Automated automobiles and ethics

What should we focus on?

Hakkala, Antti

University of Turku, Department of Future Technologies

Heimo, Olli I.

University of Turku, Turku School of Economics

Corresponding Author: Antti Hakkala, ajahak@utu.fi

Abstract:

In this paper we observe automated road vehicles via lenses of technology, hacking, society, and ethics. We analyse current ideas, technologies, and discourse around automated vehicles, focusing on security issues in current and by extension future automobiles. As the seemingly inevitable transformation from manual to automated road transportation progresses, we provide necessary and topically relevant discussion on what should be the focus when the next stage of transportation is developed.

Keywords: Autonomous systems, ethics, information security, society

Citation: The full citation information will be inserted here once the paper is accepted.

Introduction

Automated vehicles are here. Computers are replacing humans as truck drivers and Google’s AI is driving down the streets of California (Hawkins, 2017; Levin & Harris, 2017; Rushe, 2017), militaries plan using automated vehicles to support troops in dangerous areas (Magnuson, 2017), drones fly our deliveries through our skies (Amazon, 2018) and Elon Musk is predicting that no one will be allowed to drive a car in near future (Lowensohn, 2015) – because automation supposedly makes less mistakes.

Automated vehicles have the potential to solve many of problems, e.g. people driving while tired, distracted, intoxicated, or with limited skills. Human mistakes can be eliminated by taking the human out of the loop. But what are the downsides of this technology and have ethical issues been considered in sufficient depth while developing this emerging technology?

The Internet of Things (IoT) has emerged as a platform of interconnected everyday devices equipped with microprocessors and the capability to collect, process and share data with other similar devices and backend servers. As IoT solutions became more common, hacking them became more commonplace too. Nowadays a news report of a refrigerator sending junk mail (Bort, 2014) would not even get published – while similar news went viral just four years ago!

Presently, news are portraying hacking incidents on a massive scale. According to Kaspersky Lab, millions of computers are mining crypto currencies without their owners knowing (Lopatin & Bulavas, 2017). Also, major ransomware attacks have been made in recent years where the computer had been taken over and the information inside it encrypted (see e.g. (BBC, 2017a, 2017b; Constantin, 2016)).

The underlying question about autonomous vehicles and ethics lies within the need, utility and risks involved with the emerging technology. As shown before, the need and utility are unarguable but how do they cope with the risks? It is important to observe that practically everything controlled by a computer can be hacked, and modern cars are not exempt. This in turn puts all those people on the streets, roads, and alleys at a risk. The questions we must therefore pose are: can we afford the risks involved with autonomous cars? Are they a better option than cars driven by humans?

Automated Road Vehicles

The contemporary automobile already hosts multiple Electronic Control Units (ECUs) responsible for controlling various operational functions of the car, ranging from engine and drivetrain control to brakes and entertainment systems. Each ECU executes its own programming with lines of code numbering in the millions, and given that a modern car has around 50-70 individual ECUs, an estimation that a car runs on 100 million lines of code, or more, is certainly reasonable (Gaertner, 2015). Although computers are already responsible for many operational tasks, the driving responsibility is still in the hands of the human driver for most, if not all, situations. Indeed, the share of control a computer has over the actual driving task is a key metric in defining what is an autonomous vehicle.

The key terms and concepts for autonomous vehicles are defined in the Society of Automotive Engineers (SAE) standard J3016, “Taxnonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles” (SAE International, 2016).

It establishes six levels of driving automation for on-road vehicles, going from 0 (no automation) to 5 (full automation). On levels 0-2 the human driver is in charge, while on levels 3-5 the computer has the (main) responsibility for driving the vehicle.

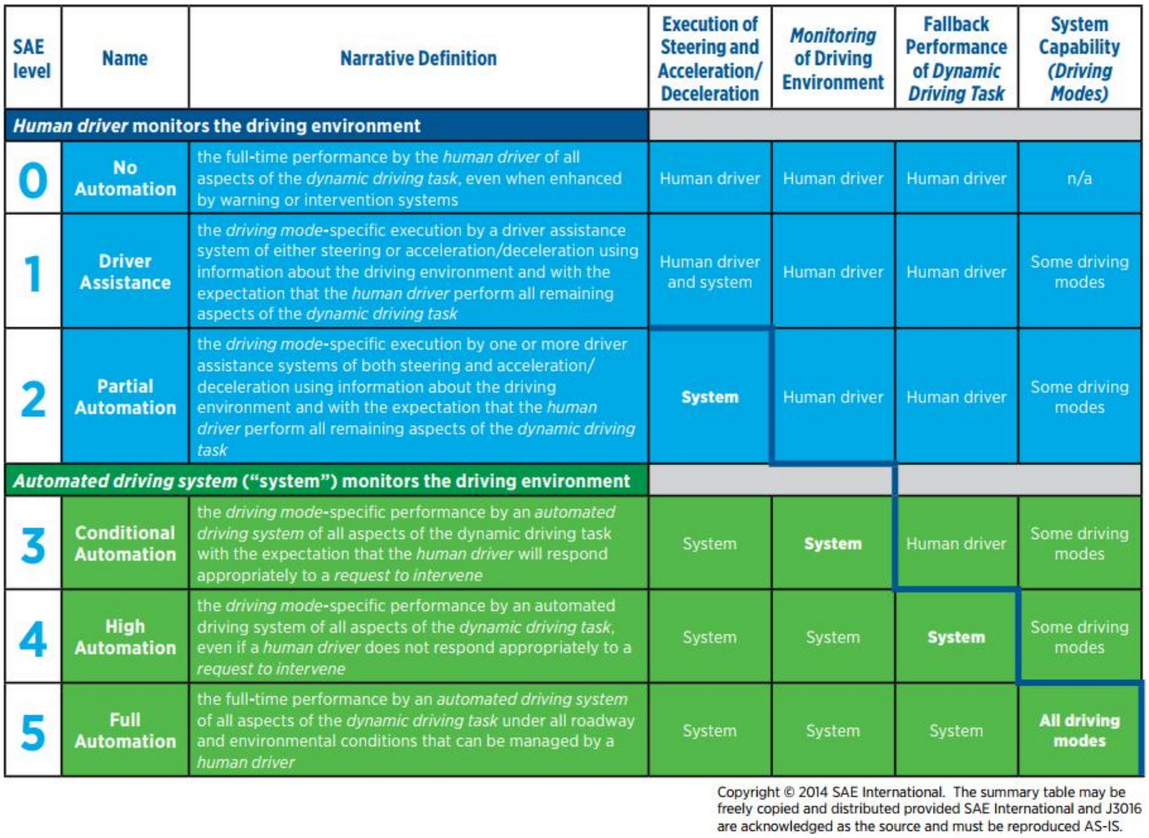

SAE level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	*Monitoring* of Driving Environment	Fallback Performance of *Dynamic Driving Task*	System Capability *(Driving Modes)*
Human driver **monitors the driving environment**						
0	No Automation	the full-time performance by the *human driver* of all aspects of the *dynamic driving task*, even when enhanced by warning or intervention systems	Human driver	Human driver	Human driver	n/a
1	Driver Assistance	the *driving mode*-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	Human driver and system	Human driver	Human driver	Some driving modes
2	Partial Automation	the *driving mode*-specific execution by one or more driver assistance systems of both steering and acceleration/ deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	**System**	Human driver	Human driver	Some driving modes
Automated driving system **("system") monitors the driving environment**						
3	Conditional Automation	the *driving mode*-specific performance by an *automated driving system* of all aspects of the dynamic driving task with the expectation that the *human driver* will respond appropriately to a *request to intervene*	System	**System**	Human driver	Some driving modes
4	High Automation	the *driving mode*-specific performance by an automated driving system of all aspects of the *dynamic driving task*, even if a *human driver* does not respond appropriately to a *request to intervene*	System	System	**System**	Some driving modes
5	Full Automation	the full-time performance by an *automated driving system* of all aspects of the *dynamic driving task* under all roadway and environmental conditions that can be managed by a *human driver*	System	System	System	**All driving modes**

Copyright © 2014 SAE International. The summary table may be freely copied and distributed provided SAE International and J3016 are acknowledged as the source and must be reproduced AS-IS.

Figure 1: Summary of SAE levels of driving automation. (SAE International, 2016)

On level 0, all aspects of the Dynamic Driving Task (DDT) (i.e. real-time operational and tactical control of the car) are controlled by the human driver. If there are automatic safety or warning systems that observe the environment, they only warn the human driver if defined safety parameters are violated, and do not interfere with the DDT. On level 1, an automated system can interfere with steering OR acceleration/deceleration, while on level 2 multiple automated systems together can interfere with steering AND acceleration/deceleration, effectively taking full control of the car from the human driver. On level 3 and onwards, the Automated Driving System (ADS) oversees the DDT. On level 3 the human driver is the backup operator should the ADS encounter a situation it cannot handle. On level 4 the ADS is expected to oversee the DDT and, for some driving modes or scenarios (e.g. driving on a freeway, parking, etc.) manage error situations where a human backup operator would be required on a level 3 system. On level 5, the ADS is in full control of the car, in all situations and all driving modes, without the need (or even the possibility of) human intervention.

As the automation level increases, the computer must make more decisions (both ethical and operational) as the human driver increasingly becomes a passenger in the vehicle due to decreasing ability to control the vehicle. Indeed, the J3016 standard explicitly states that on SAE levels 4 and 5, the ADS does not have to immediately disengage from the DDT upon human request for control (SAE International, 2016, pp. 20-21).

Technical implementation

As is shown in Figure 1, autonomous vehicles need to steer the vehicle, monitor the driving environment, and make decisions on how to perform the DDT. For monitoring the environment, an autonomous vehicle must rely on various sensors. Common sensor types include video feeds in various wavelengths (normal light, infrared), ultrasound, LIDAR and RADAR. All these sensors provide information upon which decisions on how to perform the DDT are founded. A model of the dynamic environment based on sensor observations is used to determine the surroundings of the vehicle, and to determine whether there are any obstacles on the road, where other vehicles are located, detecting pedestrians, etc.

Should the data provided by the sensors be somehow compromised, either by accident or by design, the decisions made by the vehicle on how to perform the DDT are not necessarily correct. The integrity of the sensor data is therefore critical to the correct functioning of the car. Should an adversary wish to interfere with an autonomous vehicle, a simple way is to attack the sensors, causing for example the vehicle to incorrectly perceive a pedestrian as something else, or fail to observe them altogether.

Status of vehicle automation

Few truly autonomous vehicles (SAE level 3 or higher) are on the road as of this writing, in May 2018. Google reports that their autonomous cars – currently developed under the name Waymo – have travelled over 5 million miles. In addition BMW, Nissan, Ford, General Motors, Delphi, Tesla, Mercedes Benz, and Bosch have their own projects going on (Wang, 2018). Thus far, no car company has an autonomous vehicle beyond SAE level 3. Tesla has autonomous driving capability, informing that their hardware in SAE level 5 but the software does not meet these standards – yet. They argue that the software can be updated afterwards to match the SAE 5 requirements. These autonomous cars have been found to be working in optimal conditions in places such as California or Nevada where the weather conditions are better. The limitations come when in wintertime or other similar “bad weather” conditions due to missing road markings etc. In winter conditions, an autonomous vehicle is so far forced to rely on extremely precise location service and pre-processed route information, limiting the manoeuvrability and speed of an autonomous vehicle (Tervola, 2017).

During the brief testing period on public roads, autonomous road vehicles have been involved in numerous accidents, even fatal ones. So far, four accidents that have involved fatalities have occurred, one in China and three in the United States (Horwitz & Timmons, 2016; Levin & Wong, 2018; Tesla, 2018). In some cases, the functionality of the ADS is questioned, while in some cases there is a clear indication of human operator error.

Why autonomous cars?

Yet, as an artefact, the car will change significantly in the following decade.

The reasons for developing autonomous cars are numerous and varied. Wouldn't it be nice for an individual if their car would drive your kids to school and hobbies, fetch you from the pub or be available for multitude of drivers during the day? Given properly functional sensors and software, the car will have capabilities beyond human drivers. For example, infrared vision, significantly shorter reaction time, freedom from distractions or tiredness, and more economic driving are all possible with a computer behind the wheel. To emphasise, on paper, a computer should complete the task much better than a human. But how about on the tarmac?

If a computer manages the driving task better than a human, then road safety should improve as well. Thus, the autonomous vehicles should – if done properly – reduce accidents and thereby injuries and deaths caused by human error. As the society's aim is to safekeep its' members, at a first glance it should not be a difficult choice whether to allow autonomous vehicles to the road.

Public transport could also benefit from automation. A major expense article in public transport – as in many other fields – is wages. Autonomous vehicles do not require pay or complain about long working hours, being capable of operating 24/7. Moreover, the errors made by overworked bus drivers could be mitigated by automation. Also, the private sector could benefit from the increased safety brought by autonomous vehicles, as well as the potential savings that would have a huge impact to many economies. To emphasise, the most common job in most of the U.S. states is driver (Bui, 2015). The paradigm shift brought by widespread adaption of autonomous vehicles would be comparable to the introduction of the automobile and the subsequent decline of the horse carriage driver profession.

The first phase of this major societal shift has already been in progress for a while. Cars have become increasingly automated and thus more dependent on computerised systems. Therefore, these systems must also be protected against unauthorised and unwanted access and tampering. The car therefore is not only a data storage or a system that produces data, but it also is a tool to make day-to-day life easier and safer. No one would buy a car that does not inherently promote values of safety and security. Besides the driver, the security aspects must also be considered from the point of view of passengers and other road-users as well.

Whereas a modern car is controlled by the teamwork of computer and driver, both controllers have their own separate fields of responsibility, as is illustrated by SAE automation level 0-2 definitions. Some overlapping fields, e.g. collision avoidance systems, exist in contemporary vehicles. Many new cars are equipped with systems that take control of the DDT in some predefined scenarios, such as when the car is in danger of colliding with a pedestrian. Most major car manufacturers provide SAE level 1 or 2 Advanced Driving Assistance Systems (ADASs) with collision avoidance as an option in their vehicles.

Security and automated vehicles

Whereas the driver of a car is relatively hard to hack, computer systems are more vulnerable. To protect all interest groups, hacking such critical computer systems should be made as hard as possible – as completely unhackable computer systems do not exist. Or as Eiza and Ni state about cars and computers: “if you see word ‘software’, replace it with ‘hackable’; and if you see word ‘connected’, replace it with ‘exposed’” (Eiza & Ni, 2017).

Vulnerabilities and built-in features in different cars that can be used against the driver have been around for a while. One of the clearest examples was the vulnerability found in the Jeep Cherokee, resulting in the recall 1.4 million Cherokees after researchers demonstrated they could remotely hijack the car’s system over the internet. The attack was described to be one *“..that lets hackers send commands through the Jeep’s entertainment system to its dashboard functions, steering, brakes, and transmission, all from a laptop that may be across the country.”* (Greenberg, 2015) Numerous other examples also exist; Another clear threat is tampering with the sensors of the car from the outside (Eiza & Ni, 2017).

It is important to notice that almost everything a hacker can make a computer, or a phone do, they can do to a car – and much more. The car is a computer with wheels and an engine. Therefore, it should not be hard to hack the car to mine cryptocurrencies – with the energy cost of gasoline none the less. A car could also be a target of a ransomware attack just to get the car to work or keeping the health and well-being of the people inside of it as ransom. Imagine your vehicle suddenly announcing that unless 200€ worth of bitcoin is paid to a specified address, the autonomous vehicle will deliberately crash – with the passengers still within.

We need cybersecurity where not only the information and its’ integrity nor only the communication is protected, but also the physical world and the assets that the security must protect are taken into accord. E.g. Simson Garfinkel has stated that the hackers are the real obstacle for self-driving cars (Garfinkel, 2017).

While the car has essentially become a computer (a set of networked computers, actually) with wheels and an engine, it is also approximately 2 tons of metal moving over 30 metres per second – or a bus/truck moving a tad slower but weighing 10 to 80 tons – with humans inside and moving in an environment with people in vicinity. Therefore, the possibilities for a hacker to cause harm are increased when hacking a car compared to a normal computer.

How cars can be hacked

Adding complexity to a computer system increases the attack surface available to a malicious actor. As Bruce Schneier has observed in an interview, “complexity is the worst enemy of security” (Chan, 2012). And make no mistake, cars are already complex machines. As we already discussed, being controlled by 50-70 ECUs and more than 100 million lines of code *must* leave room for errors to occur. At a high abstraction level, the attack surfaces on a modern car can be divided into *internal* and *external* (Checkoway et al., 2011). For an adversary to use the internal attack surface of a car, they must have physical access. External attack surfaces do not require physical access, but rather allow an adversary to attack the car over distance.

A well-known conventional wisdom in cyber security states that an adversary gaining physical access to the target system equals game over for the defenders, as with physical access it is possible to hack *any* device. Even tamper-proof integrated circuits and cryptoprocessors can be hacked (see, c.f., (Anderson, 2010)).

The methods for gaining access to the systems of a modern car, absent of direct physical access, can be grouped into three groups: indirect physical access, short-range wireless, and long range wireless methods (Checkoway et al., 2011). Indirect physical access methods include attacking through the ODB-II port indirectly by first compromising a computer used for diagnostics or attacking the entertainment system by a specifically crafted media files that play normally but also contain a malicious payload that exploits a vulnerability in the entertainment system and takes control of the car. Short-range wireless attack vectors include Bluetooth systems, RFID car keys, wireless tire pressure monitoring systems, or Wireless Local Area Networks, among others. Long-range wireless attack vectors include broadcast channels such as RDS radio systems and targeted radio channels such as remote telematics systems operating over cellular networks. All the attack methods discussed above can be – and have been – exploited to gain complete control of all systems of a car (Checkoway et al., 2011).

Recent examples of serious vulnerabilities using indirect physical access and wireless attack surfaces include those found in Tesla (Keen Security Lab, 2016) and BMW (Keen

Security Lab, 2018) vehicles, both discovered by the same Chinese security team. The Tesla vulnerability was one of the first practical attacks capable of taking complete control of a state-of-the-art car over wireless. Altogether 14 different vulnerabilities were found in BMW vehicles, with various levels of access required. Sis remote vulnerabilities are detailed in the report, using Bluetooth and cellular networks.

But the vulnerabilities in one car model do not directly translate to another make and model, let alone another car manufacturer, right? It would at first seem logical that vulnerabilities in car systems would be limited to a single car model or a limited subset of each model, but unfortunately this is not the case. Like the aviation industry, the automotive industry relies on standardisation to provide safety, interoperability and cost savings in the manufacturing process. This also means that the same standardized technologies, parts, controllers, and modules manufactured by automotive industry component providers are used in various car models across different manufacturers.

To give a recent example, researchers have found vulnerabilities in keyless entry systems used in cars manufactured by VW Group that affect most car models manufactured by the group between 1995 and 2016, and a vulnerability in another keyless entry system, Hitag2, affects cars of various models from ten or more different manufacturers (Garcia, Oswald, Kasper, & Pavlidès, 2016).

Even older cars that do not expose wireless attack vectors by themselves can be attacked if aftermarket entertainment or diagnostic systems are installed. Modern Android-based aftermarket entertainment systems do offer wireless connectivity over Wifi and Bluetooth, as well as access to internal car networks through the ODB-II bus. It has been shown that access to the internal network of a vehicle gives an attacker complete control over all systems of the car (Checkoway et al., 2011). This security aspect may be overlooked by many car owners who just want to upgrade their car entertainment systems. So far to our knowledge, there are no known accidents or other security issues involving these systems, but it may just be a matter of time.

Car life cycle and security

It is paramount to understand that as cars will become more computerised, they will also have more in common with other modern computerised devices such as smart phones, televisions, or other smart appliances. Like these more mundane devices, the issue of software updates must be considered also for cars. As a modern telephone, a home appliance, and now, a car, is controlled by software, the software must be kept up to date to protect them against hackers that try to exploit vulnerabilities.

Cybersecurity solutions must thus be implemented and maintained with meticulous care for the whole life-cycle of the car. A major issue in cyber security is the existence of old and even obsolete devices that are still actively used, regardless of discovered vulnerabilities that leave those devices open for exploitation. For example, the Wannacry ransomware that spread in the wild in May 2017 (BBC, 2017a) used EternalBlue, a vulnerability in the Windows operating system that was publicly exposed in April 2017 (The Shadow Brokers, 2017), to spread from device to device. After disclosure, it was immediately patched by Microsoft, even to old and deprecated versions of Windows that no longer received any updates, such as Windows XP. Multiple organizations were still heavily affected by this attack. For example, the UK National Health Service (NHS) still has hundreds of thousands of computers still running unpatched Windows XP as their operating system, and these computers were targeted by the ransomware, leading to a serious compromise of NHS systems (Clarke & Youngstein, 2017).

EternalBlue exemplifies the issue of obsolete, vulnerable devices that can be exploited with serious consequences to all stakeholders. Autonomous vehicles, due to their unique functionality and long product life cycle (not everyone can afford to drive a new car) are in danger of being used for malicious purposes if (or rather, when) vulnerabilities are discovered.

As we discussed earlier software is hackable, software providers must update their software whenever security vulnerabilities are found to ensure correct and safe operation of the system. Sometimes such vulnerabilities are patched without any public scrutiny or incidents, but we have also seen some spectacular security failures in devices. The aforementioned EternalBlue vulnerability has allegedly been used for gaining unauthorized access to computer systems long before its publication, and the swiftness of the response by Microsoft spoke volumes on how serious a threat the vulnerability was. But what happens to devices that do not receive any updates?

When a smart device reaches the end of its update cycle, it is probably still far from its actual end of product life. This leaves untold numbers of network connected devices online that are vulnerable to attacks, some of which will never be reported. An alleged CIA hacking tool, codenamed Weeping Angel, can be used to gain remote access to various Samsung Smart TV models (Wikileaks, 2017). Many Samsung devices with a vulnerable firmware version are still is active use, however, as people do not automatically replace a working device "just" because a vulnerability has been found. Such vulnerable devices can thus be found in living rooms across the globe – sometimes even creating quite the professional dilemma for a security researcher who happens to appreciate privacy, security, *and* a 65" screen.

As earlier "dumb" phones such as Nokia 3110 are "eternal" with regard to both hardware and software, they are still fully functional even 20 years later. But a smartphone bought five years ago might already have reached the end of its support life, and any subsequent vulnerabilities found in the software will compromise the security of the device and the data stored on it.

In the movies whenever “the baddies” need to get access to systems (“Hack all the cars in a five-mile radius and do it now!”) they start hacking furiously and, usually sooner than later, they get access to numerous devices all around the globe, with little to no preparation at all. In real life, though, while a similar feat *is* possible, it is extremely hard to pull off. Real life cyberattacks that compromise hundreds or thousands (or hundreds of thousands, like the Mirai botnet (Antonakakis et al., 2017) did) of devices are more of a dull affair. Such attacks are conducted over a long period of time and compromised devices are left operating normally, but with a backdoor that allows them to be taken over when necessary. This can be done by the hacker and by leaving a backdoor open to the system that the hacker can take into use or by spreading autonomously propagating malware which reports to the hacker when the backdoor is open. The latter is more dangerous, but nowadays fortunately quite rare, as operating system vendors have taken security seriously since the worst outbreaks of Sasser, MyDoom, Sobig and ILOVEYOU worms back in the early 00’s. Because the information security issues in modern cars have not been considered with similar gravity as with desktop operating systems, the potential fallout can be even worse should the similar lax attitude also extend to autonomous vehicles.

If a malicious hacker would only target, say, Bugatti or Aston Martin cars, the motive would probably be grand theft auto. Should the attacker want to hack as many cars as possible, they probably would target cars from a large manufacturer, thus giving them more attack surface. As we know, car manufacturers use similar parts and software in different models. For example a Jaguar X400 is, in fact, only “[a] little more than re-shelled Ford Mondeo” (Adams, 2011).

Therefore, if one finds a security vulnerability from one type of car, it is very likely that a similar flaw can be found from most of the cars produced by the manufacturer. Moreover the vulnerability can be very wide-spread because the largest car manufacturers have huge market shares. For example, the VW Group (Volkswagen, Audi, Skoda, Porsche, etc.) has a market share of around 23 % of all cars sold in Europe (ACEA, 2018). If one could infect just these cars and from the last 5 years, they would still have access to more than 10 % of all the cars on European roads. This gives the attacker a lot of possibilities, for example to create a large network of bots for DDoS-attacks, a network of computers for crypto-currency mining (not very effective computers, but loads of them!), or to use these to more sinister means discussed later in this paper.

Cars have become increasingly complex with the integration of computer systems. The computers are getting more operation responsibility – or all of it. Will the cars of tomorrow have sufficient software updates to stay secure and safe to use? Will they have a predetermined life cycle limited by the end of technical support lifetime? Will there be museum cars in X years? Similar concerns surface also in the case of manufacturer bankruptcy. Will a car suddenly become dangerous to its passenger should the software update cycle end, whether due to planned obsolescence or bankruptcy?

Current ethical questions and what should we focus on?

One of the problems with ethics and autonomous vehicles are the questions it raises. The current ethical research question around the autonomous cars is mostly focused on how the vehicle should behave in traffic (i.e. who or what should the vehicle collide with when a collision is otherwise unavoidable). There have been numerous theoretical and some practical (sic!) research settings where these questions have been pondered (see e.g. (Bonnefon, Shariff, & Rahwan, 2015; Hevelke & Nida-Rümelin, 2015; Lin, 2016)).

The *trolley problem*, introduced by Foot (1967), is a thought experiment where a runaway trolley is strolling down the railway tracks. Ahead, five people are tied up on the tracks, about to be ran over and killed by the trolley. You are standing by a lever, and by pulling it you can direct the trolley to a different track, where only one person is tied up. You have the option to pull the lever, and as a result, instead on five, only one person will die. The trolley problem is all about analysing what people answer to this. The problem has been rephrased in countless ways both in literature and on the internet, but all are analogous to the original problem.

Autonomous devices can face similar decisions to those illustrated by the trolley problem, and the decisions they make must be programmed by someone. Should an autonomous vehicle, in a situation where a collision is unavoidable, rather collide with a pregnant woman with a stroller, two drunk adult men, or a concrete wall (which would kill the passengers)? The discussion on responsibility, whether ethical or legal, is already intensely discussed in literature (see e.g. (Lin, 2016; MIT, 2018)).

As these questions – simultaneously with other instances of the trolley problem – are indeed important, they are just a part of the larger issue: the ethical and societal implications that this change in automotive culture will bring. First of all the issue with security is hugely more important in the discussion than the mere "shall the car run over a nurse or a priest" –discourse. Yet, if the security-issue is not brought up correctly and diminished to a single trolley problem, the automotive industry will bring their automated automobiles to the road before the discourse has finished and the requirements for safety are decided thus making the decisions and requirements harder to make.

To emphasise: the biggest ethical issue around autonomous road vehicles are security issues, as it is clearly ethically wrong to make unsafe cars. This issue is clearly not discussed enough.

We should also ponder issues of responsibility. Whereas it is obvious that when a driver makes an error, the driver is the responsible one. But when discussing about an error made by autonomous vehicle the reason for error can be e.g. faulty (or dirty) sensors, bad programming, poor quality electronics, misconducted repairs, etc. To find the responsible one from the multitude of possibilities becomes more of a task. The responsible one should none the less be found that the society can diminish the possibilities to error and therefore possibilities to serious injuries and death. Therefore there should some sort of auditing for the autonomous vehicles that they should pass and consequences for those that submit poor quality software or hardware to those vehicles.

The society will have tremendous changes with autonomous vehicles if the problems of security and functionality are solved – for better and for worse. Whereas now the parents drive their kids to play football – and therefore stand there watching the games and practices, they use time with their kids and share experiences with them. If the cars would drop the kids to their hobbies, there would not be a dozen parents participating to their kids' hobbies but just a dozen cars waiting to drive the children home. The time spent in the car while the parents drive their kids to different places is an opportunity for the family members to discuss and share experiences which strengthens their bond.

Moreover the parents could go and have a pint after the work because they have no need to be sober as driving under influence is illegal but being a passenger under influence is not. This will not support the family structure but may strengthen the bonds between workmates or friends – while bringing a possibility for increased alcoholism. As autonomous cars enable people to sleep while moving to or from work, they are more relaxed in both home and work with increased capacity to function. The cars could also be used more as there are fewer limitations (such as intoxication, being underaged etc.) to travel with the car. As the teenagers tend to move by more environment-friendly ways to school, friends', and hobbies, in future they might just take the family car – and thus use more of the planets' resources.

These all are just quick examples on how the technology can change the society, some better and some worse. We should be aware of the possibilities for both and do work on societal level to counter the bad consequences – not necessarily by limiting the automation level on traffic but with other means, e.g. promoting the values of standing next to the kids' football game!

One of the possibilities to implement automated automobiles and to counter social problem is to slow the rate of the technology taken into use. Patrignani and Whitehouse (2013) propose "slow tech" which define as following:

> *"Slow tech is a new way of looking at technology. It means designing and developing technologies that are 'slow', with the aim of being good, clean, and fair. It has, as an aspiration, the design and use of a new kind of information and communication technologies (ICT): ICT that is human-centred, and that takes into account both the limits of the planet and those of human beings."* (Patrignani & Whitehouse, 2013)

They argue that as the technology advances rapidly and changes the everyday life in increased speed and to more increased speed the people cannot keep up with it. Therefore the technology should be designed to improve human well-being and well-living, to be fair, clean, and good, not only more effective. The idea of slow tech also contains environmental values which should be taken in accord with the autonomous road vehicles, and as discussed before, the easiness of use with the autonomous vehicles may bring forth the question of “necessary trips”. According to Patrignani and Whitehouse there should be some re-thinking about the pace of the development of technology to promote the values in the society. (Patrignani & Whitehouse, 2013)

An inevitable crisis?

As discussed earlier in this paper, a prominent threat scenario is the mass hacking of autonomous road vehicles. As a single car is easily used for a terror attack – a sad story repeating itself in Europe and Canada – what damage could the attacker do with hundreds or thousands of cars?

One of the worst scenarios is using automated cars as a terror or a military first-strike weapon (Lucas, 2017). In this scenario a nation or an organisation hacks a massive amount of vehicles from another area, triggers them simultaneously to hit pedestrians, other cars, trains, bridges, or other targets causing massive amount of death and injuries simultaneously while crippling the infrastructure similar to heavy bombardment. A nation targeted by such an attack is most likely in chaos for a long time.

There have also been claims that US intelligence services (mainly CIA) uses or have used the vulnerabilities in modern cars for assassination purposes, but they are yet just claims and should be treated as such (Overly, 2017). These methods – as shown above – could be used to such described assassinations, however.

As the parameters for the safety of autonomous vehicles are a complex task to formulate, they are not the focus of this paper. Moreover we argue that we still should have those similarly as we require car manufacturers to comply with emission- or road safety standards. We should require that the cars are manufactured in such way that they do not pose hazard in new areas of road safety (such as terror attacks) while they became increasingly safe. And should this seem to be an impossible task for a car manufacturer, we still have the cars of today which (mainly) fulfil these standards.

But is the development of autonomous vehicles inevitable? We would argue so; The potential benefits of truly autonomous transportation – increased efficiency, safety, security, usability, ecological factors – are a strong driver. We also argue that instead of comparing cars to cell phones, a more apt parallel to the development of autonomous vehicles can be found in the aviation industry. This comparison between industries should be made, as the aviation industry are even more paranoid on safety issues. For example, the engine failure accident that resulted in the death of a passenger aboard a Southwest airlines flight in April 2018 (Stack & Stevens, 2018) was the first civilian aviation fatality in the US in 9 years (NTSB, 2018). Should autonomous cars reach similar safety levels to aviation, the number of road fatalities would plummet drastically.

By applying safety design best practices learned from aerospace engineering, while considering the new issues brought by computerization, safe autonomous road transportation may indeed be possible. The problem we see in this domain is not more computer-controlled systems in itself; the issue is with system connectivity and increased complexity. These issues must be addressed should we want to continue on the path eventually leading to truly autonomous vehicles.

Conclusions

We should examine the values we have embedded to automobiles. As we yearn for efficiency and ease-of-use, we should also remember the safety and security for the driver, passengers and other people alike. Whereas automated cars can react faster and do not suffer from distractions, tiredness, deceases, or intoxication, they can be hacked.

As shown earlier, the autonomous automobile is a complex set of computers working together, and the words “complex” and “computer” increase the possibility of exploitable vulnerabilities. Moreover, the reasons of hacking autonomous cars are numerous, ranging from financial gain to acts of war, and there are many who stand to benefit when the overall security levels of vehicles, whether contemporary or of the future, are low.

Ethical discourse around the subject has been somewhat limited around the trolley problem –method and therefore we call for proper discourse on both the ethical issues as well as societal challenges around the autonomous vehicles, as we argue that the biggest ethical issue around autonomous automobiles is the lack of security.

We urge the members of society to prepare themselves to yet another change in the way our society works; to diminish the possible problems such as unemployment, increased travelling by car and diminished time spent with family members, to counter the negative effects, as well as to prepare themselves to the improvements that the autonomous vehicles bring to their lives such as increased safety, time available, and possibilities to travel.

We, as a society, should also demand that these mobile computers are sufficiently secure, so that the risk for violent deeds such as terror attacks, military strikes, and assassinations is as minimal as possible. Autonomous vehicles are perhaps inevitable, but we can slow the pace of progress down by demanding that these marvels of technology follow the values of road travel: safety, efficiency, and ecology.

References

ACEA. (2018). Consolidated Registrations - by Manufacturer. Retrieved May 23, 2018, from http://www.acea.be/statistics/article/consolidated-registrations-by-manufacturer

Adams, K. (2011, November 27). The cars : Jaguar X-Type development story. *AROnline*. Retrieved from https://www.aronline.co.uk/cars/jaguar/x-type-jaguar/the-cars-jaguar-x-type/

Amazon. (2018). Amazon Prime Air. Retrieved from https://www.amazon.com/Amazon-Prime-Air/b?node=8037720011

Anderson, R. J. (2010). *Security Engineering: A guide to building dependable distributed systems*. Wiley.

Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., … others. (2017). Understanding the mirai botnet. In *USENIX Security Symposium*.

BBC. (2017a, May 13). Cyber-attack: Europol says it was unprecedented in scale. *BBC News*. Retrieved from http://www.bbc.com/news/world-europe-39907965

BBC. (2017b, October 24). “Bad Rabbit” ransomware strikes Ukraine and Russia. *BBC News*. Retrieved from http://www.bbc.com/news/technology-41740768

Bonnefon, J.-F., Shariff, A., & Rahwan, I. (2015). Autonomous vehicles need experimental ethics: Are we ready for utilitarian cars? *ArXiv Preprint ArXiv:1510.03346*.

Bort, J. (2014, January 16). For The First Time, Hackers Have Used A Refrigerator To Attack Businesses. *Business Insider*. Retrieved from http://www.businessinsider.com/hackers-use-a-refridgerator-to-attack-businesses-2014-1?r=US&IR=T&IR=T

Bui, Q. (2015, February 5). Map: The Most Common* Job In Every State. *NPR*. Retrieved from https://www.npr.org/sections/money/2015/02/05/382664837/map-the-most-common-job-in-every-state

Chan, C.-S. (2012, December 17). Complexity the Worst Enemy of Security. *Computerworld Hong Kong*. Retrieved from https://www.computerworld.com/s/article/9234815/Complexity_the_worst_enemy_of_security

Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., Savage, S., … Kohno, T. (2011). Comprehensive Experimental Analyses of Automotive Attack Surfaces. In *USENIX Security*.

Clarke, R., & Youngstein, T. (2017). Cyberattack on Britain’s National Health Service—A Wake-up Call for Modern Medicine. *New England Journal of Medicine*, *377*(5), 409–411.

Constantin, L. (2016, May 13). Petya ransomware is now double the trouble. *Network World*. Retrieved from https://www.networkworld.com/article/3069990/petya-ransomware-is-now-double-the-trouble.html

Eiza, M. H., & Ni, Q. (2017). Driving with sharks: Rethinking connected vehicles with vehicle cybersecurity. *IEEE Vehicular Technology Magazine, 12*(2), 45–51.

Foot, P. (1967). The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review, 5,* 5–15.

Gaertner, C. (2015). Trends and Lookout of the Automotive Software Industries. *Proceedings of the 6th International Conference on Software Business ICSOB 2015*. Springer.

Garcia, F. D., Oswald, D., Kasper, T., & Pavlidès, P. (2016). Lock It and Still Lose It—On the (In)Security of Automotive Remote Keyless Entry Systems. In *USENIX Security* (pp. 929–944).

Garfinkel, S. (2017, August 22). Hackers Are the Real Obstacle for Self-Driving Vehicles. *MIT Technology Review*. Retrieved from https://www.technologyreview.com/s/608618/hackers-are-the-real-obstacle-for-self-driving-vehicles/

Greenberg, A. (2015, July 21). Hackers Remotely Kill a Jeep on the Highway-With Me in It. *Wired*. Retrieved from https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/

Hawkins, A. J. (2017, October 11). Autonomous cars without human drivers will be allowed on California roads starting next year. *The Verge*. Retrieved from https://www.theverge.com/2017/10/11/16458850/self-driving-car-california-dmv-regulations

Hevelke, A., & Nida-Rümelin, J. (2015). Responsibility for crashes of autonomous vehicles: an ethical analysis. *Science and Engineering Ethics, 21*(3), 619–630.

Horwitz, J., & Timmons, H. (2016, September 20). There are some scary similarities between Tesla's deadly crashes linked to Autopilot. *Quartz*. Retrieved from https://qz.com/783009/the-scary-similarities-between-teslas-tsla-deadly-autopilot-crashes/

Keen Security Lab. (2016). Car Hacking Research: Remote Attack Tesla Motors. Retrieved May 25, 2018, from https://keenlab.tencent.com/en/2016/09/19/Keen-Security-Lab-of-Tencent-Car-Hacking-Research-Remote-Attack-to-Tesla-Cars/

Keen Security Lab. (2018). Experimental Security Assessment of BMW Cars: A Summary Report. Retrieved May 25, 2018, from https://keenlab.tencent.com/en/Experimental_Security_Assessment_of_BMW_Cars_by_KeenLab.pdf

Levin, S., & Harris, M. (2017, November 17). The road ahead: self-driving cars on the brink of a revolution in California. *The Guardian*. Retrieved from https://www.theguardian.com/technology/2017/mar/17/self-driving-cars-california-regulation-google-uber-tesla

Levin, S., & Wong, J. C. (2018, March 19). Self-driving Uber kills Arizona woman in first fatal crash involving pedestrian. *The Guardian2*. Retrieved from https://www.theguardian.com/technology/2018/mar/19/uber-self-driving-car-kills-woman-arizona-tempe

Lin, P. (2016). Why ethics matters for autonomous cars. In M. Maurer, J. C. Gerdes, B. Lenz, & H. Winner (Eds.), *Autonomous Driving* (pp. 69–85). Springer.

Lopatin, E., & Bulavas, V. (2017). Miners on the Rise. Retrieved May 16, 2018, from https://securelist.com/miners-on-the-rise/81706/

Lowensohn, J. (2015, March 17). Elon Musk: cars you can drive will eventually be outlawed. *The Verge*. Retrieved from https://www.theverge.com/transportation/2015/3/17/8232187/elon-musk-human-drivers-are-dangerous

Lucas, J. (2017, November 20). Terrorist hackers ‘could kill millions’ by remotely hacking cars, expert warns. *The Sun*. Retrieved from https://www.thesun.co.uk/motors/4951071/cars-hackers-cyber-security-threat/

Magnuson, S. (2017, March 21). Driverless Trucks Poised to Join Military Operations. *National Defense Magazine*. Retrieved from http://www.nationaldefensemagazine.org/articles/2017/3/21/driverless-trucks-poised-to-join-military-operations

MIT. (2018). Moral Machine. Retrieved May 23, 2018, from http://moralmachine.mit.edu/

NTSB. (2018). Accidents Involving Passenger Fatalities: U. S. Airlines (Part 121) 1982 - Present. Retrieved from https://www.ntsb.gov/investigations/data/Pages/paxfatal.aspx

Overly, S. (2017, March 8). What we know about car hacking, the CIA and those WikiLeaks claims. *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/innovations/wp/2017/03/08/what-we-know-about-car-hacking-the-cia-and-those-wikileaks-claims/?noredirect=on&utm_term=.c66cb11e06f7

Patrignani, N., & Whitehouse, D. (2013). Slow Tech: Towards Good, Clean and Fair ICT. In *ETHICOMP 2013 Proceedings* (pp. 384–390).

Rushe, D. (2017, October 10). End of the road. Will automation put an end to the American trucker? *The Guardian*. Retrieved from https://www.theguardian.com/technology/2017/oct/10/american-trucker-automation-jobs

SAE International. (2016). *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*.

Stack, L., & Stevens, M. (2018, April 17). Southwest Airlines Engine Explodes in Flight, Killing a Passenger. *The New York Times*. Retrieved from https://www.nytimes.com/2018/04/17/us/southwest-airlines-explosion.html

Tervola, J. (2017, December 15). VTT:n robottiauto teki epävirallisen autonomisen auton nopeusennätyksen – selviytyy lumipeitteisellä tiellä (Video). *Tekniikka & Talous*. Retrieved from https://www.tekniikkatalous.fi/tekniikka/autot/vtt-n-robottiauto-teki-epavirallisen-autonomisen-auton-nopeusennatyksen-selviytyy-lumipeitteisella-tiella-video-6692518

Tesla. (2018). An Update on Last Week's Accident. Retrieved May 21, 2018, from https://www.tesla.com/blog/update-last-week's-accident

The Shadow Brokers. (2017). "Lost in Translation" leak. Retrieved from https://github.com/misterch0c/shadowbroker

Wang, B. (2018, March 25). Uber' self-driving system was still 400 times worse Waymo in 2018 on key distance intervention metric. *Next Big Future*. Retrieved from https://www.nextbigfuture.com/2018/03/uber-self-driving-system-was-still-400-times-worse-waymo-in-2018-on-key-distance-intervention-metric.html

Wikileaks. (2017). Vault 7: CIA Hacking Tools Revealed. Retrieved May 22, 2018, from https://wikileaks.org/ciav7p1/