Columbia Business School
AT THE VERY CENTER OF BUSINESS™

Columbia CaseWorks

ID#220201

PUBLISHED ON
AUGUST 27, 2021

Autonomous Vehicles: The Technology Behind the Magic

BY C. DANIEL GUETTA*

Introduction: The History of Autonomous Vehicles

Autonomous vehicles (AVs) have a surprisingly long history, starting as early as 1500, when the artist, architect, scientist, and engineer Leonardo da Vinci designed a self-propelled cart driven by high-tension springs and preset steering.[1] Centuries later, in the 1920s, remote-controlled cars showcased radio technology that eliminated the need for drivers.[2]

In the mid 1980s, the Carnegie-Mellon Navlab 1 became the first self-driving vehicle controlled by a computer.[3] It could navigate and avoid obstacles, although it operated at a much slower speed than modern AVs. Advancements in autonomous technology continued for military and aviation uses until 2004, when the Department of Defense's research arm (DARPA) organized a competition to bring AVs to everyday use. The challenge involved teams of engineers designing a vehicle to autonomously navigate 100 miles of desert roadway. While no team successfully finished, the event was a major milestone in AV development. This concentrated focus by teams around the world set the foundation for best practices, revealed the many ethical and regulatory challenges that would later loom large, and advanced the fields of forecasting and machine learning in ways that would later prove essential to the development of today's AVs.

The Society of Automotive Engineers International (SAE International), a global nonprofit association with a mission "to advance mobility knowledge and solutions for the benefit of humanity,"[4] has developed an industry-standard scale to rate levels of autonomy, ranging between 0 and 5. Levels 0, 1, and 2 correspond to advanced driver assistance systems (ADAS) that are ultimately dependent on human monitoring, including features like parking assist. Levels 3 and above correspond to situations in which the vehicle monitors and drives itself.[5]

Author Affiliation
* *Columbia Business School*

Acknowledgments
The author is grateful to Dean's Summer Fellows William Dunn '21 and Royce Fillmore '22 for research and writing support and to Amar Shah of Wayve and Amitai Bin-Nun and Radboud Tebbens of Motional for their willingness to share their invaluable knowledge of autonomous vehicles.

The benefits of widespread AV adoption would be enormous. Safer travel, lowered insurance premiums, reduced need for traffic enforcement, to name but a few. Indeed, 94% of all serious vehicle crashes are due to human error;[6] eliminating these crashes would be transformational (although some posited it would result in a shortage of organs for transplants, most of which come from fatal car crashes[7]). There would still, however, be improvements needed before self-driving cars could become the norm.

In many ways, the progress made in the AV space over the past few years has been astonishing. For example, the American autonomous driving technology development company Waymo (formerly Google's self-driving-car project) began operating a driverless taxi service in Phoenix—with an entire fleet of vehicles with no safety operator physically behind the wheel.[8] AVs are already in use at certain retirement communities to help residents stay mobile without having to drive a car.[9])

In other ways, however, progress and widespread adoption of AVs were slower than expected, and even the most advanced cars on the road in the early 2020s (including the GM Super Cruise and the Tesla Autopilot) were typically below SAE 3, since they still required the driver to constantly monitor the driving experience. A number of more advanced SAE 3–level technologies (such as Highway Autopilot and Traffic Jam Pilot) were expected to gain extensive market approval between 2021 and 2024,[10] but hurdles remained before AVs could operate at full potential, including infrastructure issues (such as the need for more robust 5G networks to transfer high volumes of data), technical issues (ironing out problems in deep-learning vision algorithms, safety, and reliability, for example), and public perception concerns.[11] In addition, many parts of driving culture taken for granted (such as insurance requirements and regulation) would need a reimagining.

Nevertheless, the stakes were simply too high to give up, and it seemed inevitable that broad adoption of autonomous technologies would become the rule rather than the exception. It is therefore interesting to understand how the complex algorithms running these machines operate and to consider some of the different approaches various companies took to turn the dream of AVs into reality.

How an AV Sees: Sensor Technology

At the core of AV technology is its ability to sense the environment it is in. There are three main sensor types that AVs typically use to navigate: camera, radar, and lidar. Each has distinct benefits and drawbacks, and many companies use a combination of the three to create the optimum driverless experience.

CAMERA

The simplest and most obvious part of an AVs sensing apparatus is the camera. Indeed, when humans drive, they chiefly rely on vision, so it seems logical that cameras would be essential to the operation of AVs. Many common features, such as parking assist, adaptive cruise control, emergency brake assist, and lane-departure warnings, depend on cameras.

AVs have either one (mono) or two (stereo) cameras, and although stereo cameras are more expensive, they can more accurately calculate speeds and distances and create a 3D map of the world around them. AVs leverage computationally intensive deep-learning techniques to analyze camera data and navigate increasingly complex situations.[12] These deep-learning techniques have progressed in leaps and bounds over the past 10 years, but they still fall short of true human vision in many important ways. Thus, whilst it is true that humans mostly rely on vision, it is still an area of open and active debate whether cameras can provide that functionality. Cameras also suffer from a vulnerability to environmental conditions: foggy or rainy days can considerably degrade the quality of camera signals.[13] (See Exhibit 1 for an example of a signal obtained from a camera sensor, together with objects identified therein using deep-learning algorithms.)

RADAR

Short for radio detection and ranging, radar technology dates back to World War II. A radar sensor sends out a short pulse of electromagnetic waves to surrounding objects and waits for those waves to bounce back. Based on the time taken to receive these return signals and their frequency, the sensor can determine the distance and speed of surrounding objects relatively quickly and inexpensively. This technique works even in poor weather but typically results in lower-resolution images. Radar is often classified by the frequency of the signals it emits—short-range radar operates at a lower frequency and covers areas of up to 30 meters, while long-range radar uses higher frequencies and ranges up to 250 meters (AVs use both types). Radar also has a distinct advantage over cameras. Since the algorithms that interpret their signals are relatively simple, well understood, and well developed, they don't require the complex—and still somewhat brittle—deep-learning networks cameras require. (See Exhibit 2 for an example of a signal obtained from a radar sensor and note the signal's low resolution.)

LIDAR

Short for light detection and ranging, lidar is by far the most expensive of the sensor methodologies used by AVs. Much like radars, lidar sensors emit electromagnetic waves that reflect back and are picked up by a photodetector, but these waves are pulses of laser light with considerably shorter wavelengths. This gives lidar signals a very high resolution, allowing them to create an accurate 3D map of the environment around the car. These high-resolution images are then paired with extremely accurate maps and are used to help the AV localize itself. But the cost involved in creating these maps is considerable and forms a formidable barrier to entry for AV companies hoping to use them. The lidar sensors themselves are also prohibitively expensive—sometimes costing tens of thousands of dollars to outfit a single vehicle—though new technologies are reducing costs and bringing these sensors closer to mass production.[14] (See Exhibit 3 for an example of a signal derived from lidar data.)

A robust (and sometimes even furious) debate rages within the AV community as to which of these sensors should play a part in the ideal AV. One side of the debate, most vociferously championed by Tesla's Elon Musk, states that only cameras—perhaps augmented with inexpensive radar to add depth to flat pictures—are required. The secret to humans' ability to drive, the argument goes, is the brain, not extraordinary perception. So using expensive lidar sensors would simply be a futile attempt to compensate for subhuman algorithms with superhuman perception. In Musk's words:

> Lidar is a fool's errand. Anyone relying on lidar is doomed. Doomed! [They are] expensive sensors that are unnecessary. It's like having a whole bunch of expensive appendices. Like, one appendix is bad, well now you have a whole bunch of them, it's ridiculous, you'll see.[15]

Musk's point of view would certainly be attractive from a practical perspective. Lidar sensors cost a fortune, and cost is one of the hurdles in the way of mass production of AVs. In addition, for localization to work correctly, lidar requires extremely high-resolution maps of the world. Not only are these maps expensive to produce, but they are also very brittle. Indeed, any change in the environment (caused by construction, for example) can render them obsolete. (It is interesting to note that in more stable environments, such as warehouses, this limitation would be less problematic).[16]

However, many companies (including Waymo, Nuro, Uber, Aurora, Zoox, Argo, Lyft, and Motional) didn't buy into the camera-only approach, arguing that there is a deep chasm between the ability to see using a human brain versus a deep-learning algorithm. To name just one difference, these algorithms typically look at camera data as a collection of single frames, not a moving image as humans do. It turns out even humans make many mistakes based on still images. Deep-learning algorithms constantly improve, of course, but there are some situations in which even the most sophisticated deep-learning algorithm would be unable to function properly.[17] For example, a post that made the rounds on social media in the early 2020s featured a video of a Tesla Model 3 on autopilot driving behind a truck carrying three traffic lights (presumably transporting them to a repair shop). The algorithm was, understandably, confused![18]

One of the distinguishing features of an AV is the lack of room for failure—cars are dangerous, and by allowing algorithms to drive them, people are entrusting their lives to the technology. Any failure, however infrequent, would draw the ire of the public. Indeed, a number of widely publicized incidents that involved Tesla's autonomous driving system did lead to official investigations.[19] In a world in which we expect almost perfect accuracy, it is difficult to justify eschewing any sensor, however cogent the argument for doing so. This battle may continue raging—though as lidar sensors continue to become cheaper, and cameras become increasingly proficient at sensing their environment, it could eventually resolve itself.

How an AV Thinks: The Analytics Behind the Curtain

Much of the discussion around the analytics supporting AVs focuses on perception—an AV's ability to see the surrounding environment. This, however, is only one part of the AV's job, which can be divided into four distinct components[20]:

- Perception: the way the vehicle perceives the outside world. It chiefly does so by using the sensors discussed above.
- Localization: how the vehicle positions itself in the world.
- Planning or prediction: the translation of the vehicle's high-level goal (to get from Point A to Point B) into a more granular and immediate task (e.g., in the next few seconds, move 10 feet on a straight bearing). A popular method, known as model predictive control, calculates short-term goals with high accuracy and medium-term goals with low accuracy, and constantly updates these goals.
- Control: actuation of the vehicle's short-term plan. This typically involves controlling the vehicle's steering, acceleration, etc.

The multifaceted nature of this task leads to another key distinction in the way various AV players address it. More traditional companies (most famously, Waymo, but also Motional and Cruise) treated each of these tasks as distinct, with individual teams overseeing the development and perfection of each of the components, which are then operationally chained together. An entire team, for example, might oversee perception without excessive regard for how this position data might be used. The process of training such a perception model might involve taking thousands of images, human-labeled with the objects they contain, and teaching the algorithm to reproduce these labels.

Other players (Wayve, and to some extent, Tesla) took a very different approach, treating the entire task of autonomy as one process, with sensors as inputs, and actuation (steering and acceleration) as output, then attempting to train analytic models that take these sensors as inputs, and directly producing instructions for the car as outputs. The training process (at least in its most basic form) would involve feeding the algorithm data gathered during hundreds of thousands of driving hours and teaching it to reproduce the correct actuation signals taken by a human in those situations.

Both approaches could have powerful arguments in their favor. Proponents of the end-to-end approach would note that the segmented approach spends an inordinate amount of time and computational power on tasks that are ultimately useless in making the final decision. For example, if a pedestrian is crossing a street several hundred feet away, a segmented algorithm's perception task would spend valuable computational power detecting someone who will be long gone by the time the vehicle reaches that location. An end-to-end approach only uses features of the world that affect the final goal.

Another argument for the end-to-end approach says it is possible that the various components of an AV might interact in unexpected ways. For example, if the perception module fails to recognize a tree, the impact of the control module is likely to be more consequential than if the perception module failed to recognize a plastic bag. Disparate modules can, of course, be made aware of each other through skillful product management, but the process is less seamless.

On the flip side, the technical challenge of training these end-to-end models is considerably greater than each of the individual four challenges above. One reason for this is that the outputs these models seek to predict (steering and acceleration) are very low dimensional (comprising two numbers) whereas the inputs (sensor data) are extremely complex. Thus, there are potentially millions of different inputs that might lead to identical correct outputs.

Another argument cited in favor of the segmented approach is that the intermediate signals obtained can be useful for other purposes, even if they are not directly used in planning or control. For example, what metrics should be monitored and controlled to assess an AV's safety is an open question. One often-used metric controls the number of times the vehicle gets too close to a pedestrian or stationary object. A segmented algorithm with a full-fledged perception task would have the information needed to control this metric. By contrast, it is not immediately obvious how an end-to-end algorithm might be an instructor to control this specific metric. Proponents of the end-to-end approach would argue these specific metrics are somewhat arbitrary.[21]

From Cars to Robot Vacuums: The World of Autonomy

This case has thus far focused heavily on AV cars, which is perhaps unsurprising—cars are the most visible and widespread application of autonomous technology. That said, there are many other applications of autonomy, some of which are even closer to being widespread than self-driving cars.

The most obvious example is self-driving trucks. In the US alone, the freight industry was worth just under $800 billion a year in the late 2010s,[22] and the COVID-19 pandemic showed it to be brittle and sensitive to disruptions.[23] If autonomy were to tackle even part of this job, the impact would be enormous. Creating a truck that can drive itself on long stretches of homogeneous highway is also much simpler technically than creating a passenger car to drive in a dense, urban environment. Also promising, if less ambitious, is the adoption of autonomous robots in warehouses and factories—in these more controlled environments, autonomy is much easier to achieve.

Beyond these industrial applications, the advances in autonomy have trickled down to many lower-impact products, such as common household items. One specific example of this is robot vacuums—vacuum cleaners mounted on mobile robots that can automatically clean an entire room or house without guidance.

In 1956, the American science fiction author Robert A. Heinlein described the concept in his novel *The Door into Summer,* but it was not until 1990 that three roboticists, Colin Angle, Helen Greiner, and Rodney Brooks, founded iRobot, which launched the Roomba in 2002. The device was the first commercially successful robot vacuum.[24] Robot vacuums proliferated shortly thereafter, and in 2016, Angle, then iRobot's CEO, claimed that 20% of the world's vacuums were now robots,[25] and the number has doubtless grown since.

Robot vacuums are, in essence, mini AVs, and therefore need to handle all the perception, localization, planning, and control tasks described for self-driving cars—albeit on a smaller scale. The stakes are also of course much lower, which has allowed these devices to be deployed even with less-than-perfect algorithms. (Cables and errant dog turds are well known to give robot vacuums trouble, as many blog posts have humorously reported.[26])

By the early 2020s the market for robot vacuums was well developed and highly differentiated. Most vacuums followed the same basic paradigm: the robot itself contained a vacuum, wheels, and various sensors to help the machine navigate, and the system also came with a base to be plugged into a wall for recharging the vacuum. The robot started at the base, vacuumed the space, and then returned to the base for recharging and periodic emptying of detritus.

Higher-end models distinguished themselves in several ways. They typically included advanced features such as wet mopping, carpet detection, and adaptive vacuum strength, depending on how much there was to pick up. The base in these luxury models sometimes acted as a trash receptacle—the robot could empty itself at the base and resume vacuuming.

Of most relevance to this case, these robots also were differentiated by the complexity of their sensors. The most expensive of these models contained lidar and camera sensors. Cheaper models did not contain any complex sensors, instead relying on the tactile feedback received when they bumped into objects and/or various parts of a room.[27] Whatever sensors they used, all robot vacuums faced similar navigation problems. Manufacturers had to ensure that their machines covered the entire area designated for cleaning, without omitting any spots and repeating as few areas as possible. And at the end of the cycle, as the vacuum ran out of battery, it needed to find its way back to the base as quickly as possible. This step was important—if the robot failed to return before it ran out of power, it would be stuck in place, requiring a human to rescue the robot from its predicament.

Assignment: A Simple Localization Problem

In this case, we will consider a very simple version of the localization problem. In particular, we will look at the workings of the 2021 model of the CBSBot[i] (a low-end robot vacuum) near the end of its cleaning cycle. To keep the vacuum economical, it contains no advanced sensors; the CBSBot only computes current speed (measured using the rate at which its wheels are rotating) and bearing (using an ordinary compass). Engineers at CBSBot hope to use this data to determine whether the robot is getting closer to its base—or farther away.

The problem is complicated by a several factors. First, the sensors are not perfect, and there will be some uncertainty around the actual values of the variables, regardless of what the sensors say. Second, the robot will not know its position relative to the base—thus, moving at half a foot per second bearing north would bring the robot closer to the base if it is below it, and farther from the base if it is above it.

---

[i] CBSBot is a fictitious product developed for the purposes of this case.

Despite these challenges, the engineers believe they might be able to use these two basic sensors as part of the robot's localization module and have asked you to investigate whether a predictive analytic model could estimate—at any given point—how much closer to the base the robot is getting.

This problem is, of course, different in nature to the localization problem a full-fledged self-driving car might face, but it will help us develop some intuition for the process of creating these models in practice.

THE TASK AT HAND: BUILDING A PREDICTIVE MODEL

You were given a CSV file with 56,173 rows along with this case. Each row corresponds to one measurement collected in a test lab from a robot on which a lidar-based sensor was placed to accurately measure the robot's position relative to the base.[28] The file contains three columns:

- `dir_moving`: the direction in which the robot is moving, measured in degrees, clockwise, with 0 degrees indicating the robot is moving north.
- `speed`: the speed at which the robot is moving, measured in arbitrary units derived from the number of rotations of the robot's wheels per second.
- `dist_delta`: the change in the robot's distance from the base in feet, during the five seconds directly following the measurements of bearing and speed taken above. Positive numbers mean the robot got closer to the base. Negative numbers mean the robot moved farther away.

(See Exhibits 4, 5, and 6 to view plots produced using these data.)

In a real-life setting in which no lidar sensor is available, you would like to build a predictive model that will anticipate the third variable based on the first two. What kind of predictive model might you use for this task? Can you think of reasons the model may not perform as well as you might expect?

Exhibits

Exhibit 1

Signal from a Camera Sensor, Together with Annotations Produced by a Deep-Learning Algorithm, Identifying Other Objects in the Image

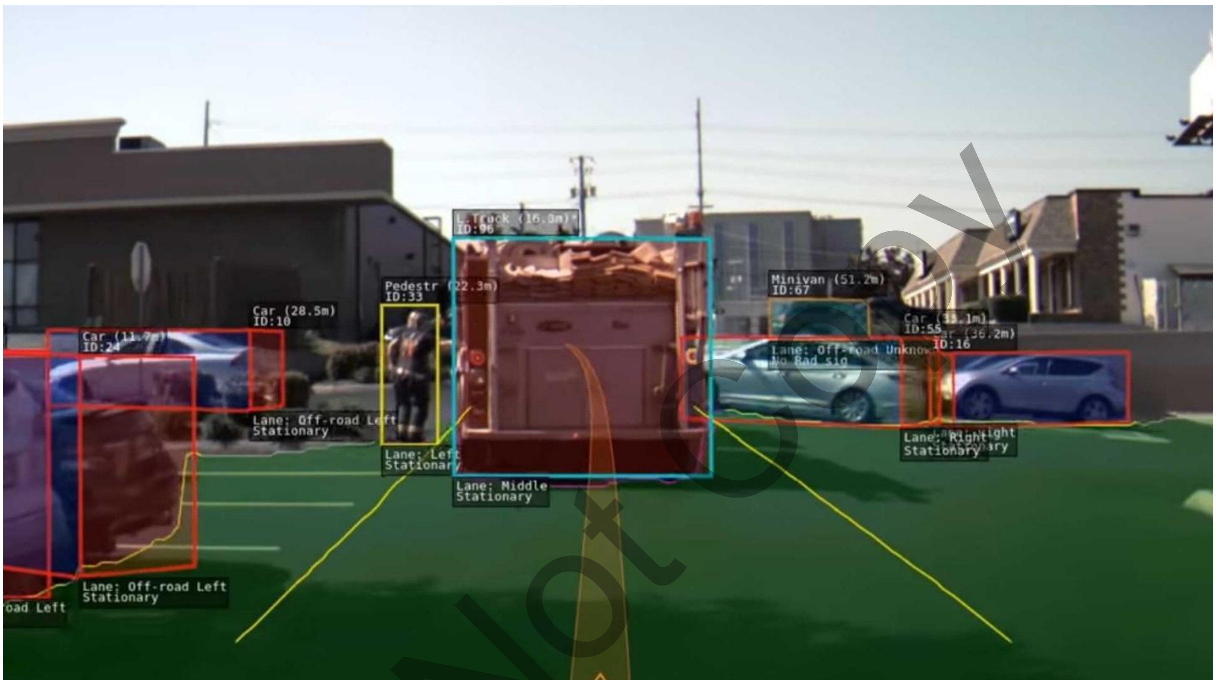

Source: Mark Kane, "Watch This Amazing Video of What Tesla Autopilot Really Sees," *InsideEVs,* April 27, 2019, https://insideevs.com/news/346873/video-tesla-autopilot-sees-fire-truck/.

Exhibit 2
Data from a Scanning Radar Sensor

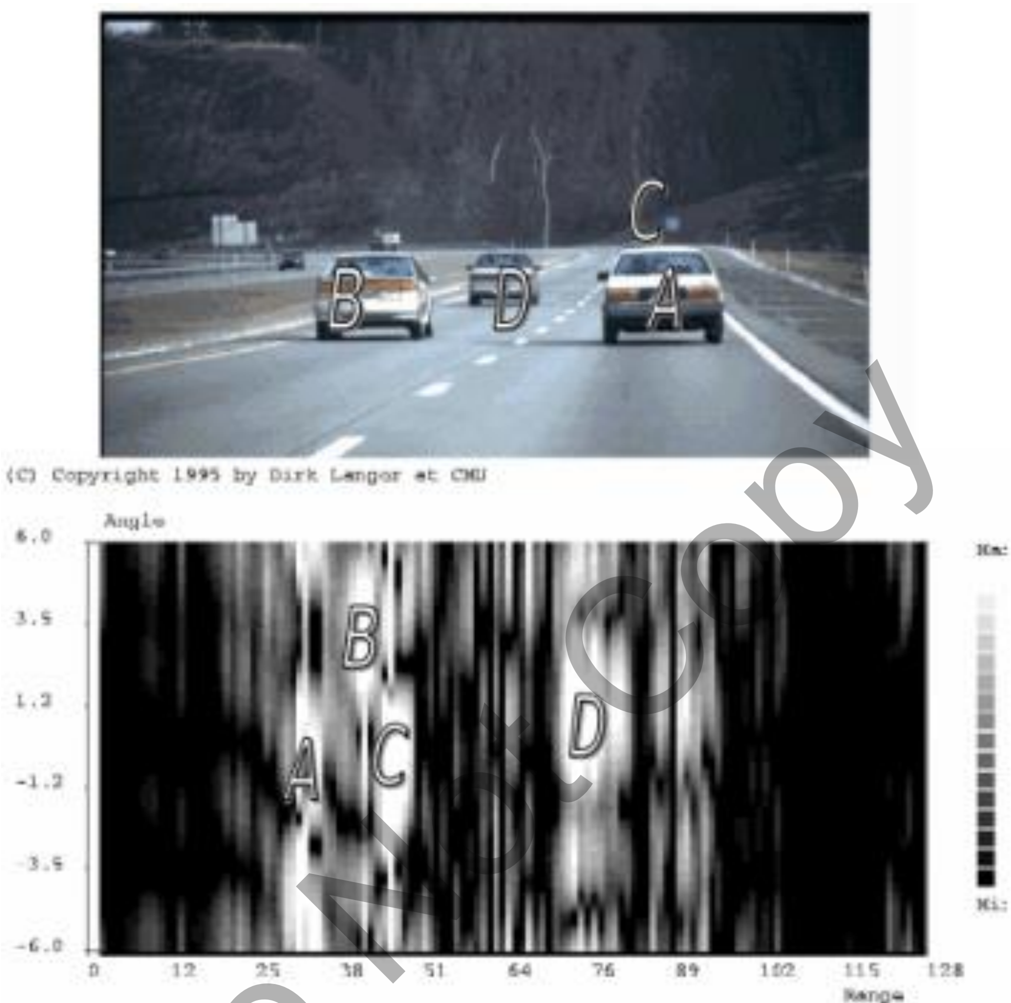

The top image is a video of the scene, and the bottom image displays the corresponding radar data (the brighter the plot, the stronger the signal returned). The x-axis represents closeness, and the y-axis represents bearing angle. Thus, objects that are very close and to the right of the image should appear at the bottom left of the signal plot, and objects that are very far and to the left of the image should appear at the top right of the signal plot.

The signal plot clearly reflects (no pun intended) the fact that vehicle A is closest, followed by vehicle B followed by C and D. Vehicle A is furthest to the right, followed by C, D, and finally B.

Source: David Kohanbash, "Lidar vs Radar: A Detailed Comparison," *Robots for Roboticists* (blog), May 4, 2017, http://robotsforroboticists.com/lidar-vs-radar/.

Exhibit 3
Data from a Lidar Sensor in Google's Self-Driving Car

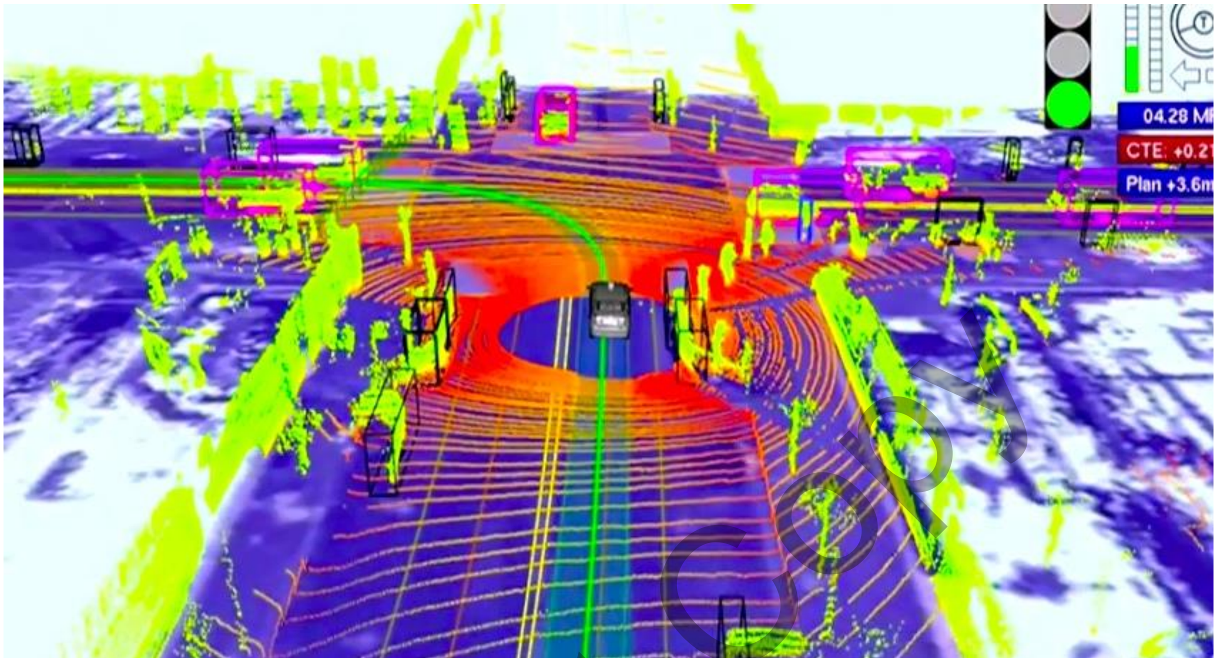

Source: Erico Guizzo, "How Google's Self-Driving Car Works," *IEEE Spectrum,* October 18, 2011, https://spectrum.ieee.org/automaton/robotics/artificial-intelligence/how-google-self-driving-car-works.

Exhibit 4

Speed Sensor Data Plotted Against Average Distance Change from the Base

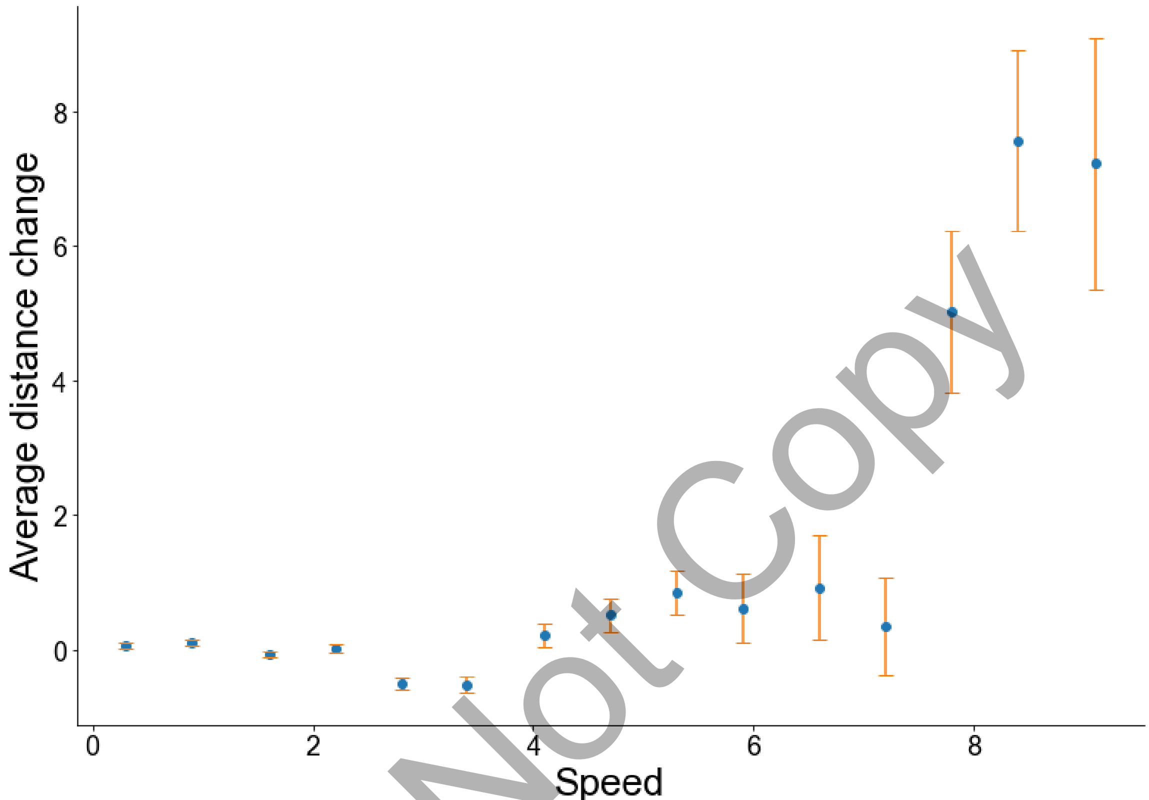

Source: Case writer's dataset, based on S. A. Pettersen et al. "Soccer Video and Player Position Dataset." From proceedings of the International Conference on Multimedia Systems (MMSys), Singapore, March 2014, https://datasets.simula.no/alfheim/.

Exhibit 5

Bearing Sensor Data Plotted Against Average Distance Change from the Base

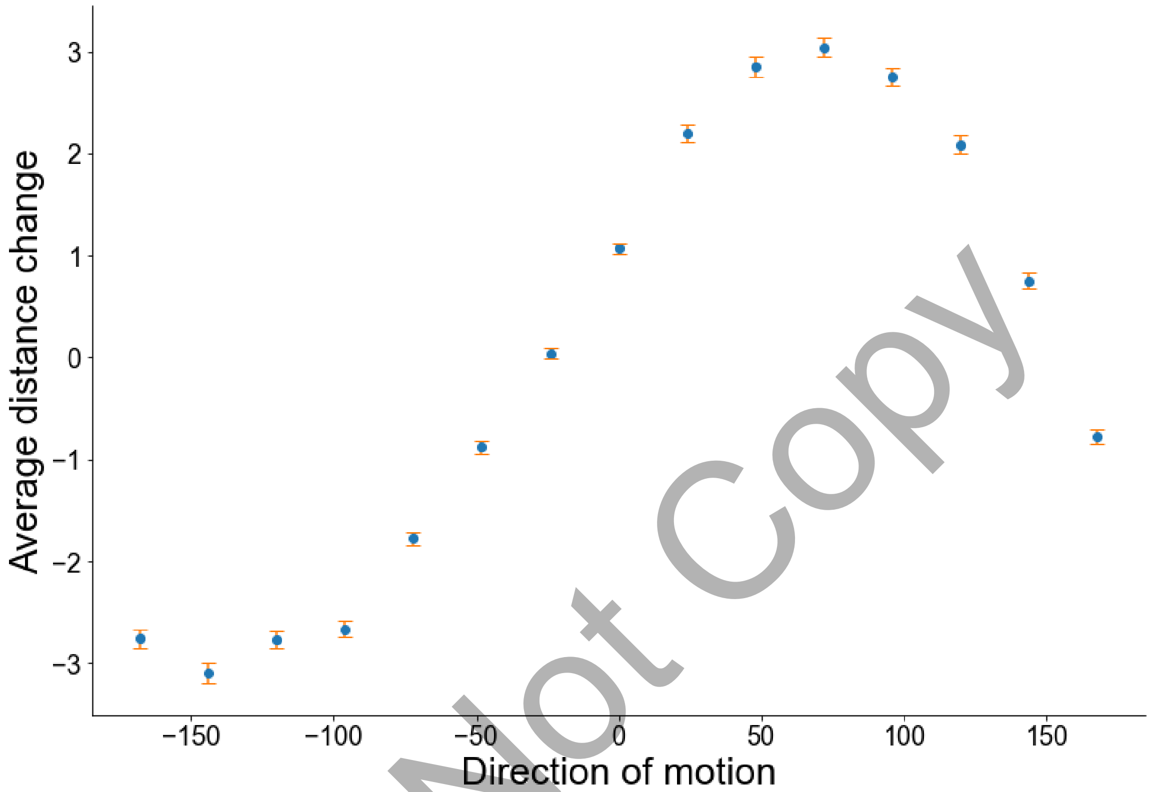

Source: Case writer's dataset, based on S. A. Pettersen et al. "Soccer Video and Player Position Dataset." From proceedings of the International Conference on Multimedia Systems (MMSys), Singapore, March 2014, https://datasets.simula.no/alfheim/.

Exhibit 6

Speed Sensor Data Plotted Against Average Distance Change from the Base, Segmented by Direction of Travel

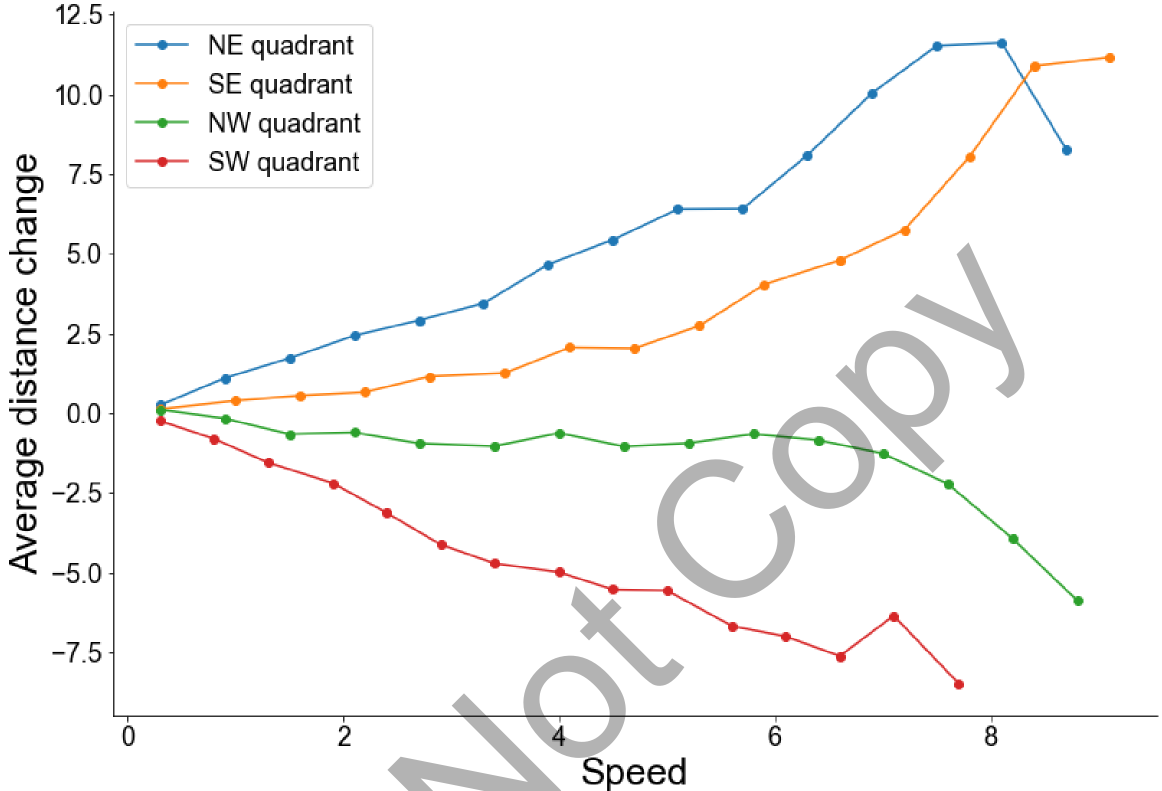

Source: Case writer's dataset, based on S. A. Pettersen et al. "Soccer Video and Player Position Dataset." From proceedings of the International Conference on Multimedia Systems (MMSys), Singapore, March 2014, https://datasets.simula.no/alfheim/.

Endnotes

---

[1] "A Brief History of Autonomous Vehicle Technology," *Wired*, sponsored content by Ford, March 2016, https://www.wired.com/brandlab/2016/03/a-brief-history-of-autonomous-vehicle-technology/.

[2] Tamas Haidegger et al., "Highly Automated Vehicles and Self-Driving Cars (Industry Tutorial)," *IEEE Robotics & Automation*, December 2018, via ResearchGate, https://www.researchgate.net/profile/Tamas-Haidegger/publication/329598453_Highly_Automated_Vehicles_and_Self-Driving_Cars_Industry_Tutorial/links/5c1754a792851c39ebf2eb77/Highly-Automated-Vehicles-and-Self-Driving-Cars-Industry-Tutorial.pdf.

[3] Haidegger, "Highly Automated Vehicles."

[4] "About Us," SAE International, accessed August 25, 2021, https://www.sae.org/about.

[5] Sandra Khvoynitskaya, "Three Types of Autonomous Vehicle Sensors in Self-Driving Cars," blog, Itransition, February 11, 2020, https://www.itransition.com/blog/autonomous-vehicle-sensors.

[6] "Automated Vehicles for Safety," National Highway Traffic Safety Administration (NHTSA), United States Department of Transportation, https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety, accessed August 26, 2021.

[7] Ian Adams and Anne Hobson, "Self-Driving Cars Will Make Organ Shortages Even Worse," *Slate*, December 30, 2017, https://slate.com/technology/2016/12/self-driving-cars-will-exacerbate-organ-shortages.html.

[8] Kristen Korosec, "Waymo's Driverless Taxi Service Can Now Be Accessed on Google Maps," *TechCrunch*, June 3, 2021, https://techcrunch.com/2021/06/03/waymos-driverless-taxi-service-can-now-be-accessed-on-google-maps.

[9] Alan Ohnsman, "Voyage Snags $31 Million as It Targets a Self-Driving Niche: Retirement Communities," *Forbes*, September 12, 2019, https://www.forbes.com/sites/alanohnsman/2019/09/12/voyage-snags-31-million-as-it-targets-a-self-driving-niche-retirement-communities.

[10] Georg Doll et al., "Private Autonomous Vehicles: The Other Side of the Robo-Taxi Story," ed. Eileen Hannigan, McKinsey & Company, December 1, 2020, https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/private-autonomous-vehicles-the-other-side-of-the-robo-taxi-story.

[11] "Automotive Revolution—Perspective Towards 2030," McKinsey & Company, January 2016, https://www.mckinsey.com/~/media/mckinsey/industries/automotive%20and%20assembly/our%20insights/disruptive%20trends%20that%20will%20transform%20the%20auto%20industry/auto%202030%20report%20jan%202016.pdf.

[12] Haidegger, "Highly Automated Vehicles."

[13] "Top Technology Challenges for Autonomous Vehicle Image Sensors," Smithers, March 2019, https://www.smithers.com/resources/2019/mar/challenges-for-autonomous-vehicle-image-sensors.

[14] Ben Dickson, "Tesla AI Chief Explains Why Self-Driving Cars Don't Need Lidar," *VentureBeat,* July 3, 2021, https://venturebeat.com/2021/07/03/tesla-ai-chief-explains-why-self-driving-cars-dont-need-lidar/; German Sharabok, "Why Tesla Won't Use Lidar," *Towards Data Science,* September 1, 2020, https://towardsdatascience.com/why-tesla-wont-use-lidar-57c325ae2ed5.

[15] Matt Burns, "'Anyone Relying on Lidar is Doomed,' Elon Musk Says," *TechCrunch,* April 22, 2019, https://techcrunch.com/2019/04/22/anyone-relying-on-lidar-is-doomed-elon-musk-says/.

[16] Ben Dickson, "Tesla AI Chief Explains."

[17] Haidegger, "Highly Automated Vehicles."

[18] FSD in 6 Months (@FSD_in_6m), "The author posted this video in the comments. It's a truck full of traffic lights," Twitter, June 2, 2021, 5:46 p.m., https://twitter.com/FSD_in_6m/status/1400207129479352323.

[19] Reuters, "Thirty Tesla Crashes Linked to Assisted Driving System Under Investigation in US," *Guardian* (US edition), June 17, 2021, https://www.theguardian.com/technology/2021/jun/18/thirty-tesla-crashes-linked-to-assisted-driving-system-under-investigation-in-us; Lora Kolodny, "Tesla Faces Another NHTSA Investigation after Fatal Driverless Crash in Spring, Texas," *CNBC,* April 19, 2021, https://www.cnbc.com/2021/04/19/tesla-faces-another-nhtsa-investigation-after-fatal-driverless-crash-in-texas.html; Sophia Ankel, "Safety Regulators Are Investigating Autopilot's Role in 30 Tesla Crashes That Killed 10 People, Report Says," *Business Insider,* June 18, 2021, https://www.businessinsider.com/tesla-autopilot-crashes-regulators-open-probes-into-30-report-2021-6; Associated Press, "Tesla Autopilot Safety under Investigation after 'Violent Crash,' NHTSA Says," *USA Today,* March 16, 2021, https://www.usatoday.com/story/money/cars/2021/03/16/tesla-autopilot-safety-nhtsa-investigation/4714393001/.

[20] Industry observations in this section, unless otherwise noted, are based on discussions the case writer had with industry experts. Interpretations are the case writer's alone.

[21] Yi Xiao et al. *Multimodal End-to-End Autonomous Driving* arXiv:1906.03199v2. October 25, 2020. https://arxiv.org/pdf/1906.03199.pdf.

[22] E. Mazareanu, "Trucking Industry in the US—Statistics and Facts," Statista, September 11, 2020, https://www.statista.com/topics/4912/trucking-industry-in-the-us/#:~:text=Trucking%20is%20responsible%20for%20most,billion%20U.S.%20dollars%20in%202019.

[23] Seana Smith, "The Shipping Industry Just Couldn't Tackle COVID-19 Disruptions, Other Shocks: Expert," *Yahoo News,* https://news.yahoo.com/the-shipping-industry-just-couldnt-tackle-covid-19-disruptions-other-shocks-expert-160850293.html.

[24] Brian Bennett, "Why are IRobot's New Roombas So Damn Expensive?" *CNet*, June 2, 2019, https://www.cnet.com/home/smart-home/why-are-irobots-new-roombas-so-damn-expensive/.
[25] Darrell Etherington, "IRobot Says 20 Percent of the World's Vacuums Are Now Robots," *TechCrunch*, November 7, 2016, https://techcrunch.com/2016/11/07/irobot-says-20-percent-of-the-worlds-vacuums-are-now-robots/.
[26] Jolie Kerr, "Our Roomba Vacuumed the House with Dog Shit," *Deadspin* (blog), November 30, 2015, https://deadspin.com/our-roomba-vacuumed-the-house-with-dog-shit-1744174290.
[27] Brian Bennett, "Samsung's New Robot Vacuum Uses Lidar and Empties Its Own Bin Like a Fancy Roomba," *CNet*, January 11, 2021, https://www.cnet.com/home/kitchen-and-household/samsung-jetbot-90-ai-new-robot-vacuum-uses-lidar-and-empties-its-own-bin-like-a-fancy-roomba/.
[28] These data are adapted from the following paper: S. A. Pettersen et al. "Soccer Video and Player Position Dataset." From proceedings of the International Conference on Multimedia Systems (MMSys), Singapore, March 2014, https://datasets.simula.no/alfheim/. They originally referred to the motion of soccer players but were adapted for this case.