



Autonomous Driving on the Thin Trail of Great Opportunities and Dangerous Trust

Sandro Mund and Patrick Glauner

Abstract

Achieving fully autonomous driving cars is a considerable technological milestone that will have significant impact on many lives and the adaption of new technologies. The question of when this milestone will be achieved is currently being debated and contradictory forecasts are increasingly being made. In this chapter, the most important components of self-driving cars are presented and different approaches are discussed. We show what makes autonomous driving so challenging and what misjudgments have been made in the past. In particular, the role of artificial intelligence will be illuminated to give a clear picture of what progress is realistic in the coming years. Next, we discuss related challenges that need to be solved in the coming years. Based on our own research, we will also show how hard it is to interpret models, like neural networks, i.e., understanding why they make the decisions they make in the context of self-driving.

1 Introduction

The conditions and degrees of autonomy under which a vehicle is described as self-propelled differ greatly. A commonly used classification for this are six levels (Society of Automotive Engineers 2014) depicted in Fig. 1.

Starting at zero, i.e., a normal car, a standard vehicle that already contains systems that can intervene briefly in the control system, for example, to prevent the wheels from locking when braking. The first level includes assistance systems

S. Mund (✉)
Trier University of Applied Sciences, Trier, Germany

P. Glauner
Deggendorf Institute of Technology, Deggendorf, Germany
e-mail: patrick@glauner.info



Fig. 1 Six levels of autonomous driving. Source: authors

that are already commercially successful but require the driver to be constantly and attentively observed. Examples are automatic speed control, lane keeping, and parking. Vehicles reaching the second level can drive most of the time on their own but require attention from the driver to catch mistakes. The following level extends this intervention by a time window and only from the fourth level the control of the driver is no longer necessary. Level four, however, is limited to the fact that this safety can only be guaranteed under certain circumstances or on certain routes. Only with level five a fully autonomous vehicle will be spoken of. The jump to the last level is the most challenging one, because all imaginable traffic situations have to be dealt with.

The social impact of the jump to level 5 is also the greatest. For example, it has been argued that every 23 s someone dies in traffic (World Health Organization 2019). In the future, many lives could be saved by this new technology. Cheap autonomous taxis will threaten jobs; however, the shared use of vehicles will in turn reduce environmental pollution. Many potential consequences are possible and whether they are positive or negative is often debatable. Autonomous driving not only affects society as a whole, but driving itself is also a social act. People give each other hand signals and disregard traffic rules to react to extraordinary situations. Equally important is interaction with passengers for vehicles that cannot drive fully autonomously to ensure that controls are not neglected. Many questions from different disciplines are therefore important to make statements about the acceptance and use of self-driving cars. In order to answer the question of whether a technical implementation is possible at all, there is the particularly important area known as artificial intelligence (AI).

This chapter is structured as follows: In Sect. 2, we discuss a number of challenges that autonomous cars need to master in order to understand their surrounding environment. We argue in Sect. 3 why AI is key to doing so. Meanwhile, in Sect. 4, we first review the history of autonomous driving and then discuss the state of the art as well as a number of predictions that have been made for the foreseeable future. In Sect. 5, we discuss how easy it has become in recent years for a large number of people to acquire the knowledge of how to build and use the complex technologies needed for building autonomous cars. Next, we look at interpretability of machine learning models in the context of autonomous cars in Sect. 6. We also present some of our research results on this topic in the framework of convolutional neural networks. Last, we summarize this chapter in Sect. 7.

2 Understanding the Environment

The architecture of a system for autonomous control of a vehicle is complex. Various components such as sensors, powerful hardware for computations, or the control of the vehicle bus must communicate with each other in real time and ensure reliability. Contradictory signals have to be handled and risks are not always avoidable. Among the most important tasks of such systems are perception, localization, planning, and control. By combining the various sensors, an overall representation of the environment of a vehicle can be obtained that is required for following steps. An intermediate step here is to bring the differently coded information into a uniform shape in order to have a consistent representation of the outside world. In the next step, the system uses this overall information to localize the vehicle. This means to include the position of the vehicle in this model of the outside world. Based on this, the effects of different control signals over a certain distance are calculated and the necessary control signals for this selection are determined.

Different sensors with their own strengths and weaknesses are used in the process. Regular cameras are cheap and have a good range but cannot measure distances. Light detection and ranging (LIDAR) systems fire millions of laser steels per second and measure how long it takes for them to jump back (Cracknell 2007). With this information an accurate 3D map can be created. The costs of a LIDAR are very high and the sensor is not robust enough to be used on a large scale for commercial vehicles. However, since this sensor is preferred for use, a lot of money is currently being invested to solve these problems. In contrast, radars use radio waves to generate images of the environment. Comparing both approaches, LIDARs are more accurate but radar sensors are much cheaper and not prone to fog, rain, or snow.

As the environment becomes more complex, the implementation becomes more challenging. If only machines were involved in road traffic, no AI would have to attempt to predict human behavior. Because of the enormous complexity, the movement in public road traffic presents research and development with a number of challenges. A further dimension in distinguishing autonomous systems is their robustness, e.g., whether they need to be intervened immediately, within defined time windows or not at all by humans. A key threshold here is the achievement of safe driving without human control, so-called fully autonomous driving.

3 The Critical Role of Artificial Intelligence

Systems for controlling autonomous vehicles consist of a multitude of components with different tasks. For example, recognizing the road, predicting the actions of other drivers, or planning the way through the next curve. These components work with various types of information, which can be provided by sensors. The information obtained in this way must be processed, evaluated, and merged in order to enable independent driving. AI is used at various points in such complex systems

and in particular machine learning, a branch of the AI that creates models from data. A particularly known method is deep learning (LeCun et al. 2015), which are neural networks that consist of multiple intermediate layers and are therefore called “deep.” For an autonomous vehicle to be able to safely participate in road traffic, many sub-tasks with different requirements have to be mastered in order to ensure the safety of people has priority. Some aspects of deep learning lack a well-founded theory (Lin et al. 2017) and it is generally challenging to verify whether a statistical model works well under all circumstances. Unlike with model-based development, where system correctness can be proven mathematically, there is no absolute certainty. Redundancy in the execution of sub-tasks and their control by testable systems helps but leads to new challenges. It is practically not possible to define abstract rules that cover all potential situations. Data-based models are more scalable because they become better with more data and can thus address situations that can hardly be described by rules.

The increase in performance (i.e., accuracy) comes with the loss of interpretability. A combination of methods of machine learning with knowledge-based systems could solve this problem, but how to do so is a contemporary research challenge. Generally, the fundamental problems of AI have not been solved yet. Human thinking is hardly understood and there is no method that is promising to simulate this intelligence. Artificial general intelligence (AGI) does not appear to be achievable in the foreseeable future without an unexpectedly large breakthrough (Shanahan 2015). Some cases in traffic require abstract thinking to understand complex situations. A person can understand when a stop sign is painted over, stolen, mirrored, or just printed on a T-shirt. Whether such a level of intelligence is necessary to control a car safely enough or whether it is possible to collect sufficient data is questionable. With the focus on methods of machine learning for autonomous driving, it is therefore crucial to have good data that contains rare special cases.

Controlling a car automatically in standard situations is a relatively simple task today. Small amounts of data are sufficient to teach a model how to stay on track and avoid objects. Mastering the remaining fraction of cases is much more challenging as it was assumed several times in the past, though. The question of when fully autonomous vehicles are roadworthy can be answered by determining when enough of these situations can be considered so that self-driving cars are statistically safer than humans. Reaching this threshold of safety is tried feverishly. The commercial success of these vehicles in turn depends on many more factors. There are laws, insurance and production costs, or the acceptance and trust, as well as the social change to mobility-as-a-service, just to name a few.

4 Ambitious Goals and Their Consequences

Today, machine learning algorithms and especially neural networks are a key component for self-driving cars. But this was not always the case. In this section, we review advances in autonomous driving, AI, and contemporary R&D challenges.

4.1 Advances in Autonomous Driving and Artificial Intelligence

Experiments with autonomous vehicles have existed since the beginning of the twentieth century. As early as 1939, General Motors sponsored radio-controlled electric cars powered by electromagnetic fields, generated by circuits embedded in the roadway. Already at that time there were optimistic estimates to have completely autonomous cars in a few decades. It was around this time that a small number of scientists from various disciplines began to discuss how artificial brains could be created, which led to the founding of the field of AI research (McCarthy et al. 1955).

At that time, one was just as optimistic to achieve good results quickly. That optimism, however, was to be paid off. In the 1970s, it became clear that many problems were much more challenging than expected and the high expectations could not be fulfilled, so that no further funding was provided. The time from 1974 to 1980 is often referred to as the first so-called AI winter (Russell and Norvig 2009). However, within this time further research was done and some progress was made. Neural networks at that time were known as perceptrons and only consisted of one unit (Minsky and Papert 1969). It was later mathematically shown that this model's learning capabilities are severely limited (Blum and Rivest 1989). The discovery of how parameters of a multi-layered perceptron (i.e., a neural network) can be trained changed the field substantially (Rumelhart et al. 1985). The broad consequence of following improvements in this learning process was later to be called deep learning and brought the networks back more attention in research (Hinton et al. 2006). Before neural networks regained importance again, however, experts systems therefore were very dominant. But again, the expectations were too high, which led to the second AI winter, which lasted from 1987 to 1993. Within this time in 1989, Carnegie Mellon University had pioneered the use of neural networks to steer autonomous vehicles forming the basis of contemporary control strategies (Pomerleau 1989).

The big start for the development of autonomous vehicles did not happen until a decade later. The second DARPA Grand Challenge was launched in 2005. To win the prize money of 1 million US dollars, more than 200 miles had to be driven autonomously. While in the previous year not a single car completed the course, this time five teams made it (Thrun et al. 2006). Again, there were optimistic voices that autonomous driving would be possible soon and big car manufacturers like BMW, Volkswagen, Audi, and many others started with their own experiments. Google also began in 2009 to secretly work on driving its own cars.

The same year, ImageNet (Deng et al. 2009), a very large and freely available database of more than 14 million labeled images, was launched. The availability of ImageNet has simplified access to data, increased the accessibility to train models with deep learning, and encouraged further research. More and more public libraries for machine learning algorithms appeared and were optimized for calculations on end-user graphics cards, making hardware much cheaper. NVIDIA is a leading supplier of hardware optimized for machine learning. In 2016, they demonstrated in a paper how a car can be controlled by a neural network in order to promote new

products for autonomous driving (Bojarski et al. 2016). Their approach was not new but inspired further projects.

In the same year, it also became known that an AI named AlphaGo had defeated one of the world's best professional Go players (Borowiec 2016). Due to the complexity of the board game, it was assumed that this would only be possible in a couple of decades later. This breakthrough strongly supported the hype of machine learning that continues to this day. AI is finding more and more economic applications and many advances in research have not yet arrived in the wider economy. Whether there will be a new AI winter is questionable, the continuous results speak against it. Further breakthroughs are not unlikely, especially due to the large investments of the automotive industry. But every success seems to be followed by ever greater expectations.

4.2 Contemporary Forecasts and Challenges

Billions are currently being invested and the entire automotive industry is taking AI seriously. Tesla wanted to launch a fully autonomous car in 2019 (Siddiqui 2019). Nissan, Honda, Toyota, and Hyundai have made announcements for 2020 and Volvo, BMW, and Ford-Chrysler for 2021 (Connected Automated Driving Europe 2019). These deadlines appear to be unrealistic, though. Other scientists doubt that it is possible in the near future or at all. Elon Musk, the CEO of Tesla, on the other hand predicts that it will be unusual soon to produce cars that are not fully autonomous. In an AI podcast at MIT, Musk said that Tesla has a big lead (Fridman 2019a). A message on Twitter that 1 billion miles were driven in autopilot mode supports this picture. Tesla is able to collect the most data due to a large fleet of sold cars, which are already equipped with cameras (Fridman 2019b). Nevertheless, Musk has previously often postponed his forecasts and many experts are critical of his statements.

The first company that did 10 million miles of autonomous driving is Waymo, which continues the Google Car project (Ohnsman 2018). Their self-propelled car program was initially led by the winner of the 2005 DARPA Challenge (Thrun et al. 2006). What is particularly interesting by comparing the two companies is that their approaches are entirely different. While Waymo uses LIDAR in combination with maps, which is by far the prevailing opinion, Tesla has a strong focus on cameras in combination with computer vision and deep learning. Musk said at an event for investors that LIDAR had no future for autonomous driving and every company focusing on it was doomed. Their approach is also viewed with a lot of skepticism because the camera lacks depth information and is sensitive to bad weather conditions (Templeton 2019).

There is currently no vehicle suitable for the mass market that allows autonomous driving without constant control of the human driver and important questions have not been answered. There is no algorithm that can understand complex traffic situations and many sensors are too expensive or unreliable. In the past, false promises have often been made, but unexpected breakthroughs have also occurred.

To have fully autonomous vehicles in a few years is a very optimistic estimation, though.

5 The Challenge of Easy Access to Complex Technologies

Easy access to new technologies is basically a good thing for a research discipline. More investments lead to more research and economic applications. This process is self-reinforcing and particularly pronounced in the field of deep learning. The flattening out of this hype does not seem predictable yet, which is why the entry into this technology is becoming easier and easier. Meanwhile, there is a multitude of online training courses and public libraries that make it possible to use complex systems within a short period of time without understanding them. This is not a problem in itself, but it can be a source of problems. What makes matters worse is that it is not obvious when gross mistakes are made. The deceptive certainty of an apparently good result can then lead to further damage. In the long run, it can also be just as harmful for a company not to use new technologies that could work really well with a little more specialist knowledge. It also makes sense to have in-house experts to check external results. This also protects against paying huge amounts of money for projects that have already been better solved with freely available software. Even large and established companies assign new employees to customer projects after only a short time. It is therefore risky not to thoroughly check the results of suppliers, but this often occurs in the AI area because the necessary knowledge is lacking.

Equally critical is the one-sided cooperation with universities. One-sided means in this context that students do not have any technical contact persons in the company. It is then dangerous that there is a discrepancy between practice and theory in many areas of machine learning. Data sets provided for teaching are usually unrealistically clean, so, for example, there are few statistical outliers, errors, and missing data points. In addition, the literature hardly deals with these practical problems, which are decisive for project success. In software development, requirements management (Nuseibeh and Easterbrook 2000) is regarded as particularly decisive, while many AI projects are still carried out in a relatively unstructured manner.

The area of autonomous driving gives the impression that these problems are not affected by the high entry barriers. But how realistic is this impression? In fact, there is an entire market for development in this area that makes it easy for small teams to get started. Hardware and software are available in different price ranges. Suppliers have entire kits in their assortment, e.g., (NVIDIA Corporation 2019). This is sensible because relatively little data is available for public research, even if this situation continues to improve. Therefore, it is necessary to conduct experiments under real conditions to develop robust models. However, it becomes problematic when investors are presented with supposed research projects that basically only consist of hardware and software that were bought, installed, and configured. Just as critical are unrealistic expectations and misjudgments.

The entry into the development of a competitive entire autonomous car should not be the goal of a small research team. Ensuring the road capability of a vehicle takes on dimensions that cannot be achieved without large teams that have a lot of expert knowledge and budget. In addition, cutting-edge research is not public and the greatest breakthroughs are therefore made in industry. In turn, companies can see which public research results are published. How far these two worlds are apart is often shown by the attempt to reproduce the results. Often only a small part of research results works under real conditions. Ironically, there are also publications that confirm these findings (Ioannidis 2005). Scientists are under great pressure to publish successful results and mistakes hardly have any consequences. Blind trust is also inappropriate for scientific publications. The evaluation of the conferences in which a paper has appeared can be an indicator for non-experts, but the evaluation of a subject expert is preferable. General authority should not have to be trusted blindly. In order to make sure that quality can also be examined in topics such as deep learning, the following section presents a number of corresponding guidelines.

6 Interpreting Deep Learning Models in Self-Driving Cars

Autonomous cars come with great risks to health and safety of drivers and pedestrians. Therefore, it is important to look into the underlying models and why they make the decisions they make based on input from the environment. Neural networks are often referred to as “black boxes,” meaning that is hardly understood why they make specific decision. However, approaches exist to gain a better insight into how they work (Samek et al. 2017). In particular, in the field of image processing there are several methods available (Erhan et al. 2009). In this section, we will analyze this topic in greater detail for autonomous driving.

6.1 Convolutional Neural Networks for End-to-End Driving

Convolutional neural networks (CNN) (Krizhevsky et al. 2012) have been regarded as one of the best methods in the field of image processing for years. CNNs extend neural networks by several layers of filters that learn from the input a hierarchy of increasingly more complex features. Their approach is in some sort inspired by how the human vision system works. The processing of camera images is an important part of autonomous driving. Therefore, CNNs have also gained popularity in this area, because success can be achieved quickly, such as letting a vehicle drive autonomously over a test track. Using only image processing is not suitable for road traffic. For example, the sirens of an ambulance also need to be recognized. Pure CNNs also lack the ability to include previous events in forecasts, i.e., each camera image is viewed individually. For industrial applications, they are therefore only suitable for partial tasks.

Nevertheless, research in this direction is being carried out. In so-called end-to-end driving (Bojarski et al. 2017), a single model controls a vehicle, such as a CNN




Fig. 2 Sinusoidal trajectory of an autonomous car. Source: authors

that has been trained with historical data. A person drives the vehicle first and the control signals are stored. These signals are the labels, i.e., solutions to the states of the outside world that have been reacted to. The technically simplest way to record these states is to use cameras. The final data set would then consist of a large number of images with the associated control signals. If the problem of driving is broken down to the fact that an action of the driver is to be assigned to every state of the world, it is a classic problem that can be solved using image processing. The final model gets the frames of a camera while driving and returns a control signal for each image. This approach works surprisingly well in simple environments. Simple here means that the weather conditions remain constant and few shadows and reflections influence the forecasts. However, there is a variance in how the car can be controlled. Since for each frame a slightly different control signal is returned, the steering wheel tends to tremble very strongly and the car drives sinusoidally, like being drunk as depicted in Fig. 2. This driving style does not give a feeling of safety. And this kind of safety would not be appropriate either. The steering movements of neural networks are ultimately only statistical statements about similar situations in the training data.

6.2 Visualizing What Deep Learning Models Learn

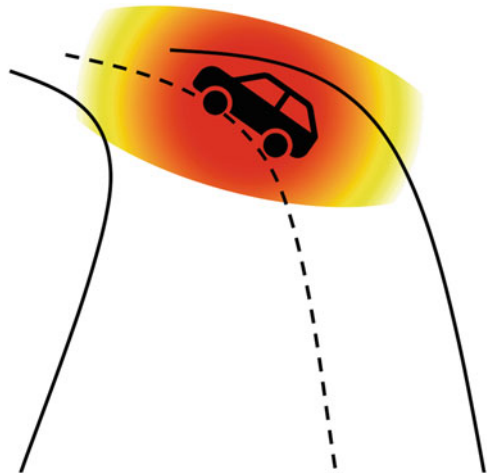
So why is it useful to deal with CNNs in the context of end-to-end driving? The properties of neural networks can be exploited here to understand how they learn from the data. What makes driving so interesting for image processing is that some elements in the images only change the output in a specific context. Most images show cars, but how a model drives the car is more influenced by where these cars are. For example, a wrong-way driver should have a significantly different effect on a control signal than a vehicle printed on a billboard. Only the context gives the picture elements a meaning that is broken down to a steering signal. In order to understand this particular context and how different models react to it, it is helpful to adapt the data sets so that better models can be trained. A further interesting aspect is the label, i.e., the direction in which the data is directed. In this way, it is possible to determine how certain sections of the input influence the overall result. In other words: Which pixels of a camera image lead to left or right steering? Areas of an image can have multiple clusters, with different influences on the overall result. How a model deals with contradictions or brings them into context is another promising research approach, also for models in other areas of image processing

or deep learning. In safety-critical systems it is particularly important to be able to make statements about how reliable a system is. Visualizing what a model sees is therefore important not only to develop better models faster, but also to understand and test them. Autonomous driving requires a lot of data to deal with as many special situations as possible. Finding blind spots in this data can be accomplished by visual analysis.

Visualizations are well understandable for humans and enable plausibility checks for complex models. For instance, a car drives particularly well on a test track. Also light and weather changes do not influence the performance. Is the model safe now? When visualizing the most important regions of the camera images by coloring using so-called heat maps, it turns out that the distinctive shape of the surrounding trees is the most important influencing factor. A simple approach to create such visualizations is to systematically modify inputs to observe their effects, such as setting the color values of pixels in an area to zero. Through many of these repetitions with different areas, each pixel can then be assigned a relative relevance to the overall image. The resulting heat maps are also called occlusion maps (OM) (Zeiler and Fergus 2014).

One question that we investigated previously was how meaningful OMs are for road traffic situations. Cutting out image areas creates new side effects, since a black pixel also matters to the model. A person would also react if there was a black, indefinable square on the road. One approach to neutralize these unwanted side effects in our research was to invert the OMs (Mund et al. 2018). This inversion is done by removing the areas that were not hidden during the creation of the maps and the other way round. The area remains the same in both cases and only whether the pixels inside or outside this window are masked changes. The two resulting maps can be combined by multiplication. The resulting map can then be used again to color areas on images according to their relevance. We sketch the outcome of the method in Fig. 3.

Fig. 3 Occlusion map sketch. Source: authors



It depicts a camera view towards the front of the car shortly before turning left. The generated occlusion map highlights the relevance of pixels for automatically deciding how to steer the car. The color should be interpreted as follows: The redder the region, the greater the importance of the corresponding pixels. The depicted occlusion map shows that the underlying neural network makes reasonable decisions as it mainly considers to region around the street turn ahead of it (where a different car is currently located) for steering left. Readers are advised to read our corresponding paper (Mund et al. 2018) for more detailed real-world visualizations.

It is particularly interesting to use this visualization for videos to highlight in real time what is currently important. It was noticeable that even with only few training data, the important regions of the images are cluster-like, even if they did not make sense at the beginning. However, the analyses are empirical and it is therefore challenging to make general statements. Especially in neural networks, observations often depend on hyperparameters, such as the network structure, or data as general insights into the learning process. We therefore repeated these visualizations and different models that were trained with different data for multiple times. With the addition of more training data, clusters started to cover more parts of the road and became more traceable. However, the visualizations became more unstable in our experiments. This means that per frame very different areas of the camera images became relevant for the model. This showed that although the learned features are sometimes logically comprehensible, the model could not generalize. A human being pays attention to different things during driving, but does not evaluate their relevance very differently many times in a second. Both the evaluation of the model predicts on the test set and the inspection of the visualizations gave a wrong impression. Once again it turned out that the previously chosen metrics were not sufficient. The critical testing of quality criteria remains a crucial and creative process that makes an important contribution.

7 Conclusions

In this chapter, we first introduced the different stages of autonomous driving and the criteria by which they can be distinguished. We then presented the most common sensors and what the most important tasks are for self-propelled cars. Different technologies also have different strengths, but a number of questions of feasibility are still open. We then highlighted the relevance of artificial intelligence and discussed its limitations and why large data sets are important for success. Next, we reviewed the relationship between autonomous vehicles and AI in a historical context. The roots of modern approaches like deep learning are older than often assumed and have a past with heights and depths. We showed that overly high expectations already led twice to so-called AI winters and how the current hype about artificial intelligence was triggered, as well as which optimistic forecasts were made up to date. We then discussed why it is important to build up expertise to maintain controls that exist for other technologies already. Finally, our own research has shown how the decisions of neural networks can be visualized in order to better

understand supposedly safe systems. The plausibility of complex systems can be visualized with simple means.

The future of self-driving cars depends on significant breakthroughs that are not yet predictable. It remains to be seen whether the enormous effort and expense involved will make this success possible.

References

- Blum, A., & Rivest, R. L. (1989). Training a 3-node neural network is NP-complete. In *Advances in neural information processing systems* (pp. 494–501). <http://papers.nips.cc/paper/125-training-a-3-node-neural-network-is-np-complete.pdf>
- Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., et al. (2016). End to end learning for self-driving cars. arXiv:1604.07316.
- Bojarski, M., Yeres, P., Choromanska, A., Choromanski, K., Firner, B., Jackel, L., et al. (2017). Explaining how a deep neural network trained with end-to-end learning steers a car. arXiv:1704.07911.
- Borowiec, S. (2016). Alphago seals 4-1 victory over go grandmaster lee sedol. Retrieved August 1, 2018, from <http://www.theguardian.com/technology/2016/mar/15/googles-alphago-seals-4-1-victory-over-grandmaster-lee-sedol>.
- Connected Automated Driving Europe. (2019). How do automakers perform with their self-driving car timeline? Retrieved December 15, 2019, from <http://connectedautomateddriving.eu/mediaroom/how-do-automakers-perform-with-their-self-driving-car-timeline/>
- Cracknell, A. P. (2007). *Introduction to remote sensing*. Boca Raton: CRC Press.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 248–255). Piscataway: IEEE.
- Erhan, D., Bengio, Y., Courville, A., & Vincent, P. (2009). Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3), 1.
- Fridman, L. (2019a). Elon musk: Neuralink, AI, autopilot, and the pale blue dot. Retrieved December 15, 2019, from <http://lexfridman.com/elon-musk-2/>
- Fridman, L. (2019b). Tesla vehicle deliveries and autopilot mileage statistics. Retrieved December 15, 2019, from <http://lexfridman.com/tesla-autopilot-miles-and-vehicles/>
- Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554.
- Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (pp. 1097–1105). <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436.
- Lin, H. W., Tegmark, M., & Rolnick, D. (2017). Why does deep and cheap learning work so well? *Journal of Statistical Physics*, 168(6), 1223–1247.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). A proposal for the Dartmouth summer research project on artificial intelligence. <https://web.archive.org/web/20070826230310/http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- Minsky, M., & Papert, S. A. (1969). *Perceptrons: An introduction to computational geometry*. Cambridge: MIT Press.
- Mund, S., Frank, R., Varisteas, G., & State, R. (2018). Visualizing the learning progress of self-driving cars. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)* (pp. 2358–2363). Piscataway: IEEE.

- Nuseibeh, B., & Easterbrook, S. (2000). Requirements engineering: A roadmap. In *Proceedings of the Conference on the Future of Software Engineering*, pp. 35–46. New York: ACM.
- NVIDIA Corporation. (2019). NVIDIA DRIVE: Scalable AI platform for autonomous driving. Retrieved December 15, 2018, from <http://www.nvidia.com/en-us/self-driving-cars/drive-platform/>
- Ohnsman, A. (2018). Waymo racks up 10 million test miles ahead of launching its robotaxi business. Retrieved December 15, 2018, from <http://www.forbes.com/sites/alanohnsman/2018/10/10/waymo-racks-up-10-million-test-miles-ahead-of-launching-its-robotaxi-business/#46de3097eb1f>
- Pomerleau, D. A. (1989). ALVINN: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems* (pp. 305–313). <http://papers.nips.cc/paper/95-alvinn-an-autonomous-land-vehicle-in-a-neural-network.pdf>
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1985). Learning internal representations by error propagation. Technical report, California University San Diego La Jolla Institute for Cognitive Science.
- Russell, S. J., & Norvig, P. (2009). *Artificial intelligence: A modern approach* (3rd edn.). Upper Saddle River: Prentice Hall.
- Samek, W., Wiegand, T., & Müller, K.-R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. arXiv:1708.08296.
- Shanahan, M. (2015). *The technological singularity*. Cambridge: MIT Press.
- Siddiqui, F. (2019). Tesla floats fully self-driving cars as soon as this year. Many are worried about what that will unleash. Retrieved December 15, 2019, from <http://www.washingtonpost.com/technology/2019/07/17/tesla-floats-fully-self-driving-cars-soon-this-year-many-are-worried-about-what-that-will-unleash/>
- Society of Automotive Engineers. (2014). Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems. Warrendale: Society of Automotive Engineers. https://www.sae.org/standards/content/j3016_201401/
- Templeton, B. (2019). Elon Musk’s war on LIDAR: Who is right and why do they think that? Retrieved December 15, 2019, from <http://www.forbes.com/sites/bradtempleton/2019/05/06/elon-musks-war-on-lidar-who-is-right-and-why-do-they-think-that/#305b7af32a3b>
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., et al. (2006). Stanley: The robot that won the DARPA grand challenge. *Journal of field Robotics*, 23(9), 661–692.
- World Health Organization. (2019). Global status report on alcohol and health 2018. World Health Organization.
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European Conference on Computer Vision* (pp. 818–833). Berlin: Springer.