Volvo Cars

Submission to the
National Transportation Safety Board (N.T.S.B.)
For the Tempe Accident involving an UBER test vehicle based on a Volvo XC90 MY2017

Gothenburg, Sweden; October 2019
HWY18MH010

TABLE OF CONTENTS

1. Introduction

UBER performed field tests with a modified model year 2017 Volvo XC90 base vehicle. The vehicle was modified in accordance with UBER specifications to be part of Autonomous Driving development by UBER.

This Submission describes Volvo Car Group from a safety culture perspective, as a developer of safety leading vehicles and functions as well as the company's development of autonomous driving technology.

With the challenge of providing autonomous driven vehicles comes a rapid growth of research needs. Many research initiatives with different stakeholders are already progressing within research areas supporting technology development and more initiatives will be required moving forward. Some findings from ongoing research activities by Volvo Cars are described in this submission supporting the investigation. In addition, a potential framework for evaluation of AD and strategies for improved safe on road testing is described.

2. Volvo Car Group

The first Volvo car rolled off the Gothenburg production line in Sweden in 1927. Since then, Volvo Car Group has been a world-leader in safety technology and innovation. Today, Volvo is one of the most well-known and respected car brands in the world with sales in about 100 countries.

Volvo Car Group (Volvo Cars) is owned by Zhejiang Geely Holding (Geely Holding) of China. Our group structure comprises Volvo Cars and related businesses: car subscription service Care by Volvo, mobility company M, electric performance arm Polestar as well as sizeable stakes in new Chinese car brand LYNK & CO and software company Zenuity. The 'Volvo' name is owned by Volvo Trademark Holding AB, which is jointly owned by Volvo Cars and Volvo Group.

Volvo cars are marketed and sold by regional market companies and national sales companies through approximately 2,300 local dealers in about 100 countries. Most of the dealerships are independent companies. As of December 2018, Volvo Cars had around 43,000 full-time employees around the globe.

Volvo Cars head office, product development, marketing and administration functions are mainly located in Gothenburg, Sweden. Since 2011, Volvo Cars has had offices in Shanghai and Chengdu (both China). The Volvo Cars China headquarters in Shanghai includes a Technology Centre and functions such as sales and marketing, manufacturing, purchasing, product development as well as other supporting functions.

Apart from the main European car production plants in Gothenburg (Sweden) and Ghent (Belgium), Volvo Cars has since the 1930s manufactured engines in Skövde (Sweden). Production of body components has been located in Olofström (Sweden) since 1969.

The company furthermore operates assembly plants in Kuala Lumpur (Malaysia) and Bengaluru (India), as well as research and development centres in Shanghai, Stockholm, Lund (both Sweden) and Sunnyvale (USA). Volvo Cars also has design centres in Gothenburg, Camarillo (USA) and Shanghai.

In 2013, Volvo Cars started production at its first Chinese manufacturing plant in Chengdu. A second plant in Daqing started production in 2014, while Volvo Cars also operates a manufacturing plant in

Luqiao and an engine factory in Zhangjiakou in China. In the US, Volvo Cars has a manufacturing facility in Berkeley County, South Carolina, which opened in 2018.

3. Safety at Volvo Cars

Volvo Cars has been working with car safety and safety features since the beginning of the company’s history. The working methodology used is called the Circle of Life. It is a successful procedure where real-life data from accidents is collected and analysed to develop cars that are even safer than the ones before.

In 1970, Volvo Cars established a comprehensive traffic accident research programme, the Traffic Accident Research Team, which collects data from traffic accidents with Volvo cars around Sweden. The R&D teams use this data when creating and evaluating new technologies and cars. Additional data is further gathered by the Traffic Accident Research Team through accident reconstructions, computer simulations and crash tests in our state-of-the-art crash laboratory in Torslanda, Sweden. The data analysis insights are then used in the design process. This is a key enabler to our success and leadership.

As we move into the era of autonomous cars, safety will be even more crucial. To facilitate the transition to autonomous cars, we place a lot of focus on the human perspective. We drive the development of global standards for communication between autonomous cars and people moving around the car and we generally believe in an open approach where we collaborate with other stakeholders in the cities, the automotive industry and adjacent industries. We want to ensure a safety perspective in all our consumer offers.

Together with a world leading active safety supplier, Veoneer (a spin-off from Autoliv), we have formed the joint venture Zenuity to develop advanced driver assistance systems and AD software technology.

Since 2016, we have been working with mobility provider Uber to develop an autonomous driving ready base vehicle based on the XC90. In 2018 we announced the intention to team up with the Chinese tech giant Baidu for the development of fully autonomous vehicles for fleet usage in China.

We will leverage our leading position in safety as well as AD technology capabilities as we commercialise our solutions. We will offer premium autonomous cars to consumers and we aim to be the partner of choice for robo-taxi fleet operators.

4. Research findings in support of the investigation and future development of AD

Two recent and important findings from our current research are coming to the forefront in discussions on how to safely develop and deploy AD capable vehicles in the light of the Tempe investigation.

The ability of drivers to be ready as a fall back

When testing AD software in a development vehicle fleet, the test driver is *by definition* put in a supervising role rather than being an active driver. In SAE terminology, the test driver assumes the role of the fall-back ready user required for Level 3 driving automation, where a "DDT [dynamic driving task] fall-back-ready user is considered to be receptive to a request to intervene and/or to an evident vehicle system failure, whether or not the ADS [automated driving system] issues a request to intervene as a result of such a vehicle system failure." (SAE J3016, p. 22).

Experience with human factors in automation over the last 50 years indicates that understanding human supervision capabilities is key to ensure safe operation of automated systems (Bainbridge, 1983; Billings, 1988; Sheridan, 1992; Endsley & Kiris, 1995; Sarter & Woods, 1995; Parasuraman & Riley, 1997; Lee et al., 2017).

It has been assumed that within the domain of driving, operators (drivers) are less affected by the underperformance problems documented in other automation domains, due to the very tight coupling between failing to respond to an automation malfunction and the consequences that follow. In short, it has been assumed that anyone with a driver's licence should be able to brake and/or steer as required whenever needed to avoid a crash, as long as they are actually monitoring the traffic environment. This assumption underlies the SAE definition above.

This assumption is now being questioned, and discussions on how to ensure that drivers do monitor traffic and are not doing something else has started to increase. This is reflected in arguments over terms such as Driver engagement, Driver-in-the-loop, Monitoring, Supervision, Fall-back readiness, Receptiveness, Availability, Over-reliance, and Complacency. However, the implicit understanding which follows with the above assumption, i.e. that if a driver is monitoring properly s/he will also respond properly, has not yet received the same attention and research focus.

To investigate this implicit understanding, Volvo recently performed a series of test track experiments at the ASTA Zero proving grounds in Sweden as part of the ADEST project (Autonomous Driving Effects on Sustainable Transportation). The aim was to further study driver behaviour during supervised automation, and in particular what would happen when the supervising operators after half an hour of successful automated driving in a research vehicle encountered an automation failure (see publications from june 2018 and onward, i.e. Victor et al. 2018, Gustavsson et al. 2018, Tivesten et al. 2019).

There were two main results from these experiments. One is that supervision reminders can be highly effective when it comes to increasing eyes on road time as well as maintaining hands on the steering wheel. The other is that eyes on road and hands on wheel may not be a sufficient safeguard to ensure proper management of conflict situations that may occur during supervised fully automated driving as tested in the experiment. The crash rates in the study were independent of how detailed instructions on system limitations the test subjects were given as well as independent of the type of supervision reminder issued and the type of conflict object encountered.

Analysis of the interview data suggests that some vehicle operators developed high trust in the research vehicle after experiencing highly reliable and stable driving performance during the 30-minute drive, and they tended to forget the initial instructions on function limitations rather quickly. Indeed, only 11 out of 76 vehicle operators were classified as “role models”, i.e. acted and explained their actions in full accordance with the content of the instructions before the drive.

These results illustrate that the concept of a fall-back ready driver needs to be questioned, and that more robust ways of ensuring test fleet safety likely need to be pursued.

The means of fleet driving will not ensure safety performance of an AD vehicle alone.

If we look at that actual exposure to crash risk, we can see that crashes in manual driving are very rare events (Nidhi & Paddock, 2016; Lindman et al, 2017). The level of crash avoidance performance that automation must meet to surpass human performance or achieve the vision of zero fatalities and serious injuries is therefore very high (Johansson, 2009; Eugensson et al., 2011). For example, in Sweden, there is 39.1 million km driven per crash with severe- or fatal injuries in passenger cars when all traffic environments are included, or 126 million km driven per crash when the analysis only includes motorways and urban highways (Lindman et al, 2017).

The implication of these numbers is that establishing the safety of a particular AD technology through mileage accumulation simply is not possible. To further exemplify, Lindman et al (2017) did a calculation of the mileage required to prove that AD is safer than human driving on a motorway in Sweden. Their finding was that to prove with statistical significance that driving in AD mode reduces the risk of a fatal crash by 50% in a Swedish driving context, 7 billion km in AD mode is required. Similar findings were shown in Kalra and Paddock (2016). Thus, the safety of AD vehicles cannot be established by accumulating a large number of accident free miles in AD mode alone. Other methods to establish adequate safety are also required.

5. Submission summary

Safety evaluation of AD:

The calculation of mileage required to prove AD safety shows that in parallel to pursuing the technological development of an AD function, it is important to work on methods for establishing that the function under development will be able to avoid or mitigate the crashes of today before deploying it on the roads. While in other types of software development it might be best practice to wait for edge cases to occur and then deal with them as they come rather than actively seeking them out, when programming for road traffic, that is not a viable approach. It is important to make sure that we can manage crashes of the types that already occur today (within the Operational Design Domain pursued) before deploying these vehicles to the public in real traffic.

Volvo Cars has spent substantial effort over the last few years in working out suitable new methodologies for how to measure the potential safety impact of AD technology, and how to set up a fair comparison with the traffic safety levels of today. The Volvo Cars approach is to work with predictive methods that estimate real-life crash outcomes based on computer simulations of crash outcomes as well as test track experimental methodologies. The simulations estimate real-world safety performance by comparing the outcome of the pre-crash phase with and without the autonomous driving system enabled for a selected set of crashes (e.g. Lindman et al, 2010; Page et al, 2015), and the test track experiments help validate various aspects of the models used in the simulations. A large part of this work is summarized in the ADEST report (Victor et al, submitted), where the holistic automation safety impact assessment framework that was developed is further described.

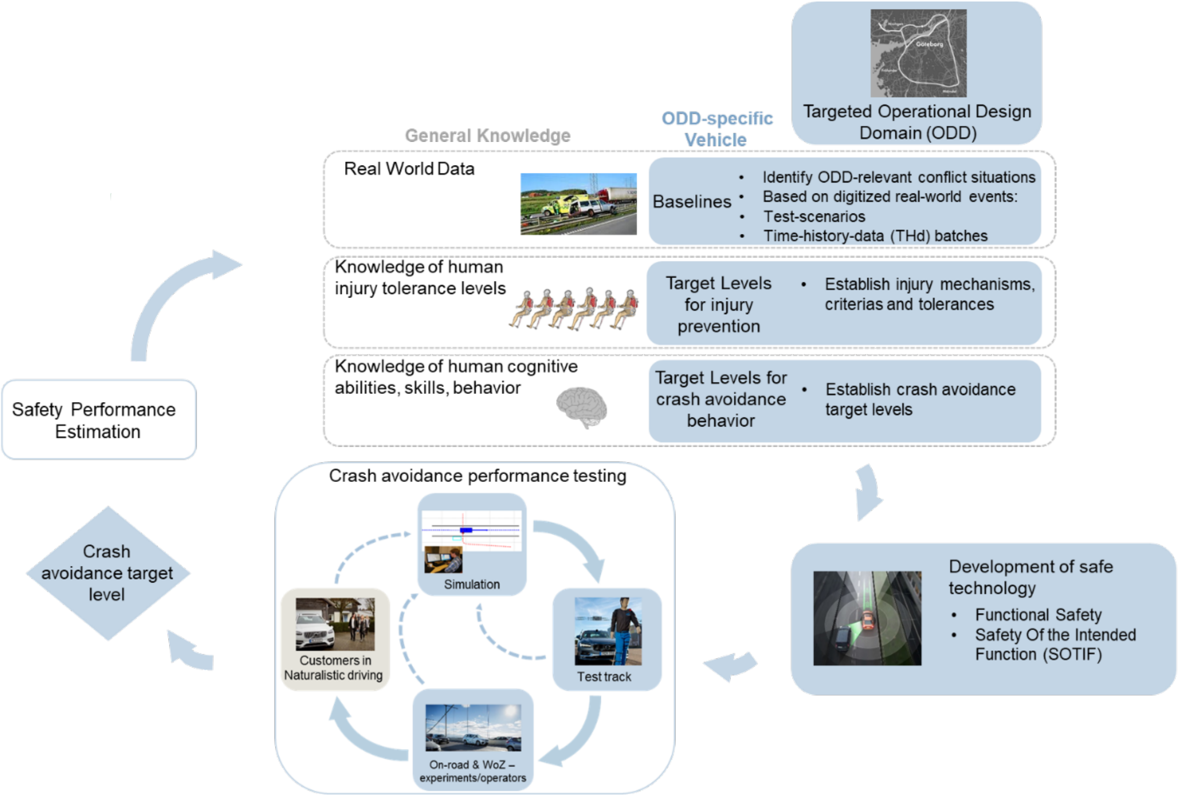

Figure 1: The holistic automation safety impact assessment framework developed by VCC in the ADEST project

On road testing during development:

The research studies of supervising capabilities in drivers from ADEST, briefly presented above, have provided novel and important insights into the challenge of the supervising role for the human in highly reliable (but not perfect) automation. They are in line what can be called the classic irony of automation (the better the automation, the less attention drivers will pay to traffic and the system, and the less capable they will be to resume control) and can influence our way of working to mitigate this aspect.

As of today, we envision an approach with some conditions to try to reduce the exposure to potential supervision failures as much as possible before test driving on public roads, such as:

- *Working rather than supervising*. Engage the supervisor. One must devise ways in which what the test driver is doing becomes meaningful and coupled to the function development. Annotation of potential conflict situations is an example of a task that could engage drivers beyond supervising the AD function.
- *Two persons in a test car*. Social aspects are important. Research, in particular within aviation safety inspections, show that people perform better when in company with others. Exactly why is not fully understood, but the effect is there and should be leveraged. A challenge in this is finding a meaningful task and role for the second person in the vehicle.
- *Short shifts*. Typically, vigilance decrements set in quite fast for routine tasks. For example, FAA recommends that airport personnel inspecting baggage should rotate every 20-30 min to keep the workload reasonable in relation to vigilance. This could be a good reference number for how long driving shifts could be for an AD function that is mature. Having two persons in the car may also simplify frequent driver rotations.
- *Adequate training*. Test drivers in an AD development fleet need training on the limits of human supervision capability, as well as on what the function being tested is designed to do, and thus by extension what type of feedback the developers will find most meaningful. As this will change over time, training is something that could be a reoccurring event with adaptation to current functionality and maturity.

Further, more co-operative and cross-functional research providing new knowledge, not only on technology, but more so to insights in human behaviour both when driving or not is required. This research need includes the development of methods and tools for analysis and verification of the safety of the AD technologies to be deployed.

6. References

Bainbridge, L. (1983). Ironies of automation. *Automatica,* 19(6), 775-779.

Billings, C.E. (1988). Toward human centered automation, in Flight deck automation: Promises and realities, S.D. Norman and H.W. Orlady, Editors., NASA-Ames Research Center: Moffet Field, CA. p. 167-190.

Endsley, M.R. and E.O. Kiris, (1995). The out-of-the-loop performance problem and level of control in automation. Human Factors. 37(2): p. 381-394

Eugensson, A., Ivarsson, J., Lie A. and Tingvall, C., “Cars are driven on roads, joint visions and modern technologies stress the need for co-operation”. International Technical Conference on the Enhanced Safety of Vehicles (ESV), 11-0352, Washington DC, USA, 2011.

Gustavsson, P., Victor, T. W., Johansson, J., Tivesten, E., Johansson, R., & Aust, L. (2018). What were they thinking? Subjective experiences associated with automation expectation mismatch. Proceedings of the 6th Driver Distraction and Inattention Conference, (October), 1–12.

Lindman, M., Isaksson-Hellman, I., & Strandroth, J. (2017). Basic numbers needed to understand the traffic safety effect of Automated Cars. In IRC-17-40 IRCOBI Conference 2017. (pp. 244–256). Retrieved from http://www.ircobi.org/wordpress/downloads/irc17/pdf-files/10.pdf

Lindman, M., Ödblom, A., Bergvall, E., Eidehall, A., Svanberg, B., & Lukaszewicz, T. (2010). Benefit Estimation Model for Pedestrian Auto Brake Functionality. In ESAR 2010.

Johansson, J., "Vision Zero - Implementing a policy for traffic safety". Safety Science, 2009:47(6), 826-831. doi:http://dx.doi.org/10.1016/j.ssci.2008.10.023.

Lee, J.D., Wickens, C., Liu, Y., Boyle, L. (2017). Designing for People: An introduction to human factors engineering, 3rd Edition, CreateSpace, ISBN: 1539808009

Nidhi, K. & Paddock, S. M. (2018) Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? Santa Monica, Calif.: RAND Corporation, RR-1478-RC. As of August 31, 2019: https://www.rand.org/pubs/research_reports/RR1478.html

Page, Y., Fahrenkrog, F., Fiorentino, A., Gwehenberger, J., Helmer, T., Lindman, M., Wimmer, P. (2015). a Comprehensive and Harmonized Method for Assessing the Effectiveness of. The 24th International Technical Conference on the Enhanced Safety of Vehicles (ESV), 1–12. Retrieved from http://www-esv.nhtsa.dot.gov/Proceedings/24/isv7/main.htm

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230-253. http://dx.doi.org/10.1518/001872097778543886

SAE. (2016). SAE J3016 Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicle. Retrieved from https://saemobilus.sae.org/content/j3016_201609

Sheridan, T. (1992). Telerobotics, automation and human supervisory control. Cambridge, MA: MIT Press.

Tivesten, E., Victor, T., Gustavsson, P., Johansson, J., & Ljung Aust, M. (2019). Out-of-the-loop crash prediction: The Automation Expectation Mismatch (AEM) algorithm. IET Intelligent Transport Systems, 1–10. https://doi.org/10.1049/iet-its.2018.5555

Victor, T. W., Tivesten, E., Gustavsson, P., Johansson, J., Sangberg, F., & Ljung Aust, M. (2018). Automation Expectation Mismatch: Incorrect Prediction Despite Eyes on Threat and Hands on Wheel. Human Factors, 60(8), 1095–1116. https://doi.org/10.1177/0018720818788164