Adaptability of Automated Driving Systems to the Hazardous Nature of Road Networks

Yrvann Emzivat[*†], Javier Ibañez-Guzmán[†], Philippe Martinet[*] and Olivier H. Roux[*]

[*]École Centrale de Nantes, LS2N – UMR CNRS 6004, France
[†]Renault S.A.S., France

Abstract—**Full driving automation refers to the sustained and unconditional performance of the entire dynamic driving task and dynamic driving task fallback by an Automated Driving System (ADS). This implies that the ADS is expected to operate the vehicle under all driver-manageable on-road conditions, without any expectation that a user will respond to a request to intervene. Such a broad operational design domain includes those areas of the road network that are inherently hazardous, which is why the ability to detect them must be an integral part of ADS functionality. The aim of this paper is to give an Automated Driving System the means to identify these hazardous areas so that it can adapt its behaviour accordingly. Intelligent agents are tasked with the evaluation of the Automated Driving System's performance during operation and record any violation of safety constraints during periodically recurring travels. If the ADS is confronted to several hazardous situations in a given area of the road network, specific rules are applied to reduce the severity and the likelihood of occurrence of similar situations during later trips. Experiments were conducted in the driving simulation software SCANeR studio to illustrate the proposed approach.**

I. Introduction

Full driving automation refers to the sustained and unconditional performance of the entire dynamic driving task and dynamic driving task fallback by an Automated Driving System (ADS). Such a system is expected to operate the vehicle under all driver-manageable on-road conditions, without any expectation that a user will respond to a request to intervene. However, designing an ADS that operates safely in a wide variety of scenarios remains an ambitious task. Due to vehicle design limitations, an ADS will ultimately be confronted to hazardous situations that it was not designed to deal with. These situations may be the result of the inherent hazardous nature of some areas in the road network. While a human driver may initially behave in a dangerous or inconsistent manner on a road whose logic or geometry is unusual, when traffic signs are occluded or when they seemingly convey contradictory information, human drivers who travel in such areas on a regular basis are able to adapt their behaviour in a way that reduces risk. Automated Driving Systems are currently unable to do so, for they lack the necessary learning capabilities. The aim of this study is to give an ADS the means to identify such areas in its operational design domain during periodically recurring travels so that it can adapt its behaviour, as a human driver would. The emphasis lies on safety, which refers to losses due to unintentional actions caused by benevolent actors [1].

Hazard perception refers to one's situation awareness capabilities for the detection of dangerous situations in the traffic environment [2]. The recognition of dangerous situations has been studied within a cooperative group of autonomous vehicles in [3], but the approach is dependant on a highly available communication system. Other inventions rely on communication to broadcast information about the fact that accidents have occurred. For example, [4] proposes a dangerous road section analysis method to warn vehicles entering a dangerous road section, while [5] presents a module that can retrieve information about the hazardous nature of a road by communicating with a server. It is assumed here that the vehicle does not rely on its communication capabilities to identify hazardous areas, for they are prone to security and safety issues.

A method and system for screening potential traffic safety hazard road sections was investigated in [6]. Identifying hazardous locations can be carried via statistical models or by incorporating the spatial configuration by means of a local indicator of spatial association [7]. To the best of our knowledge, the identification of hazardous road locations [8] is contingent on the availability of traffic accident databases, such as the ones provided in [9] and [10]. However, there is no evidence that points towards those databases covering the entirety of the hazardous areas in the road network. Moreover, new hazardous areas are bound to emerge as the road network grows and changes. Even roadworks are prone to making roads hazardous for extended periods of time. After their completion, old lane markings sometimes remain visible and can be seen alongside new ones. This can puzzle human drivers as they try to figure out what the logic of the road is.

In the approach that is proposed here, the Automated Driving System determines the location of the hazardous areas in its operational design domain by itself. The ADS also rectifies its behaviour and develops strategies to avoid finding itself in hazardous situations during later trips.

This paper complies with the taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles, as defined by SAE International [11]. The remainder of the paper is organised as follows. Section II describes the general architecture of an ADS. Section III presents the proposed approach and describes the safety module in detail. Section IV provides two relevant use cases. Section V concludes this work.

II. SYSTEM ARCHITECTURE

A. General architecture

An Automated Driving System (ADS) typically includes the following functional modules [12]:

- A *localisation module*, whose purpose is to estimate the position of the vehicle in the *embedded navigation map*. This map contains information about the road geometry, the road type, the speed limits, the presence of traffic lights and traffic signs, as well as the location of schools and hospitals. It also provides contextual information that is used for situation understanding and autonomous navigation.
- A *perception module*, whose purpose is to interpret and understand the vehicle's surroundings. Embedded sensors gather data and deliver it to the *sensor fusion component* for processing. The recovered features are used to build a *world model*, which is the explicit representation of the world as the ADS knows it. In particular, the perception system detects dynamic and static obstacles (vehicles, pedestrians, cardboard boxes, etc), attempts to identify them and measures their speed.
- A *navigation module*, whose purpose is to guide the vehicle in the driving environment. It consists of a trajectory generation component, which generates a set of obstacle free trajectories, a *decision making module*, which picks the optimal manoeuvre or trajectory, and a *control system*, which executes the chosen trajectory through the use of actuators.

B. Additional components

The invention features a *safety module* (fig. 1), which consists of two components:

- A *recording component*, whose role is to observe and record any violation of safety constraints.
- A *rectification component*, whose role is to analyse previously encountered hazardous situations in order to adapt the behaviour of the ADS during later trips.

In order to avoid issues related to user confidentiality, the monitoring component of the safety module does not operate at all times. Moreover, storing unnecessary information might clutter memory. The rectification module operates whenever the ADS drives on a road where hazardous situations have been previously recorded.

C. Human machine interface

The user defines a list of destinations called *points of interest* through the human machine interface, by supplying the system with their address or GPS coordinates and giving them a label, such as 'work' or 'home'. It is hypothesised that trips between two points of interest are instances of periodically recurring travels, which is why the recording component is activated during such trips.

In practice, the user chooses a destination (by selecting the corresponding point of interest in the list). It is for the system to determine that the vehicle is initially located in the vicinity of a point of interest. For example, the vehicle might not be parked at the same location every evening, yet the ADS must be able to determine that the ongoing trip falls within the framework of daily commuting to work and that the monitoring component must be activated.

The monitoring component can also be activated during non-recurring travels. For example, if the ADS travels on a road section that is usually taken during daily commuting or if the ADS travels in the vicinity of a point of interest, then the monitoring component is activated for good measure.
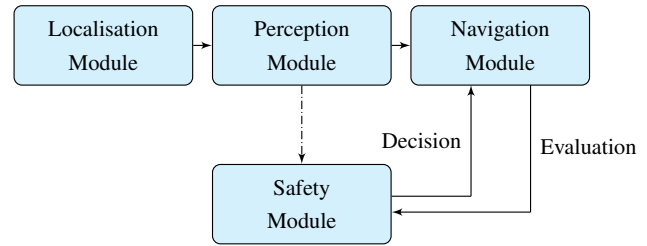

Fig. 1. Architecture of the Automated Driving System

III. OPERATIONAL BEHAVIOUR OF THE AUTOMATED DRIVING SYSTEM

A. Recording component

1) Agents: An *intelligent agent* is an autonomous entity that directs its activity towards achieving goals. It observes through sensors and acts upon an environment using actuators. Intelligent agents may also learn or use knowledge to achieve their goals. The decision making process can be interpreted as a periodic deliberation task among a set of agents, each of which supervises a set of conditions. *Safety agents* ensure that the ADS complies with the highway code (safety distances, maximum speed) and acts upon any indication that it is not behaving in a safe or correct manner (abrupt and repeated speed variations). *Non-critical agents* are responsible for the optimisation of travel time and user comfort. The decision is said to be *consonant* if the chosen trajectory is expected to verify all the safety constraints.

Besides their role in the decision making process, the agents are tasked with the evaluation of the Automated Driving System's performance during operation (fig. 1). A broken safety condition expresses a discrepancy between the confidence that the safety module placed in a chosen course of action and the subsequent hazardous situation, which occurred during its execution. In the following, it is assumed that the behaviour of the ADS is correct. In other words, the decision making process always leads to a consonant decision. An unfulfilled safety condition can thus be the result of the actions of a human driver or ADS that behaves in an unexpected or dangerous way or the indication of an inherent hazardous driving environment. It can also be linked to the transition phase of an event in traffic (fig. 2). For example, an overtaking manoeuvre that involves a vehicle moving in between the
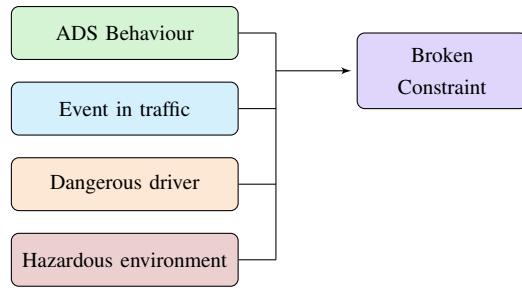
Fig. 2. Possible causes of a broken safety constraint

ADS-equipped vehicle in question and the vehicle in front of it might lead to a temporary violation of the safety distances.

2) Hazardous situations: Each safety agent is responsible for the verification of one safety condition. If the condition does not hold, a beacon is added to the map. The constraint violation c extends over a continuous area called a λ_c zone. A λ_c *zone* is characterised by the position of its endpoints, the time it took to cross, a list of relevant values that characterise the fact that the constraint did not hold along the zone, and other variables that describe the severity or magnitude of the constraint violation. For example, failing to maintain safety distances with the vehicle in front of the ADS-equipped vehicle will result in the creation of a list of relevant values containing the inter-vehicle distance measured at each verification period. Other useful information, such as time to collision values, are recorded as well. The agent also records relevant contextual data (time of the incident, season, weather and distribution of obstacles in the scene during and prior to the problem).

A *hazardous situation* involves at least one broken constraint and can be associated with multiple unfulfilled safety conditions. Therefore, it can either coincide with a single zone or it can encompass multiple, overlapping ones. Figure 3 shows a possible distribution of hazardous situations in the map after multiple trips. A recorded hazardous situation is represented by a dot whose radius corresponds to the severity of the encounter.

The perception module is tasked with tracking each obstacle in the scene and records the tactical decisions (intentions and manoeuvres) made by the dynamic obstacles during the last ten seconds. This gives the ADS *situation awareness* capabilities and helps it understand how the behaviour of each obstacle in the scene has led to the hazardous situation in the first place. Each hazardous situation has a *type*, which defines the problem. For example, 'safety distances are not verified (broken constraint) with a vehicle coming from road A after it positions itself in front of the ADS-equipped vehicle on road B (situation)' is a type of hazardous situation.

3) In practice: If a safety constraint is broken for the first time during the trip, then an instance of the class of objects 'Zone' and an instance of the class of objects 'Hazardous Situation' are created. The distribution of obstacles in the scene and their past actions are recorded, alongside relevant

information about the time of day, weather and traffic density.

During each subsequent constraint verification period, the agents check if the safety conditions hold. The order of magnitude of the verification period is that of the refresh rate of the world model. Each λ_c zone that is associated with a broken safety constraint that persists is updated. If a broken safety condition is verified once more, then the zone is saved by the recording component. If another condition is broken, then another zone is created.

A time duration is defined to discriminate hazardous situations from isolated problems, such as transient events in traffic. As soon as the safety constraints are verified once again, the hazardous situation instance is updated and archived. If it only lasted for a short period of time however, the instance is deleted as it means the ADS managed to recover from the problem rapidly. If a constraint is broken at a future date, then a new instance of the class of objects 'Hazardous Situation' is created with the correponding zone or zones. After several trips, a set of lists of hazardous situations is gathered. Each list is associated to one road section and sorted according to the direction of traffic.

B. Rectification component

1) Cover of the road network: The embedded navigation map is divided into *areas* that define a cover of the road network (as defined in mathematics). The cover is initially defined according to the road logic (intersection, roundabout, continuous road section). It is then further refined to encompass specific features of the environment (schools, hospitals, speed limit). The size of the refined areas depends on the range up to which such information remains relevant. As a result, a continuous road that joins two intersections can be broken down into multiple, potentially overlapping areas. Figure 3 shows an example of such a cover.

Areas are not to be mistaken with the previously defined λ_c zones. The cover of the road network in areas is defined a priori when the ADS is built, while hazardous situations and zones are recorded by the ADS during operation.

An inherent hazardous driving environment is expected to generate more hazardous situations on a given road section or area after multiple trips than one time encounters with random drivers that behave in a dangerous manner. Therefore, the ADS attempts to rectify its behaviour in a given area of a trip if a minimum number of hazardous situations have been recorded in that area during previous trips and during a certain time period. This number (for example 'five hazardous situations in a given area in one year' or 'two hazardous situations in a given area after one hundred trips') can be adjusted depending on the type of area and the type of hazard.

During operation, once the path between two points of interest has been computed, the safety module checks if it crosses an area that exhibits a high number of hazardous situations. If that is the case, it will invoke a *rule* from a library of fallback strategies (slow down, set local speed limit to a lower value, change lane, yield, forbid overtaking, etc), with respect to the type of hazardous situations that were
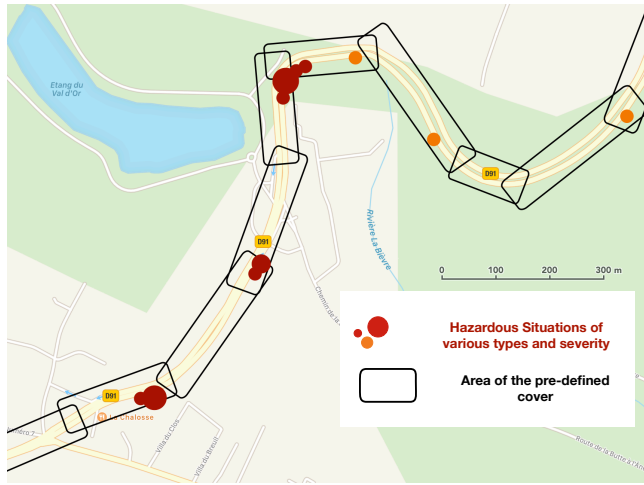
Fig. 3. Distribution of hazardous situations in the navigation map, after multiple trips

encountered. For example, if the ADS-equipped vehicle fails to maintain safety distances for a prolonged period of time despite aggressive braking, multiple times, near an area where two roads meet and if the situation awareness function of the recording component determines that the hazardous situation is related to the late detection of a vehicle that positions itself in front of the ADS-equipped vehicle, then the ADS will slow down before it reaches the intersection or change lane ahead of it. If multiple hazardous situations are linked to an overtaking attempt in a given area, then overtaking can be prohibited on that road section.

There may be several applicable rules for a given hazardous situation. These rules are sorted a priori in a certain order, according to certain features of the environment. For example, changing lane may have a higher priority near a highway ramp, while slowing down may have a higher priority near a school or hospital. A given rule can be applied in multiple ways. For example, braking can be done according to various deceleration profiles. The ADS chooses a given profile according to the number of recorded hazardous situations, their severity (given by the list of relevant values) and their distribution in the area.

A *neighbourhood* of the ADS-equipped vehicle is defined in order to measure the density of previously recorded hazardous situations in the vicinity of the vehicle. The ADS will attempt to rectify its behaviour if a minimum number of hazardous situations have been recorded in its neighbourhood. During operation, the size of the neighbourhood can vary according to the speed of the vehicle, the type of road, the distance to the next intersection, the curvature of the road, etc. The neighbourhood of the vehicle can contain several recorded hazardous situations that are of different types. A rule that applies to all of the previously encountered problems may be chosen over a rule that only applies to the most critical one.

2) Performance criterion: The architecture of the safety module is presented in figure 4. The performance criterion reflects the difference between the number of hazardous situations that occur before and after the new strategy has

Fig. 4. Inner workings of the safety module

been adopted. If the number of hazardous situations decreases significantly, then the adopted strategy is deemed satisfactory. Otherwise, the strategy is either reinforced or replaced by a lower priority rule.

The hazardous nature of an area of the road network is contingent on a number of factors. For example, the presence of patches of black ice depends on the weather and the season. A road can become more or less dangerous depending on the traffic density and related driving behaviours. Occluded traffic signs can make some roads even more hazardous during nighttime. Recording relevant contextual data allows the vehicle to add context to the occurrence of hazardous situations. Such data can also be used to check if the application of the selected rule is meaningful.

IV. CASE STUDIES

This section illustrates the practicality of the proposed approach by explicitly defining the thought process of the ADS in two simulated examples. For the sake of clarity, it is

Fig. 5. Case study A: Poor visibility at an intersection

TABLE I
EXPERIMENTAL RESULTS

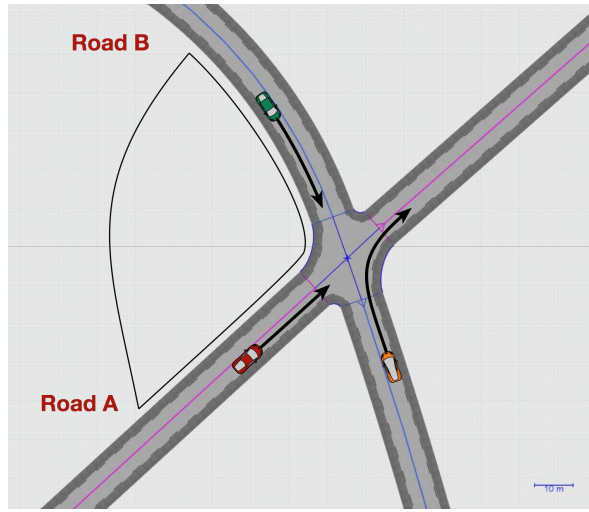| Strategy	None	32 km.h^{-1}	23 km.h^{-1}
Close calls (nb.h^{-1})	5.9	3.1	0.12
Collisions (nb.h^{-1})	1.5	0.8	0.01

assumed *for each scenario* that the hazardous situations were encountered in one area.

A. Hazards due to poor visibility near an intersection

Let us consider an Automated Driving System that travels on a priority road (denoted road A), which only has two lanes (one in both directions). The ADS is regularly confronted to hazardous situations near an intersection, which is represented in figure 5. Safety distances are not verified with vehicles coming from road B as they refuse to yield and either cross the intersection at high speed or position themselves in front of the ADS-equipped vehicle on road A. This corresponds to two different types of hazardous situations, as defined earlier. It is assumed here that these situations are due to poor visibility and a missing or occluded road sign on road B. The ADS has no knowledge of this fact and cannot possibly determine the origin of the problem. However, the ADS is able to infer the late detection of other vehicles at this intersection thanks to its situation awareness capabilities.

The corresponding scenario was created within the driving simulation software SCANeR studio as it could not be tested with a real autonomous vehicle due to the risk of damage to the prototype. Since the traffic module of the simulation software prevents such a scenario from occurring under normal conditions, the speed profiles of the vehicles driving on road B is defined artificially in order to bypass it. The speed limit is set to 50 km.h^{-1} on both roads. On average, when the other vehicles are detected, their speed is close to 45 km.h^{-1}, while the speed of the ADS-equipped vehicle nears 40 km.h^{-1}. Four vehicles were added to the simulation environment. They were given the behaviour of an ADS devoid of a safety module on road A and the behaviour of a driver that is unaware of the presence of the intersection on road B. Simulations ran for 12 hours. There were 18 collisions and 71 close calls. Most collisions are the consequence of both vehicles initiating an emergency break and colliding in the central part of the intersection.

The addition of the safety module in the proposed framework implies that a cover of the road network has been defined. In particular, the intersection is included in one of the areas of the cover. The ADS-equipped vehicle activates the rectification component after being confronted to several hazardous situations. The only applicable rule here is 'slowing down'. Time to collision (TTC) is defined as the time required for two vehicles to collide if they continue at their present speed and on the same path [13]:

$$TTC = \frac{\Delta L}{V_2 - V_1} = \frac{x_2 - x_1}{V_2 - V_1}, \qquad (1)$$

where ΔL is the distance between both vehicles, x_1 and x_2 are the lane abscissas of the vehicles, V_1 is the speed of the leading vehicle and V_2 is the speed of the following vehicle. Studies have shown that it is a fitting measure of the severity of conflicts on the road. TTC cannot be applied directly as the paths of the ADS-equipped vehicle and of the other vehicles are different. This notion can still be adapted if a vehicle positions itself in front of the ADS-equipped vehicle on road A. Assuming the speed of the vehicle is 45 km.h^{-1} when it becomes visible to the ADS-equipped vehicle, it can be determined with equation (1) that setting the speed of the ADS-equipped vehicle to 32 km.h^{-1} before that point guarantees a TTC criterion of 4 s can be achieved. If its speed is slightly lower than 45 km.h^{-1} however or if it keeps traveling on road B, then the proposed strategy is not satisfactory. Results are given per hour of simulation (Table I).

If the system is devoid of any means of elaborating a new strategy to take into account the other type of hazardous situation, the roads that lead to this intersection are prohibited by the safety module. Here, the chosen rule can be strengthened further to reduce risk. Table I shows that limiting the speed of the ADS-equipped vehicle to 23 km.h^{-1} prior to reaching the intersection allows the ADS to reduce risk significantly and can be used as a fallback strategy by the ADS.

B. Hazards near slip roads on a highway

Let us consider an Automated Driving System that travels on a highway. The ADS is confronted to hazardous situations near a slip road (fig. 6) as human drivers attempt to exit the highway, despite initially traveling on the left lane. Safety distances are not verified with respect to these vehicles, which results in aggressive braking. Such a scenario can be the result of poor signalling on the highway section that leads to this particular exit or to poor driving habits from users in this particular geographic area. Once again, the ADS is unable to infer the cause of the situation it is confronted to.

Fig. 6. Case study B: Hazardous situations near an exit ramp

TABLE II
EXPERIMENTAL RESULTS

Strategy	None	Lane change
Close calls (nb.h^{-1})	2.2	0.01
Collisions (nb.h^{-1})	1.6	0

The corresponding scenario was created within SCANeR studio. The speed limit is set to 130 km.h^{-1} on the highway. Vehicles were forced to attempt a lane change in order to reach the exit. The ADS drives near the slip road every other minute. Results are reported in Table II.

In the proposed framework, the slip road is included in one area, as part of the cover of the road network. After being confronted to the same hazardous situation multiple times, the rectification component of the safety module is activated by the ADS. Two rules can be applied here: 'changing lane' prior to the slip road and 'slowing down'. Given the context (driving on a highway), changing lane is given higher priority than slowing down (as would otherwise be the case near a school or hospital). When the vehicle should attempt to change lanes is the only remaining decision to make. By changing lanes 600 m before reaching the slip road, the Automated Driving System has more than 15 seconds to position itself. This strategy greatly reduces risk as shown in Table II. The only recorded close call is the result of the ADS being unable to change lane as too many vehicles attempted to exit the highway. Once the slip road has been left behind, the ADS can position itself in the right lane once again.

V. CONCLUSION

A system that allows an Automated Driving System (ADS) to adapt its behaviour to inherent hazardous areas of the road network during periodically recurring travels was presented. The system attempts to mimic the chain of thought of a human driver who learns from his past experience to reduce risk.

While the system was presented from a safety perspective, the method can also be adapted to improve the driving style of an Automated Driving System. Communication was proscribed as the safety module should not rely solely on the availability of such means. However, the data gathered by all the ADS-equipped vehicles could be retrieved to elaborate a hazardous behaviour map that is specific to Automated Driving Systems. The strategies of an ADS could then be shared with other vehicles that travel on inherent hazardous areas of the driving environment for the first time.

The goal of the proposed approach was to provide an ADS with situation awareness and learning capabilities to increase safety. This approach can be used in the context of full driving automation so that the ADS can adapt to situations that were not considered during its design phase. It also provides the means to understand how an ADS should behave in a hazardous areas. Future work includes a framework that takes into account uncertainties in the world model (such as the velocity and the position of objets in the environment). Other layers can be added to the situation awareness capabilities of the vehicle, such as the estimation of the intention of other drivers. Finally, further study of the classification of hazards in a given area with respect to relevant contextual hazardous information (such as the weather, time of day, season and traffic) is to be conducted as it should give the ADS more flexibility when adopting one fallback strategy over another.

REFERENCES

[1] N. G. Leveson, *Safeware: System Safety and Computers*. New York, NY, USA: ACM, 1995.
[2] M. S. Horswill and F. P. McKenna, *A cognitive approach to situation awareness: Theory and application*, ch. Drivers' hazard perception ability: Situation awareness on the road, pp. 155–175. Ashgate Publishing, Ltd., 2004.
[3] T. Batz, K. Watson, and J. Beyerer, "Recognition of dangerous situations within a cooperative group of vehicles," in *Intelligent Vehicles Symposium, 2009 IEEE*, pp. 907–912, IEEE, 2009.
[4] "Dangerous road section analysis method and dangerous road section analysis device," 2015. CN 201310731159.
[5] "System for correcting and providing real-time dangerous road-section information based on device in the car and its method," 2013. KR 20140119471.
[6] "Sections of road traffic safety hazard screening methods and systems," 2012. CN 201210426533.
[7] E. Moons, T. Brijs, and G. Wets, "Identifying hazardous road locations: hot spots versus hot zones," in *Transactions on Computational Science VI*, pp. 288–300, Springer, 2009.
[8] M. Bíl, R. Andrášik, and Z. Janoška, "Identification of hazardous road locations of traffic accidents by means of kernel density estimation and cluster significance evaluation," *Accident Analysis & Prevention*, vol. 55, pp. 265–273, 2013.
[9] I. Yamada and J.-C. Thill, "Comparison of planar and network k-functions in traffic accident analysis," *Journal of Transport Geography*, vol. 12, no. 2, pp. 149–158, 2004.
[10] I. B. Gundogdu, "Applying linear analysis methods to gis-supported procedures for preventing traffic accidents: Case study of konya," *Safety Science*, vol. 48, no. 6, pp. 763–769, 2010.
[11] SAE International, "Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles," 2016.
[12] A. Eskandarian, *Handbook of intelligent vehicles*. Springer London, 2012.
[13] R. van der Horst and J. Hogema, *Time-to-collision and collision avoidance systems*. University of Groningen, Traffic Engineering, 1993.