating perpetual safety of the system. By leveraging the structure and the ordering properties of the intersection coordination problem, the authors build their safety concept around an infinite sequence of safety inputs that each vehicle possesses, which plays the role of an infinite horizon contingency plan.

In [43] the safety of automated vehicles is formally verified by predicting the set of all possible occupancies of the automated vehicle as well as other traffic participants. The authors apply reachability analysis using mathematical models subject to uncertain inputs (e.g., sensor noise and disturbances) and uncertain initial states, and verified their results on real automated vehicles. Using identical concepts, [44] presented a method for formal collision avoidance detection on arbitrary road networks, where the existence of a collision assessed by checking if the host's occupancy set intersects with the occupancy set of any other object. State constraints are used to limit the expansion of the occupancy set and, should an object violate a specific constraint, such constraint is removed and the object's occupancy set recomputed accordingly. The same research group also described in [45] a new tool that uses reachability analysis to predict the occupancy of surrounding traffic participants, named SPOT for Set-Based Prediction Of Traffic Participants. More precisely, the authors focused on the problem of how to ensure safe motion plans and consider all possible manoeuvres (e.g., full throttle or full braking), physical constraints, as well as traffics rules that the traffic participants are assumed to obey. Given these assumptions, the obtained results are provably an over-approximation with respect to the exact set of possible occupancies, which yields that the SPOT toolbox can be used for instance to verify intended trajectories, as the underlying approach is safe. Another interesting aspect of this work is that the authors also describe how one can react to observed violations of assumptions such as, for example, when an illegal lane change is performed. Recently, the same authors have proposed in [12] a novel approach using reachability analysis to derive Time-to-React (TTR). The authors show how over-approximations of the TTR can be efficiently computed for different urban and rural traffic scenarios, and compare their approach to an optimization-based estimation technique.

Using identical concepts, [46] also dealt with safety verification of motion plans, and focused in particular on cases where (previously safe) motion plans becomes unsafe due to the violation of assumptions, i.e., a misbehavior of dynamic obstacles. To do so, the authors derived invariably safe sets in order to determine the Point of No Return (PNR) and the Point of Guaranteed Arrival (PGA). These concepts allowed the authors to divide motion plans into safe sections and safety-critical passageways, where in the latter the system is exposed to potential collisions if obstacles violate assumptions used for verification. Most interestingly, the authors developed in this paper the concept of safety costs that allows the minimization of safety-critical passageways by assigning costs to paths before their execution. By integrating the cost function into the optimization algorithm, the planner directly determines the safest trajectory possible and ultimately trajectories with a passageway of size zero, i.e., trajectories that guarantee safety for an infinite time horizon. For validation purposes, the authors used the previously mentioned SPOT toolbox to demonstrate their approach for overtaking manoeuvres on a double-sided traffic road.

In [8] and [47], a group of researchers from Chalmers University of Technology used reachability theory to determine the likelihood of a collision. Inputs and states are assumed to be bounded, e.g., constraints on steering angle, steering angle rate, lateral distance to lane markings, slip angle, etc, in order to reflect a regular driving behaviour. The goal is to determine if vehicles can reach a pre-defined safe set without violating either input or state constraints, using reachability tools to identify the safe, invariant set. Naturally, if the current state belongs to the invariant set, the situation is considered to be safe and no intervention is to be triggered, while in the other cases a suitable assisting control policy is requested. Within the same department, other researchers described in [48] a conflict resolution technique for traffic intersections that combines optimal control with sequential decision-making. In particular, the authors used reachability theory to derive model-based heuristics, characterizing for instance each vehicles degree of freedom for avoiding collisions, and present a coordination scheme that scales linearly with the number of participants. The authors exploited the notion of a decision order, firstly introduced in [49], based on which vehicles sequentially solve local optimal control problems. This yields that each vehicle avoids collisions by adapting to the already computed plans by vehicles preceding it in the order. A receding horizon implementation of the above mentioned strategy was also studied in [50], where the authors divided the decentralized solution of the local optimization problems in two parts: a finite-time problem where collision avoidance is enforced as terminal constraints, and an infinite horizon problem defining the cost-to-go that can be calculated offline.

A lot of attention has recently been given to formal verification and supervisory control for collision avoidance, and in particular for traffic intersections, see e.g. [36], [51]–[58]. In particular, a set-based approach for the design of a least restrictive supervisor has extensively been used by MIT's and Politecnico di Milano groups, see [36], [54]. A least restrictive supervisor is defined as a control algorithm that ensures the safety of the requested input (either the command of a driver or potentially from an autonomous vehicle's driving logic), while intervening as little as possible (least restrictiveness) [36]. The system takes as input the user-requested command and returns either (i) the same command if it is considered to be safe or (ii) an overriding, safe-by-design control command. In practice, identifying the set of safe manoeuvres equates with determining the largest set of states for which there exists a control input (e.g., throttle or steering input) that can avoid collisions. This set is commonly referred to as the Maximal Controlled Invariant Set (MCIS), the reader can refer to [59] and [60] for further details. The approaches proposed in [54] and [36] are based on the solution of two separate problems: (i) the *Verification Problem*, that determines if there exists an input signal that leads all vehicles safely through the intersection; and (ii) the *Supervisor Problem*, returning a safe overriding control signal when safety conditions are violated. From a control perspective, the solution to the Verification Problem can be solved by determining the membership to the MCIS, see Fig. 2 for an illustration of the MCIS concept. As illustrated in Fig. 2, the notion of MCIS is closely related to the concepts of Forbidden Set, which corresponds to all configurations representing a collision, as well as the Capture Set (sometimes also called Attraction Set in some literature), i.e., the set that incorporates all configurations from which the system will inevitably reach the Forbidden Set or, in other words, will result in a collision. It is important to notice that determining the MCIS is often a computationally difficult problem, and it has been proved to be NP-hard in the case of some collision avoidance problems, see [54] for further details. Hence, even if standard general algorithms exist in the literature, they are limited by numerical complexity to handle problems involving just a few road users (typically two).

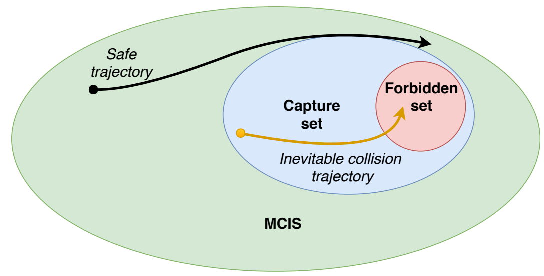

Fig. 2: Illustration of a formal verification approach for threat-assessment. Here, the Forbidden Set corresponds to a configuration corresponding to a collision, the Capture Set to all configurations from which the system will inevitably reach the Forbidden Set (i.e., set of configurations corresponding to a collision), and the MCIS denotes the Maximum control invariant Set (i.e., the set for which it exist at least one control input that allows collision avoidance for all future time instants).

A set of efficient solutions for the intersection collision avoidance problem was proposed in [54], leveraging the equivalence between the Verification Problem and a Scheduling Problem, which consists in assigning jobs to a resource while satisfying given requirements [61]. The Verification Problem can be written as a scheduling problem where the intersection represents the resource, the vehicles represent the jobs to be assigned to the resource, and the time spent by each vehicle in the intersection is the length of the job to be executed. Hence, determining the safety of the system can be posed as a control problem determining the existence of a crossing schedule, wherein the objective corresponds to minimizing the total time needed to clear the intersection. Based on the results of [36], several papers have leveraged such equivalence for the design of supervisors for more complex scenarios handling, for instance, disturbances and measurement noise [56] or uncontrollable vehicles [57]. Optimality arguments are also introduced in the design of the supervisor in [62]. The algorithm presented in [62] is able to identify the optimal corrections to a human-decided input, which allows more driver-friendly and less aggressive manoeuvres. More recently, the authors extended their results to a road-network of interconnected intersections, see [63]. The road-network verification problem is decomposed in more treatable sub-problems, and the authors show how the above mentioned results for a single intersection can be applied to handle such clearly more complex scenarios. Using similar ideas, the authors of [64] also employed formal control theoretical methods for hybrid-systems in order to guarantee collision-free intersections. In particular, the authors leverage the properties of systems that evolve on a partial order and whose dynamics preserve the ordering. For instance, vehicles approaching an intersection are guaranteed to maintain a certain order within the same lane, as vehicles cannot go over each other. The authors showed how such class of systems admits an explicit solution to the safety problem, and proposed linear complexity discrete-time algorithms with guaranteed termination. An experimental implementation of such systems has also been described in [65]. A set of works addressed the problem of multi-vehicle collisions based on abstraction techniques. For instance, [66] provided a recent overview over several works from a research group at the University of Michigan, where the authors constructed a discrete-event system abstraction, and formulate the problem in the framework of supervisory control for discrete-event systems with uncontrollable events, which allows the mitigation of the computational limitations related to the presence of continuous dynamics and infinite state spaces.

Finally, some other works have also combined formal methods with optimization approaches (discussed in the next section). For instance, [67] provided a new approach combining set- and optimization-based methods for formal path-planning for autonomous vehicles. More precisely, the authors model vehicles as differential inclusions composed of simple dynamics and a set-based uncertainty formulation, and obtain reachable sets through an interesting combination of optimization techniques and reachability analysis for control synthesis purposes. From a systems and control perspective, [68] focused on the robust model predictive control for constrained systems, proposing an approach combining optimization-based techniques and set-based concepts. In particular, the authors focused on constrained linear discrete-time systems that are subject to state and measurement disturbances, and the main idea behind this work is to consider the state estimation error as an additional unknown uncertainty, that is bounded by a simple, precomputed invariant set.

D. Probabilistic methods

The underlying idea behind probabilistic TA-approaches is to leverage available information of the system's uncertainties to make decisions with a certain level of confidence. For the sake of clearness of the discussion, the remaining of this section is organized in the following parts: subsection II-D.1 focus on probabilistic TA-methods, while subsection II-D.2 discusses probabilistic decision-making.

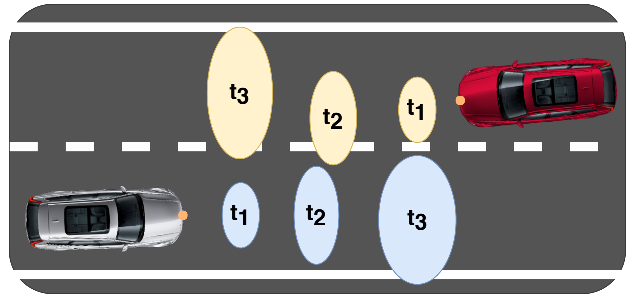

Fig. 3: Illustration of a probabilistic threat-assessment. The figure shows an oncoming scenario, with the host vehicle in silver and the oncoming vehicle in red. Vehicles are approximated as point-of-mass (solid orange circles). Illustrated by the blue and yellow ellipses is the uncertainty associated with the predicted future occupancies, that grows with the prediction time t_k.

1) Probabilistic analysis: Generally speaking, a probabilistic threat-assessment method assigns probabilities to different events, e.g., how likely it is to collide with another object in a near future given some assumptions on uncertainties. In automotive applications, some of the major sources of uncertainty include dynamical modelling errors, measurement noise and the misinterpretation of the drivers' intention. The interested reader can see Fig. 3 for an illustration of the working principles of a basic probabilistic TA.

While there is an infinite number of actions that a driver can take (and that are in general hardly predictable), each action can be linked to a small set of high-level manoeuvres such as performing a lane change, overtake, etc. This principle has been used in [69], for instance. By posing the problem as determining the probability of high-level manoeuvres, one can then reduce both the search space and the computational complexity. Here, each object is assigned a goal function based on different objectives, which is used to define a probability distribution over the set of manoeuvres for each object. The probability of a collision is computed via Monte Carlo simulations, where each object is simulated using samples from respective maneuver distribution. This method has later been improved in [70], where the authors considered the driver's awareness of other objects' location, e.g., if the object is situated behind, ahead, or at either sides of the host vehicle. It is also assumed that, for example, a driver is paying more attention to an object ahead rather than behind the ego vehicle. Based on this work, [71] improved the dynamic vehicle model in order to reduce modelling mismatch, as the model used in [70] suffered from a biased mean value for the lateral and longitudinal acceleration. More recently, [72] introduced a road-aligned curved coordinate system, aimed at simplifying the modelling when the road bends. Another work utilizing a curved coordinate system, see [73], assessed the risk in terms of a probabilistic TTC level, based on uncertainties regarding occupancy of the adjacent lanes.

Focusing on predicting the occupancy for traffic participants, [74] proposed a probabilistic approach using a Markov chain abstraction, considering measurements uncertainties as well as other traffic participants' unknown intentions. Additionally, the road geometry as well as the interaction among the traffic participants are also taken into consideration, and the probability of a crash is computed given a defined trajectory for the host vehicle. This work was later extended in [75] where Markov chain abstraction and Monte-Carlo simulations are compared in terms of computational performance and precision. According to the analysis, the Markov chain abstraction is superior for probabilistic occupancy prediction while the Monte-Carlo simulations is the best choice to derive the collision risk. A different probabilistic motion prediction algorithm, based on an Unscented Kalman Filter (UKF) combined with a reachability based decision-making protocol, has also been given in [76] to decide on an emergency intervention for intersection scenarios. A safety intervention is based on two different but naturally correlated thresholds: one defining the degree of confidence on the model based predictions; the other representing the probability of an unavoidable collision.

Other works have used Bayesian approaches. For example, [77] and [78] used a Dynamical Bayesian Network (DBN) to compute the risk that the driver is missing to stop when approaching an intersection. In scenarios with multiple objects, it is typically the case that an action of one object will cause a reaction of the other participants. Hence, each driver's action is conditioned by the behavior of the other participants, which increases the complexity of the prediction task. The authors of [79] used a DBN to model the physical relationship together with the conditional drivers behavior of the objects in the scenario. To account for changes on the drivers' actions, the DBN is embedded in a Partially Observable Markov Decision Process (POMDP), where each new action can be taken with a certain probability. The DBN model parameters are learned by utilizing an expected maximization approach using unlabeled observations. In [80] the threat level is computed by combining network level collision predictions with vehicle level collision predictions using a DBN. On the network level, the threat is assessed based on high-level traffic information, e.g., average speed and traffic flow on a road segment, while on vehicle level the threat is assessed based on individual motion predictions of the surrounding vehicles. These results show that the vehicle level assessment is improved in cases where the situation is classified as hazardous by the network level assessor. Using the principles of Bayesian Occupancy Filtering (BOF), [81] and [82] estimated the future occupancy of different participants. The environment is represented as a 2-dimensional grid, where each cell contains both static and dynamic information about occupancy and velocity. In [81], the grid indicates the occupancy by other objects and is updated via inference based on sensor measurements. However, prediction performance is poor in presence of temporarily occluded objects. This is addressed in [82] by extending the method with prior map knowledge so to increase the prediction performance for occluded objects.

Some works have focused on computational efficiency aspects, which are crucial for real-time applications. On this topic, [83] mixed set theoretical methods with a probabilistic approach, and divided the threat-assessment problem into a preliminary and a specialized part. The preliminary part searches for potential collisions with other vehicles using a geometric occupancy analysis. The occupancy of each object is efficiently computed via segmented cones representing the reachable set of states for predefined scenarios. The set of objects that can cause a collision are thoroughly analyzed in the specialized part, where the probability of collision is computed using statistical inference.

2) Probabilistic decision-making: In general, it is not obvious how to trigger an automatized intervention in a stochastic setup. However, using hypothesis testing, one can optimize the triggering threshold and therefore minimize the risk of an erroneous decision.

The authors of [7] and [84] derived a decision rule for automated braking based on hypothesis testing for the required lateral acceleration needed to avoid a collision, where the estimates of the longitudinal dynamics are uncertain. The Bayesian framework for the decision-making is generalized in [85] to deal with any stochastic threat-assessment algorithm. Furthermore, they show that the risk of a collision can be computed online with numerical Monte-Carlo integration for an automated braking system.

A similar approach has been used in [86] and [87], where the threat-assessment is divided into two levels: one for a physical system level threat-assessment, i.e., how difficult it is for a vehicle to avoid a collision; and one for the assessment on how the driver would perceive a potential intervention. The authors argue that it might be possible to trig earlier interventions based on less detailed data, if it can be be ensured that the driver will most likely not be annoyed by the intervention.

Finally, [88] has recently focused in behavior generation aspects for automated vehicles, using motion planning techniques that consider uncertain predictions of other road users. The main contribution of this work relies on an online Partially Observable Markov Decision Process (POMDP) algorithm, which is able to incorporate uncertainties and provide optimized solutions for the on intersections with an arbitrary layout and a variable number of traffic participants.

E. Data driven methods

Machine Learning (ML) methods recently became very popular for several tasks within ADAS/AD applications, in particular due to the rapid advances in computational hardware and algorithms that are strongly supported by large companies such as NVIDIA and Intel. The underlying idea behind ML is to learn an environment/behavior/property purely based on a data set, often denoted as the training set. More precisely, ML algorithms approximate the relationship from one element to another within the system, which are usually complex and not formally described. If the training set is rich and large enough, in the sense that all configurations of the system are sufficiently excited/represented, a ML-approximated function is expected to efficiently generalize with new data, e.g., verification data or real-time sampled data. In the context of ADAS applications, ML approaches are usually used for predictions where the input is the current state of the system and the output a prediction of the future state. A popular method within the ML family is the Artificial Neural Network (ANN). For further details on learning methods, the reader can refer to [89].

There are many different implementations of machine-learning concepts. For example, the authors of [90] and [91] used a technique based on Gaussian processes in order to learn the drivers' behavior at intersections and lane-change scenarios, respectively. Furthermore, ANN, Recurrent Neural Networks (RNN) and Support Vector Machines (SVM) have been considered to predict the drivers' behavior specifically for lane-change scenarios, see [92] where it is shown that SVM offers superior results with respect to RNN and ANN. In a similar work, naive Bayes, SVM, logic regression, nearest neighborhood, decision trees, extra trees and Random Forest (RF) classifiers are compared in [93] for lane-change scenarios, where it is concluded that the tree based classifiers are superior with respect to the remaining approaches. Furthermore, a RF algorithm is used in [94] to learn automated parking manoeuvres, where the solution of the RF is approximated by a General Radial Basis Function (GRBS) for enhanced real-time capabilities.

Another important application for ANN is the automated annotation of scenarios in logged data, see [95]. For example, a given scenario is automatically classified by its properties, i.e. "the scenario is an intersection with a vehicle approaching from the left side".

For AD applications, supervised ANN has already been used for End-to-End learning [96]–[98], where the objective is to derive an actuation signal directly from the raw image input obtained from a frontal camera. In other words, the goal for the ANN is to mimic the behaviour of the driver given the training data set. While End-to-End learning considers no assumptions on the underlying problem, other works proposed instead to use state estimates as inputs to an ANN. This approach is used in [99], where an ANN is learned to detect unintended lane departures.

III. Application, actuation and validation

Section II presented a description of the different threat-assessment methods and the underlying main ideas and principles. In order to relate the reviewed works to different automotive-related applications, a summary table is given in Table I. The table is organized with respect to three arguments: i) what type of application the works are aimed for; ii) the actuation principle; and iii) how the results have been validated. In particular, we identified five different collision-prone scenarios: at intersections, frontal and rear-end collisions (a vehicle is traveling in same direction as the host), lane departures (LKA), and collisions with Vulnerable Road Users (VRU). We also classified works with respect to the actuation principle, i.e., relying on braking and/or steering actions, and the validation approach separated in two categories, simulation only and experimental (or simulations based on real-data). Note that broad scope publications such as dissertations and review papers, as well as literature related to automotive applications that do not fit with the above mentioned structure are gathered at the end of the table. Furthermore, non automotive-oriented literature focusing on general aspects of sensing, control theory, or computer science, for instance, are excluded.

From Table I, one can see that a great part of the papers in this literature review used probabilistic and formal methods. Moreover, one can also observe that, among the cited literature, a lot of attention has been given to conflict resolution at traffic intersections, as such scenarios incorporate the same technical challenges such as round-abouts or on-ramp merging, for instance. A significant part of the considered publications provided experimental results, even if most of the works rely on simulation-only results. It is important mentioning that, naturally, the distribution of Table I and the discussion (see next section) is not necessarily representative of all the existing literature, but rather an overview of the papers analyzed in this review.

IV. Discussion

In the rest of the paper, we will discuss the different approaches in terms of system performance, robustness, real-time properties and also challenges and implications of increased autonomy.

A. System performance

While all methods can benefit from an efficient data-filtering and sensor-fusion algorithms, the different approaches differ on how environment data is used within the algorithms. For instance, SBTMs usually only consider the most likely state estimates, discarding in many cases the associated distribution of the states. Instead, robustness with respect to uncertainties are gained by putting margins in the decision-making stage, that are typically tuned by empirically testing to maximize the system performance.

Probabilistic methods, on the other hand, tend to make use of the entire uncertainty model in order to estimate the probability of a collision. Nevertheless, it is generally not obvious how to derive the uncertainty model, especially in cases of time-varying uncertainties. This is, e.g., the case for uncertainties related to driver's intentions. From this literature review, it follows that probabilistic methods are preferred in scenarios where the uncertainties can be modeled by a few random variables, or in cases where the prediction horizon is sufficiently short.

Regarding formal methods, the underlying objective is usually to verify and formally guarantee whether a situation can evolve to an unavoidable dangerous situation. These methods are best suited for applications where formal guarantees of performance are required, or when the system is subject to complex combinations of safety, comfort, or logic requirements, for example.

In data-driven approaches, the only underlying assumption is that there exists a relationship between the inputs and the outputs, i.e., the current system's state can be mapped

TABLE I: An overview of the reviewed articles. Here, CA stands for Collision Avoidance, LKA for Lane Keep Assistance, and VRU for Vulnerable Road Users.

	Single-behavior TMs	Optimization methods	Formal methods	Probabilistic methods	Data driven methods
	Application				
CA with VRU	[4] [13] [17]			[69] [72] [86] [87] [83]	
CA at intersections	[5] [13]	[27] [28] [29] [30]	[12] [35] [36] [42] [45] [48] [49] [50] [53] [54] [55] [56] [57] [58] [62] [64] [65] [66]	[74] [75] [76] [77], [78] [79] [82] [83] [88]	[79] [90]
Frontal CA	[4] [5] [14] [15] [20] [21]	[24] [25] [26]	[12] [35] [39] [43] [44] [45] [46] [52]	[69] [70] [71] [72] [73] [74] [75] [86] [87]	[91] [92] [93]
Rear-end CA	[5] [6] [7] [9] [10] [11] [16] [17] [18] [19] [21]	[24] [25] [26]	[12] [35] [36] [37] [38] [39] [40] [43] [44] [45] [46] [63] [67]	[7] [69] [70] [71] [72] [73] [74] [75] [84] [85] [86] [87]	[91] [92] [93]
LKA	[8]		[8] [35] [47]		[96] [97] [98] [99]
	Actuation				
Steering control	[9] [10] [11] [13] [16] [17] [18] [19] [20]	[24] [25] [26]	[8] [12] [38] [39] [40] [43] [44] [47] [52] [67]	[72] [79]	[91] [94] [96] [97] [98]
Throttle/Braking control	[6] [7] [9] [10] [11] [13] [14] [15] [19] [21]	[27] [28] [29] [30]	[12] [36] [37] [38] [39] [40] [42] [43] [44] [48] [49] [50] [53] [54] [55] [56] [57] [58] [36] [62] [63] [64] [65] [66] [67]	[7] [76] [77] [79] [85] [86] [87] [88]	[90] [91] [94]
	Validation approach				
Simulation only	[14] [16] [17] [18] [19] [20]	[26] [28] [29] [30]	[12] [36] [37] [38] [39] [40] [42] [45] [46] [48] [49] [52] [54] [56] [57] [62] [63] [64] [66]	[69] [70] [71] [79] [83] [84] [88]	[79] [92] [99]
Experimental/ with real-data	[4] [5] [6] [7] [8] [9] [10] [11] [13] [15] [16] [21]	[24] [25] [27]	[8] [35] [43] [44] [47] [53] [55] [58] [64] [65] [67]	[7] [70] [72] [73] [74] [75] [76] [77] [78] [79] [82] [85] [86] [87]	[90] [91] [93] [94] [96] [97] [98]

	Related literature
Correlated topics and broad-scope publications	[1] [2] [3] [4] [7] [35] [41] [79] [80] [81] [95] [100] [101] [102] [103]

to a threat level. In the case of supervised learning (e.g., deep learning), a practical challenge is that the data used for training the algorithm needs to be classified a-priori, either by humans or by some automated logic. Furthermore, training data-set must be persistently exciting (i.e., exciting all possible modes of the system), which may not always be possible in collision avoidance (i.e., there is often a discrepancy between the amount of data on regular driving behaviors and collision-prone behaviors). On the other hand, one of the strong benefits of using machine-learning based methods is their ability to capture complex nonlinear behaviors, which makes it, at least in theory, suitable to find relationships that may be hard to model using classical engineering principles.

B. *Real-time properties*

From an industrial perspective, the systems performance is naturally not the only factor under consideration. Indeed, any safety system needs to be cost-efficient, both during the development phase but also upon deployment (e.g., real-time computational requirements). There is therefore an emphasis on adopting efficient but not unnecessarily complex methods.

Generally speaking, SBTMs are computationally cheap by design. For set based approaches, recent works have proposed a toolbox that is able to compute set based predictions in under 20 ms, see [45]. Other solutions, though, aim at leveraging the problem structure to reduce the online computational load or set and reachability computation, by solving sub-parts of the problem offline during the development phase and later deployed in vehicles with lower computational demands. Regarding optimization based approaches, their computational complexity usually scales poorly with the size of the problem, which mean that only small problems can be efficiently solved in real-time. Unfortunately, constrained MPC-problems tend to be non-convex, which together with long prediction horizons usually yields prohibitive complex problems.

Probabilistic methods may also lead to computational-efficient solutions in some special cases where the noise is strictly Gaussian and the dynamics linear, for example. However, that is rarely the case in practice and the solution often needs to be numerically approximated by Monte-Carlo simulations or particle filtering, which are known to be accurate but computationally expensive.

For data-driven algorithms, it is crucial that the training data is sufficiently rich and large in size. However, large networks with many layers and neurons, combined with complicated activation functions, offer great challenges in terms of computational power and memory allocation for real-time deployment. Moreover, the training process itself is normally computational expensive but can usually be done offline. Leveraging the large database that automakers gathered over the years from field tests and test track verification tests, deep-learning has been attracting an increasing amount of interest for ADAS and autonomous driving systems, in particular for image segmentation[1] as illustrated in Fig. 4. But while the first applications have shown promising results for lane-markings and road sign detection, for example, their applicability to core safety features is not straightforward and will require several years of development.

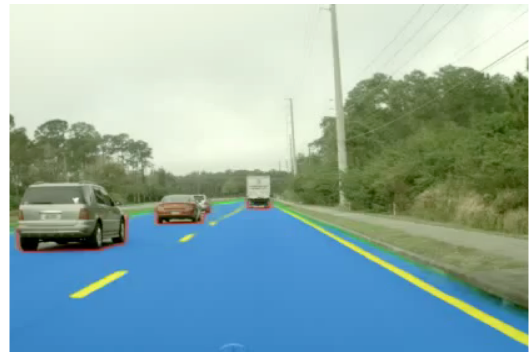

Fig. 4: Image segmentation with deep-learning algorithms, which are able to identify areas of interest in front of the vehicle: the drivable-road area is highlighted in blue, lane markers in yellow, road edges in green and vehicles' delimiters in red.

C. Robustness

Another important aspect for any safety system is robustness, here defined as the measure of the system's ability to handle uncertainties due to e.g., measurement noise or deviations from other assumptions such as the modeled driver behaviour. Formal methods have an intrinsic ability to guarantee robustness, where the level of robustness can be adjusted by over- or underestimating the uncertainties in the set definitions. SBTM- and optimization-based methods, on the other hand, fall short in terms of robustness, as a result of over-simplifications made on the system modelling and problem statement, see [104]. Imagine a solution to the optimization problem that tangents the constraints for some variables. In this case, the solution is feasible with respect to the objective function, constraints and vehicle model. But what if the model is uncertain and the state measurements are noisy? To handle such issues, it is common practice to measure the state after some time period and to solve the dynamic optimization problem again starting from the new measured state. By using feedback on the measurement information, this provides the whole procedure with a robustness that is typical for closed-loop systems, and usually referred to as Model Predictive Control (MPC) that we discussed earlier in the paper. To enhance robustness properties, extensions to the MPC framework have been made in [105] and [106]. The underlying idea of Robust MPC is that the problem contains one or several sources of uncertainties and several techniques have been proposed to solve the problem such as stochastic MPC [107] or scenario-tree MPC [108] and [109]. In the first set of works the problem is defined as a typical MPC problem, but where the constraints are only fulfilled with a certain probability. In the second approach, for problems with a finite number of realizations of uncertainties, one can formulate a problem for each uncertainty realization and structure them in a tree formation. Hence, the overall goal with scenario-tree MPC is to find a control sequence that can handle all realizations of the uncertainties and potentially using a distributed computational scheme.

D. Challenges, implications and opportunities

Despite the substantial technical progress on sensing, computation and communication capabilities in the last few years, the development and deployment of autonomous vehicles will raise great technical and societal challenges in the decades to come. A discussion on the such challenges but also future opportunities is given in, e.g., [100] or [101]. In [100], the authors discuss different necessary technologies, covering a wide spectrum of aspects including in-car computing and data management, road side infrastructures and cloud solutions. Most interestingly, the authors particularly discuss the need of high-definition maps, identified as core technology for autonomous vehicles. The opportunities and future implications of autonomous vehicles for transportation policies are also tackled in [101], where the authors provide a review on relevant literature for aspects going from safety to machine ethics. Finally, the interdisciplinary nature of the challenges raised by increasingly autonomous safety systems is also discussed in [102]. Here, the authors highlight the need for a multi-disciplinary approach across software and hardware engineering, testing or social and legal fields, and particularly point out the challenges to validate machine-learning based systems with respect to the ultra-dependable levels required for autonomous vehicle fleets. The same group of researchers also discussed in [103] the challenges in testing and validation of autonomous vehicles, a corner stone for a large-scale deployment of such technology.

[1]Image segmentation consists of a process of partitioning a digital image into multiple segments that share certain characteristics. Image segmentation is particular interesting for automotive systems in order to simplify and/or change the representation of an camera image into something that is more meaningful and easier to analyze.

V. Conclusions and future perspectives

In this paper we provide a literature review and a technical discussion on the different methods used for threat-assessment for automotive active-safety systems. We have covered a wide range of publications focusing on intelligent and automated vehicles, and classified them in comprehensive categories based on the underlying threat-assessment and decision-making process. It is worth highlighting that the discussions driven in this paper are articulated around a decision dilemma, according to which there exists a clear trade-off between high performance (i.e., the ability to avoid a collision) and precision (i.e., performing an avoidance maneuver only when needed). While long prediction horizons and high intensity interventions are important to properly handle very critical situations, they may, when combined with uncertain estimates, increase the number of unwanted/false interventions. Short prediction horizon, on the other hand, may provide high precision perception and situation awareness, but the driver's time to react is reduced and the effectiveness of an ADAS may be limited. Consequently, making use of methods in such a way that they best handle this intrinsic conflict, is likely one of the main challenges for future generations of ADAS systems.

Despite the many comprehensive works reviewed here, there is clearly still room for substantial improvements for future research. For example, one of the take-home messages is that the methods covered here have different strengths and weaknesses that are dependent on the context. Hence, it is tempting to think that it might be possible to combine different methods to maximize the over all performance on a broader set of situations, even though it is not obvious how this can be achieved. Moreover, the drivers' intention is also a crucial aspect of all types of methods, since it determines a substantial part of the future outcome. Techniques such as driver state estimation using in-vehicle mounted cameras and vehicle-to-vehicle communication, likely to enter the market soon, will allow for new possibilities in terms of threat-assessment and decision-making strategies.

Acknowledgments

This research was supported by Zenuity and the Swedish Governmental Agency for Innovation Systems/FFI through Contract 2014-05621. The authors are also grateful to the reviewers and the editors for their efforts in reviewing and improving this paper.

References

[1] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *Robomech Journal, Springer*, no. 3, 2014.

[2] D. González, J. Pérez, V. Milanés, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 4, pp. 1135–1145, 2016.

[3] B. Paden, M. Cap, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[4] A. Van der Horst, *A time based analysis of road user behaviour in normal and critical encounters*, 1990, no. HS-041 255.

[5] J. C. Hayward, "Near-miss determination through use of a scale of danger," *Pennsylvania Transportation and Traffic safety Center*, 1971.

[6] R. van der Horst and J. Hogema, *Time-to-collision and collision avoidance systems*, 1994.

[7] J. Jansson, *Collision Avoidance Theory with Application to Automotive Collision Mitigation*, 2005, no. 950.

[8] P. Falcone, M. Ali, and J. Sjöberg, "Predictive threat assessment via reachability analysis and set invariance theory," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1352–1361, 2011.

[9] R. Kiefer, M. Cassar, C. Flannagan, D. LeBlanc, M. Palmer, R. Deering, and M. Shulman, "Refining the camp crash alert timing approach by examining" last second" braking and lane change maneuvers under various kinematic conditions," *Accident Investigation Quarterly*, no. 39, 2004.

[10] R. Kiefer, D. LeBlanc, and C. Flannagan, "Developing an inverse time-to-collision crash alert timing approach based on drivers last-second braking and steering judgments," *Accident Analysis Prevention*, vol. 37, no. 2, pp. 295 – 303, 2005.

[11] S. Noh and W. Y. Han, "Collision avoidance in on-road environment for autonomous driving," *International Conference on Control, Automation and Systems*, 2014.

[12] S. Sontges, M. Koschi, and M. Althoff, "Worst-case analysis of the time-to-react using reachable sets," in *IEEE Intelligent Vehicles Symposium*, 2018.

[13] A. Tamke, T. Dang, and G. Breuel, "A flexible method for criticality assessment in driver assistance systems," *IEEE Intelligent Vehicle Symposium*, no. Iv, pp. 697–702, 2011.

[14] M. Brännström, E. Coelingh, and J. Sjöberg, "Threat assessment for avoiding collisions with turning vehicles," *IEEE Intelligent Vehicles Symposium*, pp. 663–668, 2009.

[15] ——, "Model-based threat assessment for avoiding arbitrary vehicle collisions," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 658–669, 2010.

[16] A. Eidehall, "Multi-target threat assessment for automotive applications," *IEEE International Conference on Intelligent Transportation Systems*, 2011.

[17] D. Madås, M. Nosratinia, M. Keshavarz, P. Sundstrom, R. Philippsen, A. Eidehall, and K.-M. Dahlen, "On path planning methods for automotive collision avoidance," *IEEE Intelligent Vehicles Symposium*, 2013.

[18] A. Eidehall and D. Madås, "Real time path planning for threat assessment and collision avoidance by steering," *IEEE International Conference on Intelligent Transportation Systems*, 2013.

[19] J. Nilsson, A. Ödblom, and J. Fredriksson, "Worst-case analysis of automotive collision avoidance systems," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 4, pp. 1899–1911, April 2016.

[20] S. Hosseini, N. Murgovski, G. Rodrigues de Campos, and J. Sjöberg, "Adaptive forward collision warning algorithm for automotive applications," *American Control Conference*, pp. 5982–5987, 2016.

[21] Y. Zhang, E. Antonsson, and K. Grote, "A new threat assessment measure for collision avoidance systems," *IEEE Intelligent Transportation Systems Conference*, pp. 968–975, 2006.

[22] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control : Theory and Design*, 2012.

[23] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*, 2017.

[24] S. Anderson, S. C. Peters, T. E. Pilutti, and K. Iagnemma, "Design and development of an optimal-control-based framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios," *Springer Tracts in Advanced Robotics*, vol. 70, pp. 39–54, 2011.

[25] S. Anderson, *A unified framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles*, 2009.

[26] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, and F. Borrelli, "Stochastic predictive control for semi-autonomous vehicles with an uncertain driver model," *IEEE Conference on Intelligent Transportation Systems*, pp. 2329–2334, 2013.

[27] B. Gutjahr and M. Werling, "Automatic collision avoidance during parking and maneuvering — An optimal control approach," *IEEE Intelligent Vehicles Symposium*, pp. 636–641, 2014.

[28] R. Hult, G. Rodrigues de Campos, P. Falcone, and H. Wymeersch, "An approximate solution to the optimal coordination problem for autonomous vehicles at intersections," *American Control Conference*, 2015.

[29] R. Hult, M. Zanon, S. Gros, and P. Falcone, "Primal decomposition of the optimal coordination of vehicles at traffic intersections," *IEEE Conference on Decision and Control*, pp. 2567–2573, 2016.

[30] M. Zanon, S. Gros, P. Falcone, and H. Wymeersch, "Asynchronous algorithm for optimal vehicle coordination at traffic intersections," *World Congress of the International Federation of Automatic Control*, 2017.

[31] P. Tabuada, *Verification and Control of Hybrid Systems*. Springer, 2008.

[32] C. Belta, *Formal Methods for Discrete-Time Dynamical Systems*. Springer, 2017.

[33] C. Baier and J.-P. Katoen, *Principles of model checking*. Springer, 2008.

[34] A. Aswani, H. Gonzalez, S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.

[35] M. Althoff, "Reachability analysis and its application to the safety assessment of autonomous cars," Ph.D. dissertation, Technische Universitat Munchen, 2010.

[36] A. Colombo and D. Del Vecchio, "Least restrictive supervisors for intersection collision avoidance: A scheduling approach," *IEEE Transactions on Automatic Control*, vol. 60, no. 6, pp. 1515–1527, 2015.

[37] S. Loos, A. Platzer, and L. Nistor, "Adaptive cruise control: Hybrid, distributed, and now formally verified," *Lecture Notes in Computer Science, Springer*, vol. 6664, pp. 42–56, 2011.

[38] M. Hilscher, S. Linker, R. Olderog, and A. Ravn, "An abstract model for proving safety of multi-lane traffic manoeuvres," *Theories of Programming and Formal Methods. Lecture Notes in Computer Science, Springer*, vol. 6991, pp. 404–419, 2011.

[39] M. Hilscher, S. Linker, and E. Olderog, "Proving safety of traffic manoeuvres on country roads," *Theories of Programming and Formal Methods. Lecture Notes in Computer Science, Springer*, vol. 8051, pp. 196–212, 2013.

[40] W. Damm, H.-J. Peter, J. Rakow, and B. Westphal, "Can we build it: formal synthesis of control strategies for cooperative driver assistance systems," *Mathematical Structures in Computer Science*, vol. 23, no. 8, pp. 676–725, 2013.

[41] A. Rizaldi and M. Althoff, "Formalising traffic rules for accountability of autonomous vehicles," in *IEEE Conference on Intelligent Transportation Systems*, 2015.

[42] H. Kowshik, D. Caveney, and P. R. Kumar, "Provable systemwide safety in intelligent intersections," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 3, pp. 804–818, March 2011.

[43] M. Althoff and J. M. Dolan, "Online verification of automated road vehicles using reachability analysis," *IEEE Transactions on Robotics*, vol. 30, no. 4, pp. 903–918, 2014.

[44] M. Althoff and S. Magdici, "Set-Based Prediction of Traffic Participants on Arbitrary Road Networks," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 2, pp. 187–202, 2016.

[45] M. Koschi and M. Althoff, "Spot: A tool for set-based prediction of traffic participants," in *IEEE Intelligent Vehicles Symposium*, 2017, pp. 1686–1693.

[46] C. Pek, M. Koschi, M. Werling, and M. Althoff, "Enhancing motion safety by identifying safety-critical passageways," pp. 320–326, 2017.

[47] M. Ali, P. Falcone, and J. Sjöberg, "Model-based threat assessment for lane guidance systems," *American Control Conference*, pp. 4586–4591, 2011.

[48] G. Rodrigues de Campos, P. Falcone, R. Hult, H. Wymeersch, and J. Sjöberg, "Traffic coordination at road intersections: Autonomous decision-making algorithms using model-based heuristics," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 8–21, 2017.

[49] G. Rodrigues de Campos, P. Falcone, and J. Sjöberg, "Autonomous cooperative driving: a velocity-based negotiation approach for intersection crossing," *IEEE Conference on Intelligent Transportation Systems*, pp. 1456–1461, 2013.

[50] G. Rodrigues de Campos, P. Falcone, H. Wymeersch, R. Hult, and J. Sjöberg, "Cooperative receding horizon conflict resolution at traffic intersections," *IEEE Conference on Decision and Control*, pp. 2932–2937, 2014.

[51] C. Tomlin, G. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A study in multi-agent hybrid systems," *IEEE Transactions on Automatic Control*, vol. 43, pp. 509–521, 1998.

[52] F. Fadaie and M. E. Broucke, "On the least restrictive control for collision avoidance of two unicycles," *International Journal of Robust Nonlinear Control*, vol. 16, pp. 553–574, 2006.

[53] R. Verma and D. Del Vecchio, "Semi-autonomous multi-vehicle safety: A hybrid control approach," *IEEE Robotics & Automation Magazine*, vol. 18, pp. 44–54, 2011.

[54] A. Colombo and D. Del Vecchio, "Efficient algorithms for collision avoidance at intersections," *ACM Conference on Hybrid Systems: Computation and Control*, 2012.

[55] M. Hafner, D. Cunningham, L. Caminiti, and D. Del Vecchio, "Cooperative collision avoidance at intersections: Algorithms and experiments," *IEEE Transactions on Intelligent Transportation Systems*, 2013.

[56] L. Bruni, A. Colombo, and D. Del Vecchio, "Robust multi-agent collision avoidance through scheduling," *IEEE Conference on Decision and Control*, 2013.

[57] H. Ahn, A. Colombo, and D. Del Vecchio, "Supervisory control for intersection collision avoidance in the presence of uncontrolled vehicles," *American Control Conference*, 2014.

[58] H. Ahn, A. Rizzi, A. Colombo, and D. Del Vecchio, "Experimental testing of a semi-autonomous multi-vehicle collision avoidance algorithm at an intersection testbed," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.

[59] J. Lygeros, C. Tomlin, and S. Sastry, "Controllers for reachability specifications for hybrid systems," *Automatica*, vol. 35, pp. 349–370, 1999.

[60] E. C. Kerrigan, "Robust Constraint Satisfaction : Invariant Sets and Predictive Control," Ph.D. dissertation, University of Cambridge, 2000.

[61] M. Pinedo, *Scheduling: Theory, Algorithms, and Systems*. Springer, 2008.

[62] G. Rodrigues de Campos, F. Della Rossa, and A. Colombo, "Optimal and least restrictive supervisory control: Safety verification methods for human-driven vehicles at traffic intersections," *IEEE Conference on Decision and Control*, 2015.

[63] A. Colombo, G. Rodrigues De Campos, and F. Della Rossa, "Control of a city road network: Distributed exact verification of traffic safety," *IEEE Transactions on Automatic Control*, vol. 62, no. 10, pp. 4933 – 4948, 2017.

[64] M. R. Hafner and D. Del Vecchio, "Computational tools for the safety control of a class of piecewise continuous systems with imperfect information on a partial order," *SIAM Journal on Control and Optimization*, vol. 49, no. 6, pp. 2463–2493, 2011.

[65] M. Hafner, D. Cunningham, L. Caminiti, and D. Del Vecchio, "Cooperative collision avoidance at intersections: Algorithms and experiments," *IEEE Transactions on Intelligent Transportation Systems*, 2013.

[66] E. Dallal, A. Colombo, D. Del Vecchio, and S. Lafortune, "Supervisory control for collision avoidance in vehicular networks using discrete event abstractions," *Discrete Event Dynamic Systems*, vol. 27, pp. 1–44, 2017.

[67] B. Schürmann, D. Hess, J. Eilbrecht, O. Stursberg, F. Köster, and M. Althoff, "Ensuring drivability of planned motions using formal methods," in *IEEE International Conference on Intelligent Transportation Systems*, 2017.

[68] D. Mayne, S. Rakovic, R. Findeisen, and F. Allgöwer, "Robust output feedback model predictive control of constrained linear systems," *Automatica*, vol. 42, no. 7, pp. 1217–1222, 2006.

[69] A. Broadhurst, S. Baker, and T. Kanade, "Monte Carlo road safety reasoning," *IEEE Intelligent Vehicles Symposium*, pp. 319–324, 2005.

[70] A. Eidehall and L. Petersson, "Threat assessment for general road scenes using monte carlo sampling," *Intelligent Transportation Systems*, pp. 1173–1178, 2006.

[71] S. Danielsson, L. Petersson, and A. Eidehall, "Monte Carlo based Threat Assessment: Analysis and Improvements," *2007 IEEE Intelligent Vehicles Symposium*, pp. 233–238, 2007.

[72] 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.

[73] J. Kim, K. Jo, W. Lim, M. Lee, and M. Sunwoo, "Curvilinear-coordinate-based object and situation assessment for highly automated vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 3, pp. 1559–1575, 2015.

[74] 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.

[75] 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.

[76] G. Rodrigues de Campos, A. H. Runarsson, F. Granum, P. Falcone, and K. Alenljung, "Collision avoidance at intersections: A probabilistic threat-assessment and decision-making system for safety interventions," *IEEE International Conference on Intelligent Transportation Systems*, 2014.

[77] S. Lefèvre, "Risk Estimation at Road Intersections for Connected Vehicle Safety Applications," Ph.D. dissertation, Universite de Grenoble, 2012.

[78] A. Armand, D. Filliat, and J. Ibañez-Guzman, "A Bayesian Framework for Preventive Assistance at Road Intersections," *IEEE Intelligent Vehicles Symposium*, pp. 1128–1134, 2016.

[79] T. Gindele, S. Brechtel, and R. Dillmann, "Learning driver behavior models from traffic observations for decision making and planning," *IEEE Intell. Transp. Syst. Mag.*, vol. 7, no. 1, pp. 69–79, 2015.

[80] C. Katrakazas, "Developing an advanced collision risk model for autonomous vehicles," Ph.D. dissertation, Loughborough University, 2017.

[81] C. Chen, C. Tay, C. Laugier, and K. Mekhnacha, "Dynamic environment modeling with gridmap: A multiple-object tracking application," *9th Int. Conf. Control. Autom. Robot. Vision, 2006, ICARCV '06*, 2006.

[82] T. Gindele, S. Brechtel, J. Schröder, and R. Dillmann, "Bayesian occupancy grid filter for dynamic environments using prior map knowledge," *IEEE Intell. Veh. Symp. Proc.*, pp. 669–676, 2009.

[83] 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. 12, no. 4, pp. 942–953, 2011.

[84] R. Karlsson, J. Jansson, and F. Gustafsson, "Model-based statistical tracking and decision making for collision avoidance application," *American Control Conference*, pp. 3435–3440, 2004.

[85] J. Jansson and F. Gustafsson, "A framework and automotive application of collision avoidance decision making," *Automatica*, vol. 44, no. 9, pp. 2347–2351, 2008.

[86] F. Sandblom and M. Brännström, "Probabilistic threat assessment and driver modeling in collision avoidance systems," *IEEE Intelligent Vehicles Symposium*, pp. 914–919, 2011.

[87] M. Brännström, F. Sandblom, and L. Hammarstrand, "A probabilistic framework for decision-making in collision avoidance systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 2, pp. 637–648, 2013.

[88] C. Hubmann, J. Schulz, M. Becker, D. Althoff, and C. Stiller, "Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction," *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 1, pp. 5–17, 2018.

[89] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation Learning: A Survey of Learning Methods," *ACM Computing Surveys*, vol. 50, no. 2, pp. 1–35, 2017.

[90] A. Armand, D. Filliat, and J. Ibanez-Guzman, "Modelling stop intersection approaches using Gaussian processes," *IEEE Conference on Intelligent Transportation Systems*, pp. 1650–1655, 2013.

[91] R. Huang, H. Liang, P. Zhao, B. Yu, and X. Geng, "Intent-Estimation- and Motion-Model-Based Collision Avoidance Method for Autonomous Vehicles in Urban Environments," *Applied Sciences*, vol. 7, no. 5, p. 457, 2017.

[92] D. Urun, J. Edelbrunner, and I. Iossifidis, "Autonomous driving: A comparison of machine learning techniques by means of the prediction of lane change behavior," *IEEE International Conference on Robotics and Biomimetics*, pp. 1837–1843, 2011.

[93] N. Motamedidehkordi, S. Amini, S. Hoffmann, F. Busch, and M. R. Fitriyanti, "Modeling tactical lane-change behavior for automated vehicles: A supervised machine learning approach," *IEEE International Conference on Models and Technologies for Intelligent Transportation Systems*, pp. 268–273, 2017.

[94] G. Notomista and M. Botsch, "A machine learning approach for the segmentation of driving maneuvers and its application in autonomous parking," *Journal of Artificial Intelligence and Soft Computing Research*, vol. 7, no. 4, pp. 243–255, 2017.

[95] R. Gruner, P. Henzler, G. Hinz, C. Eckstein, and A. Knoll, "Spatiotemporal Representation of Driving Scenarios and Classification using Neural Networks," *IEEE Intelligent Vehicles Symposium*, pp. 1782–1788, 2017.

[96] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, "End to End Learning for Self-Driving Cars," *arXiv:1604*, pp. 1–9, 2016.

[97] M. Bojarski, P. Yeres, A. Choromanska, K. Choromanski, B. Firner, L. Jackel, and U. Muller, "Explaining How a Deep Neural Network Trained with End-to-End Learning Steers a Car," *arXiv*, pp. 1–8, 2017.

[98] S.-G. Jeong, J. Kim, S. Kim, and J. Min, "End-to-end learning of image based lane-change decision," *IEEE Intelligent Vehicles Symposium*, pp. 1602–1607, 2017.

[99] J. M. Ambarak, H. Ying, F. Syed, and D. Filev, "A neural network for predicting unintentional lane departures," *IEEE International Conference on Industrial Technology*, pp. 492–497, 2017.

[100] H. G. Seif and X. Hu, "Autonomous driving in the icity - hd maps as a key challenge of the automotive industry," *Engineering*, vol. 2, no. 2, pp. 159 – 162, 2016.

[101] S. A. Bagloee, M. Tavana, M. Asadi, and T. Oliver, "Autonomous vehicles: challenges, opportunities, and future implications for transportation policies," *Journal of Modern Transportation*, vol. 24, no. 4, 2016.

[102] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 90–96, 2017.

[103] ——, "Challenges in autonomous vehicle testing and validation," *SAE International Journal of Transportation Safety*, vol. 4, no. 1, pp. 15–24, 2015.

[104] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, "Examples when nonlinear model predictive control is nonrobust," *Automatica*, vol. 40, no. 10, pp. 1729 – 1738, 2004.

[105] P. J. Campo and M. Morari, "Robust model predictive control," *IEEE American Control Conference*, pp. 1021–1026, 1987.

[106] A. Bemporad and M. Morari, *Robust Model Predictive Control : A Survey*. Springer, 2007.

[107] B. Kouvaritakis and M. Cannon, *Model predictive control: Classical, robust and stochastic*. Springer, 2015.

[108] M. C. Steinbach, "Tree-sparse convex programs," *Mathematical methods of operations research*, vol. 56, no. 3, pp. 347–376, 2003.

[109] E. Klintberg, J. Dahl, J. Fredriksson, and S. Gros, "An improved dual newton strategy for scenario-tree mpc," *IEEE Conference on Decision and Control*, pp. 3675–3681, 2016.

John Dahl received the M.Sc. degree in Systems, Control and Mechatronics in 2013 from Chalmers University of Technology, Gothenburg, Sweden, where he is currently working toward the Ph.D. degree at the Department of Electrical Engineering. He is currently doing research with the Department of Electrical Engineering and Zenuity AB, Gothenburg. His research interests include active safety and estimation techniques, in particular, for threat-assessment and decision-making in collision-avoidance systems.

Gabriel Rodrigues de Campos received the Ph.D. degree in Automatic Control in 2012 from Grenoble University/Grenoble INP, France. He is currently a researcher with Zenuity in Gothenburg, Sweden. Prior to joining the Zenuity, he was a postdoctoral fellow with the Department of Signals and Systems, Chalmers University of Technology, Sweden and the Politecnico di Milano, Italy. His research interests include cooperative and distributed control, multi-agent systems, and intelligent transportation systems.

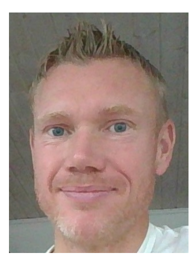

Claes Olsson received his PhD in 2005 in Automatic Control from Uppsala University, Sweden. The initial research focus was within the field of active vibration isolation applied to automotive design. Claes joined Volvo Car Corporation in 2005 where he carried out development and research for active safety systems and received the role as a Technical Expert on Collision Avoidance. He is currently with Zenuity AB acting as a technical expert within the field of ADAS and, in particular, collision avoidance technology.

Jonas Fredriksson received his M.Sc. in Computer Science Engineering from Luleå University of Technology, Sweden in 1997 and Ph.D. in Automatic Control from Chalmers University of Technology, Sweden in 2002. Currently, he is Professor in Mechatronics at the Department of Electrical Engineering at Chalmers. His research activities include modeling, control and simulation, with a special interest in automotive applications.