



Self-driving cars find their way in the world

Colin McCormick

Sensors, sophisticated algorithms, and powerful onboard computing are the essential ingredients.

Self-driving cars, otherwise known as autonomous vehicles (AVs), are best thought of as rolling robots. By continuously sensing their surroundings, choosing a set of actions, and then implementing them, AVs can navigate various environments. That cycle of “sense, decide, and act” depends on a suite of sensors that feed data to various algorithms and an onboard computing platform to run the system in real time. To understand the trade-offs being explored by AV developers, it helps to start with the sensors.

Lights, camera, action

An AV is easy to spot on the road because of the unusual sensor package that juts up from its roof. A key part of that package is the light detection and ranging (lidar) unit. Automotive lidar, such as the system on the Ford Argo shown on the next page, is composed of an array of semiconductor lasers and optical detectors mounted on a rotating platform in an enclosure. By emitting pulses of near-IR light and measuring the return time of reflections, the system calculates a “point cloud” of objects surrounding the vehicle. The point cloud is updated at roughly 10 Hz, at a range of about 50–100 m, and with a spatial resolution of ± 2 cm. But that density of high-resolution data comes at a price: Lidar units can cost tens of thousands of dollars.

AVs augment their lidar data with radar measurements. Operating at either 24 GHz or 77 GHz for short- or long-distance detection, respectively, radar has much lower spatial resolution than lidar. However, because it is easy to measure the Doppler shift of returned radar pulses, automotive radar can also determine the radial velocity of objects; that measurement, in turn, improves the AV's ability to track pedestrians and other vehicles. Radar works in rain, snow, and other weather conditions that can blind lidar. Another advantage: Radar units can be had for a few hundred dollars.

AVs use lidar and radar data for two main tasks. First, they determine their own position and orientation in space, a process known as localization. Coarse localization is possible with GPS—often augmented with gyroscopes and accelerometers—but autonomous driving requires much higher precision. AVs use algorithms to combine lidar and radar data in order to identify landmarks such as walls, trees, and signposts. The need for high-resolution “base maps” that locate those landmarks means that AVs are usually limited to operating in previously mapped areas. Generating, updating, and distributing those maps to

AVs is an important operational challenge and one of the fastest-growing segments of the AV industry.

The second main task for lidar and radar aboard AVs is the identification and tracking of moving objects, including other vehicles and pedestrians nearby. Statistical algorithms such as Kalman filters maintain estimates of the current position and velocity of all tracked objects and update them using lidar and radar data. Sensors that can detect objects at long distances are particularly valuable for high-speed driving because every 50 m of additional range provides one additional second of warning about oncoming vehicles.

Besides lidar and radar, AVs use the digital video camera as a primary sensing technology. Multiple cameras are on the job—all using standard CMOS technology but with some focused on nearby objects and others focused farther away. Deep convolutional neural networks—a sophisticated form of artificial intelligence—analyze the images to detect and classify objects such as pedestrians, bicycles, stoplights, and other vehicles; to identify lane lines and open space on the road; to read street signs; and to perform other tasks. The information can be used to validate or veto proposed trajectories calculated by the AV based on lidar and radar data.

Path planning

AVs use information about their own position and velocity, the status of objects around them, and lane markings to make decisions about exactly where they should drive. The process, known as path planning, includes a higher-level, strategic component—for instance, whether the AV should switch lanes to pass a vehicle in front of it—and a lower-level, tactical component that determines the optimal steering angles and acceleration or braking to accomplish a maneuver with minimal jerk.

Path-planning algorithms attempt to meet several prioritized goals, from avoiding collisions to completing trips in the least time, as constrained by speed limits. They may also include higher-level decisions, such as determining the best overall street-by-street route to get from the start of a journey to its end.

An important aspect of the AV operational model is the different update rates of various sensory inputs and data processing. The sensors (lidar, radar, and cameras) update at high frequency (typically tens of hertz) to provide the AV with the most recent measurements of nearby objects and road conditions. The localization algorithm updates its estimate of the vehicle's



A FORD ARGO VEHICLE in action. Visible on the roof are two lidar units (top) and multiple video cameras (bottom).

position and velocity at medium frequency (typically a few hertz), and the path-planning algorithm updates its proposed near-term trajectory at the lowest frequency (typically around one hertz) to reflect the fact that physical maneuvers should be relatively smooth and consistent for safe driving.

Because most of those functions must be conducted in real time, all the data processing must be done aboard the vehicle. To maximize the processing speed, especially for the deep neural networks analyzing video-camera images, AVs typically use a graphics processing unit (GPU). The fact that the processors consume hundreds of watts has led AV designers to explore trade-offs between complex algorithms and power draw to avoid compromising the vehicle's overall fuel efficiency. GPUs are also used extensively in the offline training of neural networks; the data collected from many operating vehicles are fed into the networks and used to periodically update AV software.

The future of lidar

The use of lidar for AVs dates back to a 2005 Defense Advanced Research Projects Agency grand challenge: the first successful demonstration of autonomous driving in an uncontrolled environment. The top-finishing vehicles all used lidar devices for high-resolution spatial information about their surroundings. Since that time, almost all AV developers have built their systems around lidar devices, despite the high cost. A notable exception is Tesla, whose Autopilot semiautonomous system relies only on radar and video cameras but continues to face questions about its safety performance. By contrast, companies like Ford and Cruise Automation are experimenting with AVs using multiple lidar units that provide redundancy and enhanced spatial data.

As the single most expensive AV component, lidar has been the focus of intense R&D by industry. One area of interest is replacing the potentially failure-prone rotation architecture. Several companies are working to develop lidar based on micro-electromechanical mirrors, optical phased-array techniques, and wide-field flash illumination. Those methods steer laser beams without bulk mechanical motion and may enable "smart" beam steering that would increase the density of lidar beams near identified objects to acquire more spatial information about them.

A second area of interest is eye safety. Wavelengths near 900 nm, which are typical in automotive lidar, can cause retinal damage. To minimize the hazard, the optical power is kept low. But that limits the sensitivity and range of the time-of-flight method. Some lidar developers have therefore begun using lasers at 1550 nm, which are less dangerous for eyes. The wavelength allows for higher power but requires more expensive indium gallium arsenide detectors.

Industry has also explored alternatives to time-of-flight detection. One of them is the frequency-modulated continuous-wave method, in which the lidar frequency is continuously changed. Using heterodyne detection, the lidar device measures the distance to a reflecting object by sensing the beat frequency between emitted and returned light. The technique is sensitive to the Doppler shift, so it potentially allows lidar to also detect radial velocity.

Another technique being developed is amplitude modulation, which works much like an optical lock-in detector. Amplitude-modulation devices could use lower optical power and might also provide a solution to lidar cross talk, a potential problem if multiple AVs operate in the same area.

Autonomous vehicles combine a wide array of cutting-edge technologies to create a system that can navigate under diverse conditions. AVs already demonstrate impressive technical performance in laboratories and on public roads, and the pace of innovation is rapid. You may find yourself riding in an autonomous vehicle sooner than you think.

Additional resources

- W. Schwarting, J. Alonso-Mora, D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control Robot. Autonomous Sys.* **1**, 187 (2018).
- F. Rosique et al., "A systematic review of perception system and simulators for autonomous vehicles research," *Sensors* **19**, 648 (2019).
- J. Wallace, "Photonics products: Lidar systems: Automotive lidar draws heavily on photonics industry," *Laser Focus World*, 1 November 2018.

PT