12th ITS European Congress, Strasbourg, France, 19-22 June 2017

Paper ID TP0832

Operational Safety for Autonomous Vehicles

Sven Maerivoet[1*] and Lars Akkermans[2]

1. Transport & Mobility Leuven, Belgium, sven.maerivoet@tmleuven.be

2. Transport & Mobility Leuven, Belgium, lars.akkermans@tmleuven.be

Abstract

Autonomous vehicles (AVs) become more and more prevailing in today's world. There is hence a clear need for researching both functional and operational safety. In this paper, we study the latter, implying the full context of a vehicle embedded in a surrounding with other traffic participants, and the related accidents risks. We set up an accident typologies framework, based on available accident analyses (sketches) and literature information, adapted to Belgian accident statistics. This is then validated by means of software simulations, by defining certain scenarios for which we set up various experimental conditions. The results are then interpreted in terms of surrogate safety measures (SSMs) which are then converted into calculted accident and consequently injury risks. This allows to estimate the impact of autonomous vehicles on traffic safety, as well as allowing us to compare them to non-autonomous vehicles (or vehicles with a lower level of autonomy).

Keywords:

Autonomous vehicles, operational safety, traffic safety, surrogate safety measures

Introduction

Functional and operational safety

Autonomous, automated and driver assisting vehicles are highly positioned on the international research agenda. Analysts predict a very large market potential in different sectors. Therefore different companies within the automotive industry are preparing for this future market and develop solutions that enable them to enter this market. Autonomous vehicles (AV) are significantly on the rise, with some OEMs even promising up to Level 5 autonomy [1] by the end of 2017. Whereas the boundaries set by legislation are slowly being dissolved, there is still a large amount of work left concerning the area of (traffic) safety. In this respect, our paper proposes a mechanism to define the operational safety of AVs, as opposed to the already widespread research done on their functional safety. In this paper, we understand the latter to be concerned with the actions to be taken in case of failing sensory input or failing sensors altogether. Operational safety on the other hand considers the vehicle as an element in a full situational context, including other traffic participants and behavioural reactions.

Our main goal, in relation to safety performance, lies in the development of a framework for the modelling and validation of the safety performance of (semi- and fully) autonomous vehicles within a regular mobility context. Currently, the safety question in relation to the use and combination of autonomous vehicles in a closed or open environment is mostly guided by a strict process of stipulating technical requirements for the sensory part and the stipulation of a limited and probably exhaustive set of escape and/or avoidance algorithms on the vehicle control side (that is, the functional safety). Or, in other words: the sensors need to make sure that all and everything that can be monitored is monitored and the vehicle controls need to be operated in such a way that either an immediate standstill or limited evasive action can be taken. However, in itself this does not prove safety, it rather proves technical aptitude to react to a limited input set. Given the nature of accidents (high portion of human error, often with a combination of a combination of crucial errors until an accident no longer can be avoided), this remains a weak point.

Furthermore, no real measuring instrument is currently available that compares “human” error with “machine” error. There is no real test that presents users and developers with an idea of how safe an autonomous vehicle is in comparison to a non-autonomous vehicle. This is both a problem for engineers (who want to know how to improve on their design), funding (which is potentially very sensitive to public opinion and safety), end-users (who would hesitate to buy a vehicle for which each component can be tested, but for which the overall performance remains somewhat shady) and other stakeholders (such as insurances, law makers, etc.). Because of this, our primary is to develop a framework that allows for a true objective and comparative analysis of the safety performance of autonomous vehicles. This can for example be done by looking at an identified set of accidents: so-called fundamental accident diagrams. In a simplified version, these can be found on insurance claims but they are also used by police forces and accident investigators to describe how accidents occurred and under what conditions. By and large, a limited set of diagrams can be identified that represent the majority (95%+) of accidents [ETAC]. These diagrams can be allocated with an occurrence likelihood. In order to reach the primary goal, the autonomous vehicles' technical properties, as well as possibly performance in a testing situation, simulation or real-world situation would need to be analysed in translated into similar occurrence likelihoods, both for individual as well as grouped accident diagrams.

Additional benefits from the development of such a framework would also be the possibility to provided targeted/specific information on the largest improvement potentials in terms of safety for autonomous vehicles, as well as the potential to further develop the technical requirements that are currently being posed on the sensory equipment and vehicle-control algorithms.

Longitudinal modelling in EMDAS

Our research is conducted within the EMDAS (Environmental Modelling for automated Driving and Active Safety) project (led by Flanders' MAKE), which started from a bus that is already guided laterally (thus automated steering), and which is supported by an extensive safety case for operation in the public space based on longitudinal guidance. This project brings human driving notions into an automated vehicle, adding a plurality of sensors, including visual light, infrared light, and radar to the guided bus. By applying high definition image processing techniques and carefully mixing the images as function of weather conditions, we are able to measure the surroundings of the vehicle in most weather circumstances. In similarity to a human driver and in contrast to state-of-the-art automated vehicles we model the surroundings in visible and non-visible spaces and take appropriate action when the vehicle remains in search for information from the non-visible spaces during driving. Within the surroundings we find the traffic signs, lights, pedestrians, and other road users. These road users are classified and probabilities are given for their interaction in the perceived surroundings. The momentary sum of this information is called the environmental model. This environmental model will be used to give the bus its longitudinal guidance by means of an automated speed-profile guaranteeing safety, energy conservation, and comfort.

Modelling system

The whole longitudinal guidance system is modelled using a simulator and a genuine vehicle. While the simulator allows us to do parameter space exploration and demonstrate safety in common accident situations, the real vehicle allows us to demonstrate the correspondence to the simulator. Both the simulator and the real vehicle are created as flexible open platforms and will be a longer term valuable asset for the project's partners and the Flemish community to perform this and subsequent research in automated vehicles. The software system is based on the coupling of MATLAB and Simulink, together with the detailed PreScan software that allows the detailed setup of various scenarios.

Setting up an accident typologies framework

In order to define the various situations in which an AV is expected to drive and interact with other vehicles, possibly leading to accidents, we need to set up an accident typologies framework. To that end, we initially thought to access the GIDAS (German In-Depth Accident Study) database (see also http://www.vufo.de/forschung-und-entwicklung/gidas/?L=1). It was however difficult to obtain access to this database, given the high fees (relative to the original project's budget). Our second attempt therefore focused on the CaDaS (Common Accident Data Set) database: this defines a standard model to report road accidents data for each accident on a common way, thereby ensuring high quality and homogeneity of the input data. We applied extensive efforts to request access to the database, but decided after a long waiting period to fall back on the publicly available information, as well as the information from the SafetyNet project.

Framework background

The framework is based on accident sketches:

- 62 CADAS sketches

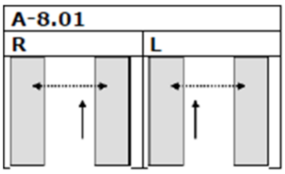

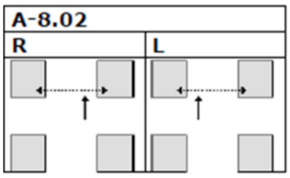

- SafetyNet 8 main accident events (18 detailed ones) and 16 main causes (131 detailed causes in 45 groups)

An example of such a detailed context is shown in the following tables:

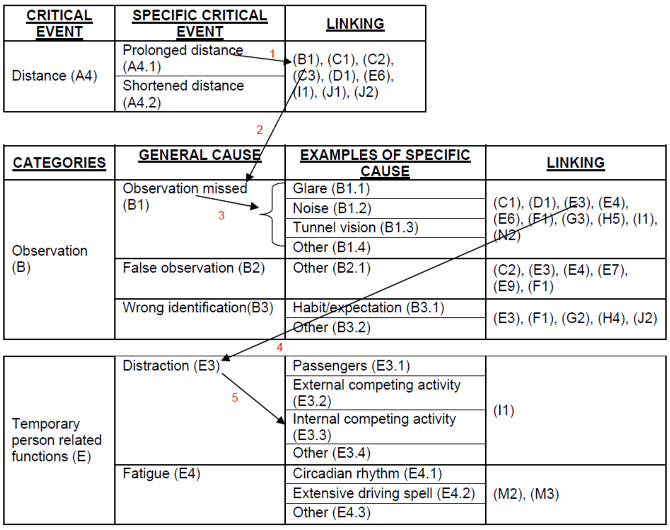

CRITICAL EVENT	SPECIFIC CRITICAL EVENT	LINKING
Distance (A4)	Prolonged distance (A4.1)	(B1), (C1), (C2), (C3), (D1), (E6), (I1), (J1), (J2)
	Shortened distance (A4.2)	

CATEGORIES	GENERAL CAUSE	EXAMPLES OF SPECIFIC CAUSE	LINKING
Observation (B)	Observation missed (B1)	Glare (B1.1) Noise (B1.2) Tunnel vision (B1.3) Other (B1.4)	(C1), (D1), (E3), (E4), (E6), (F1), (G3), (H5), (I1), (N2)
	False observation (B2)	Other (B2.1)	(C2), (E3), (E4), (E7), (E9), (F1)
	Wrong identification(B3)	Habit/expectation (B3.1) Other (B3.2)	(E3), (F1), (G2), (H4), (J2)
Temporary person related functions (E)	Distraction (E3)	Passengers (E3.1) External competing activity (E3.2) Internal competing activity (E3.3) Other (E3.4)	(I1)
	Fatigue (E4)	Circadian rhythm (E4.1) Extensive driving spell (E4.2) Other (E4.3)	(M2), (M3)

Identified accident formats

The different accident sketches in turn allowed us to:

- Assign occurrence percentages to the most prevailing cases
- Advantage: we also obtain literature information on the influence of the biomechanical delays
- That gives us a better estimation of the impact on deaths and wounded

An example of this is shown in the following table:

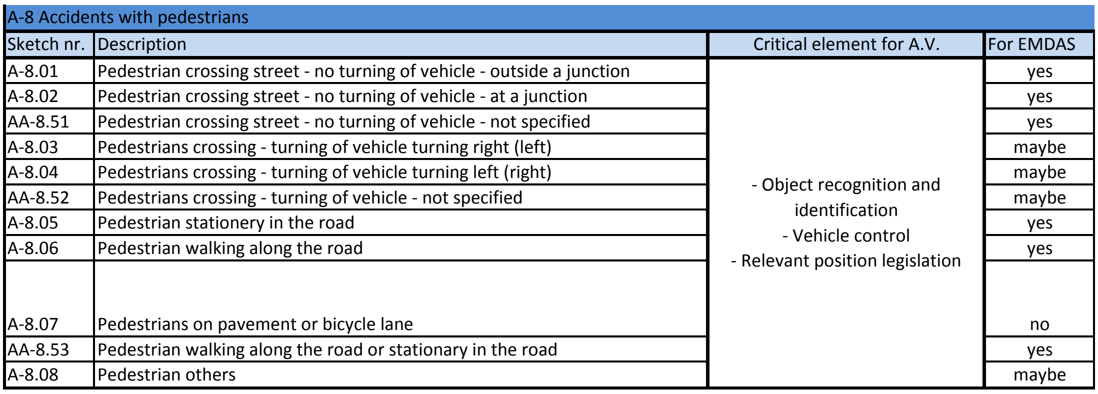

A-8 Accidents with pedestrians			
Sketch nr.	Description	Critical element for A.V.	For EMDAS
A-8.01	Pedestrian crossing street - no turning of vehicle - outside a junction	- Object recognition and identification - Vehicle control - Relevant position legislation	yes
A-8.02	Pedestrian crossing street - no turning of vehicle - at a junction		yes
AA-8.51	Pedestrian crossing street - no turning of vehicle - not specified		yes
A-8.03	Pedestrians crossing - turning of vehicle turning right (left)		maybe
A-8.04	Pedestrians crossing - turning of vehicle turning left (right)		maybe
AA-8.52	Pedestrians crossing - turning of vehicle - not specified		maybe
A-8.05	Pedestrian stationery in the road		yes
A-8.06	Pedestrian walking along the road		yes
A-8.07	Pedestrians on pavement or bicycle lane		no
AA-8.53	Pedestrian walking along the road or stationary in the road		yes
A-8.08	Pedestrian others		maybe

Once this information is known, we can start linking accident sketches to events and causes, as shown in the following scenario matrix:

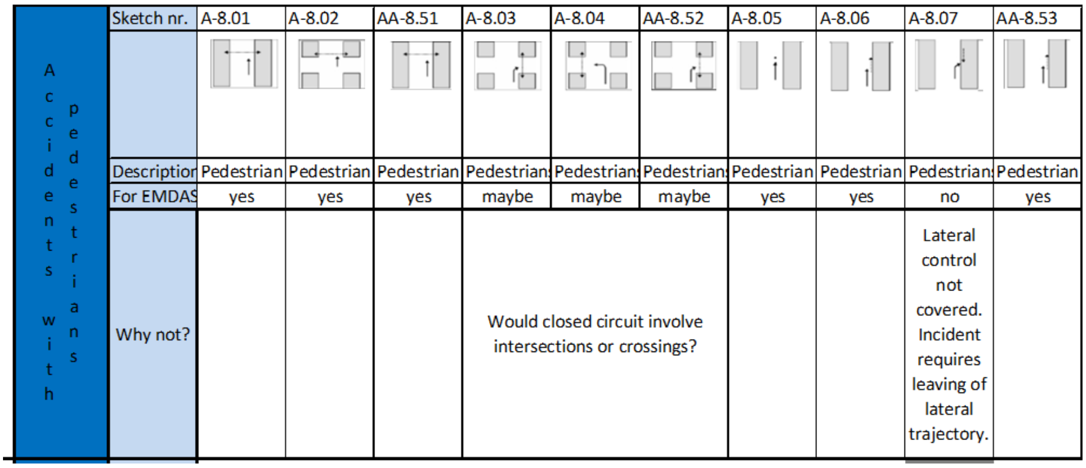

Accidents with pedestrians	Sketch nr.	A-8.01	A-8.02	AA-8.51	A-8.03	A-8.04	AA-8.52	A-8.05	A-8.06	A-8.07	AA-8.53
	Descriptior	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian	Pedestrian
	For EMDAS	yes	yes	yes	maybe	maybe	maybe	yes	yes	no	yes
	Why not?				Would closed circuit involve intersections or crossings?					Lateral control not covered. Incident requires leaving of lateral trajectory.	

Application to the Belgium case

The previous information is then checked and applied for specific observed traffic safety in Belgium, based on available accident statistics on the one hand, and Google Maps identification of existing local infrastructure on the other hand. This gives us the following frequency table, based on information from the Belgian Institute for Traffic Safety (2012 and 2013).

		overall	built-up area	rural area	highway
Frequency based table	road type	est. Victim frequencies	est. Victim frequencies	est. Victim frequencies	est. Victim frequencies
	overall		34.0%	48.1%	6.9%
regular infrastructure	straight	64.4%	20.4%	31.7%	6.0%
	slight curve				0.8%
	normal curve				
	sharp curve				
	crossing (x, perpendicular)	34.5%	13.3%	15.4%	▨
	crossing (x, non-perpendicular)				▨
	crossing (y)				▨
	crossing (t)				▨
	roundabout (3 roads)	1.2%	0.7%	1.0%	▨
	roundabout (4 roads, perpendicular)				▨
	roundabout (4 roads, non-perpendicular)				▨
	Ovonde				▨
	Road splitting/merging (highway ramp)	▨	▨	▨	0.6%

Validation of accident typologies

Our previously described framework captures around 90% to 95% of all accident cases in Belgium. These are in turn validated in the PreScan simulation software (in cooperation with Flanders' MAKE, based on scenario selections). The idea for the latter is to have a common use case to evaluate functional and operational safety hazards. It is important to note that this validation is not exhaustive, nor complete. It does however provides a sound methodology to assess traffic safety of AVs versus non-AVs (in comparison). We furthermore focus on accidents whereby interactions with a bus prevail. This way, we build trust in our methodology, so that it can be applied to other traffic safety cases as well.

Creating PreScan experiments

All our road typologies are aligned with those of Flanders' MAKE's HARA framework. Examples of selected scenarios are:

- Straight road: unimpeded and car-following
- Intersection: without and with traffic lights
- Pedestrian crossing: at zebra and elsewhere
- …

The idea here is to start with one common scenario, and then build the validation on a case by case basis. The MATLAB + Simulink + PreScan software combination allows detailed simulations of various traffic interactions, as graphically depicted in the following figures:

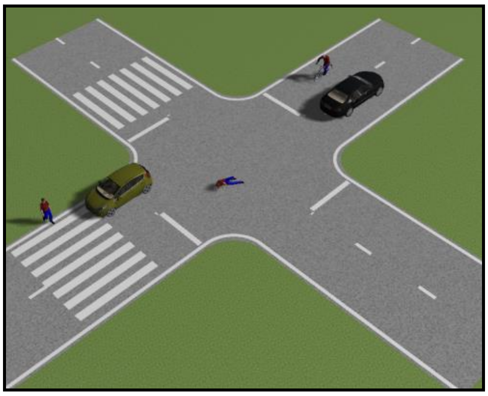

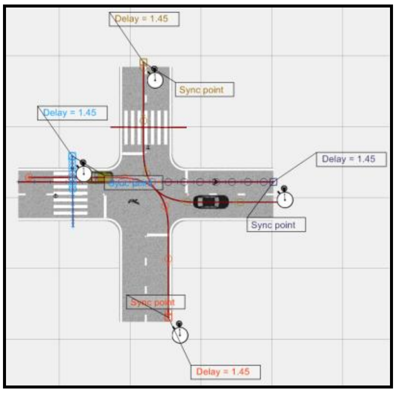

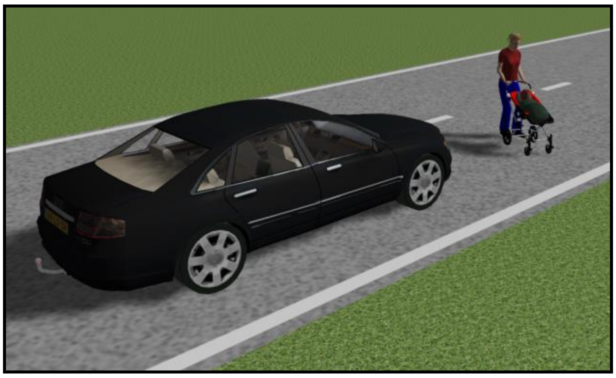

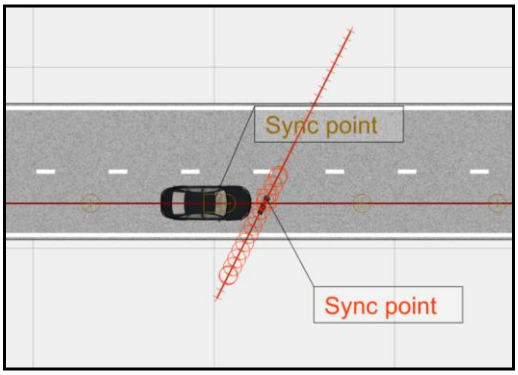

In our specific test setup, we create the following experiment where bus is in car-following mode on a straight road:

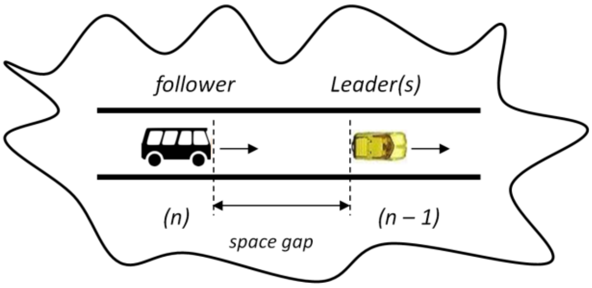

- Initial conditions:
    - Leading vehicle
        - Position ($\mathbf{x_L}$)
        - Speed ($\mathbf{v_L}$)
    - Our following bus
        - Position ($\mathbf{x_B}$)
        - Speed ($\mathbf{x_B}$)

        ➔ Space gap $\mathbf{g_s} = | \mathbf{x_L} - \mathbf{x_B} |$
- Two groups of scenarios:
    - The leading vehicle gradually decelerates to stop ($\mathbf{v_L} \rightarrow 0$)
    - The leading vehicle suddenly decelerates (emerg. brake) ($\mathbf{v_L} \rightarrow 0$)

Experimental setup:

- Perform experiments in ($\mathbf{v_L}$, $\mathbf{v_B}$, $\mathbf{g_s}$) parameter space
    - $\mathbf{v_L}, \mathbf{v_B} \in \{15, 30, 50\}$ km/h
    - $\mathbf{g_s} \in \{1, \ldots, 30\}$ m (@ 50 km/h → ~14 m/s; 2 sec leadway)
- Bus is equipped with a (longitudinal) controller (WP4)
    - The controller gets input from the sensors in the bus model
    - We consider this bus model as black-box
    - Is the controller stochastic?

      ➔ If so, we increase the number of simulations
    - The speed profile of the lead vehicle is non-stochastic
- Case ➔ Scenarios ➔ Experiments ➔ Simulation runs

Introduction of surrogate safety measures

Surrogate Safety Measures (SSM) are events that can be correlated with crash rates. SSMs could be used as indicators of accidents in safety evaluations. SSMs are in particular useful when testing for situations where no real or not enough accident data is available. SSMs can be used in the development of intelligent driver support systems (such as collision avoidance systems) bus also more advanced systems such as Automated Vehicles. SSMs can provide a very useful insight when mixed traffic occurs (not all vehicles are AVs).

SSMs function as indicators and are linked with associated likelihoods to have accidents (collision risk) and accident outcomes (collision severity), given a number of assumptions (such as human driver, deceleration ratios, etc.).

Examples of SSM as indicators for Collision Risk are presented in the following [2, 3]. It should be noted that different indicators are suitable for different types of conflicts: head-on, rear-end, sideswipe, intersections/crossing traffic, etc. [4] provides a good overview of some of the SSMs which could be applied in EMDAS.

- Time-based measures:
    - Gap time (GT)
    - (Minimum) time to collision (TTC)
    - Time to accident (TTA)
    - Encroachment time (ET)
    - (Minimum) post encroachment time (PET)
    - Initially attempted post encroachment time (IAPT)
    - Sideswipe collision risk (SSCR).
- Required braking power measures:
    - (Initial) deceleration rate (DR)
    - Deceleration rate to avoid collision (DRAC)
    - Proportion of stopping distance (PSD)
    - Crash potential index (CPI)
- Safety indices:
    - Time exposed time to collision (TET)
    - Time integrated time to collision (TIT)
    - Difference between TET and TIT
- Examples of SSM as indicators for Collision Severity:
    - Unsafety density parameter (UD)
    - Max speed (MaxS)
    - Relative speed (DeltaS)
    - Kinetic energy
    - Maximum “post collision” DeltaV (MaxDeltaV)

Using these SSMs we can describe two types of conflicts: on a single location in time and space (conflict point), or during a range of times and locations (conflict line). A special case are rear-end conflict lines. This allows for a clear grouping:

- Crossing flows – conflict point events
- Merging crossing flows – conflict line events
- Adjacent flows – lane-changing conflict line events
- Following flows – rear-end conflict line events

Note that additional collision types such as pedestrian collisions and U-turn related conditions do pose some difficulties for SSMs. In addition, evasive maneuvers are mostly not represented (changing lanes, swerving, accelerating, …) and not all conflict event contributors are directly integrated in the (estimation of) SSMs. However the methodology to estimate SSMs can be adjusted up to a certain extent to allow for differences in these contributors: visual obstructions and occlusion, sunlight blinding, weather conditions, road signage, …

Experimental application of SSMs

In our software similations, we ‘continuously’ measure:

- SSM_1 = space gap ($\mathbf{g_s}$)
- SSM_2 = speed difference ($\mathbf{v_B} - \mathbf{v_L}$)

Hence, each SSM leads to a {**safe**, **unsafe**, **accident**} interpretation based on literature.

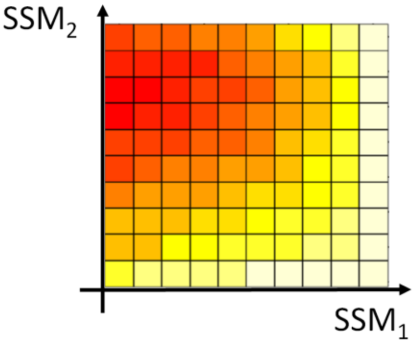

The next step is then to analyse the SSMs in light of AVs:

- Is the controller safe or not?
  [➔ **accident risk**]
- If not: what is the result?
  [➔ **injury risk**]

Depending on what these SSMs represent (fatal / severe / slightly / almost wounded), we calculate the injury risk and possibly modify the interpretation of the SSMs specifically in light of AVs.

It needs to be noted that this use of SSMs is different from more traditional uses. Current uses of SSMs in the context of vehicle automation are limited to the analysis of an overall safety effect of the use of AVs in a mixed automation context or similar. For example, while AVs might accept small time gaps between vehicles to cross vehicle streams, this may cause problems for non-AVs which are within such a vehicle stream. Currently, the SSMs that are described are somewhat limited in truly assessing the safety impact of the introduction of AVs in (mixed) environments.

Within this project, we solely created and tested a methodology which uses SSMs and modelling for assessing the safe operation of the AV, and only the AV, in the same way that one would test a non-AV. In a sense, we tested the proposed methodology against (1) the sensors, their monitoring of the environment and the interpretation of the data they provide, and (2) the response the AV poses to such an environment. And this, for a preselected set of driving conditions. Note that if this method on the longer term proves to be feasible and can be validated, this can be extended to more driving conditions and more complicated measurements.

Example application

Based on the results from the simulation model (carried out under controlled conditions), we will be able to for example assess what the benefits of autonomous vehicles are compared to a human driver, taking into account the operation safety of the vehicle. Given a (simulated) human driver and a computer controller, we use the previously described framework to understand where and how an autonomous vehicle outperforms a human, given the accident typology associated with the former. An example of such an analysis is shown in the following table, where the speed difference between both ‘drivers’ after an impact of the vehicle with a static object:

Impact speed difference A.V. - Human driver														
10	20	30	40	50	60	70	80	90	100	110	120	130	140	150
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2.84	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*
10.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
10.63	*3.11*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*
9.59	9.92	0.59	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
8.41	*14.20*	*5.25*	*-0.56*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*
7.03	14.58	9.98	3.69	-1.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5.30	*15.68*	*12.17*	*9.14*	*2.98*	*-1.13*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*
2.56	15.00	15.58	10.51	7.09	2.76	-0.92	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2.17	14.27	17.47	12.11	9.96	6.69	2.86	-0.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.83	*13.51*	*17.83*	*14.08*	*11.13*	*9.75*	*6.68*	*3.20*	*0.08*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*	*0.00*
1.39	12.71	19.22	16.77	12.44	10.71	9.71	6.93	3.72	0.80	-1.89	0.00	0.00	0.00	0.00
0.75	11.84	18.91	20.04	13.95	11.76	10.54	9.75	9.20	4.39	1.64	-0.92	0.00	0.00	0.00
0.00	10.91	18.34	20.23	15.75	12.90	11.43	10.50	9.85	9.37	5.18	2.57	0.13	-2.18	0.00
0.00	9.90	17.76	20.58	18.05	14.16	12.37	11.28	10.52	9.97	9.55	6.08	3.59	1.25	-0.98
0.00	8.76	17.15	21.34	21.59	15.59	13.40	12.10	11.23	10.60	10.11	9.74	9.43	4.68	2.43

An autonomous vehicle as such thus therefore not have to drive millions of kilometres in order to gain appreciation and trust from the community and homologation bodies, so as to be commissioned on the road. This hugely impacts legislation, field operational tests (FOTs), and wide adoption of AVs.

Going further

In principle an AV drives millions of vehicle kilometres as so-called validation of its functioning. Examples of reporting are then Google and Tesla that check/incorporate exceptional cases in their continuous updates. Our proposed framework already covers 80% to 90% of the infrastructure, based on traffic safety analyses. As we incorporate more context, we can also identify small mistakes leading to innocent versus critical situations with respect to opening up the public roads for AVs. We can furthermore also create a tool to list “unexpected” or “impossible” situations to find the limits of AVs related to our safety scenarios. This is very useful as, e.g., the UK is also following this approach. Some of the key ideas here are (1) to create a checklist for supporting the modifications of existing roads to accommodate AVs, and (2) the identification of points of attention related to AVs’ traffic safety. This furthermore allows then to answer the following research questions:

(1) Are AVs, equipped with the right technology, suitable for daily use in existing infrastructure?

(2) Is the existing infrastructure suitable for AVs, albeit in mixed traffic?

(3) Which modifications are necessary to make a road or vehicle suitable?

A possible approach to answer these is then to setup three modules:

- Default infrastructure: Assuming ‘ideal’ conditions within a simulation environment, to test sensors, data collection and analysis, vehicle control…
- Exceptional infrastructure: Aimed at testing the interaction between infrastructure, AVs, and the various phases in autonomous driving, based on ‘real’ conditions and looking at potential conflicts, defunct infrastructure, … Goal is to make infrastructure future-proof
- Checklist for modifications To be used when supporting field trials or real deployment of AVs on (Belgian) roads.

Conclusions

As autonomous vehicles (AVs) become more and more prevailing, and not just remaining a topic of hot conversation, we clearly see an arising need for researching both functional and operational safety. In this paper, we study the latter, implying the full context of a vehicle embedded in a surrounding with other traffic participants, and the related accidents risks. Our paper proposes a methodology for defining and accident typologies framework, based on available accident analyses (sketches) and literature information. These are then subject to a statistical verification by means of available data in of accidents on the Belgian road infrastructure. Our framework is then validated by means of software simulations, giving more credibility to its use and the related operational traffic safety. The software defines certain scenarios for which we set up various experimental conditions. The results are then interpreted in terms of surrogate safety measures (SSMs) which are then converted into calculted accident and consequently injury risks. This allows to estimate the impact of autonomous vehicles on traffic safety, as well as allowing us to compare them to non-autonomous vehicles (or vehicles with a lower level of autonomy). Thus our paper created and validated a methodology which allows for the comparison of safety-related aspects, such as the safe functioning (i.e., road safety), between autonomous vehicles and non-autonomous vehicles.

Acknowledgements

We wish to thank the members of the EMDAS consortium (Flanders' DRIVE, Xenics, VDL, Universiteit Gent, TomTom, Transport & Mobility Leuven, and Siemens—LMS). EMDAS is funded by the Flemish IWT (VLAIO), reference 140647.

References

1. SAE International (2014) Automated Driving: Levels of Driving Automation Are Defined in New SAE (Society of Automotive Engineers) International Standard J3016, https://www.sae.org/misc/pdfs/automated_driving.pdf
2. Gettman, D.; Larry, H. (2003) Surrogate Safety Measures from traffic simulation models, final report. FHWA-RD-03-050
3. Behbahani, H.; Nadimi, N. (2015) A framework for applying surrogate safety measures for sideswipe conflicts. International Journal for Traffic and Transport Engineering, 2015, 5(4): 371-383. http://dx.doi.org/10.7708/ijtte.2015.5(4).03
4. Cunto, F. J. (2008) Assessing Safety Performance of Transportation Systems using Microscopic Simulation. PhD Thesis, University of Waterloo, Waterloo, Ontario, Canada