A new integrated collision risk assessment methodology for autonomous vehicles

Dr. Christos Katrakazas*
Chair of Transportation Systems Engineering,
Department of Civil, Geo and Environmental Engineering
Technical University of Munich
Arcistrasse 21, Munich, 80538, Germany
Tel: +49 (0) 89 28910463
Email: c.katrakazas@tum.de

Professor Mohammed Quddus
School of Architecture, Building and Civil Engineering
Loughborough University
Loughborough LE11 3TU, United Kingdom
Tel: +44 (0) 1509228545
Email: m.a.quddus@lboro.ac.uk

Professor Wen-Hua Chen
Department of Aeronautical and Automotive Engineering
Loughborough University
Loughborough LE11 3TU, United Kingdom
Tel: +44 (0)1509 227 230
Email: w.chen@lboro.ac.uk

*Corresponding author

Abstract

Real-time risk assessment of autonomous driving at tactical and operational levels is extremely challenging since both contextual and circumferential factors should concurrently be considered. Recent methods have started to simultaneously treat the context of the traffic environment along with vehicle dynamics. In particular, interaction-aware motion models that take inter-vehicle dependencies into account by utilizing the Bayesian interference are employed to mutually control multiple factors. However, communications between vehicles are often assumed and the developed models are required many parameters to be tuned. Consequently, they are computationally very demanding. Even in the cases where these desiderata are fulfilled, current approaches cannot cope with a large volume of sequential data from organically changing traffic scenarios, especially in highly complex operational environments such as dense urban areas with heterogeneous road users. To overcome these limitations, this paper develops a *new risk assessment methodology* that integrates a network-level collision estimate with a vehicle-based risk estimate in real-time under the joint framework of interaction-aware motion models and Dynamic Bayesian Networks (DBN). Following the formulation and explanation of the required functions, machine learning classifiers were utilized for the real-time network-level collision prediction and the results were then incorporated into the integrated DBN model for predicting collision probabilities in real-time. Results indicated an enhancement of the interaction-aware model by up to 9%, when traffic conditions are deemed as *collision-prone*. Hence, it was concluded that a well-calibrated collision prediction classifier provides a crucial hint for better risk perception by autonomous vehicles.

1. Introduction

Existing transport systems are not as economically efficient, as environmentally benign, nor as safe as they should be, and one key cause of this is due to the 'human element'. Human drivers are responsible for a 94% of the critical pre-collision events according to a recent survey from the National Highway and Traffic Safety Administration (Singh, 2015). Recent advancements in artificial intelligence, sensor fusion, vehicle technology and software algorithms have brought about the introduction of semi- or fully-autonomous vehicles closer to reality, especially in commercial fleets. Autonomous Vehicles (AVs) can learn, adapt, take decisions and act independently of human control and are, therefore, envisaged to make a profound impact on the economy, safety, mobility and society as a whole. Nonetheless, the most important advantage offered by AVs relates to improved road safety that is promised by researchers and manufacturers worldwide (Campbell et al., 2010).. A large number of traffic collisions and the related casualties could, therefore, potentially be reduced by removing the human involvement from the task of driving through the rapid uptake and penetration of AVs. Although AV technologies could deliver a *step change* in safety and mobility, they create new translational research challenges

In order to ensure the safety of its occupants and other traffic co-participants, an AV has to perform the sense-plan-act methodology in which *sensing* relates to understanding the surrounding environment, *planning* is the decision making and *acting* is actually moving the vehicle according to the planning (*Katrakazas et al, 2015)*. Possibly, ensuring safety in the planning module is the most complex in which a motion model generates a trajectory in the face of uncertainties at all levels. Two major challenges of the planning module prevail: (1) sensors may fail to detect what is happening around the vehicle and this may have a serious impact on the planning module and (2) vehicle software cannot plan for all the situations that the vehicle will possibly encounter. Consequently, addressing safety remains a pivotal challenge for AVs for both academia and industry worldwide. This is confirmed by recent incidents that resulted in three fatal collisions in the US and 60 collisions in the State of California according to their Department of Motor Vehicles as of April 2018. Examining their casual factors reveals that

AVs should be taught to understand not only what the surroundings are but also the context in order to enrich their situational awareness and decision making. Therefore, a planning module will take the circumstances or context into account rather than consider a vehicle as an independent entity, especially during the transition period from the fully manual to the fully autonomous driving era. Cases of contextual and circumferential aspects include: AVs drive through dense urban traffic, complex road settings, construction zones, residential streets where children suddenly appear and disappear by filtering through parked vehicles, segments with unstable traffic dynamics and hard-to-predict traffic co-participants, roads with traffic incidents such as vehicle breakdowns, traffic bottlenecks, network deficiencies and collision hot-spots. Even when AVs are doing everything they are supposed to, the underlying safety challenge would be how these factors could be taken into account in the collision-risk assessment of AVs.

Currently, a motion model is used to predict the intended trajectories of other vehicles and surrounding objects in a specific traffic environment and compare them with the trajectory of the interested AV in order to estimate the collision risk. Computational complexity, however, emerges when searching for an efficient trajectory representation in which vehicles are assumed to move independently (Agamennoni et al., 2012; Lefèvre et al., 2014). Recent approaches (e.g. Agamennoni et al., 2012; Gindele et al., 2015; Lefèvre, 2012) try to address the problem of risk assessment of AVs by taking into account contextual information (i.e. information on the traffic scene and the motion of other vehicles) as well as human-like reasoning about vehicles' interaction without predicting the trajectories of all other vehicles. The main method for making such predictions is the use of probabilistic models, especially Dynamic Bayesian Networks (DBNs) which are a robust framework for drawing an inference from the vehicle dynamics and the contextual information and can handle missing or erroneous data while maintaining real-time tractability (e.g. Murphy, 2012; Lefèvre et al., 2014). Nonetheless, perfect sensing or communications between vehicles are often assumed (Katrakazas et al., 2015; Paden et al., 2016).

The inherent limitations of robotics-based approaches on risk assessment in the context of organically changing dynamic road environments indicate that alternative methods should be sought as supplements for building a robust and comprehensive risk assessment module for an AV.

Over the past years, the estimation of the probability of a traffic collision occurring in real-time has also been studied by many researchers working in the traffic safety and traffic engineering perspective of Intelligent Transportation Systems (ITS). Real-time collision prediction for ITS is formulated on the basis that the probability of a collision's occurrence could be estimated from traffic dynamics during a short-time prediction horizon from data retrieved online (Abdel-Aty and Pande, 2005). The predominant technique of evaluating collision risk relates to comparing traffic measurements (e.g. speed, flow, occupancy) on a specific road segment just before a reported collision with traffic measurements from the same segment and time at normal situations (Pande et al., 2011). It can be understood that the traffic engineering perspective addresses the macroscopic problem of identifying a location with high-risk collision occurrence. This spatio-temporal risk could potentially provide a broader picture of the road network in terms of hazardous traffic conditions as an additional safety layer to AVs. An approach to bridge vehicle-level and network-level risk assessment is yet to be fully understood and utilised.

In order to realise the full benefits of AVs and to ensure that society is satisfied with this disruptive vehicular technology, its underlying safety challenge needs to be properly addressed. This paper directly tackles this challenge through a unique world-leading activity that incorporates fundamental concepts from the two schools of thought - *robotics* (vehicle-based) and *traffic engineering* (segment-based). The incorporation of this macroscopic spatio-temporal collision risk (henceforth termed as "network-level risk") into microscopic vehicle-level risk, therefore, forms the motivation of this current paper. This study offers a methodological expansion to existing DBN-based risk assessment of AVs with the aim of increasing their perception of the environment and easing online computations by exploiting real-time safety information for the road segment on which the ego-AV travels on. Such a risk assessment module can be embedded in the path or manoeuvre planning routines of autonomous vehicles, assuring a safe navigation of the ego-vehicle.

The rest of the paper is organised as follows: first, the existing literature and its main findings are synthesised. An analytic description of the proposed DBN for collision risk estimation in real-time is described next. This is followed by a presentation of the data needed for such an analysis and the methods used to estimate the risk of a collision. Results from machine learning classifiers (i.e. *k*-Nearest Neighbours, Neural Networks, Support Vector Machines, Gaussian Processes), used for network-level collision prediction and integrated with simulated and real-world vehicle-level data, are then presented. Finally, scenarios where the proposed model and network-level information in general could assist the safe navigation of AVs are given.

2. Literature Review

Risk assessment of AVs has been primarily addressed in the literature by utilizing different motion models (i.e. models that describe the movement of vehicles with regards to their surroundings). Lefèvre et al. (2014) presented a detailed survey to compare and contrast recent research on traffic environment modelling and prediction and introduced several risk estimators for intelligent vehicles. According to their work, motion models are classified into: (i) physics-based, (ii) manoeuvre-based and (iii) interaction-aware models. The first category of the motion models describes according to the laws of physics while the second one relies on estimating the intentions of the other traffic participants based on either clustered trajectories or manoeuvre estimation and execution. These two categories of motion models do not take the environment into account but rather consider vehicles as independent entities. Interaction-aware motion models exploit inter-vehicle relationships as to easily identify any dangerous situations in real-time.

Because of the incorporation of contextual information when modelling the motion of the vehicles in a traffic scene, interaction-aware models with regards to risk assessment is the focus of this literature review. It should, however, be noted that there is a dearth of research that integrate vehicle-level risk

assessment with the context-aware risk assessment in order to derive a more comprehensive risk assessment of AVs (Agamennoni et al., 2012).
As noted in the survey of Lefèvre et al. (2014), the vast majority of interaction-aware motion models are built using DBN models due to their capability of handling missing data efficiently, the simplistic representation of the relationship between the variables and the real-time tractability of the model for drawing an online inference.

Lefèvre (2012) pointed out that if an ego-vehicle has to predict all the future trajectories of the vehicles in its vicinity and to analyse them for any potential collisions, the whole process would become intractable for real-time applications. Her work exploited the power of interaction-aware models by the application of DBNs for the purpose of risk assessment at road intersections. Elegantly, instead of predicting the trajectories of all nearby vehicles, only vehicles which were found to disobey traffic rules or gap acceptance models were analysed for any potential collisions. It was however assumed that vehicular communications were enabled so as for the vehicles to exchange their spatial, speed and turning measurements through appropriate message delivery protocols. Nevertheless, an important observation was that collision risk does not only need intersecting trajectories but also behavioural or infrastructural information in order to enhance risk estimation for AVs. In the same principle, Worrall et al., (2012) showed the real-time efficiency of an interaction-aware model with the aid of DBNs. They constructed a fully probabilistic model based on a DBN using an improved calculation of the Time-to-Collision (TTC) variable for risk assessment. Their approach was, however, failed to handle complex traffic scenarios; for instance, “give-way” at non-signalised junctions. Moreover, communications were again assumed to be available and the approach was actually tested on mining facilities which could not efficiently represent traffic dynamics on real-world road networks.

Recent approaches were formulated to better describe the traffic environment by including network-related information. Gindele et al. (2015), for instance, included information on car-following models and the interactions among the vehicles in the adjacent lanes so as to faster recognise the intention of each vehicle and assessed risk using the TTC metric. Their DBN approach requires many variables which consequently need to be trained to efficiently describe, for example, the relationship between traffic participants, the influence of traffic rules to traffic participants, the influence of the geometry of the road on the actions. In order to address some of these issues, Kuhnt et al. (2015) proposed to use a static street model in order to provide an extra hint to a motion model. Their approach, however, fails to provide an efficient description of the inter-vehicle dependencies. Recently, Bahram et al. (2016) showed that even without vehicular communications, if the knowledge of the road geometry and traffic rules is available, the prediction time for anticipating the manoeuvres of other vehicles can be significantly improved. Nevertheless, network-level knowledge was limited to train classifiers that have the capability of detecting any manoeuvre associated with the acceleration and deceleration of vehicles as well as lateral offsets in relation to the centre-line of a lane.

It can be concluded from the literature that interaction-aware motion models have gained attention in modelling the inter-relationship between the participants of a traffic scene explicitly. However, complex traffic scenarios are difficult to tackle and learning specific manoeuvres of the drivers and classifying them as *safe* or *dangerous* are time-consuming due to the massive datasets needed. In order to address these challenges, traffic–related information is starting to become part of these models but their complexity and assumptions may hinder a comprehensive but simple representation of the traffic environment. Last but not the least, although network-level collision prediction has been researched over the years, an approach to bridge vehicle-level and network-level risk assessment is yet to be fully understood and utilised.

The overriding objective of this paper is, therefore, to address this methodological gap by extending typical DBN-formulations based on the principles of interaction-aware motion models aided by network-level collision risk prediction as an additional safety layer. The purpose is to enhance the overall risk assessment method of AVs with a particular focus on faster predictions and more comprehensive reasoning. The work builds on previous research (i.e. Lefèvre, 2012 and Worrall et al. 2012) which showed that such methods can be efficiently implemented in real-time while keeping the complexity of the DBN motion model as low as reasonably practicable.

3. Methodological background

The focus of this study is to integrate network-level collision prediction with interaction-aware motion models under a Bayesian framework for risk assessment of AVs. Time-varying traffic scenes have to be modelled appropriately allowing an ego-AV to reliably estimate the collision risk from the presence of surrounding vehicles as well as the interactions between these vehicles that are deemed to pose the greatest threat. Therefore, an appropriate framework for modelling dynamic systems must be applied.

Data acquisition for AVs is dependent on the temporal frequency of their built-in sensor unit. As a result, input data to the risk assessment algorithm are inherently sequential.

Murphy (2002) indicated that state-space models such as Hidden Markov Models (HMMs) and Kalman Filter Models (KFMs) perform better in sequential data problems associated with finite-time windows, discrete and multivariate inputs or outputs and they can be easily extended. A known drawback of HMMs is that they suffer from high sample and high computational complexity. This means that learning the structure of the model and inferring the required probability may take longer to accomplish. Furthermore, simple HMMs require a single discrete random variable which cannot cope with the description of a constantly changing environment such as a traffic scene. Factorial HMMs and coupled HMMs enable the use of multiple data streams but the former has problems related to the correlation between the hidden variables and the latter needs the specification of many parameters in order to perform an inference (Murphy, 2012). KFMs rely on the assumption that the system is jointly Gaussian which makes it inappropriate to jointly accommodate both discrete and continuous variables (Murphy, 2002).

In order to overcome the above limitations in handling sequential data, Murphy (Murphy, 2002) proposed the use of DBNs. DBNs are an extension of Bayesian networks which is a graphical representation of a joint probability distribution of random variables to handle temporal sequential data (e.g. Koller and Friedman, 2009). DBN representation of the probabilistic state-space is straightforward and requires the specification of the first time slice, the structure between two time slices and the form of the Conditional Probability Distribution (CPDs). A crucial part in defining a DBN is the declaration of hidden (i.e. latent) and observed variables.

When applied for the anticipation of the motion of the vehicles and risk assessment for automated driving, a typical DBN layout that takes the inter-vehicle dependencies into account is shown in Figure 1 (Lefèvre, 2012). The DBN requires the definition of three layers:

Layer 1: the highest level corresponds to the context of the vehicle's motion. It can be seen as a symbolic representation of the state of the vehicle (Agamennoni et al., 2012). It can contain information about the manoeuvre that the vehicle performs (as seen in Lefèvre, 2012) or the geometric and dynamic relationships between vehicles (as seen in Agamennoni et al., 2012). The variables contained in this level are usually 'discrete' and 'hidden' (e.g. manoeuvre undertaken or compliance with traffic rules).

Layer 2: this level corresponds to vehicle's physical state such as kinematics and dynamics of the vehicle. It usually includes information about the position, the speed and the heading of the vehicle but can also accommodate information coming from a dynamic model for the motion of the vehicle (e.g. the bicycle model). The variables contained in this level are usually 'continuous' and 'hidden' (e.g. speed, position, acceleration)

Level 3: the lowest level corresponds to the sensor measurements that are accessible (e.g. measured speed of the ego-vehicle). The measurements are processed in order to remove noise and create the physical state subset. The variables at this level are always 'observable'.

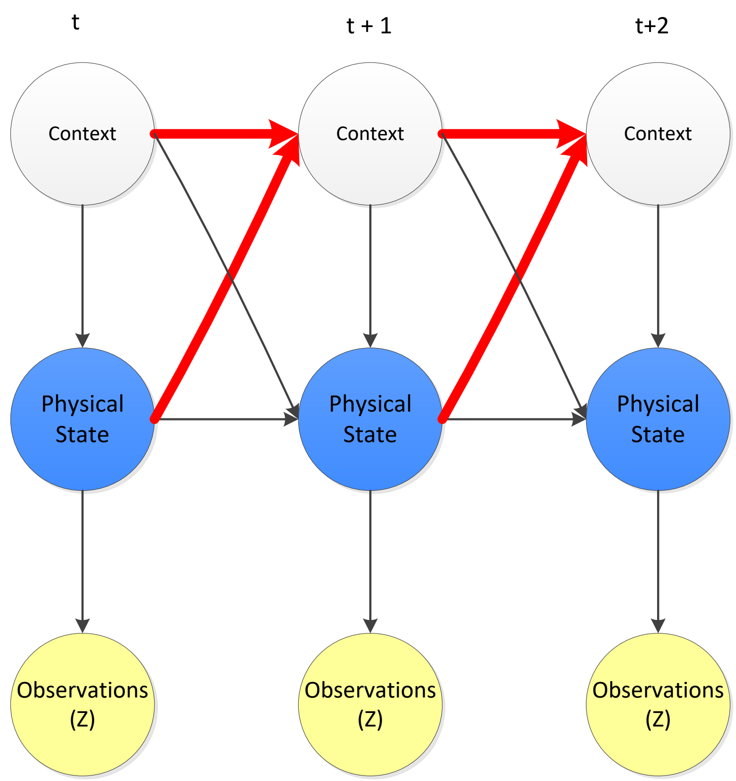

Figure 1: Graphical representation of a typical DBN-based interaction aware

In Figure 1, it is noticeable that for every time moment the specific context of each vehicle influences the physical state of the vehicle and consequently the physical state is depicted on the observations from the sensors. Accordingly, it is apparent from the thick solid arrows that the context of each vehicle at a specific time slice is dependent on the context and the physical state of every vehicle in the traffic scene at the previous time slice. This means that the probability of a vehicle belonging to a specific context in the next time slice requires the estimation of the union of probabilities which describe the context for each of the vehicles in the scene along with the probability distributions of variables related to their physical states. For more clarity, assume that an ego-vehicle is travelling in the middle lane of a motorway and senses that a lead vehicle on the left lane intending to change its lane. Based on the traffic rules, it is logical to assume that the ego-vehicle would slow down or change its lane to the right. If there is a vehicle in the right lane, then the context of "slowing-down" would have a higher probability than the context of "change its lane to the right" or "change its lane to the left" and the differences in the context would depend on the physical measurements of all vehicles in the scene (i.e. the position and speed of the ego-vehicle and the other two vehicles).

To enhance risk assessment for automated driving without increasing the complexity of such DBN-based interaction-aware motion models, a new structure is developed in this paper by incorporating an additional layer that deals with network-level collision risk.

4. Developed DBN model for motion prediction and risk assessment

In order to include the network-level collision prediction in the motion prediction and risk assessment routine, a new layer along with its relationship with other layers are introduced as depicted in Figure 2.

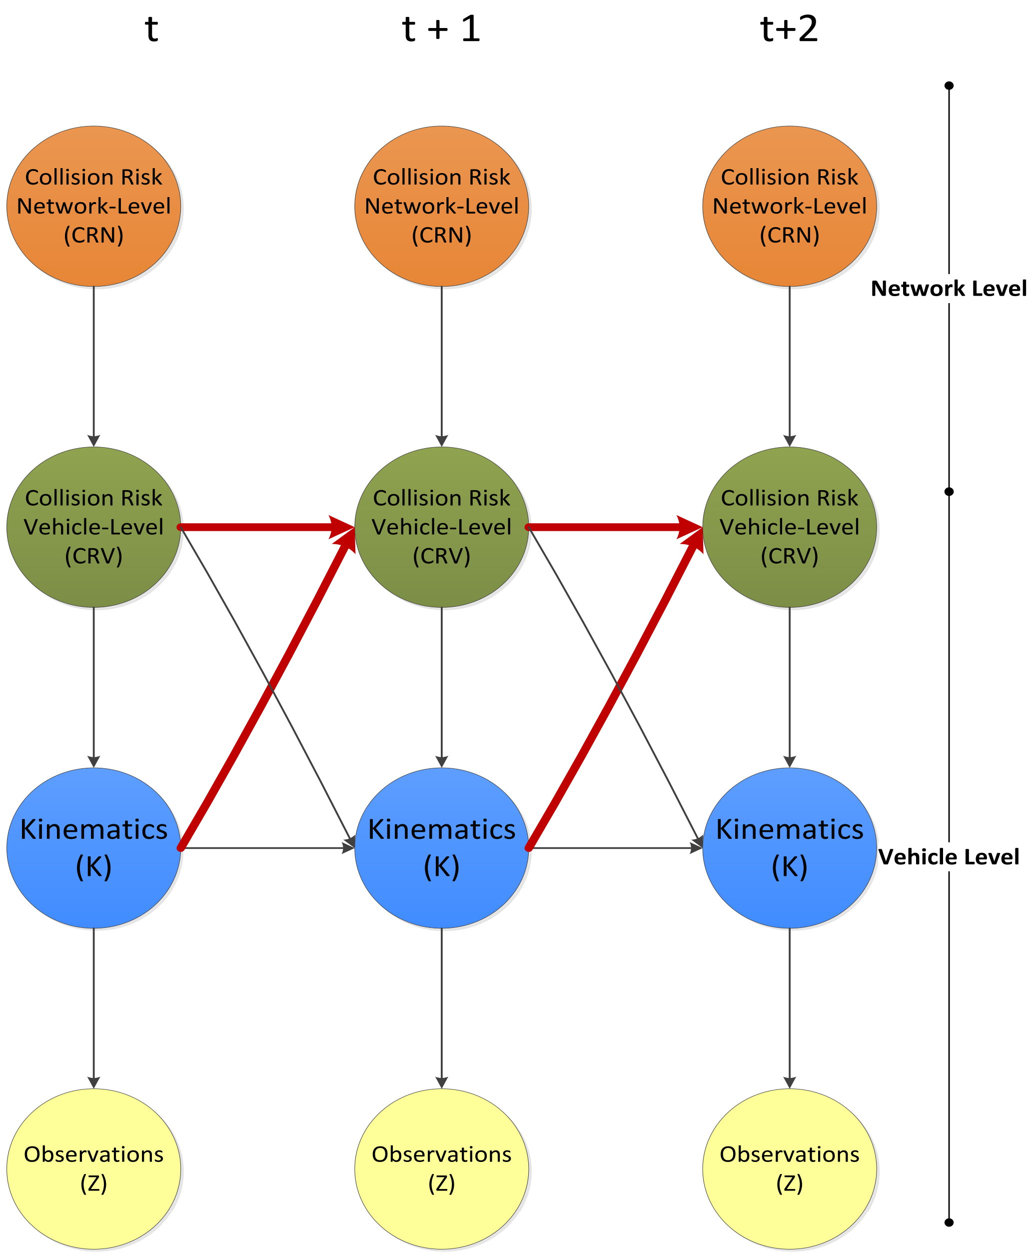

Figure 2: Developed DBN Network

Comparing Figures 1 and 2, it can be observed that the context layer is broken into two interacting safety-related domains: (i) network-level collision risk and (ii) vehicle-level risk. The topology of the DBN is designed in such a way that it accurately represents the dependencies between the layers: i) if any safety risk is identified at a network-level, it should be depicted in the vehicle-level; ii) the vehicle-level safety risk is depicted on the motion of the vehicles, and iii) the motion of the vehicles is depicted on the observations from the sensors. The model presented above could, in theory, be applied to any traffic situation by defining the variables CRN, CRV, K, and Z accordingly.

4.1. Variable definitions

Network-level real-time collision risk (CRN): Represents the safety context of the road segment on which the ego-vehicle is travelling on (i.e. whether the traffic conditions on the road segment are *collision-prone* or *safe*). The variable in this layer is '*discrete*' taking only two values:

1. *Safe traffic conditions*
2. *Collision-prone traffic conditions*

As a result, $(\mathrm{CRN}_{\mathrm{n}}^{\mathrm{t}})$ indicates the probability that the traffic conditions on the road segment (with lengths 300-500m as indicated by Pande et al., 2011) on which a vehicle n travels at time t are "*collision-prone*" or "*safe*" based on traffic dynamics. The input variables for estimating network-level collision risk consist of aggregated traffic conditions data (e.g. the mean speed of the vehicles, the mean number of the vehicles, the mean occupancy). Because many vehicles are travelling on a road segment, it is assumed that once the network-level collision risk is estimated for the segment, then its value is the same for all the vehicles in this specific segment.

Vehicle-level risk (CRV): Represents the safety context of one vehicle in a traffic scene (i.e. whether a vehicle can potentially cause a collision with the ego-vehicle). The variable in this layer is also 'discrete' but takes four values describing the safety context of each vehicle depending on the network-level safety context:

1. *Safe driving on a road segment having safe traffic conditions*
2. *Safe driving on a road segment having collision-prone traffic conditions*
3. *Dangerous driving on a road segment having safe traffic conditions*
4. *Dangerous driving on a road segment having collision-prone conditions*

"*Safe*" and "*Dangerous*" driving can be a user-defined function and indicate the characterization of the manoeuvres undertaken by the vehicles in the traffic scene. Safe driving does not pose a threat to another vehicle, while dangerous driving indicates that the motion of one vehicle could be considered *unsafe* by another vehicle in the traffic.

From Figure 2 it can also be observed that the estimation of the vehicle-level safety context depends on the network-level safety context as well as the union of safety contexts and kinematics of all the vehicles in the vicinity of the ego-vehicle. Consequently, network-level collision prediction provides a hint to the estimation of vehicle-level collision probabilities in which the multi-vehicle dependencies are taken into account.

<u>Sensor measurements (Z):</u> Represents the available observations from the sensors of the ego-vehicle. Z_n^t denotes the available measurements that describe the state of the vehicle *n* at time *t*. The variables in this layer are 'continuous'.

The measurements for each vehicle are assumed to include:

$\mathrm{Pm}_n^t = (\mathrm{X}_n^t \mathrm{Y}_n^t, \theta_n^t) \in \mathbb{R}^3$: the measured lateral and longitudinal position (X_n^t, Y_n^t) and heading of the vehicle (θ_n^t)

$\mathrm{Vm}_n^t \in \mathbb{R}$: the measured speed of the vehicle

<u>Kinematics of the vehicles (K)</u>: Represents the physical state of a vehicle. $\boldsymbol{K_n^t}$ denotes the conjunction of all the variables that describe the physical state of the vehicle *n* at time *t*. The variables in this layer are continuous as they are referring to continuously measured quantities such as position and speed.

Based on the available measurements described previously, the following variables are selected to represent the physical state of a vehicle:

$\mathrm{P}_n^t = (X_n^t Y_n^t, \theta_n^t) \in \mathbb{R}^3$: the real values of the position and heading of the vehicle

$V_n^t \in \mathbb{R}$: the real value of the speed of the vehicle

4.2. Joint Distribution

For the proposed DBN depicted in Figure 2 the joint distribution of all the vehicles is estimated as (Bessiere et al., 2013):

$$\begin{aligned}
&\boldsymbol{P}\left(\boldsymbol{CRN^{0:T}}, \boldsymbol{CRV^{0:T}}, \boldsymbol{K^{0:T}}, \boldsymbol{Z^{0:T}}\right) \\
&\quad = \boldsymbol{P}\left(\boldsymbol{CRN^{0}}, \boldsymbol{CRV^{0}}, \boldsymbol{K^{0}}, \boldsymbol{Z^{0}}\right) \prod_{\boldsymbol{t=1}}^{\boldsymbol{T}} \prod_{\boldsymbol{n}}^{\boldsymbol{N}} P(\mathbf{CRN}_{\boldsymbol{n}}^{\boldsymbol{t}}) \times P(\mathrm{CRV}_{\boldsymbol{n}}^{\boldsymbol{t}} | \mathbf{CRV}_{\boldsymbol{N}}^{\boldsymbol{t-1}} \mathbf{K}_{\boldsymbol{N}}^{\boldsymbol{t-1}} \mathrm{CRN}_{n}^{t}) \\
&\quad \times P(\mathrm{K}_{\boldsymbol{n}}^{\boldsymbol{t}} | \mathrm{CRV}_{n}^{t-1}\ \mathrm{K}_{n}^{t-1} \mathrm{CRV}_{n}^{t}) \times P(\mathrm{Z}_{n}^{t} | \mathrm{K}_{n}^{t}) \qquad (1)
\end{aligned}$$

where n is the vehicle ID number in the vicinity of the ego-vehicle, t is the time moment, T is the total time duration of the measurements and N is the total number of vehicles that are observed in the traffic scene. Bold letters indicate that the indicated layers are calculated for all the vehicles. For example, $\mathbf{CRV}_{\boldsymbol{N}}^{\boldsymbol{t-1}}$ indicates the vehicle-level risk context for time *t-1* for all the vehicles in the traffic scene.

4.3. Estimating the risk of collision by using a hint from network-level risk prediction

Modelling the motion of the vehicles with regards to network- and vehicle-level risks requires a new estimation framework to be developed. In order to quantify the influence that network-level risk estimation has on estimating vehicle-level collision risk, it is essential to infer the probability that there is a vehicle-level “unsafe” situation, given the hint from the network and the measurements from the sensors.

In the majority of recent studies on network-level collision prediction (e.g. Sun and Sun, 2015), traffic conditions at 5-10 minutes before the collision are deemed to be the most suitable to identify collision events timely and initiate an intervention by the responsible traffic agencies. However, 5 to 10-minute aggregation may not suitable for the real-time safety assessment of AVs where sensor information is available at a higher sampling frequency (e.g. 1 Hz, 0.1 Hz). It is, however, a reality that traffic agencies aggregate traffic data at pre-defined time intervals (e.g. 30-second or 1-minute, 5-minute and 15-minute). Because of the difference at the temporal horizon between network-level collision prediction and vehicle-level measurements, it is assumed that the CRN layer is an observable layer. CRV and K are hidden layers because the variables in these layers are inferred through the vehicle's sensor measurements. The sensor measurements layer (Z) is obviously an observable layer.

Exact inference in such non-linear and non-Gaussian models is not tractable. Therefore, in order to estimate the probability of a "dangerous" vehicle-level context given the traffic situation and the sensor measurements the use of particle filters (Merwe et al., 2000) is proposed as they have been proven to work well in similar situations (Lefèvre, 2012; Murphy, 2002).

If an inference algorithm is chosen, then the probability to be inferred is:

$$P\left([CRV_n{}^t \in \{dCP, dSA\}] | CRN_t, Z_{0:t}\right) > \lambda \qquad (2)$$

where:

- $\boldsymbol{CRV_n{}^t}$ denotes the vehicle-level safety context of vehicle *n* at time *t*;
- $\boldsymbol{dCP}, \boldsymbol{dSA}$ denote a "**d**angerous" vehicle travelling on a road segment with **C**ollision-**P**rone traffic conditions and a "**d**angerous" vehicle travelling on a road segment with **SA**fe traffic conditions respectively;
- $\boldsymbol{CRN_t}$ denotes the network-level collision risk for all the vehicles on a specific road segment;
- $\boldsymbol{Z_{0:t}}$ denote the sensor measurements until time moment t;
- $\boldsymbol{\lambda}$ is a threshold to identify "dangerous" encounters between the surrounding traffic participants and the ego-vehicle.

Equation 2 indicates that given a hint for the safety assessment of a road segment, the motion of the vehicles in that specific segment is affected. This resembles the fact that human drivers are also affected when the information of traffic incidents such as a broken-down vehicle on the roadway or a queue formation in the downstream is displayed via Variable Message Signs.

4.4. Note on the similarities and differences with other probabilistic models

The model depicted in Figure 2 bears resemblance to a Switching State Space Model (SSSM) with regard to explaining the dynamics of the traffic scene by switching between a discrete numbers of contexts. In SSSMs the switching process is regulated by a discrete Markov process which indicates which context is active at every time step. However, in the proposed model, this switching process is conditionally Markov, because the context variable in the vehicle level (CRV) depends not only on the discrete variable of the previous time step but on the continuous kinematics of the vehicles of the previous time step.

The structure of the proposed model also resembles a Coupled Hidden Markov Model (CHMM) (Brand et al., 1997) because of the way the different time slices connect. In CHMMs the current hidden layer depends on the hidden layer in the previous time step as well as the hidden layer of a neighbouring Markov Chain. However, CHMMs are usually intended for maximum likelihood estimation, while this work emphasizes on prediction. The obvious difference with CHMMs is that the proposed model accommodates continuous nodes, whereas CHMMs only work with discrete-valued variables. Furthermore, the use of CHMMs for solving the problem this work tackles introduces computational complex, as a different CHMM should be constructed for each interaction between two vehicles.

4.5. Parametric forms

In order to estimate the joint distribution of the network for inference, the functions that calculate each of the probabilistic distributions of each layer need to be defined. since the focus of the approach is the incorporation and enhancement of network-level collision prediction into existing motion models for automated driving a brief description of the parametric forms for vehicle-level risk and there are a large number of variables for the problem, kinematics and sensor measurements are presented.

4.5.1. Vehicle-level risk $P({CRV_n}^t)$

The content of vehicle-level risk is derived from the previous vehicle-level risk context and kinematics of all the vehicles on the scene, and is influenced by the current network-level collision prediction. The estimation of the probability that the motion of one vehicle is considered “dangerous” or “safe” is derived through a feature function that takes as inputs the current network-level risk, the previous vehicle-level risk context of the vehicle and the previous vehicle kinematics:

$$P\left(\mathrm{CRV}_{\boldsymbol{n}}^{\boldsymbol{t}} \middle| \mathbf{CRV}_{\boldsymbol{N}}^{\boldsymbol{t-1}} \mathbf{K}_{\boldsymbol{N}}^{\boldsymbol{t-1}} \mathrm{CRN}_{n}^{t}\right) = \boldsymbol{f}(\mathbf{CRV}_{\boldsymbol{n}}^{\boldsymbol{t-1}}, \mathbf{K}_{\boldsymbol{n}}^{\boldsymbol{t-1}}, \mathrm{CRN}_{n}^{t}) \tag{3}$$

In order for this feature function to be defined, three steps need to be considered:

a) Using a Kalman Filter (Murphy, 2012), the physical state of the vehicles in the traffic scene is estimated. For example, after applying a Kalman filter algorithm, the elements $\{X_{ego}^{t}, Y_{ego}^{t}, \theta_{ego}^{t}, v_{ego}^{t}\}$ and $\{X_{n}^{t}, Y_{n}^{t}, \theta_{n}^{t}, v_{n}^{t}\}$ will be known. v_{ego}^{t} *and* v_{n}^{t} denote the speeds of ego-vehicle and vehicle-n respectively.

   If Δp_t denotes the relative position between ego-vehicle and vehicle-*n*, and Δv_t denotes the relative speed between ego-vehicle and vehicle n then the time-to-collision (TTC) and the distance-to-collision (δ) between the ego-vehicle and vehicle-n are expressed as follows (Agamennoni et al., 2012):

   Time to collision: $\mathrm{TTC}_{\mathrm{n}}^{\mathrm{t}} = \frac{\Delta p_t^T \Delta v_t}{\Delta v_t^T \Delta v_t}$ (4)

   Distance to collision: $\delta_{\mathrm{n}}^{\mathrm{t}} = \sqrt{\Delta p_t^T \Delta p_t - \mathrm{TTC}_{\mathrm{n}}^{\mathrm{t}} \Delta p_t^T \Delta \mathrm{v}_{\mathrm{t}}}$ (5)

   If $P_n^t = (X_n^t, Y_n^t, \theta_n^t)$ denote the position and heading of vehicle *n* at time moment *t* and v_n^t denotes the speed of the vehicle, an indicator function (f_K) can indicate if vehicle-*n* brakes dangerously, changes lane dangerously or drives safely with regard to the ego-vehicle. For rear-end collisions TTC-based thresholds could be of use (e.g. Toledo et al., 2003):

$$f_K = f(\mathrm{TTC}_{\mathrm{n}}^{\mathrm{t-1}}) = \begin{cases} 1\text{: dangerous } if\ \mathrm{TTC}_{\mathrm{n}}^{\mathrm{t}} < Critical\ TTC \\ 0: safe; otherwise \end{cases} \tag{6}$$

b) If a vehicle in the previous time epoch was indicated as "dangerous" in the road segment that the ego-vehicle is driving on, then it is assumed that the CRV context was "dangerous". Otherwise, it is assumed that the motion of all the vehicles was "safe". Thus, another indicator function to take the previous vehicle-level risk of all vehicles into account can be defined as:

$$f_{CRV_N} = \begin{cases} 1 \; if \; \sum_{n=1}^{N} CRV_n^{t-1} > 0 \\ 0, otherwise \end{cases} \tag{7}$$

where N is the total number of vehicles that the ego-vehicle can sense.

c) In order to take network-level collision risk into consideration and easily identify dangerous traffic participants, the network-level classification metrics are considered as a coefficient:

d) $$f_{CRN_n} = \begin{cases} \frac{Accuracy+Recall}{2} \; if \; \mathrm{CRN}_N^t = dangerous \; and \; {f_{CRV_N}}^{t-1} = 1 \\ 1 - \frac{Accuracy+Specificity}{2} \; if \; \mathrm{CRN}_N^t = safe \; and \; {f_{CRV_N}}^{t-1} = 0 \\ 1 - recall \; if \; \mathrm{CRN}_N^t = safe \; and \; {f_{CRV_N}}^{t-1} = 1 \\ 1 - specificity \; if \; \mathrm{CRN}_N^t = dangerous \; and \; {f_{CRV_N}}^{t-1} = 0 \end{cases} \tag{8}$$

By that definition, if a vehicle is detected to pose a threat (i.e. dangerous) and the traffic conditions are collision-prone, a compromise between the accuracy of the classifier and its recall is boosting the identification of a hazardous road user. If traffic conditions are indicated as safe, then the compromise is made between the accuracy and the specificity of the classifier which shows its ability to correctly classify safe traffic conditions. Afterwards, this compromise is subtracted from 1 to indicate the probability of a vehicle being dangerous. When the network-level classifier indicates safe traffic but a vehicle is sensed to be posing a "threat" to the ego-vehicle, then the prediction is boosted by the false negative rate (given by the formula: $1 - recall$). Lastly, when traffic conditions are indicated as dangerous but there is no vehicle posing a threat, then the vehicle-level risk is boosted by the false alarm rate (i.e. $1 - specificity$).

Having all three indicative functions, the probability of the current vehicle-level collision risk context could be calculated as shown in the following example:

$$P(\mathrm{CRV}_{\boldsymbol{n}}^{\boldsymbol{t}} = "dCP\ or\ dSA" \big| \mathrm{CRV}_N^{t-1}\mathrm{K}_N^{t-1}\mathrm{CRN}_n^t) = \frac{\sum_{n=1}^{N}(f_{K_n}=1)+\sum_{n=1}^{N}(f_{CRV_n}=1)+f_{CRN_N}}{3N} \tag{9}$$

where *N* is the total number of vehicles that the ego-vehicle can sense. *3N* is chosen as a normalising factor in order for the probability to be within [0,1] even when one vehicle is posing a threat (i.e. $\sum_{n=1}^{N}(f_{K_n}) = 1, \sum_{n=1}^{N}(f_{CRV_n}) = 1\ and\ f_{CRN_N} = 1$). It is assumed that the sampling and risk estimation frequencies will be adjusted as soon as a risk is estimated.

4.5.2. Kinematics $P(K_n^t | CRV_n^{t-1}\ K_n^{t-1} CRV_n^t)$

The variables describing the kinematics layer must contain all the information needed in order to characterise the contexts. In this work, it was explained that the physical state vector will contain information on the position of a vehicle (in an absolute reference system, its heading and its speed). It is assumed that vehicles move according to the bicycle model as shown in Figure 3 (Snider, 2009). The kinematic bicycle model merges the left and right wheels of the car into a pair of single wheels at the centre of the front and rear axles as seen in Figure 3. It is assumed that wheels have no lateral slip and only the front wheel is steerable.

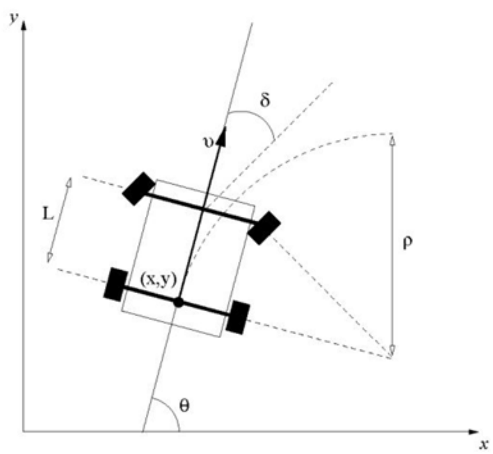

Figure 3: Bicycle model kinematics

The equations of motion for all vehicles in the traffic scene can be integrated over a time interval Δt using a simple forward Euler integration method (Press et al., 1993) in order to acquire the evolution of kinematics over time.

In the proposed model in Figure 3 and in its joint distribution as shown in Equation (1), it is observed that the current kinematics depend on the previous and current vehicle-level risk context as well as on the current kinematics of the vehicle. It is assumed that vehicles moving in a specific context will follow kinematics according to that context. As a result, the parametric forms of the position, heading, and speed of each of the vehicles should be defined according to the current vehicle context and the previous kinematics only. For example:

$$\mathrm{P}(\mathrm{P}_n^t|\mathrm{CRV}_n^{t-1}\ \mathrm{K}_n^{t-1}\mathrm{CRV}_n^t) = P(\mathrm{P}_n^t|\mathrm{CRV}_n^t\mathrm{K}_n^{t-1}) \tag{10}$$

In order to expose the dependency of current kinematic measurements on the previous vehicle-level safety context, context-specific constraints (e.g. constraints on the TTC between ego-vehicle and another vehicle) should be defined to distinguish between contexts. For example, if the derived TTC is below 1 second, this could indicate a “dangerous driving” in a road segment with safe or collision-prone traffic conditions. The parametric forms of the probability distribution of position and speed of the vehicles can be assumed to follow normal distributions (Lefèvre, 2012).

For example, the likelihood of the position and heading of a vehicle is defined as a tri-variate normal distribution with no correlation between x, y, and θ

$$P(\mathrm{P}_n^t|[\mathrm{CRV}_n^{t-1} = C_i][\mathrm{P}_n^{t-1} = X_n^{t-1}Y_n^{t-1}, \theta_n^{t-1}][\mathrm{V}_n^{t-1} = v_n^{t-1}]) = N\big(\boldsymbol{\mu}_{\boldsymbol{xy\theta}}(X_n^{t-1}Y_n^{t-1}, \theta_n^{t-1}, C_n), \boldsymbol{\sigma}_{\boldsymbol{xy\theta}}\big) \tag{11}$$

where $\boldsymbol{\mu}_{\boldsymbol{xy\theta}}(X_n^{t-1}Y_n^{t-1}, \theta_n^{t-1}, C_n)$ is a function which computes the mean position and heading of the vehicle $(\mu_x, \mu_y, \mu_\theta)$ according to the bicycle model and the context-specific constraints, $C_n$denotes the context of vehicle-n and $\boldsymbol{\sigma}_{\boldsymbol{xy\theta}} = (\sigma_x, \sigma_y, \sigma_\theta)$ is the standard deviation which can be acquired from the covariance matrix of the Kalman Filter algorithm.

4.5.3. Sensor measurements $(Z_n^t|K_n^t)$

The sensor model used is adopted from (Agamennoni et al., 2012) because of the use of the Student *t*-distribution which performs better with outlier data. The sensor model can be defined as:

$$P(Z_n^t/K_n^t) \sim Student(C^T K_n^t, \sigma^2 I, \nu) \qquad (12)$$

where *C* is a rectangular matrix that selects entries from the kinematic (physical state), *v* are the degrees of freedom, I is the identity matrix and σ is related to the accuracy of the sensor system.

4.5.4. Network-level collision risk $P({CRN_n}^t)$

In theory, every technique which can be utilised for real-time collision prediction can be applied to estimate the probability of a road segment having collision-prone traffic conditions in the proposed DBN. As the problem of identifying if the traffic conditions at a specific road segment are collision-prone or note is a binary classification problem, the outcome of every technique would be a binary indication (e.g. 1 for collision-prone conditions and 0 for safe traffic).

Binary classifiers are usually evaluated through the following performance metrics:

$$Accuracy = \frac{\mathrm{T_{conflict}} + \mathrm{T_{safe}}}{\mathrm{T_{conflict}} + \mathrm{T_{safe}} + \mathrm{F_{safe}} + \mathrm{F_{conflict}}};$$

$$Recall = \frac{\mathrm{T_{conflict}}}{\mathrm{T_{conflict}} + \mathrm{F_{safe}}};$$

$$Specificity = \frac{\mathrm{T_{safe}}}{\mathrm{T_{safe}} + \mathrm{F_{conflict}}}$$

where $T_{conflict}$ represents a correct detection of conflict-prone traffic conditions identified as conflict-prone, $F_{conflict}$ represents an incorrect detection of conflict-prone traffic conditions identified as safe, $\mathrm{T_{safe}}$ is a safe traffic condition instance correctly identified as safe, and $\mathrm{F_{safe}}$ is a safe traffic condition instance falsely identified as conflict-prone.

In order to transform the classification result, a probability of a road segment having collision-prone traffic conditions can be estimated as:

$$P({CRN_n}^t = "dangerous") = (\frac{Acc+Rec}{2}), \text{if CR} = 1 \tag{13}$$

where *CR* is the classification result for the aggregated traffic conditions in real-time (i.e. 0 or 1), *Acc* and *Rec* are accuracy and recall of the calibrated classifier. It can be observed that if the classifier indicates a collision-prone situation then the probability of the road segment being "dangerous" is estimated by taking into account the overall accuracy of the classifier and its performance in identifying conflict-prone conditions (i.e. recall). It goes without saying that when *CR=1* the probability of the road segment being safe is:

$$P({CRN_n}^t = "safe") = 1 - P({CRN_n}^t = "dangerous") \tag{14}$$

Accordingly, for CR=0: $P({CRN_n}^t = "safe") = (\frac{Acc+Spec}{2})$ (15)

$$P({CRN_n}^t = "dangerous") = 1 - P({CRN_n}^t = "safe") \tag{16}$$

where *Spec* is the specificity of the classifier (i.e. the classifier's performance in identifying safe traffic conditions).

From equations (13) - (16), the importance of building robust classifiers with less false alarms and solid identification of both normal and collision-prone traffic is observable.

5. Data Description

In order to demonstrate how a network-level hint on collision risk can be employed in real-time risk assessment for autonomous driving, the necessary network and vehicle-level data need to be acquired.

As disaggregated traffic data are more useful for the purposes of this study, traffic microsimulation software - PTV VISSIM (PTV, 2013) is used along with the Surrogate Safety Assessment Model (SSAM) (Pu and Joshi, 2008) which extracts conflicts using the simulated vehicle trajectories from VISSIM. A 4.52-km section of motorway M62 between junction 25 and 26 in England was used as the study area. 15-minute traffic data obtained from the UK Highways Agency Journey Time Database (JTDB) corresponding to every day of the years 2012 and 2013 were used as input to the microsimulation software. For the simulated network the vehicle composition is given in Table 1.

Table 1: Vehicle composition for the studied link segment (M62 motorway, junctions 25-26)

Year	2012		2013	
Vehicle category	**Number of vehicles**	**Ratio**	**Number of vehicles**	**Ratio**
Cars and LGV	57136	0.84100209	62591	0.85727
HGV	10643	0.156657541	10238	0.140224
Buses	159	0.002340369	183	0.002506
Total	67938	1	73012	1

Four simulation runs (i.e. one for identifying conflicts and three for the identification of normal traffic conditions) were utilized. The number of additional runs was chosen in order to cope with the imbalance between conflict and safe conditions which can prove essential for classification purposes (He and Garcia, 2009). The simulations were calibrated using the GEH statistic (Transport For London, 2010) and travel-time measurements. The conflicts were identified in SSAM if the TTC between two vehicles was below 1.3 seconds and Post-Encroachment Time (PET) was below 1 second. That is because TTC below 1.3 seconds is lower than the average human reaction time (Triggs and Harris, 1982) and PET values close to zero show imminent collisions (Pu and Joshi, 2008). For every conflict, the nearest upstream detector on the road segment was identified by comparing the time of the conflict with the time the vehicles passed from every detector. This specific detector was marked as “conflict detector”. Traffic data aggregated at 30-seconds intervals were extracted for every conflict detector, the corresponding upstream and downstream detectors on the same lane and the detector in the adjacent lane. In order to obtain the non-collision cases for every conflict detector, the conflicts for the other three simulation runs were assessed to see if any conflicts occurred in their vicinity. If there was no conflict, the traffic measurements obtained from that detector represented ‘safe’ conditions. Otherwise, the detector was discarded. As four simulations were run, having used one simulation for the extraction of conflict-prone conditions and the three other simulations for the extraction of collision-free conditions, the procedure was repeated an additional three times so that every simulation run was used for the extraction of both ‘conflict-prone’ and ‘safe’ conditions. In total the final simulated dataset consisted of 7,800 conflict events and 23,400 non-conflict cases.

According to the guidelines from the Federal Highway Administration (FHWA) (Dowling et al., 2004), the GEH-statistic (Transport For London, 2010) and the link travel time were used. The GEH statistic correlates the observed traffic volumes with the simulated volumes as shown below:

$$GEH = \sqrt{\frac{(V_{sim} - V_{obs})^2}{\frac{V_{sim} + V_{obs}}{2}}}$$

where V_{sim} is the simulated traffic volume and V_{obs} is the observed traffic volume.

After a number of trial simulations, the best GEH values were obtained by using the following parameters for the Wiedemann 99 car following model:

- Standstill distance: 1.5 m
- Headway time: 0.9 sec
- Following variation: 4 m

For the simulation to efficiently resemble real-world traffic it is essential that (Dowling et al., 2004):

1. GEH statistic < 5 for more than the 85% of the cases
2. The differences between observed and simulated travel times is equal or below 15% for more than 85% of the simulated cases.

The validation results are summarized in Fig. 4 and 5, and the comparison between traffic flow and travel time in simulation and reality are depicted in Fig. 6 and 7. The calibration was performed using the entire simulated dataset (from all four periods) and the observed traffic conditions and conflicts so as to have a unified dataset.

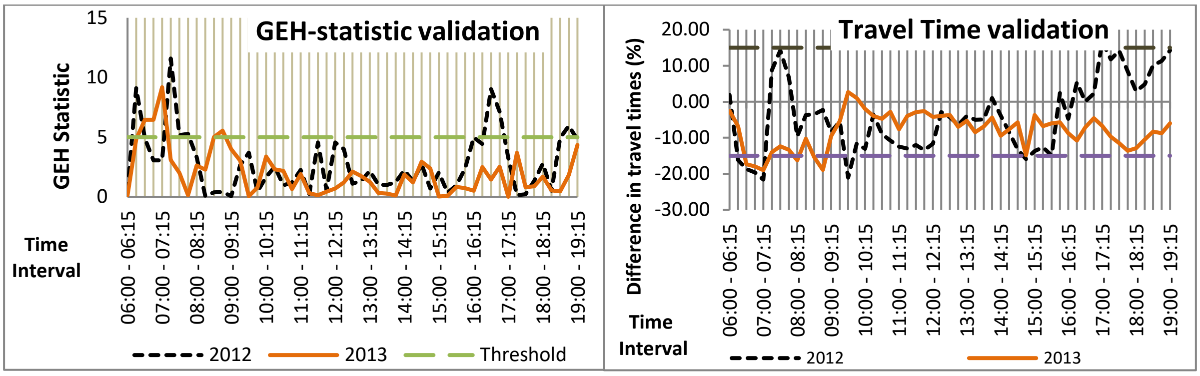

Fig. 4. GEH statistic and Travel time validation for each time interval and year.

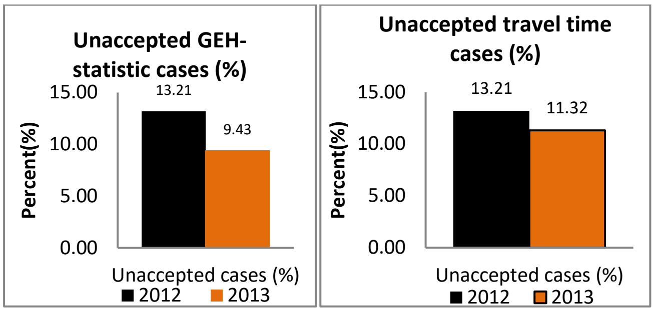

Fig. 5. Percentage of unaccepted cases for each year regarding the GEH statistic and travel time.

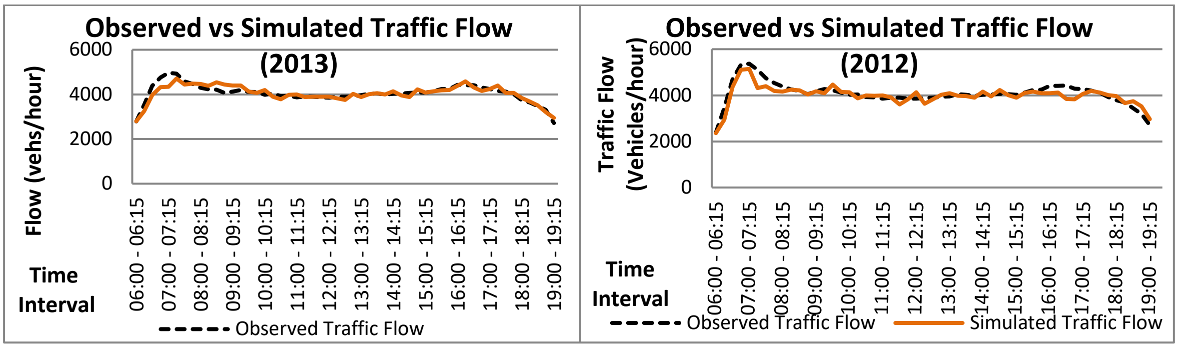

Fig. 6. Observed vs Simulated Traffic flow for each year

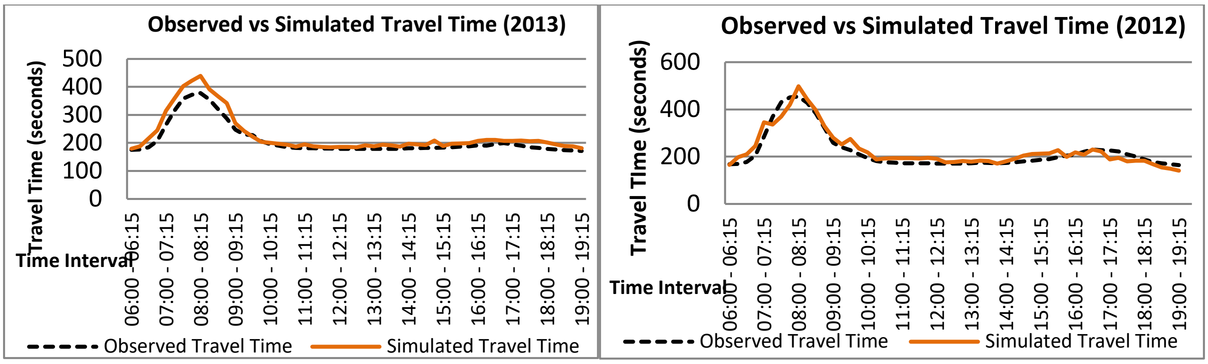

Fig. 7. Observed vs Simulated travel time for each year

In the simulations that were undertaken, the GEH values for most of the time intervals were found to be less than five. However, there were intervals where GEH values were found to be between 5 and 10. These values indicated either a calibration problem or a data problem. Because of the large number of simulations undertaken (~1000 for every scenario) it was assumed that the bad GEH values related to the highly aggregated traffic data (i.e. 15-minute by road-level). Therefore, it was decided to keep the simulation results for the intervals with GEH values between 5 and 10

In order for the conflicts to be validated, the Crash Potential Index (CPI) was used as suggested by Flavio (Cunto, 2008). CPI is calculated through the following equation:

$$CPI_i = \frac{\sum_{t=t_{i_i}}^{tf_i} (P(MADR^{(a_1,a_2,\dots,a_n)} \le DRAC_{i,t}) \cdot \Delta t \cdot b}{T_i} \qquad (17)$$

where CPI_i is the CPI for vehicle i, $DRAC_{i,t}$ is the deceleration rate to avoid the crash (m/s^2), $MADR^{(a_1,a_2,\dots,a_n)}$ is a random variable following normal distribution for a given set of environmental attributes, t_{i_i} and t_{f_i} are the initial and final simulated time intervals for vehicle i, Δt is the simulation time interval (sec), T_i is the total travel time for vehicle i and b is a binary state variable denoting a vehicle interaction. For MADR according to (Cunto, 2008) a normal distribution with average of 8.45 for cars and 5.01 for HGVs with a standard deviation of 1.4 was assumed for daylight and dry pavements. The results for the calibration of the conflicts are shown in Fig.8

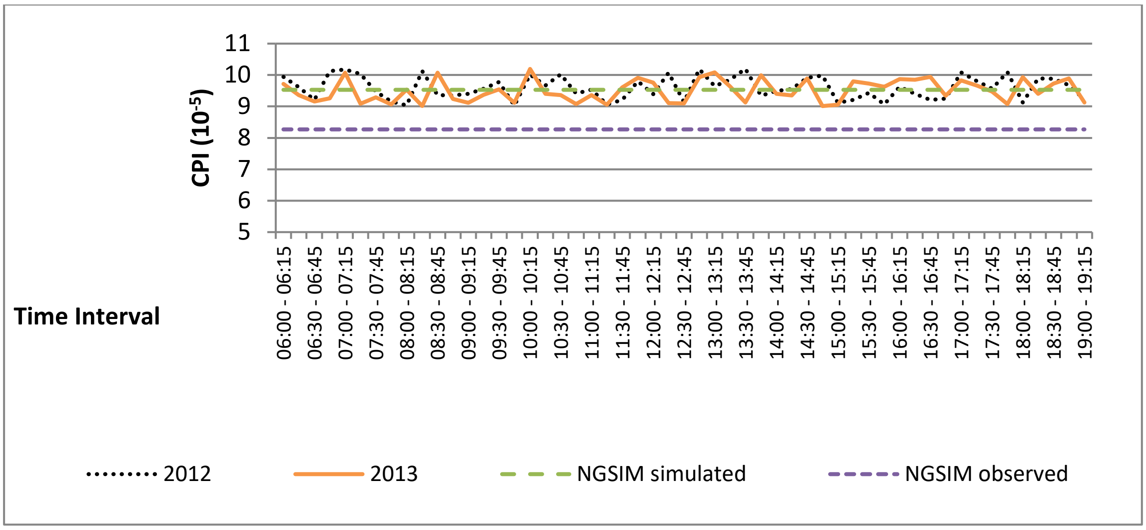

Fig. 8. Conflicts validation

In Fig. 8 it is shown that for the majority of the time intervals, CPI is similar to the simulated CPI of the NGSIM dataset and close to the values of the observed NGSIM CPI. Therefore, it can be concluded that the simulated conflicts resembled realistic hazardous scenarios

It should be noted here, that the sole purpose of the simulation, was to extract highly disaggregated traffic data and corresponding conflicts between vehicles, in order to be used for the proposed DBN model. The simulated dataset does not contain any AVs and therefore the Wiedemann motorway model was used, to replicate car-following behavior. The DBN model was not run within the simulation environment, but the traffic data created from simulation were used to test the proposed AV real-time safety assessment model.

In addition to the simulated traffic data, 5-minute aggregated traffic and the corresponding accident data were provided by the Department of Transportation planning and Engineering of the National Technical University of Athens. The data contained traffic and collision information during a 6-year period (2006-2011). Collision and traffic data concerned two major roads of the metropolitan area of Athens (i.e. Mesogeion and Kifissias avenues). In total the Athens dataset contained 472 collision cases and 917 non-collision cases.

The collision database that was provided included the following variables:

- Collision: 0 for non-collision cases and 1 for collision cases
- Average of speed, occupancy and volume upstream and downstream of the accident location (3 * 2 locations= 6 traffic variables) in 5-minute intervals for 1-hour before the accident time

It should be noted that the 5-minute average correspond to the closest upstream detection from the location of the accident. As disaggregated traffic data are within the scope of this paper, only the 5-minute prior to the accident were extracted and used for the development of the models. For more information on the Athens dataset the reader is prompted to Theofilatos, (2015).

For the estimation of the vehicle-level risk, data were collected using the instrumented vehicle of the School of Civil and Building Engineering of Loughborough University. The vehicle is equipped with the following sensors:

- a Near InfraRed (NIR) Camera
- a short and long-range automotive radar
- a GNSS and 3D Dead Reckoning system
- a lane-departure and forward collision warning camera system

All the sensors are aligned along the centre of the longitudinal axis of the car. The position of the sensors and the experimental vehicle are depicted in Figure 9.

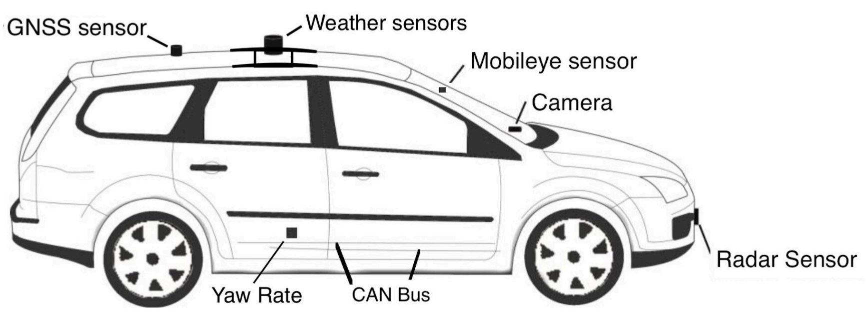

Figure 9: The experimental vehicle along with its sensors

For the purposes of this paper, only data from the GNSS system and the automotive radar have been used. The vehicle data were collected on April 23rd 2017, between 10:53 am and 11:51 am on the M1 motorway (J23-J18) from Loughborough to the Watford Gap service station. Regarding the radar sensor, it identifies targets and objects with a sensor cycle of 15.15 Hz. A target can be anything which reflects radar waves, whereas an object is a target which has been traced by the software used by the radar sensor over a few measurements. Only the object measurements have been used, as they are more representative of the vehicles and obstacles surrounding the ego-vehicle. The speed of the ego-vehicle as measured by the GNSS module during the driving trip and the total number of vehicles sensed by the ego-one during the driving trip are depicted in Figure 10. For each of the vehicles sensed and according to the GNSS ego-vehicle position as well as the radar object readings, a TTC metric was derived in order to identify dangerous traffic participants.

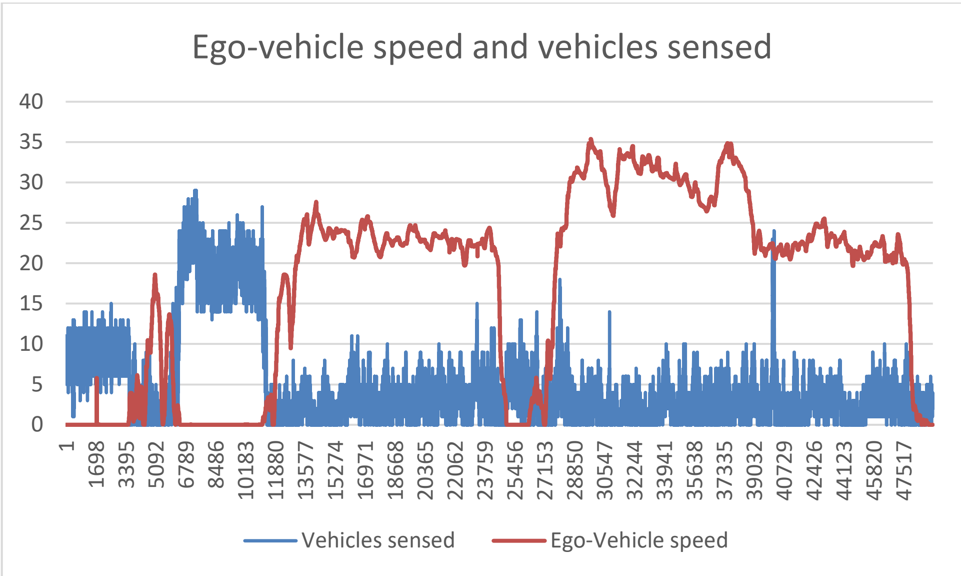

Figure 10: Ego-vehicle speed during the driving trip

6. The impact of network-level collision prediction on vehicle-level risk assessment

The developed DBN network which integrates network-level and vehicle-level collision prediction was given in Figure 2. The part that is of interest for this work is the top part of the graph as shown in Figure 11. More specifically, the estimation will be related on how a good prediction by a network-level classifier enhances or decreases the identification of a dangerous road user given that the measurements about vehicle-level and kinematics in a previous time epoch are known.

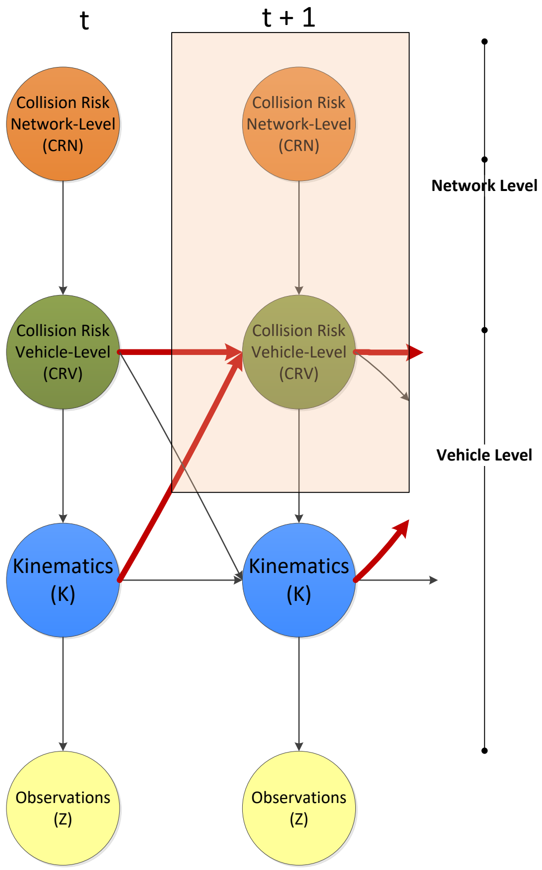

Figure 11: Variables of interest in the developed DBN

In this section, the vehicle-level risk is estimated with and without the network-level risk. For that purpose, the results from two machine learning classifiers are going to be initially utilized for the estimation of vehicle-level risk. These are:

- The k-Nearest Neighbour (kNN) classifier using the imbalanced learning technique of Synthetic Minority Oversampling Technique (SMOTE) along with Edited Nearest Neighbours (ENN) utilized with the 30-second simulated data.
- A Gaussian Processes (GP) classifier using traffic data aggregated at 5-minute intervals from Athens, Greece, which are classified using the imbalanced learning technique of Neighbourhood Clearing (NC).

These classifiers were chosen in order to estimate vehicle-level risk with as little prediction horizon as possible using disaggregated traffic data after a comparison with other classifiers such as support vector machines, neural networks and k-nearest neighbours. Imbalanced learning (He and Garcia, 2009) was chosen to assist with classification results because of the difference in the proportion between collision and non-collision cases which is a known problem of real-time collision prediction datasets(Xu et al., 2016).

6.1. Estimation of vehicle-level risk using simulated data

Assuming that vehicle-level measurements were not available, the following artificial scenarios are formulated for the estimation of the vehicle-level risk:

6.1.1. Traffic data aggregated at 30-second intervals

It is assumed that once traffic conditions are classified, the prediction is broadcasted for a time interval equal to the traffic data aggregation. Therefore, if the traffic data aggregation is 30-seconds, every CRN prediction lasts for 30 seconds. In this scenario, it is assumed that traffic conditions are classified as conflict-prone and, at time t_1=10 seconds after the beginning of the CRN prediction, there is a traffic participant that poses a threat to the ego-vehicle. Furthermore, it is assumed that this “dangerous” vehicle has kinematics that indicate an imminent danger for the ego-vehicle. Hence, according to equations (6) and (7): ${f_{K}}_{N}^{t=10} = 1$ and ${f_{CRV}}_{N}^{t=10}$=1. It should be noted here that 10 indicates the time moment occurring ten seconds after the network-level prediction and hence 20 seconds remain for the end of the temporal aggregation interval.

The kNN classifier under SMOTE-ENN with 30-seconds temporal aggregation resulted in 77.56% accuracy, 77.14% recall and 77.71% specificity.

<u>Scenario 1: Traffic conditions are predicted as conflict-prone</u>

According to equation (13):

$$P({CRN_n}^t = "dangerous") = (\frac{Acc+Rec}{2}) = \frac{0.7756+0.7714}{2} = 0.7735 = 77.35\%$$

Furthermore, as the traffic conditions are estimated as dangerous and ${f_{CRV}}_N^{\,t=10}$=1, the boosting parameter for the vehicle-level safety context ${f_{CRN}}_N$ is equal to $P({CRN_n}^t = "dangerous")$. Consequently, ${f_{CRN}}_N^{\,t=10} = 0.7735$.

Figure 12 illustrates the estimation of vehicle-level risk context when the ego-vehicle is sensing 1, 3, 5 and 10 vehicles in its vicinity, with and without the network-level hint.

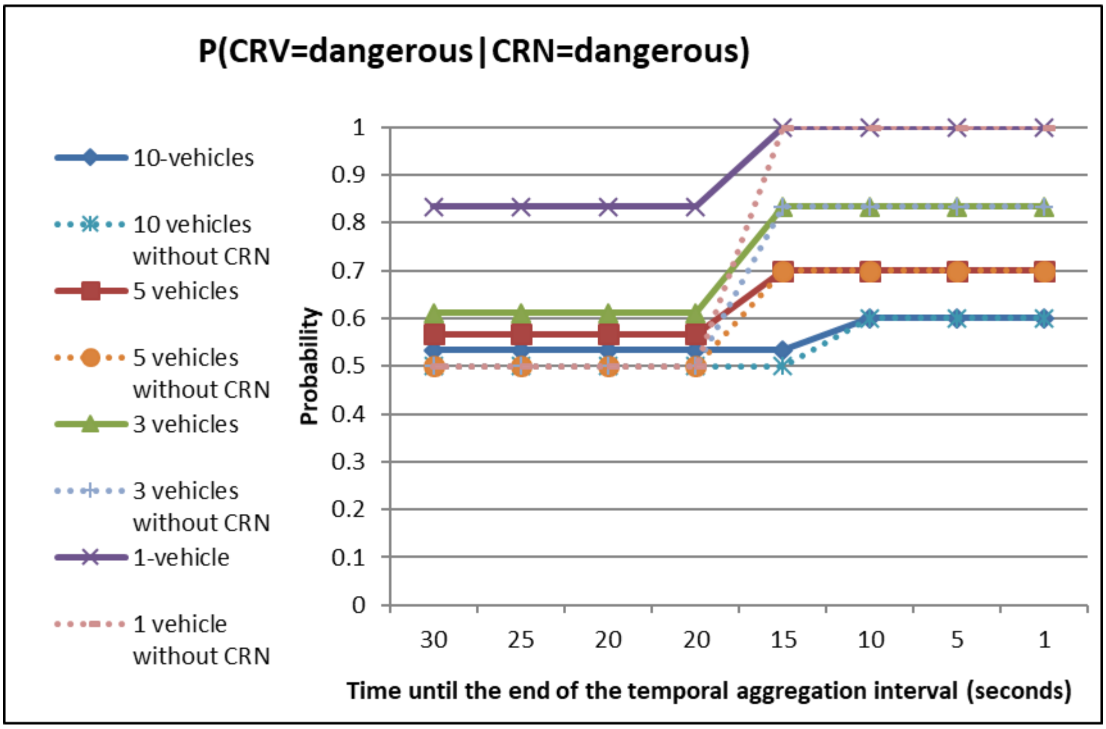

Figure 12: Estimation of P(CRV=dangerous|CRN=dangerous) for a multiple vehicle scenario

From Figure 12, the potential enhancement of the vehicle-level safety context could be observed. First of all, if network-level safety information is available, the probability of a vehicle being considered as a threat is higher, which may be conservative as an approach but induces a hint to the ego-vehicle that a danger is imminent. Moreover, it is shown that this extra hint results in a faster increase of probability when a vehicle is sensed to be performing a dangerous manoeuvre, which could lead to the faster identification of a dangerous road user and an earlier initiation of the manoeuvre to avoid the danger. If, for example, a threshold is defined (e.g. if probability is over 65%) in order to raise a warning to the risk assessment module of the AV, then Figure 12 demonstrates that the threshold is raised faster if network-level information is available.

To further demonstrate how vehicle-level safety is affected, a second artificial scenario was investigated. This relates to the probability of a vehicle driving dangerously, given that the network-level collision risk is predicted as safe.

<u>Scenario 2: Traffic conditions are predicted to be “safe”</u>

According to equation (15):

$$P({CRN_n}^t = "safe") = (\frac{Acc + Spec}{2}) = \frac{0.7756 + 0.7771}{2} = 0.77635$$

Because in this scenario the traffic conditions are estimated as safe and ${f_{CRV_N}}^{t=10}$=1, the boosting parameter for the vehicle-level safety context f_{CRN_N} is equal to $f_{CRN_N} = 1 - recall$ in order to represent the false negative rate i.e. the probability that the traffic conditions are falsely identified as safe.

Hence, ${f_{CRN_N}}^{t=10} = 1 - recall = 1 - 0.7714 = 0.2286$%=22.86%.

Figure 13 illustrates the estimation of the probability of the vehicle-level risk context being dangerous when the ego-vehicle is sensing 1, 3, 5 and 10 vehicles in its vicinity with and without the network-level hint.

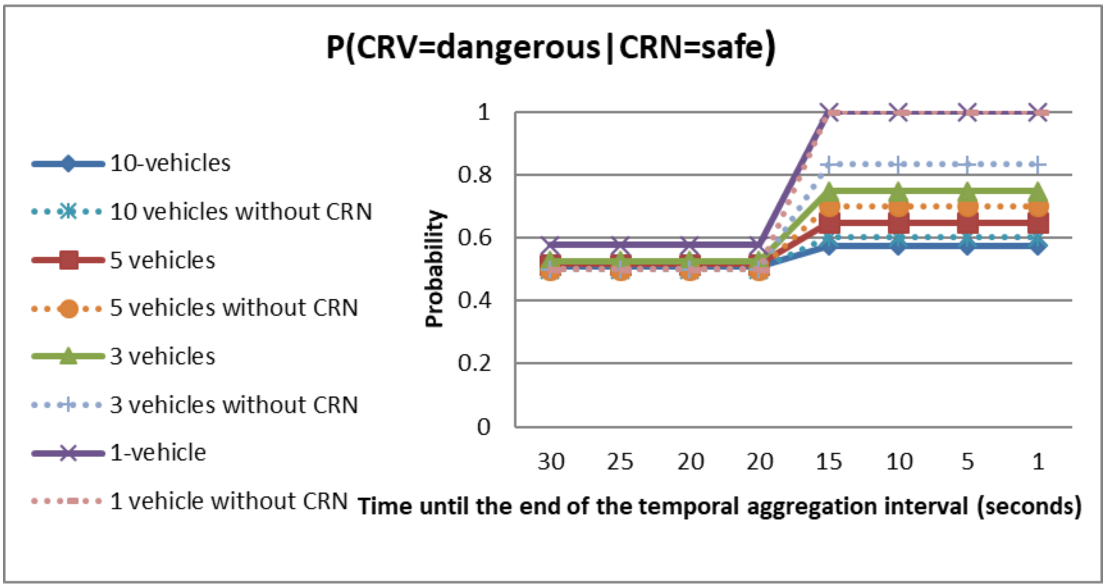

Figure 13: Estimation of P(CRV=dangerous|CRN=safe) for a multiple vehicle scenario

From Figure 13, it is shown that the estimation of the probabilities without the network-level hint results in higher rates and a faster identification of the dangerous road user. Only when just one vehicle is in the vicinity of the ego-one and the dangerous road user is obvious, the two approaches (i.e. with and without network-level information) yield similar results. This indicates that when NLCP indicates safe traffic conditions, more trust should be given to the vehicle measurements rather than the network traffic information.

6.1.2. Traffic data aggregated at 5-minute intervals

In order to further test the impact of network-level collision information on vehicle-level collision risk, the classifier developed on the 5-minute aggregated data from Athens was utilized. The classifier achieved 83.95% accuracy, 91.71% specificity and 68.86% recall. For this scenario, the number of vehicles was randomly sampled for each time moment. It was also assumed that a vehicle performs dangerous manoeuvres starting from t=180 before the end of the temporal aggregation to t=100 seconds before the end of the temporal aggregation interval. Hence, ${f_{K_N}}^{t=180:100} = 1$ and ${f_{CRV_N}}^{t=180:100}$=1.

Scenario 1: Traffic conditions are predicted as *collision-prone*

According to equation 13:

$$P({CRN_n}^t = "dangerous") = (\frac{Acc+Rec}{2}) = \frac{0.8395+0.6886}{2} = 0.7641 = 76.41\%$$

Furthermore, for the time intervals *t=300:180* and *t=100:0*, the traffic conditions are estimated as dangerous but there is no vehicle performing dangerous manoeuvres. Therefore, the boosting parameter for the vehicle-level safety context during these intervals is:

$$f_{CRN_N}{}^{t=300:180\ \&\ t=100:0} = 1 - \frac{Accuracy+Specificity}{2} = 0.1217$$

For the time interval *t=180:100,* traffic conditions are estimated as collision-prone and there is only one vehicle performing a hazardous manoeuvre. Therefore, the boosting parameter for the vehicle-level safety context during these intervals is:

$$f_{CRN_N}{}^{t=180:100} = \frac{Accuracy + Recall}{2} = 76.41\%$$

Figure 14 illustrates the estimation of the probability of a vehicle being dangerous during the 5-minute traffic data temporal aggregation interval in a multiple vehicle scenario.

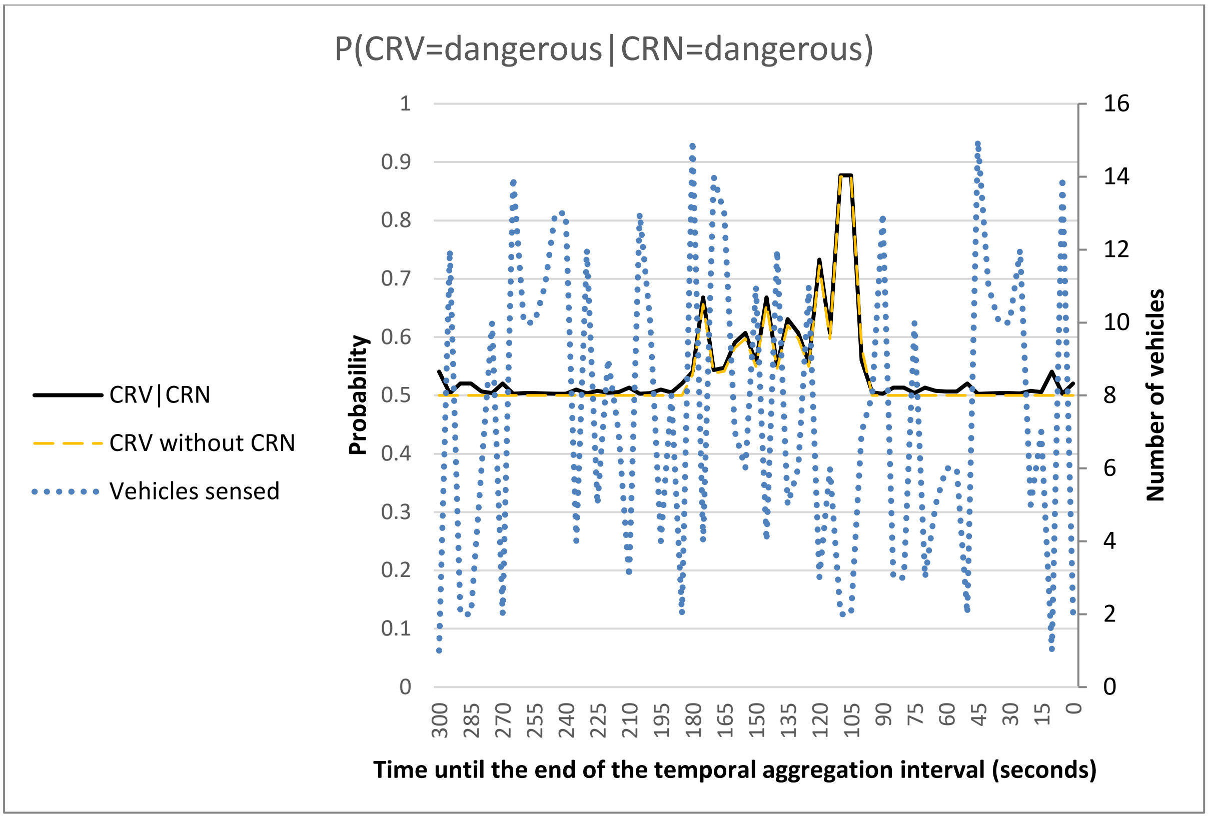

Figure 14: Estimation of P(CRV=dangerous|CRN=dangerous) for a 5-minute traffic data aggregation interval

From Figure 14, it is further justified that the use of CRN estimation enhances the probability of identifying whether another vehicle driving dangerously with respect to the ego-vehicle. From *t=180* seconds until *t=100*, when a nearby vehicle is assumed to perform dangerous manoeuvres, the probability of the vehicle being dangerous given the network-level hint is relatively higher than the corresponding probability without the network-level information. Moreover, it is demonstrated that the lower the number of vehicles, the more obvious it is to recognize the vehicle which is driving "dangerously". This is normal because with fewer vehicles, the one responsible for triggering a collision is easier to detect. Nevertheless, it is advantageous that the line representing the probability P(CRV|CRN) is above the corresponding probability graph which does not take into account network-level collision information. It is also observed that at a time moment when a danger is not imminent the probability is increased, which is a potential drawback. However, this can be utilized as an extra caution by an AV's planning module.

<u>Scenario 2: Traffic conditions are predicted as *safe*</u>

Given that the traffic conditions are predicted to be *safe*, the network-level collision risk can be estimated by using equation 15:

$$P({CRN_n}^t = \text{"}dangerous\text{"}) = 1 - (\frac{Acc+Spec}{2}) = 1 - \frac{0.8395+0.9171}{2} = 0.1217 = 12.17\%$$

Furthermore, for the time intervals t=300:180 and t=100:0, the traffic conditions are estimated as safe without a vehicle perceived as a threat. Therefore, during these intervals:

$${f_{CRN_N}}^{t=300:180\ \&\ t=100:0} = P({CRN_n}^t = \text{"}dangerous\text{"}) = 0.1217$$

For the time interval t=180:100 traffic conditions are estimated as safe but there is one vehicle performing hazardous manoeuvres. Therefore, the boosting parameter for the vehicle-level safety context during these intervals is:

$$f_{CRN_N}{}^{t=180:100} = 1 - Recall = 1 - 0.6886 = 0.3114$$

Figure 15 illustrates the estimation of the probability of the vehicle-level risk context being dangerous during the traffic data temporal aggregation interval and according to the vehicles sensed.

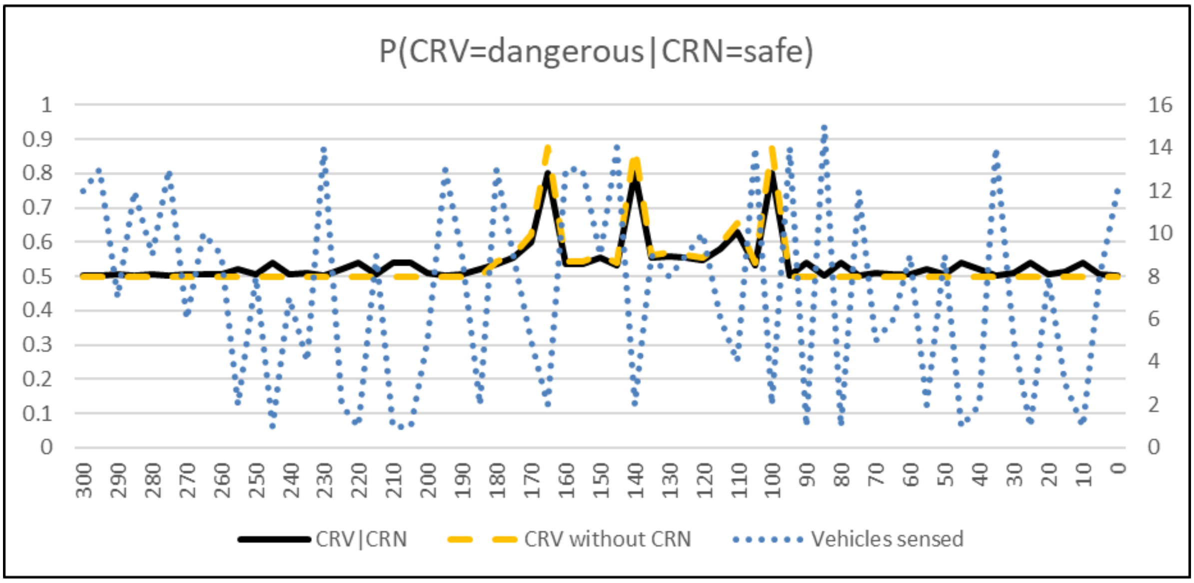

Figure 15: Estimation of P(CRV=dangerous|CRN=safe) for a 5-minute traffic data aggregation interval

Like the case when traffic data were aggregated in 30-seconds intervals and the traffic conditions were assumed to be safe, Figure 15 illustrates that, when a danger is sensed by the ego-AV, network-level information does not contribute to the enhancement of the corresponding probability.

6.2. Estimation of vehicle-level risk using real-world data

It is common knowledge that traffic data are mostly available for motorways where magnetic loop detectors and automatic vehicle identification devices exist. Therefore, the developed method is demonstrated for the case of motorway driving. Risk assessment of AVs at junctions is not considered as an example because it has been the focus of previous research (Agamennoni et al., 2012; Lefèvre, 2012).

In order to validate the credibility that network-level information has on the estimation of vehicle-level collision prediction, the vehicle-level data as described in Section 5 were utilized.

More specifically, the available TTC measurements were filtered in order to identify hazardous road users. According to the same principle as the one used in SSAM to derive conflicts, TTC values below 1.5 seconds were flagged as "hazardous" because 1.5 is the average human reaction time (Triggs and Harris, 1982). The number of hazardous vehicles during the trip is given in Figure 16.

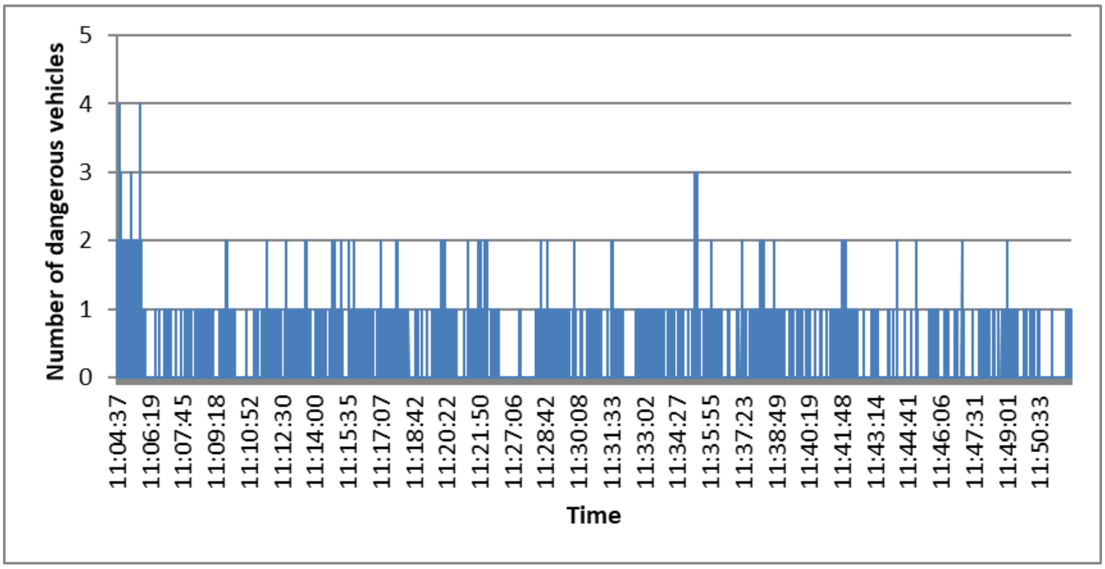

Figure 16: Number of dangerous vehicles with respect to the ego-vehicle

The time interval from 11:05:37 to 11:06:25 was used in the analysis as the highest number of "hazardous" road users was observed during that one minute. The analysis took place only during this interval so as to imitate "dangerous" driving behaviour from other traffic participants.

The classifiers that were tested for the estimation of CRV based on the network-level information and their characteristics are described in Table 2. More specifically, a kNN classifier along the imbalanced technique of SMOTE-ENN was utilized for classifying traffic data aggregated at 30-seconds intervals, a Support Vector Machine (SVM) classifier along with the imbalanced technique of Repeated Edited Nearest Neighbours (RENN) was utilized for classifying 1-minute and 3-minute traffic and conflict data and a Neural Network (NN) classifier along with SMOTE-ENN was utilized for classifying 5-minute traffic and conflict data. These are the classifiers that yielded the best classification result for every temporal aggregation interval, after a comparison of different classification and imbalanced learning techniques. For each of the classifiers the probability that a vehicle drives dangerously was estimated given that the CRN points towards collision-prone and safe traffic. For the estimation of vehicle-level risk context the formulas (13) -(16) were used. For every vehicle with TTC<1.5 seconds, it was assumed that the vehicle's kinematics were also dangerous so as to have f_{K_N}=1.

Table 2: CRN classifiers used for vehicle-level risk estimation

Traffic data aggregation	Classifier	Accuracy	Recall	Specificity	False Alarm Rate
30-seconds	kNN with SMOTE-ENN	0.7756	0.7714	0.9171	0.2229
1-minute	SVM with RENN	0.9219	0.6886	0.9996	0.0004
3-minute	SVM with RENN	0.9222	0.6891	0.9999	0.00001
5-minute	NN with SMOTE-ENN	0.8006	0.8285	0.7913	0.2087

6.2.1. Estimation of vehicle-level risk given traffic conditions are collision-prone

Figures 17-20 illustrate the results for the probability that a vehicle poses a threat to the ego-vehicle, given the available network-level information and the vehicle-level data.

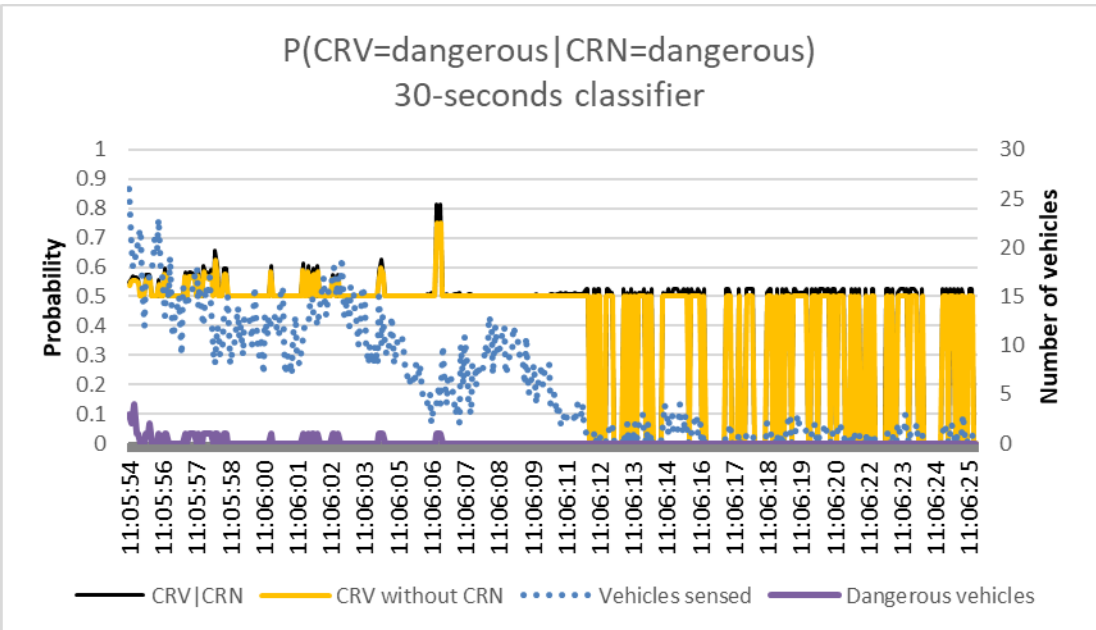

Figure 17: Estimation of vehicle-level risk using 30-seconds network-level information for conflict-prone traffic conditions

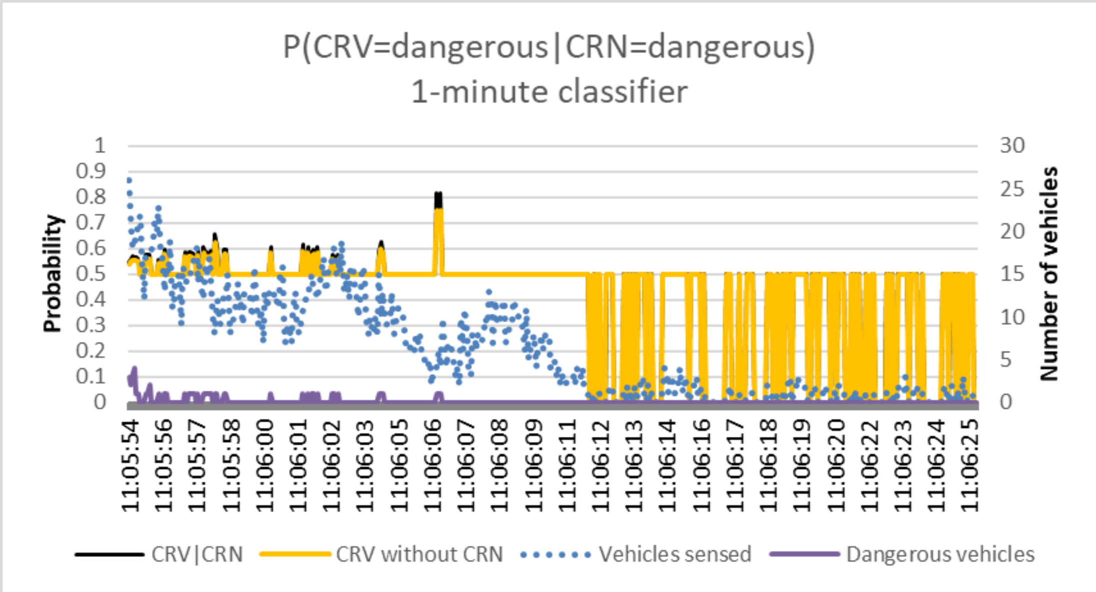

Figure 18: Estimation of vehicle-level risk using 1-minute network-level information for conflict-prone traffic conditions

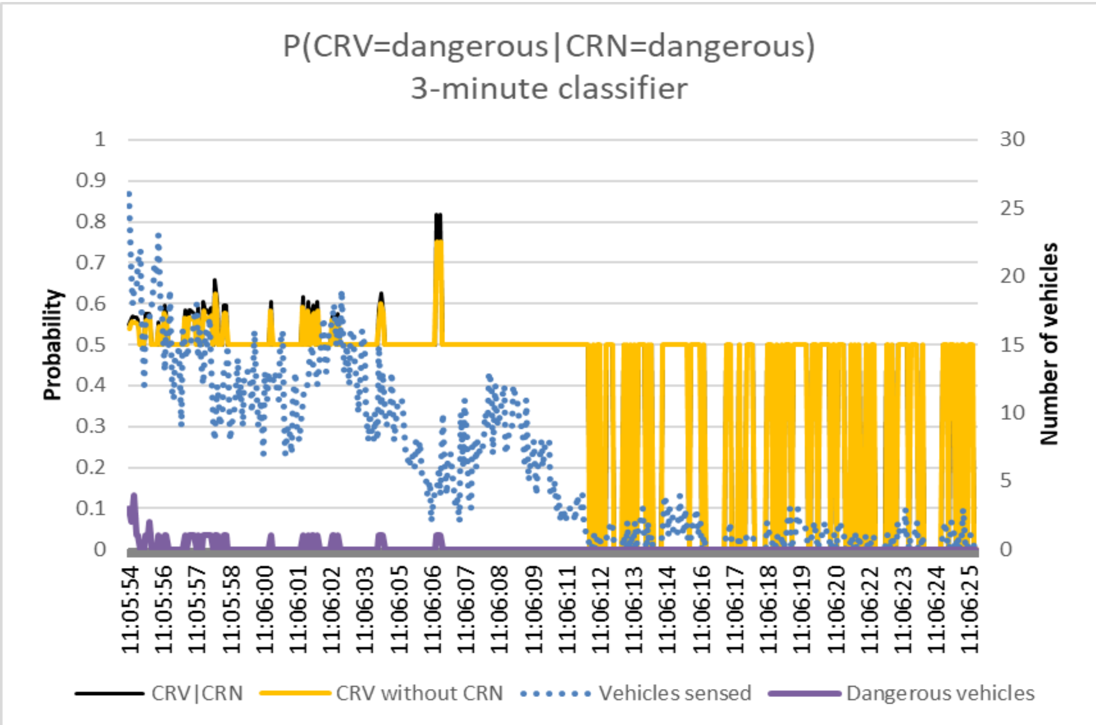

Figure 19: Estimation of vehicle-level risk using 3-minute network-level information for conflict-prone traffic conditions

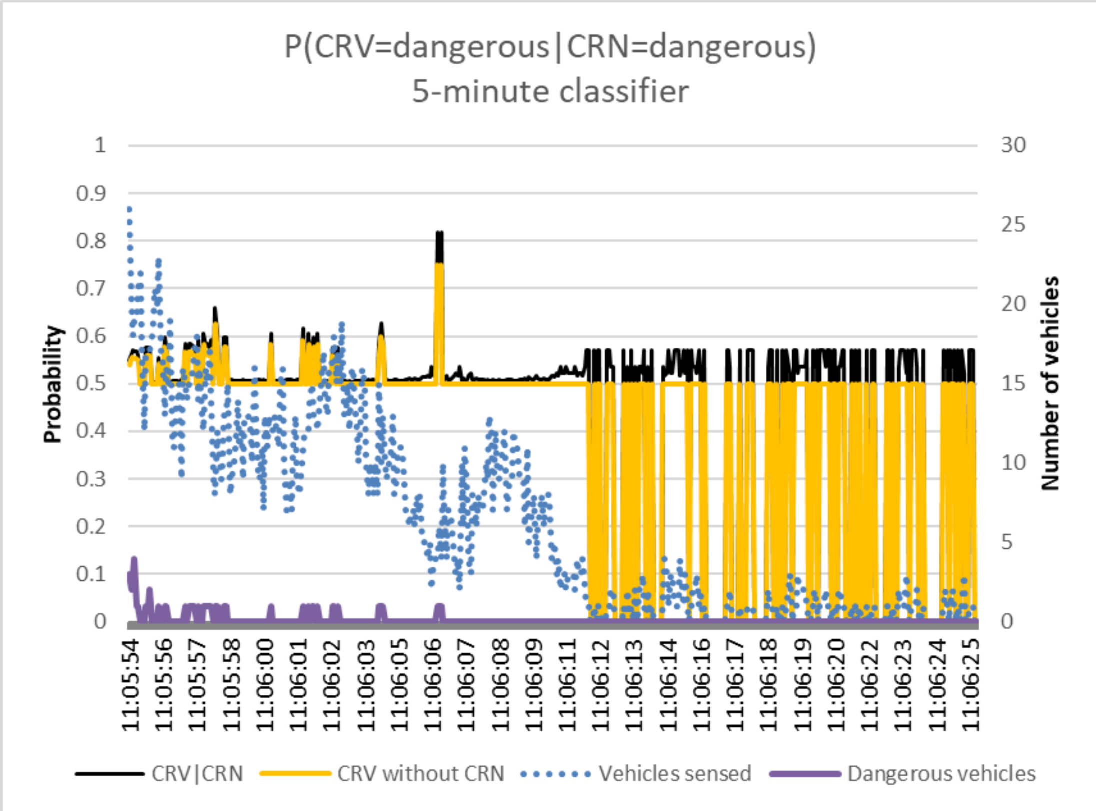

Figure 20: Estimation of vehicle-level risk using 5-minute network-level information for conflict-prone traffic conditions

After observing Figures 17-20, it is further validated that, when traffic conditions are predicted as conflict-prone, it is easier to identify if there is an imminent danger for the ego-vehicle. Even when highly disaggregated traffic data are utilized, the probability of a dangerous vehicle being dangerous is enhanced when compared to the probability obtained only from vehicle-level measurements. When the number of vehicles sensed is high, the enhancement in the probability is lower. However, the plot of CRV|CRN is always higher than the one of CRV without network-level information, assuring a greater level of safety for the ego-vehicle.

To illustrate the effect of network-level information on vehicle-level risk estimation, Figure 21 presents a plot of the percentage difference between the estimation of the probability that a vehicle drives in a “hazardous” way with regards to the ego-vehicle with and without CRN.

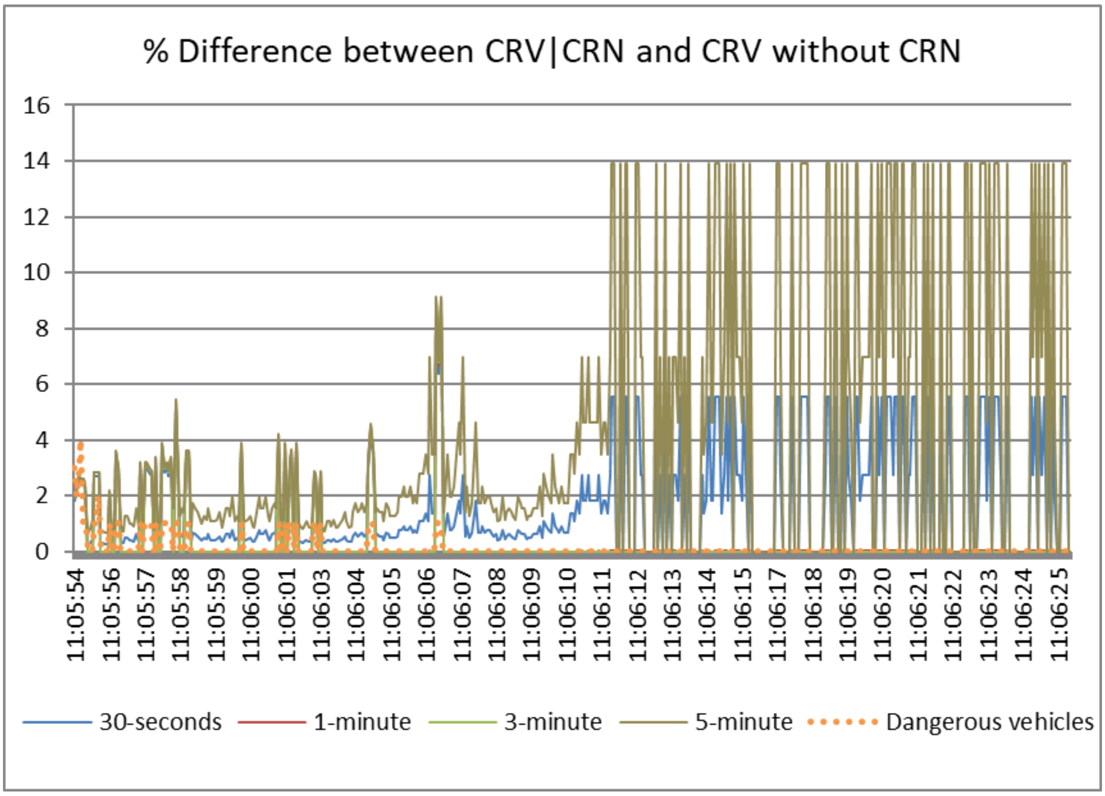

Figure 21: Difference (%) between vehicle-level risk estimation with and without network-level information for conflict-prone traffic conditions

From Figure 21 it can be concluded that the greater influence came from the 5-minute classifier. This is probably due to the ability of the classifier to better detect conflict-prone and safe traffic efficiently as observed from its recall and sensitivity statistics. When there is at least one dangerous vehicle, the estimation of a dangerous vehicle-level safety context is enhanced by up to 9%, ensuring safer navigation. When no dangerous vehicles are detected, the difference can reach up to 14%. This shows that, when traffic conditions are predicted as dangerous, the ego-vehicle can adjust to a more cautious behaviour as a conflict or collision might occur.

Overall, when traffic conditions are predicted as hazardous, the ego-vehicle can better estimate if a vehicle is driving dangerously, even when highly disaggregated traffic data information is available. Furthermore, the fact that, a small probability of a dangerous vehicle is assigned even when no dangerous vehicles are around, can be exploited in an AV risk assessment module.

6.2.2. Estimation of vehicle-level risk given traffic conditions are safe

Figures 22-25 illustrate the results for the probability that a road user is driving dangerously towards the ego-vehicle, given the available network-level information and the vehicle-level data if the traffic conditions are indicated as safe.

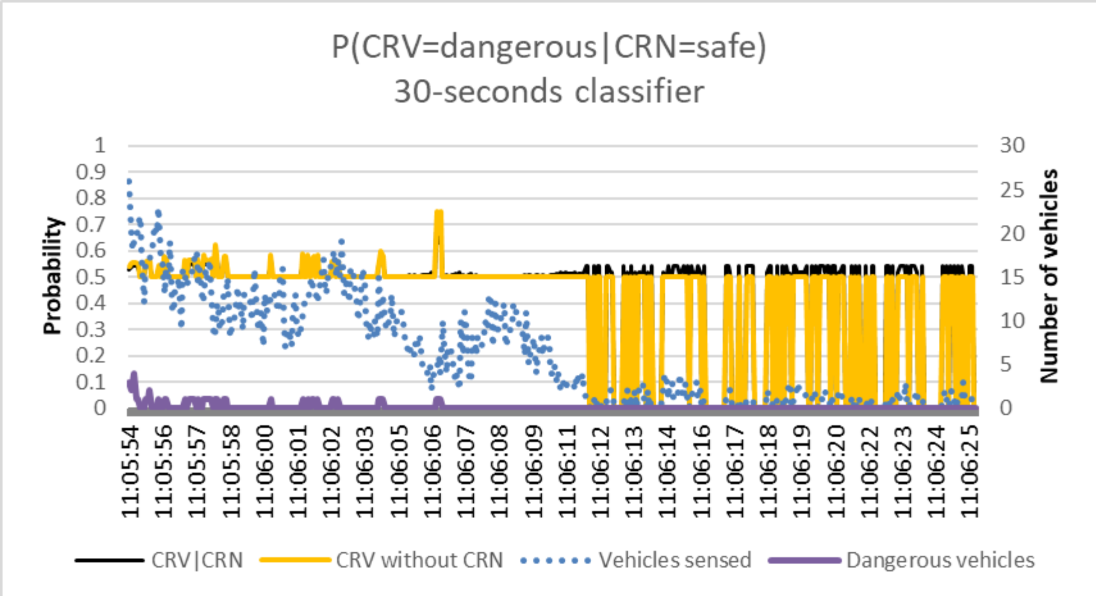

Figure 22: Estimation of vehicle-level risk using 30-seconds network-level information for safe conditions

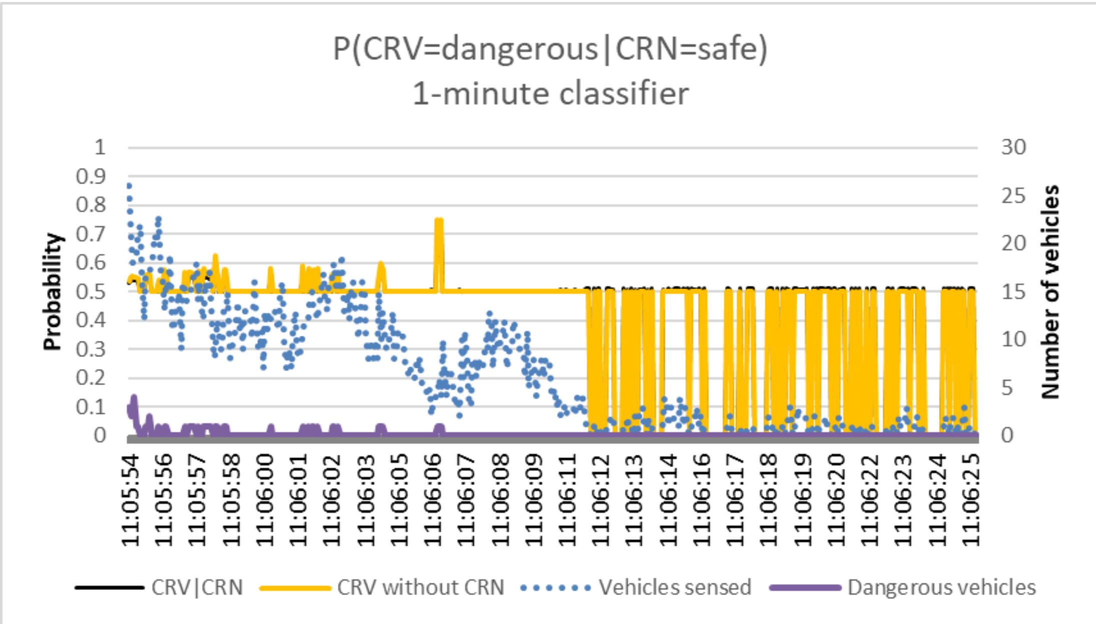

Figure 23: Estimation of vehicle-level risk using 1-minute network-level information for safe conditions

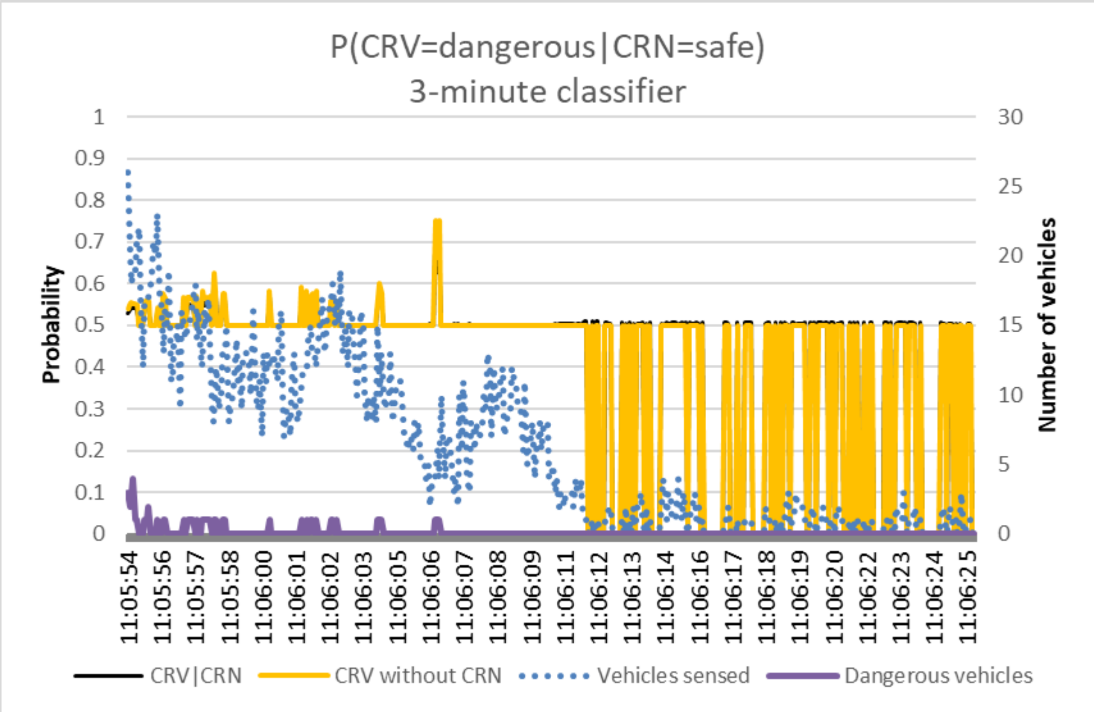

Figure 24: Estimation of vehicle-level risk using 3-minute network-level information for safe conditions

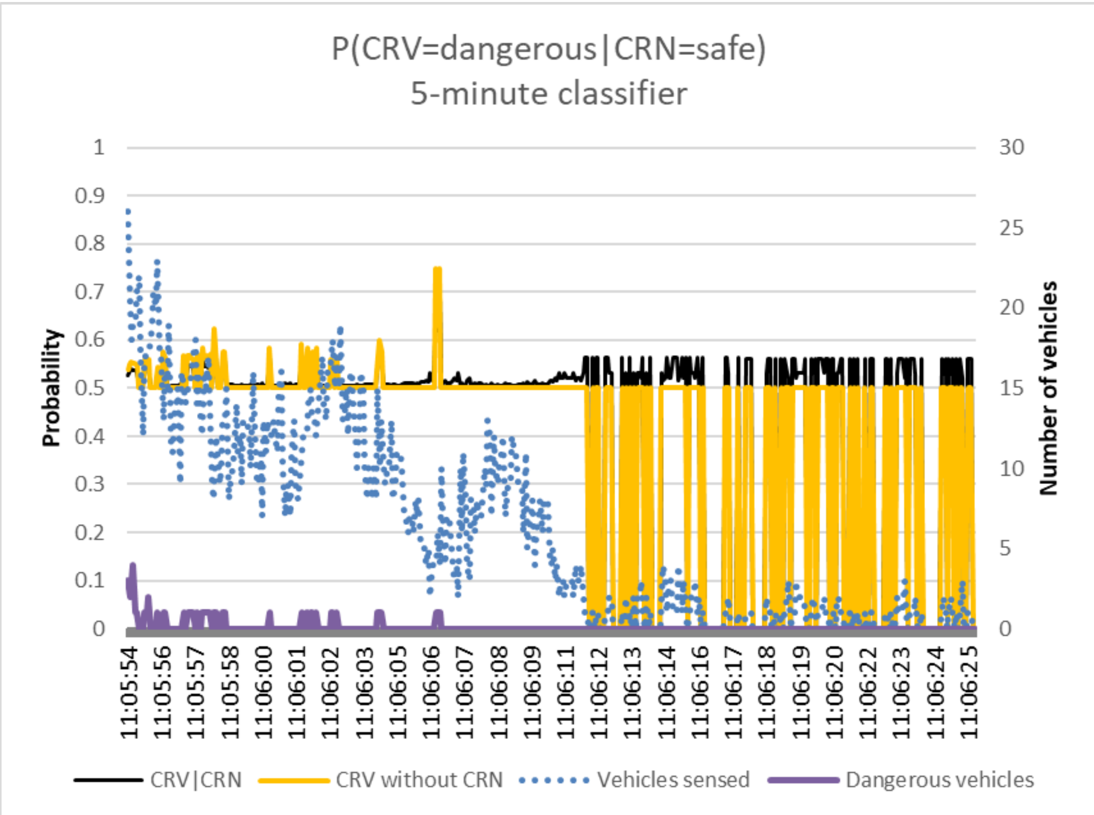

Figure 25: Estimation of vehicle-level risk using 5-minute network-level information for safe conditions

Similar to the case of simulated data, Figures 22-25 demonstrate that, if real-time network-level information points towards safe traffic conditions, then the measurements from the sensors of the ego-vehicle are more reliable to detect dangerous traffic participants. The differences between the two different ways to estimate the vehicle-level safety context probabilities are more obvious when better CRN classifiers are used, such as the 5-minute classifier demonstrated in this paper. Even when no dangerous vehicles are detected and traffic conditions are predicted as safe, the probability that a vehicle could be dangerous is elevated due to the possibility that the network-level information is falsely classified.

As with the conflict-prone conditions, Figure 26 demonstrated the percent difference between the two different approaches to estimate the probability that a vehicle is driving dangerously towards the ego-one.

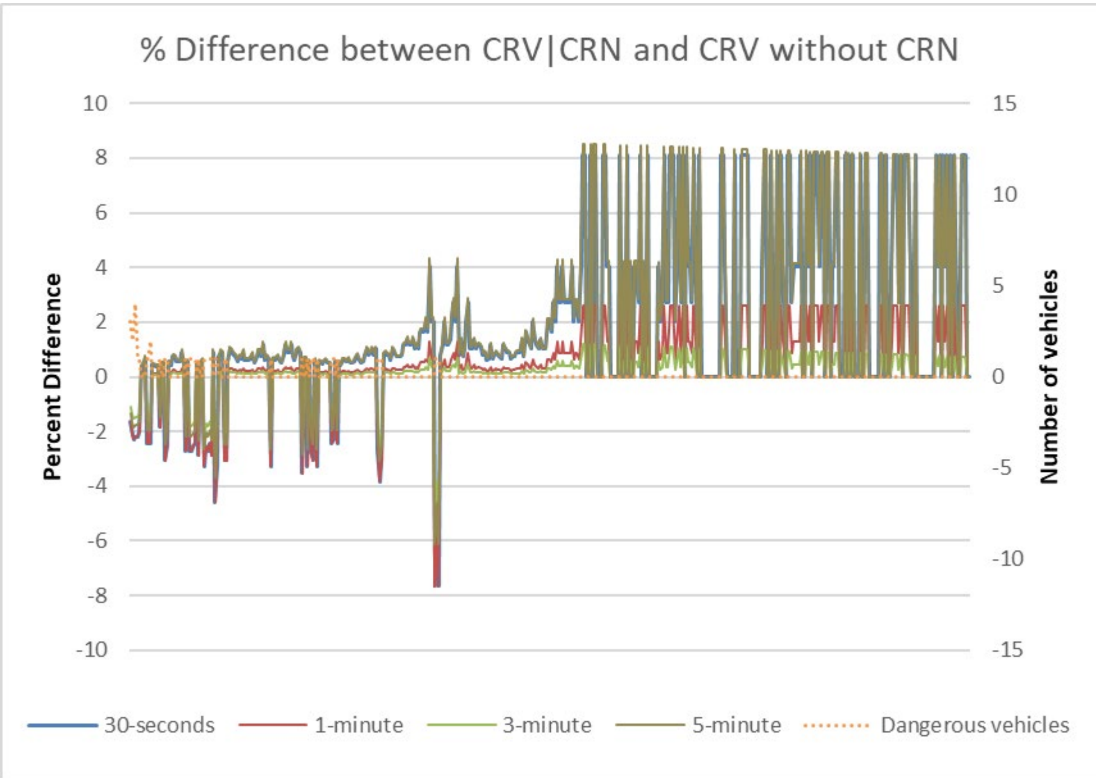

Figure 26: Difference between vehicle-level risk probability with and without network-level information for safe conditions

From Figure 26 it is noticeable that network-level information does not enhance AV risk assessment when traffic conditions are predicted as conflict-prone. As mentioned before, network-level information induces a slight probability that the network-level prediction is wrong when no vehicle is detected as dangerous. On the other hand, in cases when there is an imminent danger, utilizing vehicle-level information only, results in a better hazard recognition than the proposed methodology, reaching up to 8% more confidence in estimating a dangerous traffic participant.

It should be noted that the extracted probabilities for all the scenarios are not high enough. The scenarios developed in this paper were built on some assumptions and without highly detailed vehicle-level data. For the scenarios where traffic conditions were indicated as collision- or conflict-prone, the probability of another vehicle being dangerous was higher when CRN was available, however, further work is needed to calibrate the proposed DBN model in the cases when CRN indicates safe traffic. Nevertheless, the enhanced probability for the dangerous road user when collision-prone traffic was predicted shows that the method has potential for utilization in AV risk assessment.

7. Implementation challenges and recommendations

AVs require a plethora of data from multiple sensor platforms to generate a collision-free trajectory (Huang et al., 2013; Polychronopoulos et al., 2007). Most of AVs utilize cameras (Bertozzi et al., 2000) and laser scanners (Jiménez et al., 2012; Mertz et al., 2013) to scan the surroundings and estimate a safe path for the vehicle. However, it is still unknown how AVs would identify the optimal course of action in the face of a system failure (Dixit et al., 2016; Koopman and Wagner, 2016). In that perspective, the integrated modelling framework developed in this paper could address this challenge. As network-level collision prediction utilizes more macroscopic data compared to the data received by the sensor systems of AVs which have high frequency, the network-level prediction would act as *a-priori* for specific time periods. Consequently, if the majority of the sensing systems fail, then according to the network-level information, the AV can resolute the problem by slowing down as in the case of collision-prone traffic conditions, until it reaches a safe point or the system error is fixed. This also applies to cases where the sensor system, especially the vision-based systems, become obstructed (e.g. due to the presence of a big truck in front of the ego-vehicle or due to adverse weather conditions). Consequently, network-level collision information could assist not only the identification of "*dangerous*" road users but could act as a safety net for all the motion planning levels, i.e. from routing to manoeuvre planning. Finally, if traffic conditions are classified as *collision-prone*, then warning messages could be presented through VMS or broadcasted to the AVs communication system by traffic management agencies, prompting the passenger to take control until safety is ensured. Obviously, the proposed model is not limited to AVs only but could also be applied for Connected and Autonomous Vehicles (CAVs).

8. Conclusion

This paper developed a new methodology for the integration of two interacting domains (i.e. network-level and vehicle-level collision prediction) to enhance the risk assessment of AVs. An interaction-aware model based on Dynamic Bayesian Networks was developed to take into account not only the dependencies between the vehicles in a traffic scene but also a hint from network-level collision risk (CRN) so as to increase comprehensive reasoning about unsafe behaviour during automated driving on a road segment. Results from machine learning classifiers (i.e. *k*NN, Neural Networks, Support Vector Machines, Gaussian processes) were presented with regards to network-level collision prediction and were used as an example to show the influence of this prediction on vehicle-level risk estimation. The potential impact that network-level classifiers would have on the identification of the presence of "*dangerous*" road users was estimated using both artificial and real-world data collected from an instrumented vehicle. Both the artificial dataset and the real-world dataset revealed that the probability of identifying whether another vehicle poses a threat to an AV was increased by up to 9% if CRN indicated *conflict-prone* traffic. On the other hand, when traffic conditions were indicated as *safe*, the prediction did not enhance the probability that a road user was a "threat" for the ego-vehicle. This enhancement is greater when 5-minute traffic data are utilized for predicting network-level collisions. Nevertheless, even when highly disaggregated traffic data (i.e. 30-seconds) were used, the probability of a traffic participant posing a threat to the ego-vehicle was enhanced by approximately 6%. Since network-level predictions utilize data at a higher temporal interval than the sampling frequency of the sensors of an AV in order to provide a broader perception horizon, the developed method would allow AVs to reduce speeds, change their trajectory or prompt a passenger to take the control in order to ensure a safe journey, even when other sensor systems fail. The algorithms and techniques developed in this paper will set the “rules of the game” in advance and will significantly contribute to the ambition that self-driving vehicles should never cause any traffic collisions.

9. Acknowledgement

This research was funded by a grant from the UK Engineering and Physical Sciences Research Council (EPSRC) (Grant reference: EP/J011525/1).

References

Abdel-Aty, M., Pande, A., 2005. Identifying crash propensity using specific traffic speed conditions. J. Safety Res. 36, 97–108.

Agamennoni, G., Nieto, J.I., Nebot, E.M., Member, S., 2012. Estimation of Multivehicle Dynamics by Considering Contextual Information. IEEE Trans. Robot. 28, 855–870.

Bahram, M., Hubmann, C., Lawitzky, A., Aeberhard, M., Wollherr, D., 2016. A Combined Model- and Learning-Based Framework for Interaction-Aware Maneuver Prediction. IEEE Trans. Intell. Transp. Syst. 17, 1538–1550.

Bertozzi, M., Broggi, A., Fascioli, A., 2000. Vision-based intelligent vehicles : State of the art and perspectives ଝ. Rob. Auton. Syst. 32, 1–16.

Bessiere, P., Mazer, E., Ahuactzin, J.-M., Mekhnacha, K., 2013. Bayesian Programming, CRC Press.

Brand, M., Oliver, N., Pentland, A., 1997. Coupled hiddenMarkov models for complex action recognition. Int. Conf. Comput. Vis. Pattern Recognit. 1–6.

Campbell, M., Egerstedt, M., How, J.P., Murray, R.M., 2010. Autonomous driving in urban environments: approaches, lessons and challenges. Philos. Trans. A. Math. Phys. Eng. Sci. 368, 4649–72.

Dixit, V. V, Chand, S., Nair, D.J., 2016. Autonomous Vehicles : Disengagements , Accidents and Reaction Times 1–14.

Gindele, T., Brechtel, S., Dillmann, R., 2015. Learning Driver Behavior Models from Traffic Observations for Decision Making and Planning. IEEE Intell. Transporation Syst. Mag. 7, 69–79.

He, H., Garcia, E.A., 2009. Learning from imbalanced data. IEEE Trans. Knowl. Data Eng. 21, 1263–1284.

Huang, E., Antoniou, C., Wen, Y., Ben-Akiva, M., Lopez, J., Bento, L.C., 2013. Real-Time Multi-Sensor Multi-Source Network Data Fusion Using Dynamic Traffic Assignment Models. In: Proceedings of the 12th International IEEE Conference on Intelligent Transportation Systems, St. Louis, MO, USA, October 3-7, 2009.

Jiménez, F., Naranjo, J.E., Gómez, O., 2012. Autonomous manoeuvring systems for collision avoidance on single carriageway roads. Sensors (Basel). 12, 16498–521.

Katrakazas, C., Quddus, M., Chen, W.-H., 2017. A simulation study of predicting real-time conflict-prone traffic conditions. IEEE Trans. Intell. Transp. Syst..

Katrakazas, C., Quddus, M., Chen, W.-H., Deka, L., 2015. Real-time motion planning methods for autonomous on-road driving : State-of-the-art and future research directions. Transp. Res. Part C Emerg. Technol.

Koller, D., Friedman, N., 2009. Probabilistic Graphical Models: Principles and Techniques, Journal of Chemical Information and Modeling. The MIT Press.

Koopman, P., Wagner, M., 2016. Challenges in Autonomous Vehicle Testing and Validation Driver Out of the Loop in: 2016 SAE World Congress.

Kuhnt, F., Kohlhaas, R., Schamm, T., Marius, J.Z., 2015. Towards a Unified Traffic Situation

Estimation Model – Street-dependent Behaviour and Motion Models –. Int. Conf. Inf. Fusion 1223–1229.

Lefèvre, S., 2012. Risk Estimation at Road Intersections for Connected Vehicle Safety Applications. PhD Thesis, University of Grenoble & INRIA, France.

Lefèvre, S., Vasquez, D., Laugier, C., 2014. A survey on motion prediction and risk assessment for intelligent vehicles. ROBOMECH J. 1, 1–14.

Mertz, C., Navarro-serment, L.E., Maclachlan, R., Rybski, P., Steinfeld, A., Supp, A., Urmson, C., Vandapel, N., Hebert, M., Thorpe, C., Duggins, D., Gowdy, J., 2013. Moving Object Detection with Laser Scanners 30, 17–43.

Merwe, R. van der, Doucet, A., Freitas, N. de, Wan, E., 2000. The Unscented Particle Filter, Technical Report CUED/F-INFENG/TR 380, Cambridge University.

Murphy, K., 2002. Dynamic Bayesian Networks: Representation, Inference and Learning. PhD Thesis, University of California, Berkeley.

Murphy, K.P., 2012. Machine Learning: A Probabilistic Perspective, MIT Press.

Paden, B., Cap, M., Yong, S.Z., Yershov, D., Frazzoli, E., 2016. A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles 1–27.

Pande, A., Das, A., Abdel-aty, M.A., Hassan, H., 2011. Real-Time Crash Risk Estimation: Are All Freeways Created Equal? Transp. Res. Rec. 2237, 60-66, 2011.

Polychronopoulos, a., Tsogas, M., Amditis, a. J., Andreone, L., 2007. Sensor Fusion for Predicting Vehicles' Path for Collision Avoidance Systems. IEEE Trans. Intell. Transp. Syst. 8, 549–562.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1993. Numerical recipes in Fortran (The art of scientific computing), Mathematics and Computers in Simulation.

PTV Planung Trasport Verkehr AG, 2013. PTV VISSIM 6 User Manual.

Pu, L., Joshi, R., 2008. Surrogate Safety Assessment Model (SSAM): Software User Manual 96.

Singh, S., 2015. Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey. Natl. Highw. Traffic Saf. Adm. 1–2.

Snider, J.M., 2009. Automatic Steering Methods for Autonomous Automobile Path Tracking, Robotics Institute, Carnegie Mellon University.

Sun, J., Sun, J., 2015. A dynamic Bayesian network model for real-time crash prediction using traffic speed conditions data. Transp. Res. Part C Emerg. Technol. 54, 176–186.

Theofilatos, A., 2015. An advanced multi-faceted statistical analysis of accident probability and severity exploiting high resolution traffic and weather data.PhD Thesis National Technical University of Athens, Greece.

Thrun, S., 2010. Toward Robotic Cars. Commun. ACM 53, 99–106.

Toledo, T., Koutsopoulos, H.N., Ben-Akiva, M., 2003. Modeling Integrated Lane-Changing Behavior. Transp. Res. Rec. 1857, 30–38.

Transport For London, 2010. Traffic Modelling Guidelines TfL Traffic Manager and.

Triggs, T.J., Harris, W.G., 1982. Reaction Time of Drivers to Road Stimuli. Med. Pregl. 62, 114–9.

Worrall, S., Agamennoni, G., Nieto, J., Nebot, E., 2012. A context-based approach to vehicle behavior prediction. IEEE Intell. Transp. Syst. Mag. 32, 32–44.

Xu, C., Liu, P., Wang, W., 2016. Evaluation of the predictability of real-time crash risk models. Accid. Anal. Prev. 94, 207–215.

Zhang, S., Deng, W., Zhao, Q., Sun, H., Litkouhi, B., 2013. Dynamic Trajectory Planning for Vehicle Autonomous Driving. In: 2013 IEEE International Conference on Systems, Man, and Cybernetics. Ieee, pp. 4161–4166.