Lesson_10_Autonomous_vehicles

Autonomous Vehicles - Fabio Fossa

Autonomous driving technology: some basic elements

Before discussing the ethical aspects of a technology, we must first have an idea of how it works and what it is about. The autonomous vehicles that are currently in testing phase look very much like normal cars even if they have some additional sensors.
These might be:

- **Lydars** and **radars** to detect *obstacles* and *objects* to build a map of the surrounding area of the vehicle and perform **obstacle avoindance**
- **GPS antennas** to register the position of the vehicle on a map
- **Cameras** to perfom machine vision

Among the changes that autonomous driving would imply, there is also the design of cars, as we would have *new available spaces* (since we could ideally remove the steering wheel and the pedals) that could be used.

https://www.youtube.com/watch?v=B8R148hFxPw&ab_channel=Waymo (Waymo demonstration of automous driving)

The **degree of automation** in autonomous driving is very important (wow.) and it has various degrees. In order to build vehicles that are able to drive in every situation and weather conditions, we need to start from the easy functions and embed all the others progressively.

The term *autonomous vehicles* is used to describe very different objects and technologies.

Levels of automation

We can distinguish **6 separate levels** of automation and the vehicles that are currently on the market comply to the level 1 description, maybe some to level 2.

These levels are introduced in a paper from the Society of Automated Engeneers (SAE) J3016 "*Levels of Driving Automation".*

Level 0: no automation

Level 5: fully automated vehicle, that is able to drive itself under **every possible** circumstance. Only here there won't be a driver anymore and we will be just passengers (I'm the passenger and i ride and i ride...).

This means that all other levels from 1 to 4, the driver still **has to mantain** some control on the car and he/she will look more and more like a passenger as the we level up, but *never entirely.* The result will be a mix of the two roles.

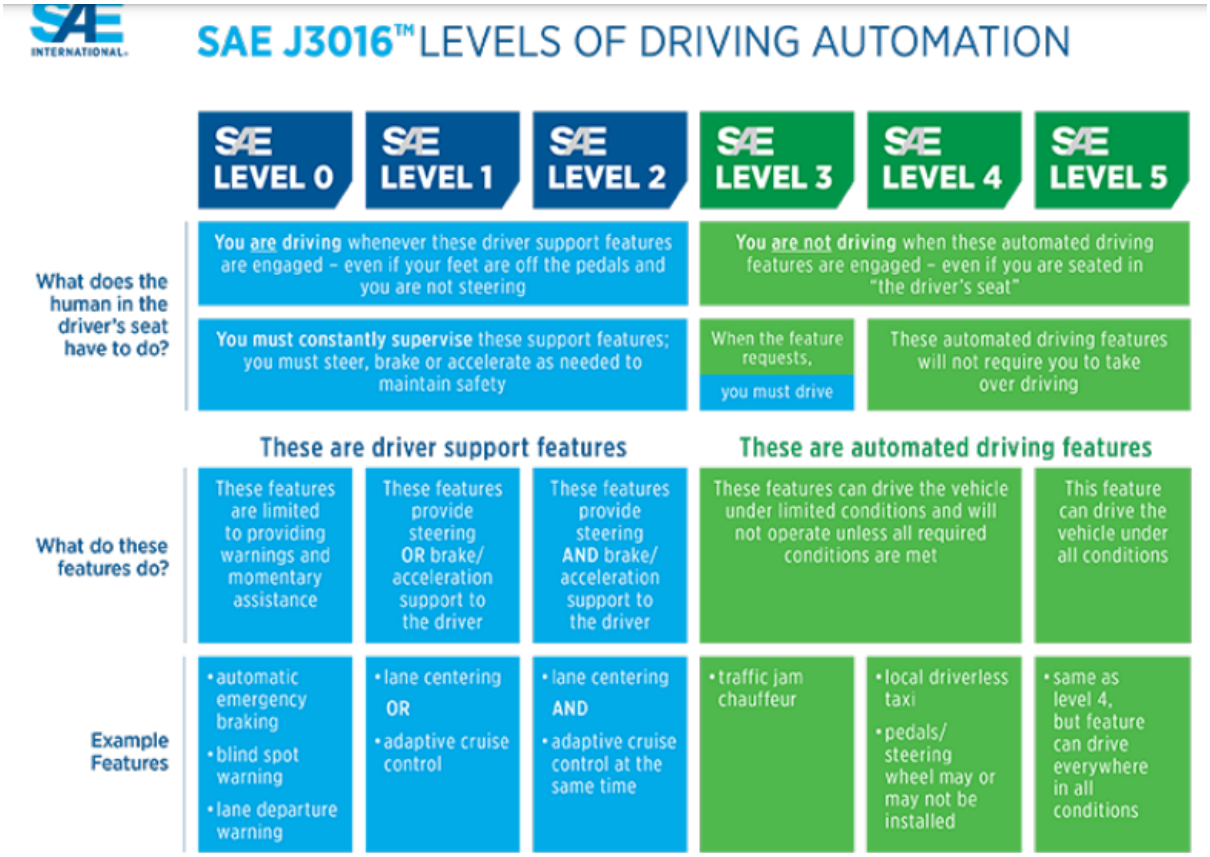

We can make a further distinction between the levels, separating them in two macro groups:

- *Human Supported Driving (0-2):* the passenger sitting in the driving seat is still the driver in all senses.
- *Automated Driving (3-5):* the driver looks more like a passenger, but the problem is to understand to what extent.

We can see that the more the automation increases, the more we are able to automate the driving functions. SHOCKING.

The difference between level 1 and level 2 is the **integration** of the automated features.

In level 3 there are some *very specific situation* where the driver can seat back and let the vehicle drive for itself. For example, if you are stuck in a traffic jam, the car might be able to handle the situation by itself even though the driver *has* to get the control back if its necessary.

In level 4, the car is functioning inside an **operational designed domain** (i.e, road condition and type, weather situation, time of the day, etc..) and is able to handle all driving

tasks without the help of a human driver as long as is inside that domain, it can even perform certain safety manouvers. The problems rise up in the **transitions** between these operational domains, where the car is no longer able to drive correctly and the control is given back to the human driver.

This is an important node from the ethical point of view, as there is the need to gather a lot of data on the psychology of the passengers like their **reaction time** (how long they need to take back the control of the car).

The driver is taken out (dead.) only in level 5. But this might be only a **technological goal**, because we might not want to reach this point from a **moral** or **ethical standpoint**.
This shows us the intimate connection between social and ethical problems and technology.

If we consider as a product of technological research, social interest and ethical concerns, we understand immediatly that *Level 5 autonomy* **is not the true goal** we want to accomplish.

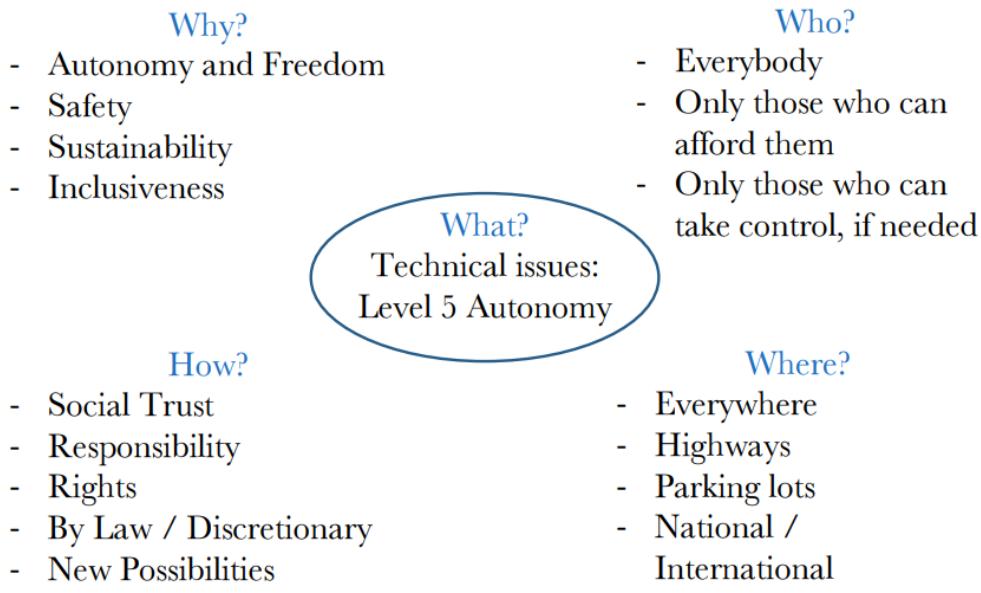

Why should we want Autonomous Vehicles?

- **Autonomy and freedom**: we would be able to use the time we spend driving to do other activities.
- **Safety**: from a more *social oriented prospective* is **dominant**. If designed correctly, we would be able to remove human errors out of driving, which are the most significant cause of accidents.
- **Sustainability and Inclusiveness :** these are different points of view that might be relevant depending on the narrative.

Who would drive this cars? (wait what?)

- **Everybody:** from the ***inclusiveness*** point of view, we might want to develop these vehicles to provide them to those who are **unable** to drive their cars like blind or physically incapacitated people. This reconnects to the *autonomy and freedom* point of the list above: these people would gain more **independency** and **freedom of movemen**t entailed by the autonomous vehicles of level 5, that would be **necessary** in this kind of scenario.

How?

Let's say that we want to ensure **responsability**: in situation where we have to deal with the moral side of driving, we want to be humans to be in charge of taking decisions. If people cannot make their own decisions, they might not trust the technology, hence, they won't buy it. This would prevent the other benefits, e.g. safety, to happen.

→ We want people to **feel in charge** and therefore we want to involve them in the moral decisions that may arise. In this situation, *level 5 autonomy* would be no longer an objective as we would have mixed controls both automatic and human.

This is an example of how ethical concerns can impact the technological ideas.

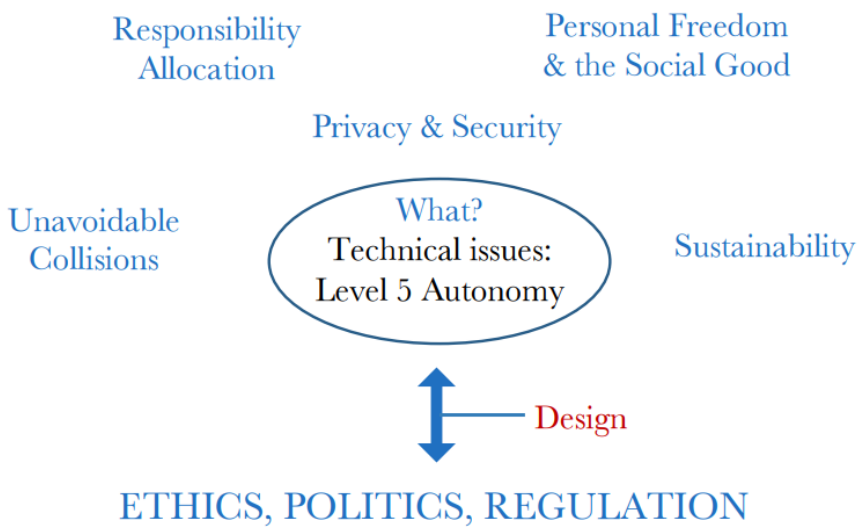

Ethical and social problems influence the technical problem of building autonomous cars. The mean to obtain alignment between this two worlds is **Regulation,** while **Design** is the venue in which values can be implemented into our technological products. So, an effort should be made to give clear and easily **applicable guidelines** to align technology design with ethical and social desires.

This is the direction that the European Union is moving towards; at the moment we are at the level of *Soft Law:* guidelines and recommendations are given to stakeholders so that can align to certain values deemed important.

Ethics of Connected and Automated Vehicles is the most important document they have right now and its aim is to enstablish a baseline for future policies.

Regulation & Policy

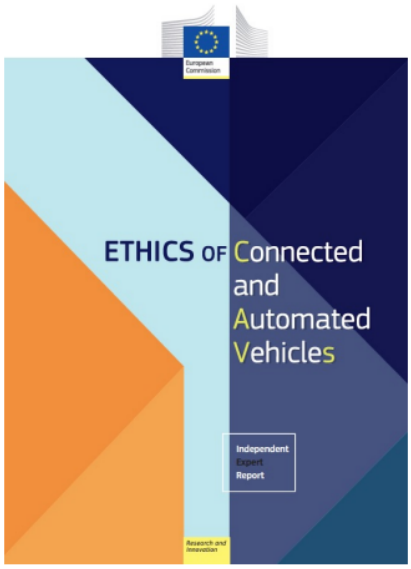

- Released on Sept. 17th, 2020
- Authored by an Independent Expert Group – 14 members, mostly academic (philosophy, law, engineering)
- Establishes a baseline for future European policy on connected and automated vehicles
- Ethical & social issues are widely accounted for

The expert group of this document presents an ethical framework based on 8 principles (*responsability, human dignity, non-maleficients, beneficients and so on)* from which are drawn 20 recommendations, that everyone, from stakeholders to users, should adhere to and, if this were to happen, we would able to build technology that would acceptable from the ethical point of view.

Problems

Unavoidable Collisions

Unavoidable Collisions

- Unavoidable collisions have been a primary worry in the ethical debate on autonomous driving
- Q: how should the system handle morally laden situations – i.e., situations where harm is unavoidable *but* can be distributed in different ways? ☞ Accident-algorithms

Classic example of unavoidable collision:
An autonomous car cannot brake while approaching a cross-road where people are crossing the street; we have to program the machine to do something about this, but what? Should it run over the pedestrians or harm the people inside the car?

This topic is relevant because we have the possibility to program up to a point what to do ***systematically*** what we believe to be the right decision, meanwhile nowadays is the driver that ***istinctually*** decides on the moment.

So we implement the *Accident-algorithms* that have to deal with all the issues reported in the slide above.

This technology brings in the domain of ethics a situation that in regular driving is not entirely there to begin with.

Unavoidable Collision

- Many issues:
  - Which value/ethical theory to implement?
  - How could we do that?
  - Who gets to decide?
  - How should this choice be made?
  - What about personal autonomy?
  - What about the rights of bystanders?
- Many problems:
  - Sometimes a bit detached from vehicle dynamics
  - Sometimes a bit detached from available/foreseeable tech
  - Sometimes a bit detached from realistic scenarios

☞ Still the problem *remains*!

Issues (nothing to add to the first in the slide):

1. How can we translate the **ethical theory** into computer language?
2. Who gets to decide how it is implemented in the vehicle? Regulators, Developers, Users, Producers, ...?
3. Should this choice be done individually or democratically?
4. What about **personal autonomy**? Should we, drivers/passengers, have a say on what the vehicle decides to do? It is builded that way by someone else, who decided which values were important...
5. **Right of Bystanders**: can we be sure that their rights will be respected?

There are many theoretical and methodological debates, even a debate on the debate itself (*metadebate*): does it make sense to think so much about unavoidable collision? Are these situations real like?

Privacy & Security

- For autonomous vehicles to function properly, a huge quantity of data must be collected, shared, and stored

  ☞ privacy issues:

- Definitions of privacy and sensible data as involved in AD
- Privacy protection throughout the entire infrastructure
- Informed consent (…)

- Autonomous vehicles pose risks proper of both usual vehicles *and* information systems

  ☞ double challenge:

- "Mechanical" vehicle safety standards
- Digital infrastructure liabilities: external attacks (security), software issues (robustness), data thefts and leaks (…)

What are the data needed by autonomous vehicles that are *sensitive*? Which are the data that need to be protected? How privacy is *envolved* in autonomous driving (AD)?

The entire infrastructure reported in the slide has to be intend as the whole *Internet of Things* technology since the car is connected to a whole bunch of devices.

We must always ask for consent when gathering data. How can we do it in a **significant/meaningful way** when we are using cameras, that gather information of **all the road users?**

Cybersecurity

This is an important topic because it is possible to *confuse and tamper* (alterare/sabotare) the sensors of the cars, like changing road signs or gps reads.

Privacy & Security

Bogus Satellite Nav Signals Send Autonomous Cars Off the Road

At the Black Hat security conference, a researcher demonstrated how making tweaks to navigation signals could send a self-driving car careening off the road.

Hacking street signs with stickers could confuse self-driving cars

Subtle or camouflaged optical hacks can change a stop sign into something else.

Researcher Hacks Self-driving Car Sensors

$60 lidar spoofing device generates fake cars, pedestrians and walls

Researchers Fool Autonomous Vehicle Systems with Phantom Images

There are many threats cybersecutity-wise that lead to new safety problems related to the attacks to the information systems of the autonomous cars.

Responsibility Allocation

- Who is to be held responsible for harm caused by accidents where autonomous vehicles are involved?
  - NOT the systems themselves, then:
  - passengers?
  - owners?
  - designers/developers?
  - producers?
  - nobody, just insurance system?

- Meaningful Human Control Approach:
  - Autonomous vehicles must be *designed* and *deployed* in a way that assures a satisfying exercise of human moral responsibility + a clear and fair distribution of legal liability

! Huge impact on Level 5 Automation !

The main question to be asked is as long as we have shared control between the car and the driver is: Who is to be held **responsible** for harm caused by accidents where autonomous vehicles are involved?

It does make little sense to hold the system responsible...

According to some scholars, there should be some mechanisms and also technological solutions that allow to assess responsability and share control in a very clear way.

Control and responsability are **linked**, of course, so we should understand which functions can be delegated to the system and which ones should be remain under human control. In this way, we would distribute responsability fairly.

As we said, this would change the concept of Level 5 automation as it would integrate human control with the automated one when it is meaningful for that to happen.

Then again, responsabilty allocation is important to allow the public to trust the new technology.

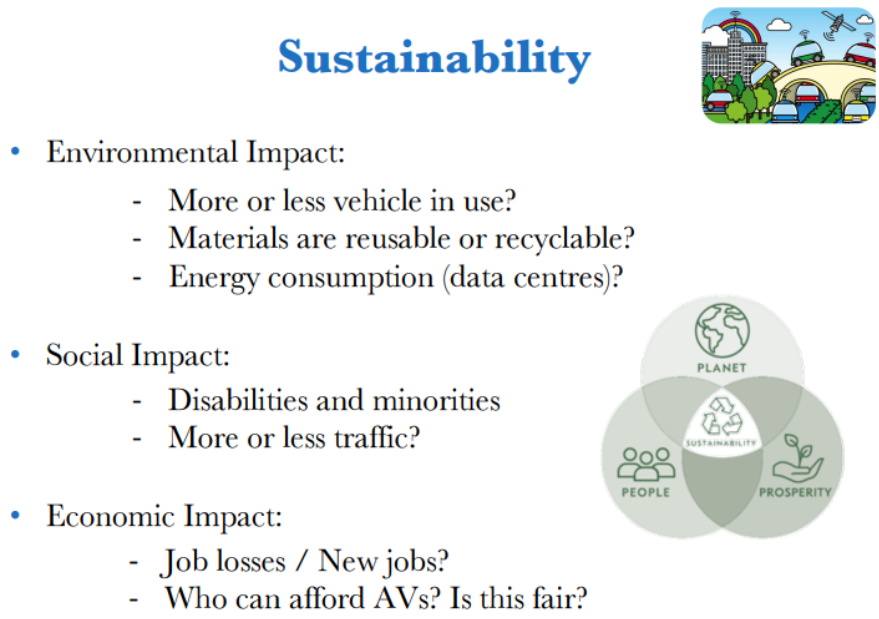

Sustainability is a composed notion of at least **three elements**:

- **The enivronment**: are the technology introduced environmentally friendly?

  There is a shared belief that we will be able to **reduce pollution** through autonomous vehicles because we would be able to handle traffic in a more efficient way. But if this will actually happen depends on many factors and some scholars believe that there would be some kind of **rebounds** (e.g, the number of vehicles in use could increase because more people would be able to drive). Also the **energy consumption** of data centres could be an issue due to the high amount of data that need to be stored.

- **The social impact of a technology**: does it introduce a desireable social impact?

  It is unclear whether AV would reduce traffic or not, but we can be confident on the advantage of including people with disabilities into the dimension of mobility.

- **The economic impact of a technology**: is it economically manageable?

  Again, the impact of new AI technology on the job market is hard to assess before hand.
  If AV were to become a reality, they should be available not only for the richest part of

the community but to everybody as it would be unfair that only rich people could have access to their benefits (uffa cattivi, la disparità sociale è brutta.).

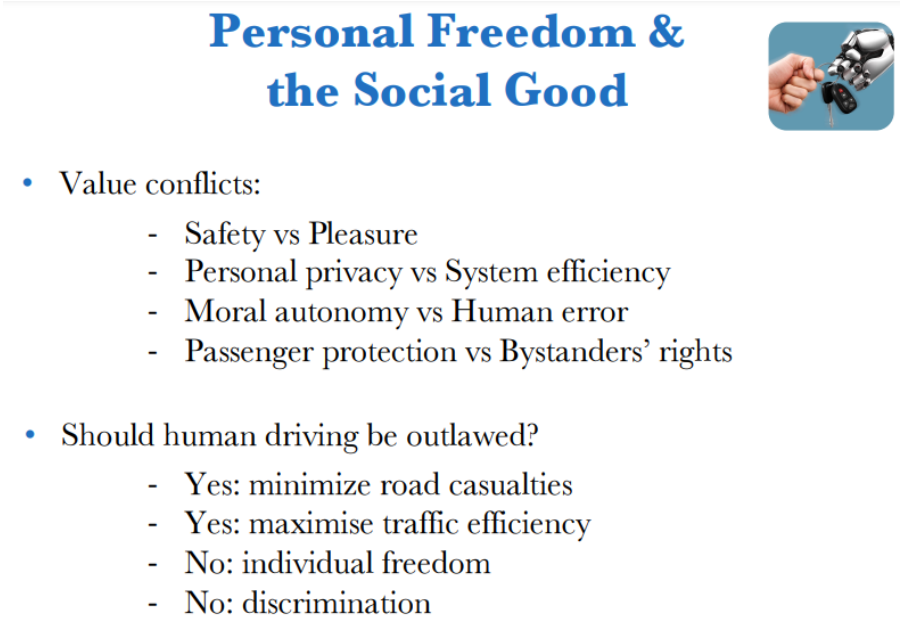

- **Safety vs Pleasure**: is it right to take away the pleasure of driving and its sense of freedom/autonomy to have safer roads?
- **Personal Privacy vs System efficiency:** imagine that to handle traffic in the most efficient way possible, you must share all the information concerning your movements and an authority is aware of them and might infer where you will be in the future, would it be okay?
- **Moral autonomy vs Human error:** here we face again the trade-off between the personal value of being in charge of moral choices and the social value of reducing the number of victims from a system prespective.
- **Passenger protection vs Bystanders' rights:** similar trade-off to the previous one that should be evaluated.

Bonus Sartor Question

Why can't we have both worlds, human and machine driven, cohexisting at the same time?

From a social point of view, we would have more accidents in a mixed scenario because if there are more human drivers there is also higher chances of human errors.

Sartor replies:" Okay, but who cares if people die, it's worth to have the freedom to choose how you want to move around"