

AUTONOMOUS VEHICLES

Abstract

Very soon, cars are likely to be fully autonomous, and driving themselves. Although the technical challenges involved in making this possible are immense, there is another side that is often overlooked. Any time machines like cars become more autonomous, there are ethical and moral implications when something goes wrong. It is not only the technical that is the domain of the engineer. The engineer also has a responsibility to delve into these potential ethical problems, and determine a satisfactory solution.

Robotics and Control Systems: Autonomous Vehicles

What's the problem?

Cars are getting smarter every year. Soon, cars will be driving themselves completely without human intervention. Although the technology that drives this progress is all very exciting, there are important questions that are almost always overlooked.

- ♣ Do we even want cars to drive themselves?
- ♣ What happens when something goes wrong? Who is responsible?
- ♣ Who will regulate this new breed of vehicle? Does it become a legal problem?

Some of these I will take for granted right now. I will assume that we do want autonomous vehicles, and that they are on their way. I will also ignore regulation and legality. Although I believe this is a very important subject, and one that should be addressed soon, it is a mess that I do not want to get lost in.

In my opinion, the biggest problems involve what happens when something goes wrong. Think about the following scenario:

While driving down a main road, a car begins to back out of a driveway twenty yards ahead of your vehicle. You slam the breaks, but can't stop in time. A car in the lane to your left prevents you from swerving out of the way.

In this situation there are effectively two options. You could hit the car backing out, or swerve and hit the car to your left. A human driver does not have the time to think about this and reacts instinctively. However, an autonomous vehicle has no instinct, and has to make a choice. What should it do? Is there a 'correct' decision?

Who cares?

Everyone who has ever been in a car should care about this problem. Self-driving cars are coming, and some are already here. Google has developed technology that allows their vehicles to largely drive themselves, and have already been testing in some parts of the country. Tesla's newest vehicles will be able to come pick you up (on private property), *with nobody in the car*. We cannot all afford to drive a Tesla, but features like adaptive cruise control, lane keeping, and park assist are coming to more and more cars. These are all steps toward full autonomy. Like it or not, this is something that everyone will have to deal with, and probably very soon.

What technology exists?

Existing systems generally try to imitate humans. This makes sense; engineers can only teach machines what the engineers already know. Usually these systems can be viewed very simply, as shown by Figure 1.




Figure 1

Simplified Systems Diagram. Source: XXXX (YEAR).

Humans use this simple system constantly. Our 'sensors' are sight, hearing, smell, taste, and touch. Each of these delivers a signal to our brains, which does the 'processing.' After coming to a decision, we move, speak, cough, blink, or do any of the other actions people are capable of performing.

For cars, a lot of advanced technology is required to make this all possible. Systems to control acceleration and deceleration have existed for years (like in cruise control), and controlling steering is no harder. The 'brains' need to figure out where to go (likely using GPS), and make decisions based on that information, like which street to turn on. This is a problem that has been explored in the field of artificial intelligence for a long time, and is usually implemented as a decision tree.

Figure 2 is an example of a simplified decision tree, specific to the problem of autonomous vehicles. The circles represent sensors, GPS and radar are shown, but there will be many more, including cameras, laser ranging, and accelerometers. The diamonds are the processing stages, where the computer analyzes the data from the sensors and makes a decision. Finally, the squares represent the possible decisions. Like

the sensor stage, the decision stage includes dozens of possibilities, which are not all listed. In Figure 2, the green items represent a single level of the tree. The blue items represent part of the second level (omitting the possible decisions). As the tree grows downward in time, it also grows wider, representing the increased variability further in the future. Note that the second level has three redundant sets of sensors and processing, one for each possible decision from the first stage. The computers in autonomous cars will need to be quite powerful, and try to predict many levels ahead in the tree. Over time, bad choices can be removed, and new ones added based on what the sensors indicate.

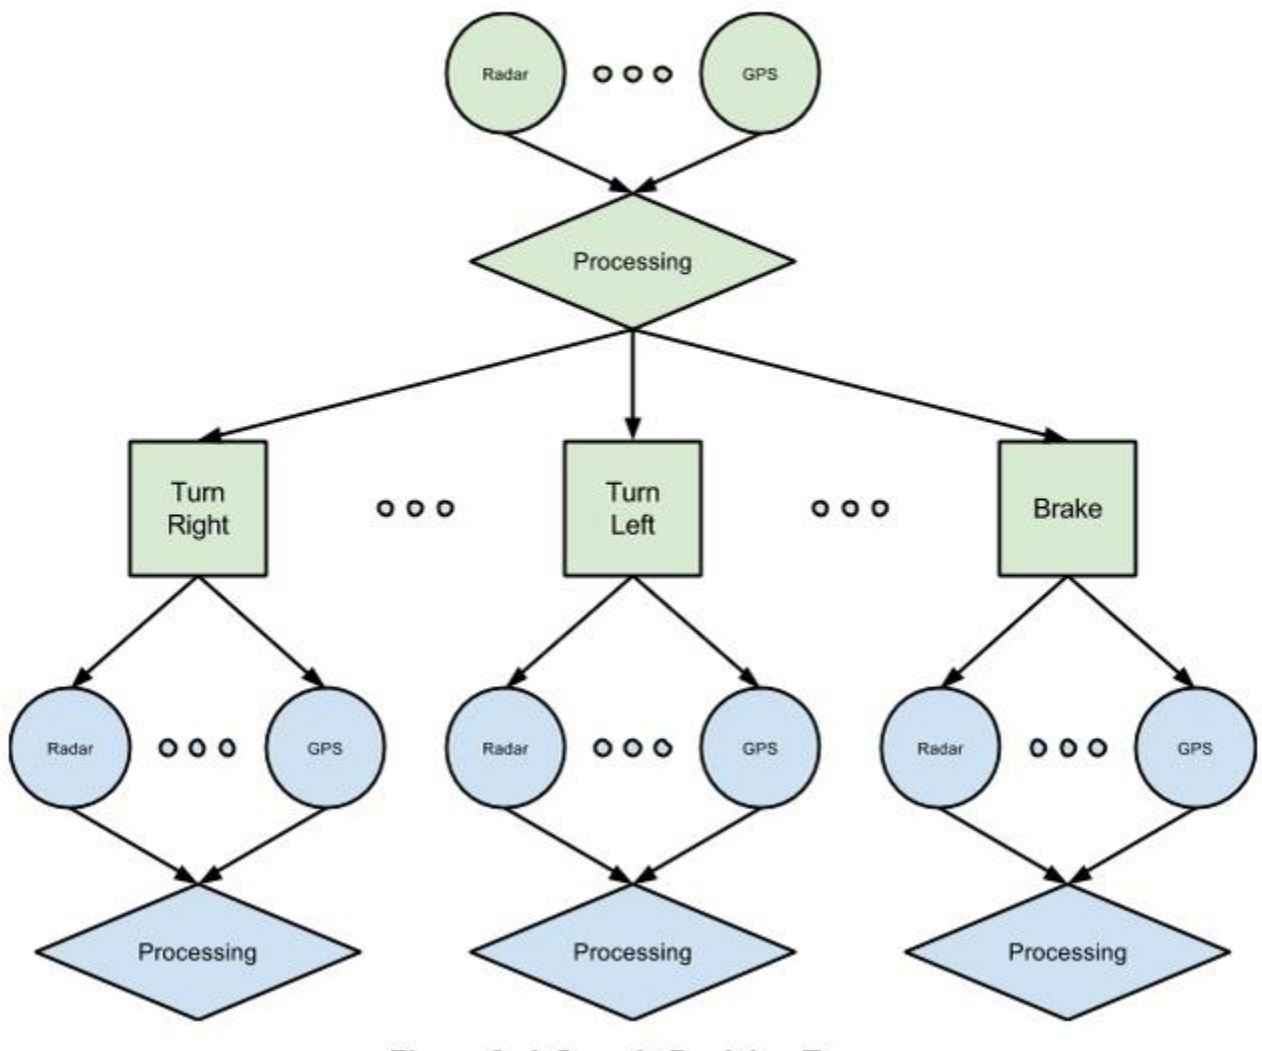


Figure 2

A Sample Decision Tree. Source: XXXX (YEAR).

Some of the biggest new technologies involve the sensors that feed information to that computer ‘brain.’ A combination of radar, lidar and sonar can provide a lot of information about surroundings, and cameras with advanced image recognition also help to allow autonomous vehicles to be aware of the immediate vicinity. These are constantly getting more accurate, and size reductions have made it possible to mount them on cars.

Why don't existing solutions work for cars?

The artificial intelligence aspects of the system are where the hard decisions must be made. In the past, most intelligent systems have dealt with hard (or impossible) choices in one of three ways.

Some of the earliest uses of artificial intelligence were to play computer games. Usually when there was no good move to make, or several moves were equally good, the computer could choose randomly. The worst possible outcome was that the game was lost.

Today, intelligent robotic systems are much more common. For example, many kinds of manufacturing assembly lines are robotic, and controlled by computers. When something goes wrong (perhaps a part is broken, a robot malfunctions, or the power goes out) typically these systems can simply shut down, and wait for a technician to fix whatever went wrong. The worst possible outcome in this case is some lost revenue due to the stopped time, or possibly some broken parts.

One of the best parallels to our autonomous cars is the autopilot system that helps fly planes. If there is some kind of emergency during flight, the system alerts a human pilot to take over until a correction can be made. Here the worst case scenario has a much more significant impact, as the lives of the passengers and crew can be at stake.

However, I claim none of these solutions will work for cars. Prompting a human to take over cannot work, because in most emergency situations there would not be time for a human to recognize the need to assist, much less actually do something to help. Stopping and waiting for the problem to be fixed does not work for similar reasons. In a moving vehicle, shutting down the system could actually lead to a more dangerous position. Imagine being brought to a full stop in the middle of a busy highway. Finally, choosing a random option is also not acceptable. These choices will not just lose a game, but could involve thousands of dollars in damage and injury to one or more people. Some people may find this acceptable, but I know that I would not like my health to be left up to chance.

Why are engineers the key to solving these problems?

In philosophy, there is a classic thought problem, which goes something like this:

An out-of-control train is barreling down the tracks toward an unsuspecting group of workers. You stand next to a switch that could divert the train down another path, saving the workers. However, on the other track stands your grandmother, who would be hit instead. Do you pull the switch? What if the person is not your grandmother? What if the person is someone you hate instead, or a serial killer?

This question has no correct answer. A reasonable person could make an argument either way. This is similar to the first example, with a car backing out and a car in the next lane. Again, a reasonable argument could be made for either. Does it matter if one is an SUV? Does it matter if one has a child in the back seat?

When making these autonomous vehicles, these questions are no longer philosophical. Real people, most likely the engineers creating the cars, will have to make real choices that could have implications of life or death. That is truly the job of an engineer – to bring math and physics and philosophy into the real world, and use them as tools to solve problems. Clearly this does not only involve math and science. It is the responsibility of any engineer to evaluate the impact of what they are creating. It is the responsibility of any engineer to determine risks involved with their project, and their responsibility to alleviate them. In the case that a risk cannot be suitably limited, it is then the responsibility of that engineer to raise their concerns, and not build blindly or ignore the risk.

Engineers are faced with real problems. Inevitably there will be risks, sometimes very dangerous ones. Every engineer must take the time to evaluate these ethical and moral problems; they are as much a part of engineering as math, and science, and problem solving.

Source : <https://sites.tufts.edu/eesenior/designhandbook/2015/autonomous-vehicles/>