berylls
STRATEGY ADVISORS

BERYLLS INSIGHTS

BEYOND ELECTRIC – HOW TESLA COULD BECOME THE APPLE OF AUTOMATED DRIVING

By Berylls Strategy Advisors

WHY TRADITIONAL OEMS SHOULD THOROUGHLY ANALYSE TESLA'S AD STRATEGY

By now, large parts of the established automotive industry have recognised Tesla's achievements as a producer of high-quality battery electric vehicles (BEVs). Although it took the newcomer years, and (often appropriate) criticism prevails, it speaks for itself that VW CEO Herbert Diess recently called Tesla a "role model" and that Porsche's first fully electric model, the Taycan, released in 2019, is constantly benchmarked against Tesla's Model S - which has been on the market since 2012. And although the company's current market capitalisation might partly be based on hype, the newcomer currently is the second most valuable car manufacturer in the world, the company's performance is simply stunning.

But Tesla is much more than just a pioneer for electric vehicles. According to CEO Elon Musk's second "Master Plan", a high-level plan depicting the company's main strategic goals, the next big innovation that Tesla will drive to market is automated driving (AD). Standing in stark contrast to many other OEMs who currently scale back their AD ambitions due to the high associated investments, Elon Musk continues to aggressively push the issue. One strong indication being that Musk recently made AD one of Tesla's top-priorities, overlooking it personally.

So why should traditional automotive care? Alongside most players in the automotive industry, we at Berylls Strategy Advisors are convinced that AD is a major pillar of future mobility. At the same time, we do recognise the tremendous challenges that still need to be overcome – be they regulatory, technical or societal. While tackling those hurdles requires huge upfront investments, we additionally acknowledge that slowed economic growth in core markets, looming trade wars, the switch to electrification and deepened digitalization, and now COVID-19 in addition, leave most traditional automotive players with decreasing financial manoeuvrability. The imperative becomes clear, AD-development must be as cost-efficient as possible while keeping customers safe. The big question is how to solve this conundrum?

While most of the automotive industry seems to focus on Waymo, a close analysis of Tesla seems much more promising. Afterall, traditional automotive has very little in common with Waymo, whose R&D is subsidized by Google's advertising business, but faces in large parts the same challenges as Tesla does. Berylls Strategy Advisors therefore took a closer look at Tesla's AD strategy, which can be discussed along the four high-level building blocks of AD intelligence: sensors, data, software and computer hardware.

THE MAIN BUILDING BLOCKS OF AD INTELLIGENCE

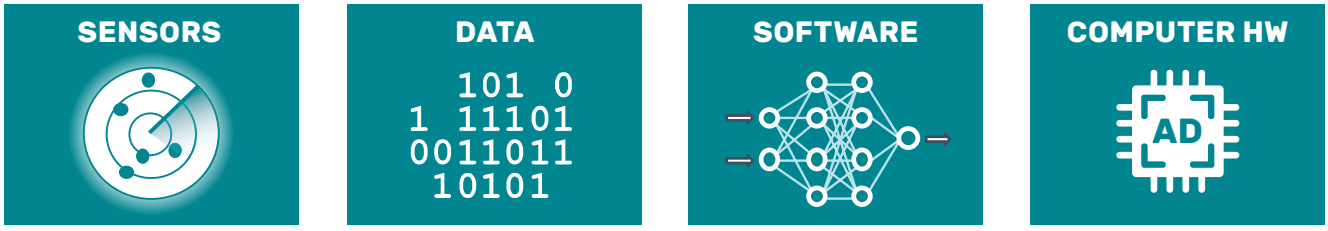

Source: Berylls Strategy Advisors

SENSORS: TESLA AGAINST THE REST OF THE INDUSTRY

When it comes to the sensor suite that gathers the data about an automated vehicle's environment, Tesla pursues a very different strategy than most of the industry. This difference mainly concerns the question whether laser sensors in form of LIDARs (Light Detection and Raging) and high-definition (HD) maps are required to solve the perception, localisation and path planning components of the AD challenge.

LIDARs are sensors emitting pulsed laser waves, which reflect on moving or static obstacles back to sensors, providing a detailed three-dimensional model (i.e. a 3D point cloud) of a vehicle's surroundings. Besides the fact that LIDARs outperform alternative sensors like cameras in range, their greatest advantage lies in their ability to accurately measure depth. This property makes it a lot easier for the self-driving software to distinguish between driveable space and obstacles. While HD maps are not conventional sensors, they equally provide information about an AV'S surroundings. A detailed map of an AV's operational design domain (ODD) provides fine-grained information about topology, drivable lanes etc. and allow the vehicle to "see beyond the horizon".

Most of the AD industry still relies on using LIDARs and many also employ HD maps. Such approaches follow the paradigm of using many different sensors, overlapping and complementing each other, to increase robustness and hence safety. LiDAR for example often outperforms cameras in very bright light or at night – the reason why Waymo has been manufacturing them inhouse since 2011 and VW recently invested "significantly" in the promising LiDAR start-up Aeva. Both companies' AD pilots also rely on HD maps.

Tesla, however, does not use either, with Musk calling both technologies "unnecessary crutches". At Tesla's "Autonomy Day" in April 2019, Musk even proclaimed that "Anyone relying on LiDAR is doomed!". Tesla's sensor suite encompasses "only" eight cameras, twelve sonars, a radar, GPS with a standard map, as well as various sensors relating to vehicle speed, steering angles, acceleration etc. While there are additional advantages like reduced power consumption, the main two advantages of eliminating LIDARs and HD maps from the sensor suite are a) a strong reduction in cost and b) a universal AD stack that is not bound to an ODD.

VISUALISATION OF TESLA'S SENSOR REDUNDANCY

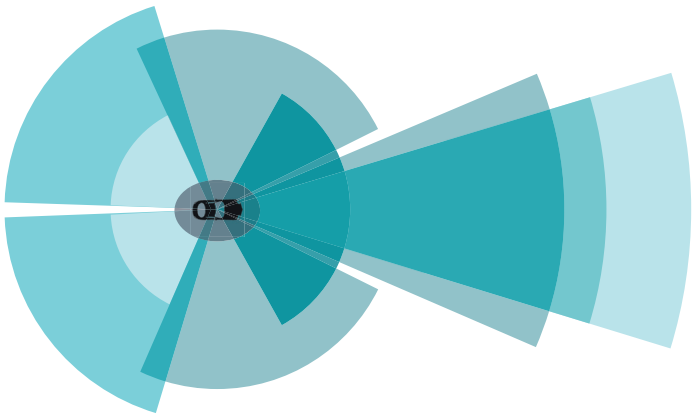

Source: Berylls Strategy Advisors, based on Tesla Inc., 2019

DATA: MORE REAL-WORLD DRIVING DATA THAN ANYONE ELSE … BY FAR

Arguing that humans also do not possess laser vision yet drive reasonably well, Tesla's core assumption here is that computer vision, in particular computer object recognition and prediction, can be trained to match and exceed this specific competency of the human brain. While there are many factors deciding whether this is possible in the near term, the crucial ingredient here is a lot (high quantity) of interesting (high quality) data from the real world.

The need for a large quantity of data stems from the use of inference learning, more precisely of deep convolutional neural networks. Similar to how humans learn throughout life, these particular tools of AI learn from examples – e.g. the software is provided with a digital image of a cyclist, along with the corresponding label. The major difference between such neural nets and us humans is, however, that we need far less examples to correctly generalise. While a child might need to see only one cyclist to then recognise any cyclist as such, AI still needs to be trained with thousands of samples to perform robustly.

TESLA VEHICLE DELIVERIES AND AUTOPILOT HW

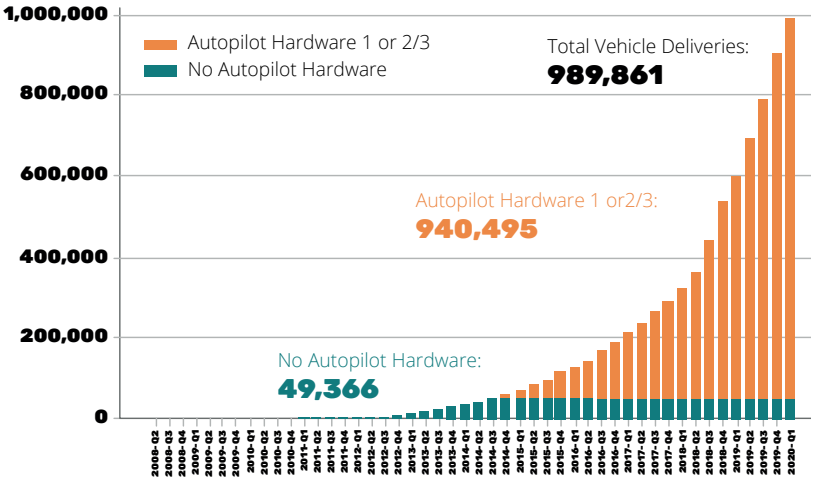

JUMP IN DATA UPLOAD DUE TO DATA ENGINE

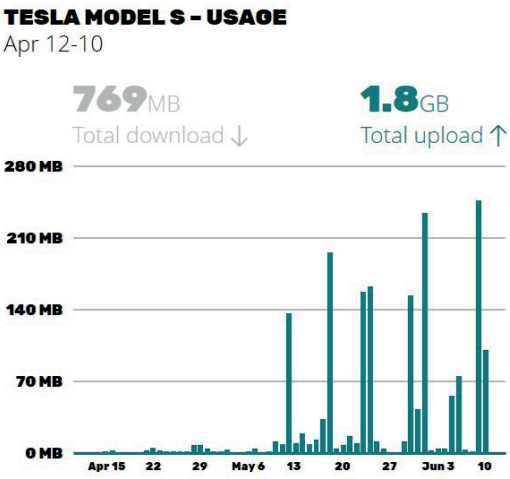

Source: Berylls Strategy Advisors, based on Lex Fridman, 2020 (left) and www.electrek.com, 2017 (right)

Although the debate continues about how much data is enough to guarantee safe deployment (refer to the Berylls Insight "The Future of Automated Driving will be decided by its Perceived Safety"), it is clear that Tesla has gathered more real-world data than any competitor by far, through a process the company calls "Fleet Learning". Enabled by the low cost of its sensor suite, Tesla has been equipping all its delivered vehicles with its Autopilot sensor suite since late 2014. Up to the end of March 2020, this amounts to more than 940k vehicles on the road. To then upload data from these vehicles, Tesla asked its customers permission to upload video snippets recorded by their vehicles' outside cameras in May 2017 – shortly after a marked increase of uploaded data could be recorded.

SOFTWARE: CUSTOMERS PAY TO TRAIN TESLA'S ALGORITHMS

But data quantity is nothing without data quality. Neural network training data also needs to be varied, as automated vehicles can in theory encounter an unlimited number of different situations. Lots has been written about how it is relatively trivial to train an AI model to make the correct decisions 95% of the time, while it is very hard to do the same for the remaining 5%. This is due to the famous "edge cases", meaning the very rare but equally varied "interesting situations". Think of bikes carried on cars, people wearing dinosaur costumes in the streets or a plastic tarp falling of a truck on the highway – automated vehicles must basically deal with anything.

Tesla's clever response to this challenge is to truly include their customers and their vehicles into the software development process. In fact, very often they pay for it; the current price of the Autopilot lies at 6,300 USD. At its Autonomy Day last year, the company shared insights into this process of continuous integration, to which chief of AI Andrej Karpathy referred to as the "data engine".

TESLA'S DATA ENGINE

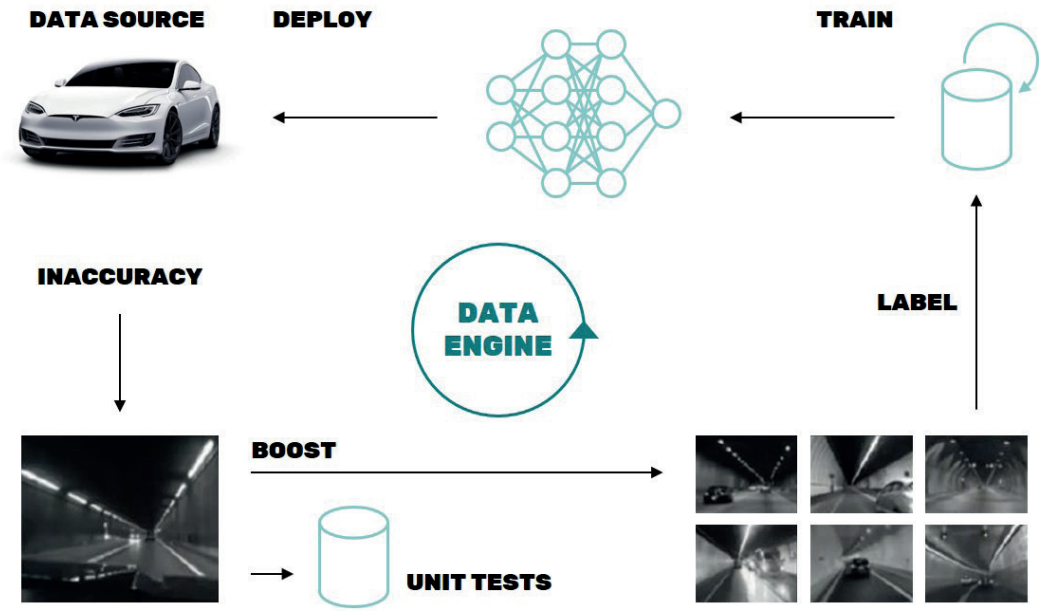

Source: Berylls Strategy Advisors, based on Tesla Inc., 2019 bottom

The process starts with training a first version of the software, potentially at random, before deploying it to a number of vehicles in the fleet. Through the use of their vehicles, customers then start to gather data on the performance of that version – either by activating Autopilot directly, or by teaching a software that is running in the background. The latter mode of data collection is called "shadow mode", a form of AI imitation learning, where the software constantly compares its predictions to the real actions of a human.

But regardless whether Autopilot is activated or only runs in shadow mode, sensor footage of a huge variety of "interesting situations", instances in which the software made a mistake, can be sent back to the Tesla servers.
Tesla engineers then analyse the data before sending a request to the whole fleet to source many more examples of this kind of situation. By doing this, Tesla even claims to label data automatically, an expensive process if done manually by humans. Equipped with many thousand examples of a single challenging real-world situation, the deep neural networks can then be retrained, before the cycle starts again. If this only worked half as well as CEO Elon Musk and his senior engineers proclaimed last year, this development process is unrivalled in increasing speed and decreasing cost.

"Tesla is going to win level 5. (...) Tesla is gathering data on a scale that none of them are. They are putting real users behind the steering wheels. I think it's the only strategy that works."

George Hotz, US-American hacker and founder of comma.ai (US self-driving startup)

We do want to issue a stark word of warning, however: Although the approach seems highly compelling, we are sceptical about letting customers use an immature version of a safety-related software. Numerous incidents where Tesla's Autopilot failed have been reported, sometimes with fatal consequences for their "drivers". Putting a clear emphasise on imitation learning rather than actually letting the software control the vehicle, should alleviate this concern.

ESTIMATED TESLA AUTOPILOT MILES

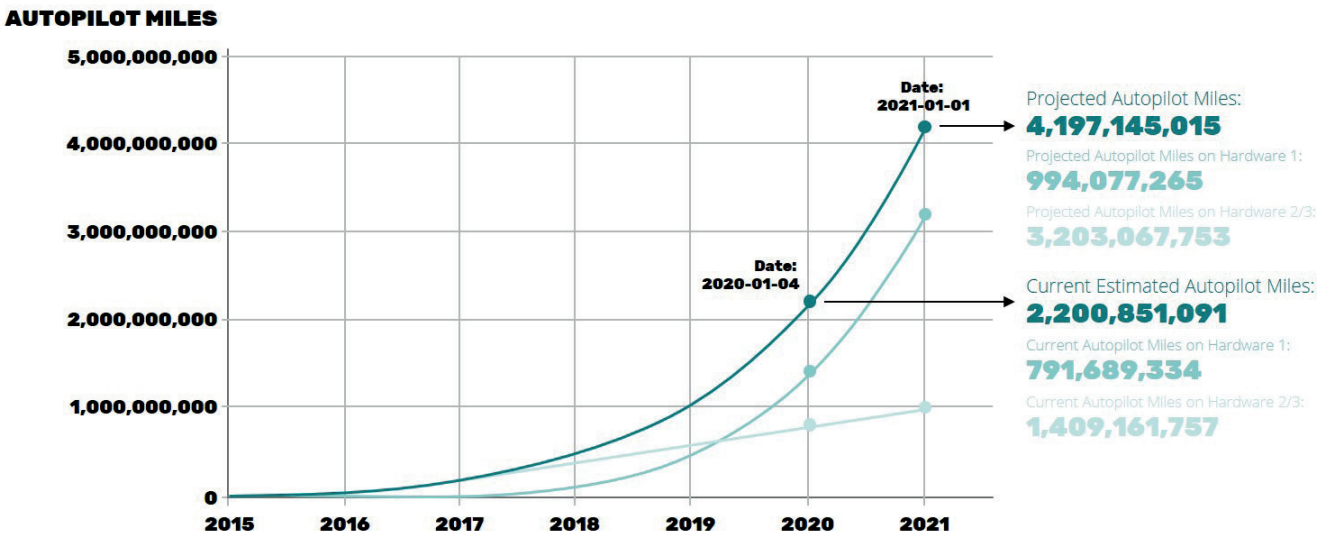

Source: Berylls Strategy Advisors, based on Lex Fridman, 2020

HARDWARE: FOLLOWING APPLE'S EXAMPLE

When it comes to the processors necessary to operate its computer vision software, Tesla's approach shows some similarities to Apple Computers. Just like Apple once decided to design both, software and hardware for PCs, smartphones, tables or watches, Tesla started working on its own full AD computer in 2016. Although the company initially worked with NVIDIA hardware, as many other car manufacturers still do, Tesla decided to hire a team to design a computer that perfectly fulfils the OEMs AD-requirements. The assembled team was first headed by Jim Keller then by Pete Bannon, both highly respected figures in the chip industry, having worked for companies like Intel, Apple, and AMD. During last year's Autonomy Day, Bannon eventually presented the Full Self Driving (FSD) Computer, which has been incorporated into every Tesla model produced from that date on. While it remains debatable whether or not Tesla's solution is truly more powerful than latest NVIDIA versions, the alleged car company did create a powerful computer, which "raised the bar for self-driving carmakers", as the new rival NVIDIA respectfully commented on its blog.

IF ELON MUSK SUCCEEDS, TESLA WILL TRULY BECOME "THE NEXT BIG THING"

We can conclude that what we are currently witnessing is nothing less than a battle of two AD philosophies, of which only one will survive. If on the one hand Tesla has overestimated the capabilities of its world-class AI team, the FSD kit will not get safety approval by regulators, and the humiliated Californians will have lost invaluable time and a lot of money. If on the other hand, however, Tesla's huge bet pays off, the rest of the industry will not be able to compete with their much cheaper AD hardware and software. While the more expensive LiDAR-based kits might still be used in commercial vehicles, Tesla might well become the only AD contestant able to sell a fully functional AD-kit at a price point that is low enough for private passenger cars – and this finally is a scenario justifying the company's current tremendous evaluation.

ACTION PLAN FOR OEMS: PICKING THE CHERRIES WHILE AVOIDING THE PITTFALLS

1. **Weigh costs and benefits of integrating AD sensors into new production vehicles for free** to make use of "fleet learning". Any approach must include relatively affordable sensors, like cameras or optimised new generation LIDARs (like promised by Aeva).

2. **Strengthen own capabilities of handling big data**, either by targeted hiring of data scientists or by building up partnerships with reputed software companies.

3. **Find a SAFE (!) way to include real-world customers in continuous software development process.** Valuing customer safety higher than time to market, the focus should lie on using imitation learning. Mass connectivity and the ability to perform high-volume over-the-air updates are of course prerequisites.

4. **Integrate this process of data sourcing intelligently into existing AD development approaches.** This particularly aims at integrating data on edge cases into physical or virtual simulations.

5. **Beware of data protection laws!** All gathering of data must respect local laws like the European GDPR – Tesla's data engine does pose a high risk to personal and potentially sensitive data.

6. **Don't be afraid to vertically integrate core components of your AD stack if necessary.** Just as Waymo is designing its own LIDARs and Tesla its own computers, traditional OEMs should equally remove bottlenecks by building up new competences if no better alternative is available on the market.

SEEKING ADVICE? MEET BERYLLS.

Berylls Strategy Advisors is a top management consulting firm specialized in the automotive industry, with offices in Munich and Berlin, in China, in Great Britain, in South Korea, in Switzerland and in the USA. Its strategy advisors and associated expert network collaborate with automotive manufacturers, automotive suppliers, engineering services providers, outfitters, and investors to find answers to the automotive industry's key challenges. The main focus is on innovation and growth strategies, support for mergers & acquisitions, organizational development and transformation, and profit improvement measures across the entire value chain.

In addition, together with our clients, experts at Berylls Digital Ventures develop solutions for digitizing and transforming the business models of OEMs, suppliers and engineering services providers.
Longstanding experience, well-founded knowledge, innovative solutions, as well as an entrepreneurial mindset distinguish Berylls consulting teams. Through partnerships with experts, Berylls can draw on in-depth technology expertise, a comprehensive understanding of the market, and powerful networks in order to develop workable solutions.

YOUR CONTACT PERSONS

Matthias Kempf
Partner
matthias.kempf@berylls.com

Robert Ziffling
Associate
robert.ziffling@berylls.com

t +49 89 710 410 400
info@berylls.com

Munich, April 2020