Machine Learning Takes Automotive Radar Further

Some may think that the greatest challenge to automating vehicles is in developing the algorithms that tell a vehicle where and how to drive – the planning and policy. It is not. The greatest challenge lies in sensing and perception, in building a perception system that can reliably create the most accurate and robust environmental model for the planning and policy functions to act upon. In this way, perception systems are fundamental to enabling higher levels of automation.

As OEMs look for the best perception systems to deploy in their vehicles to enable lifesaving, active safety capabilities, radar offers a multitude of benefits, including low system cost and resiliency through a wide range of weather and lighting conditions.

These attributes make radar an ideal foundation for building any vehicle's environmental model, and they become especially critical as vehicles move beyond basic warning functions and into assistance and automation functions. Centralizing the intelligence and applying machine learning in just the right way can turbocharge the performance, ensuring that vehicles capitalize on radar's strengths while fusing its data with that of other sensing modalities. In doing so, OEMs can create the best canvas on which to design and implement planning and policy functions that provide advanced features and solve the most challenging corner cases.

Machine Learning and Radar

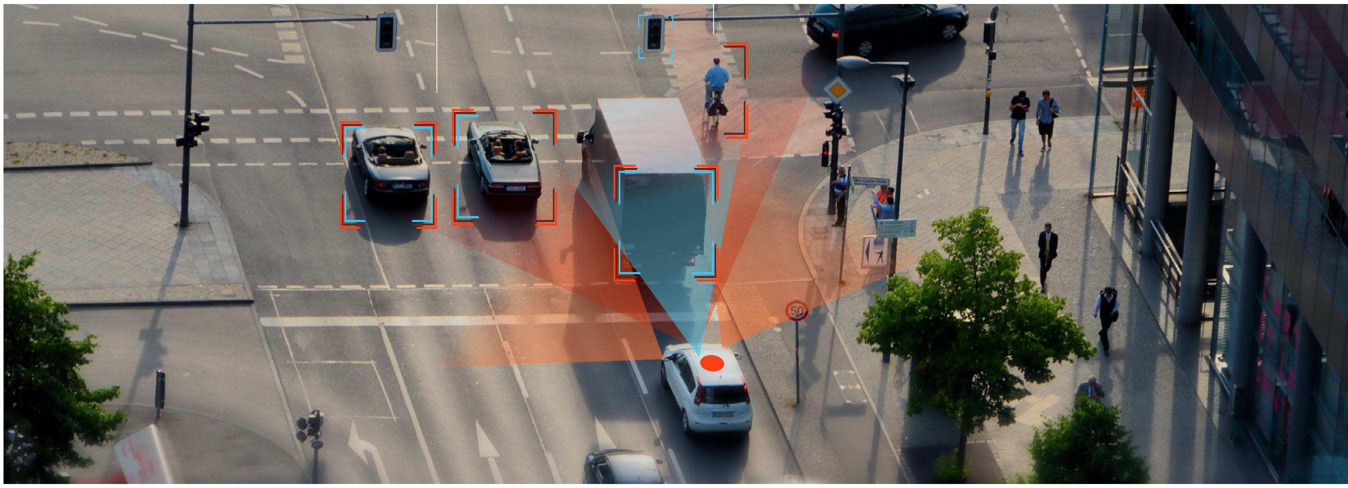

Active safety capabilities save lives and prevent accidents. For example, forward collision warning with automatic emergency braking reduces rear-end collisions by 50%, according to the Insurance Institute for Highway Safety. In a 2019 Consumer Reports survey, 57% of vehicle owners said an advanced driver-assistance feature in their vehicle had prevented them from getting into an accident. These solutions typically employ a forward-facing radar or camera – or ideally, both.

The challenge for OEMs in the coming years will be to bring more advanced active safety features to the market in a cost-effective way, allowing OEMs to offer the capabilities on more models and bring them to more consumers – while at the same time laying the groundwork for higher levels of automation, which will have to address the most difficult sensing challenges.

Success depends on two primary functions: the quality of the information provided by the sensors, and the ability of the compute to interpret that data. On the sensor side, radar-centric solutions provide an excellent foundation for this path. On the compute side, a machine learning system can use the data coming from radar sensors and combine it with data from other sources to create a very robust picture of a vehicle's environment.

BENEFITS OF RADAR

The main sensors in use today on vehicles are radar and cameras, with ultrasonics playing a role in short distances at low speeds and lidar used in autonomous driving.

Part of the reason radar is widely used is that it can reliably indicate how far away an object is. Typical long-range automotive radars can provide range measurements on objects that are as much as 300 meters to 500 meters away. Cameras, by contrast, have to try to estimate how far away an object is based on the size of the object in the camera's image and other factors. Even leveraging a stereoscopic approach, this can be challenging. Further, resolution becomes an issue, as a single pixel in a camera image is very broad at long range, making it harder for a camera to discern those objects. Focusing optics can help, but they limit the field of view, leading to a challenging compromise typical of camera-based perception systems.

At the same time, radar makes inherent measurements of relative speed, so at the same time it is providing a range measurement, it can also tell how quickly something is moving toward the vehicle or away from it. Cameras and lidars may need to take multiple images over time to estimate relative speed.

Because radar uses radio waves instead of light to detect objects, it works well in rain, fog, snow and smoke. This stands in contrast to optical technologies such as cameras – or in the future, lidar – which are generally susceptible to the same challenges as the human eye. Consider the last time you were blinded by direct sunlight while driving, or tried to see clearly through a windshield covered with dirt and grime. Optical sensors have the same challenges, but radars can still see well in those cases. And unlike cameras, radar does not need a high-contrast scene or illumination to sense well at night.

Figure 1. Radar can perceive its environment in a variety of weather and lighting conditions.

Radar also provides an OEM significant packaging flexibility, thanks to its ability to work when placed behind opaque surfaces. Optical technologies need to be able to "see" the road, which requires them to be visible from the outside of a vehicle – preferably at a high point so they can have good line of sight and stay clear of road dirt and grime. Radar, by contrast, can be placed behind vehicle grilles, in bumpers, or otherwise hidden away, giving designers significant flexibility to focus on vehicle aesthetics.

WHERE TO USE OPTICAL SENSORS

Cameras are well suited for object classification. Only a camera can read street signs, and a camera is best at telling if an object is another vehicle, a pedestrian, a bicycle or even a dog. Each of those objects is going to behave differently, so the vehicle's system will be better able to anticipate movements if it knows exactly what it is looking at.

Lidar has drawn attention because it offers some unique strengths. It can take direct range measurements at high resolution and form a grid, where each grid cell has a particular distance associated with it. Because lidar operates at a much higher frequency, it has a much shorter wavelength than traditional radar – and that means it can provide higher angle resolution than radar, allowing lidar to identify the edges of objects more precisely.

One downside of lidar is that it needs to have a clean and clear surface in front of it to be effective, which of course can be especially problematic on a moving vehicle. One unfortunate yet well-placed beetle could render a vehicle sightless.

An equally significant issue is that lidar is a less mature technology than radar, which means it's much more expensive. The expense limits how widely lidar can be used in today's high-volume automotive marketplace.

To ensure a reliable and safe solution, a vehicle should have access to a combination of different sensing technologies and then use sensor fusion (**see sidebar→**) to bring those inputs together to gain the best possible understanding of the environment. But even if that isn't possible – if the cameras are smudged and the lidar is having bug-splatter issues – the radars in the vehicle can deliver excellent information, especially when paired with the right machine learning algorithms.

Figure 2. Radar can be located behind the outer body of a vehicle.

SENSOR FUSION

Sensor fusion is the ability to bring together inputs from multiple radars, lidars and cameras to form a single model or image of the environment around a vehicle. The resulting model is more accurate because it balances the strengths of the different sensors. Vehicle systems can then use the information provided through sensor fusion to support more-intelligent actions.

Of course, the more sensors on a vehicle, the more challenging fusion becomes, but also the more opportunity exists to improve performance.

In the past, the processing power to analyze sensor data to determine and track objects had been packaged with the cameras or radars. With Aptiv's Satellite Architecture approach, the processing power is centralized into a more powerful active safety domain controller, allowing for data to be collected from each sensor and fused in the domain controller.

RADAR

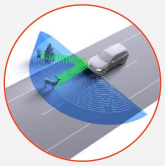

- Long-range sensing
- Object movement
- All-weather performance

LIDAR

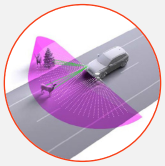

- Precise 3D object detection
- Range accuracy
- Free-space detection

CAMERA

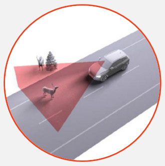

- Object classification
- Object angular position
- Scene context

	RADAR	LIDAR	CAMERA	FUSION
Object detection	+	+	O	+
Pedestrian detection	—	O	+	+
Weather conditions	+	O	—	+
Lighting conditions	+	+	—	+
Dirt	+	O	—	+
Velocity	+	O	O	+
Distance - accuracy	+	+	O	+
Distance - range	+	O	O	+
Data density	—	O	+	+
Classification	—	O	+	+
Packaging	+	—	O	+

+ = Strength O = Capability — = Weakness

MACHINE LEARNING

Machine learning is a subset of artificial intelligence that refers to a system's ability to be trained through experience with different scenarios. As vehicles become more automated, developers can use machine learning to train systems to identify objects and to better understand their environment with less data.

One challenge machine learning helps address with radar is edge detection. Radar's longer wavelengths produce lower resolution that can lead to under-resolved targets, making it difficult to tell where a target's edges are. When that happens, it becomes challenging to interpret the data and resolve the scene. Engineers are working on ways to improve the resolution of radar, such as moving up from the common 77 GHz frequency used in today's automotive applications to 120 GHz or higher, with a corresponding reduction in wavelength. That allows a much higher resolution for the same size sensor. Even with today's radars, however, machine learning can help to characterize different scenarios when the data is difficult to describe through standard algorithms.

Developers can present many examples of objects in a particular category to a machine learning system, and it can learn how signals are scattered by complex objects with many reflection points. It can take advantage of contextual information. And it can even learn from simultaneous data provided by cameras, lidars or HD maps to classify objects based on radar signals.

Further benefits are possible if we use machine learning judiciously. Instead of taking a brute-force approach and applying machine learning to all of the raw data provided by a radar, we can do some classical preprocessing and then apply machine learning just to those portions that make sense.

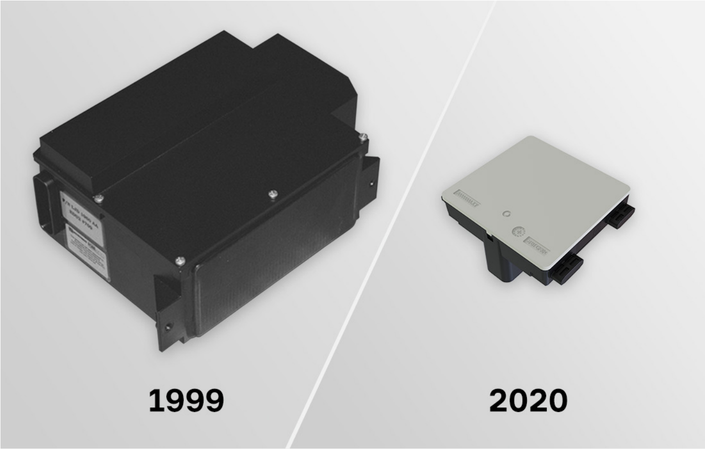

AN AUTOMOTIVE FIRST

Aptiv pioneered advanced driver-assistance systems (ADAS) technologies in 1999 with an adaptive cruise control system for the Jaguar XKR. Using a microwave radar in the front of the vehicle, the adaptive cruise control (ACC) system measured the distance and relative speed of the preceding vehicle and used throttle and braking to ensure that the Jaguar stayed 1 to 2 seconds behind it.

The technology won a PACE Award, but the radar was expensive, so the capability was aimed narrowly at luxury vehicles. Engineers joked that if you bought the radar, you got the car for free. Many generations of hardware later, the technology is smaller, lighter and less than one-tenth the cost. Radar has proven successful through decades of harsh use, and vehicles of all levels now rely on the technology to provide active safety features to consumers.

Many automotive radars utilize an array of antennas to measure angle. In classical radar signal processing, the digitized signals from each antenna are converted to range and speed. The signals are compared across the antenna array to make angle measurements. An example of preprocessing would be to use classical signal processing to isolate regions of interest, to focus on objects with certain ranges and speeds. The signals from each antenna with a common range and speed can then be used to train a system.

Common radars can utilize up to 12 antennas, and five or more radars can be employed on a single vehicle. Those antennas allow digital beamforming, where the signals from each individual antenna are digitized and then combined digitally. The result is that the radars sample the signals one time and then form beams in as many different directions as necessary. By looking across these arrays and analyzing the places where the radars overlap, the system can deduce the angles of different objects.

This kind of analysis gives the system a rich basis of information to feed into a neural network, which in turn can apply machine learning to produce an even clearer picture of the scene. Without this interim step, an AI system would have to determine the scene from the raw digitized signals themselves in real time, which means it would need to be extremely powerful and therefore more expensive and resource-intensive, and it would require long training sequences to figure out what to make of the data. Plus, such a system would be difficult to troubleshoot – if the vehicle detected an object that was not there, for example, it could be difficult to figure out where the processing went wrong. Combining classical processing with machine learning can provide some orthogonality in the data processing, which increases the robustness of the system.

While the data provided by a radar is more complex than what comes in from vision systems – providing range and range rate in addition to location of objects – it is also quite valuable. It is well worth the effort to intelligently sift through the data to extract meaning. Aptiv's 20-year history of working with automotive radar – we were the first to put a radar in a Jaguar in 1999 to enable adaptive cruise control – has given us the expertise needed to pull out the relevant data in the most efficient way.

COST AND POWER ADVANTAGE

Emerging architectures have satellite radars distributed throughout a vehicle, connected via Ethernet to a central system-on-chip with a machine-learning accelerator. Aptiv is using this kind of Satellite Architecture to process data from five radars or more and keep costs down. The approach is highly data efficient, and the machine-learning models can run on processors that cost less and consume less power than alternatives.

For example, an implementation that processes data from six short-range radars would use about 1W, whereas an implementation processing data from six cameras could consume 10W to 15W, and a high-end graphical processing unit consumes around 100W.

In another example, machine learning can glean information on range and free-space detection on radar-generated data to deliver results that are close to lidar, but at radar's lower cost.

Potential savings come from not having to build parallel implementations of processing, RAM and communications in every sensor, and from the efficiencies gained from centralizing software in a domain controller. The lower cost means that even standard or entry-level vehicles can be equipped with this lifesaving technology.

CHALLENGING SCENARIOS

There are many scenarios that human drivers encounter every day that do not lend themselves to easy solutions when it comes to advanced driver assistance systems. If there is an object in the road, is it safe to drive over? How should the vehicle adjust its driving if an adjacent truck creates a blind spot? Machine learning coupled with radar can address these and many more concerns. Here are a few examples.

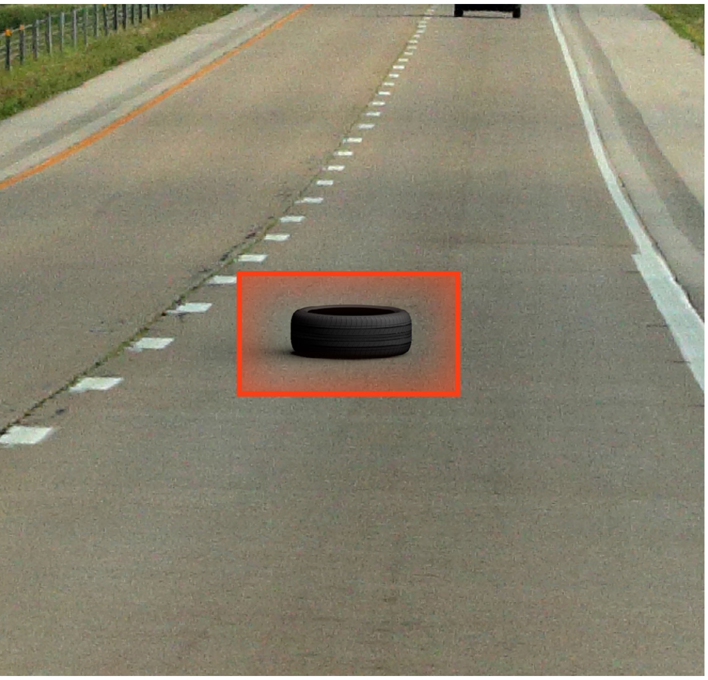

Debris in the road

Small objects or debris in the road can pose a challenge, particularly at high speeds. Radar with machine learning has been shown to improve range by more than 50% and enable it to track small objects at 200 meters, which allows plenty of time for the vehicle to either change lanes or come to a safe stop.

Objects that are safe to drive over

Human drivers often take for granted their ability to gauge whether an object on the road is something they could drive over. They do not estimate that the object is 5 cm high or 10 cm. They tend to act on intuition – a feeling, perhaps, based on their past experience. A machine learning system can also be trained with objects that are safe to drive over and those that are not. Programmers can create a portion of the overall processing chain focused on this question as a special subset of object classification – "over-drivable," yes or no? – and pass the answer on to software that can take action if needed.

Vulnerable road users

Vulnerable road users include bicyclists and motorcyclists. This has been a particular area of focus for regulatory and rating agencies because these users have little protection in the event of a collision, and they can be more difficult to identify than other vehicles. Machine learning reduces misses by 70% compared to classical radar signal processing, and sensor fusion with other sensing modalities can improve detection further.

Pedestrians

Detecting pedestrians can present unique challenges to any kind of sensors, particularly in a cluttered urban environment when many pedestrians could be crossing a street and walking in different directions. By using all dimensions of radar data as described earlier, however, advanced machine learning techniques can help the vehicle see the pedestrians in the cluttered environment. It can even spot them behind a parked car or other obstruction that may hide them from view.

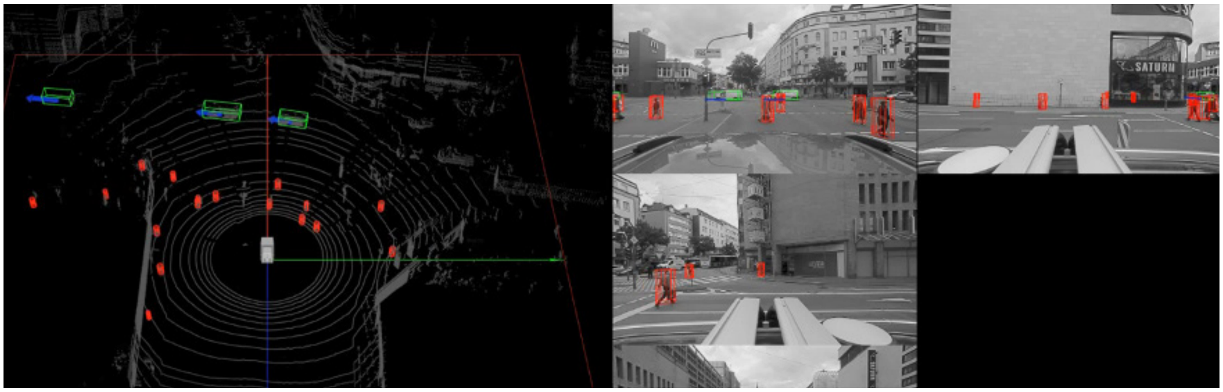

Occluded pedestrians:

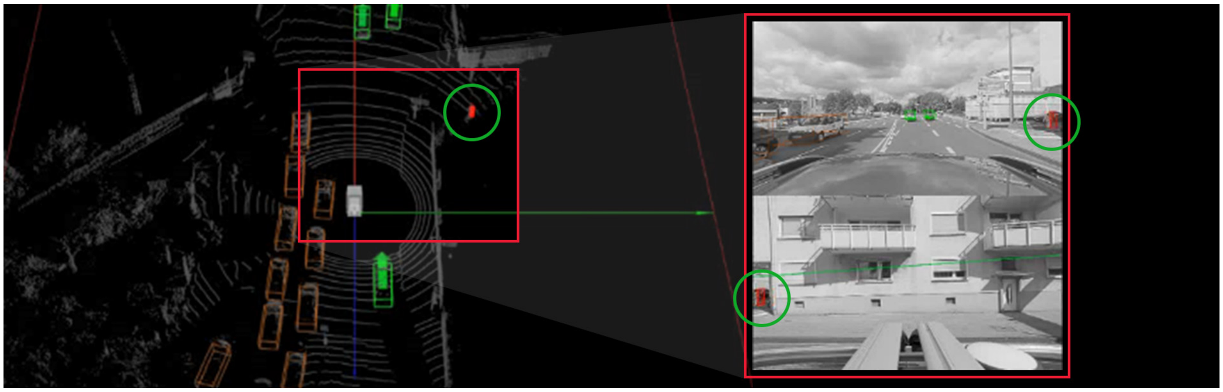

Pedestrian near path and parked vehicle:

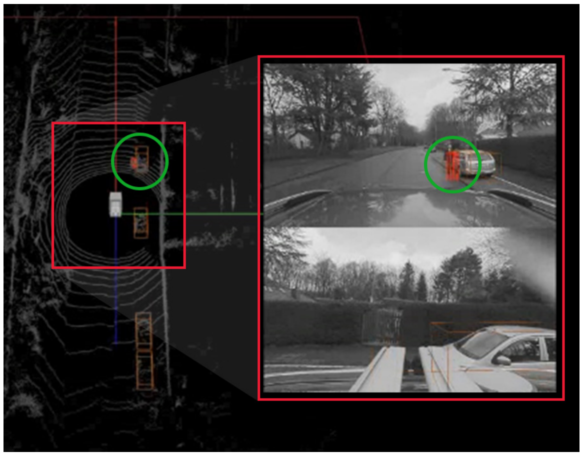

Pedestrian alert during rear parking maneuver:

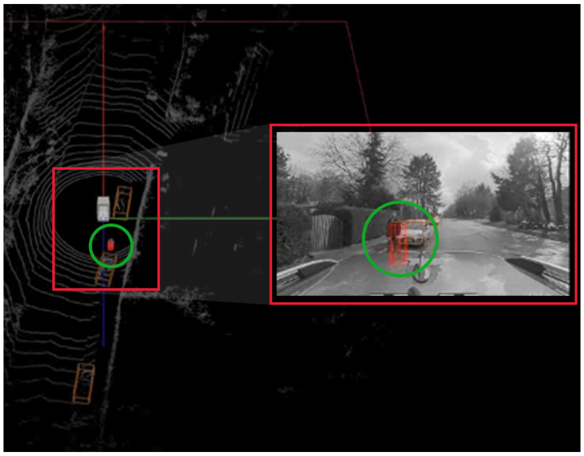

Low-reflectivity road boundaries

Some road boundaries, such as flat concrete walls seen from acute angles, do not reflect radar strongly. Machine learning can use robust segmentation and signal processing across range, Doppler and antenna response over time to figure out where those boundaries are.

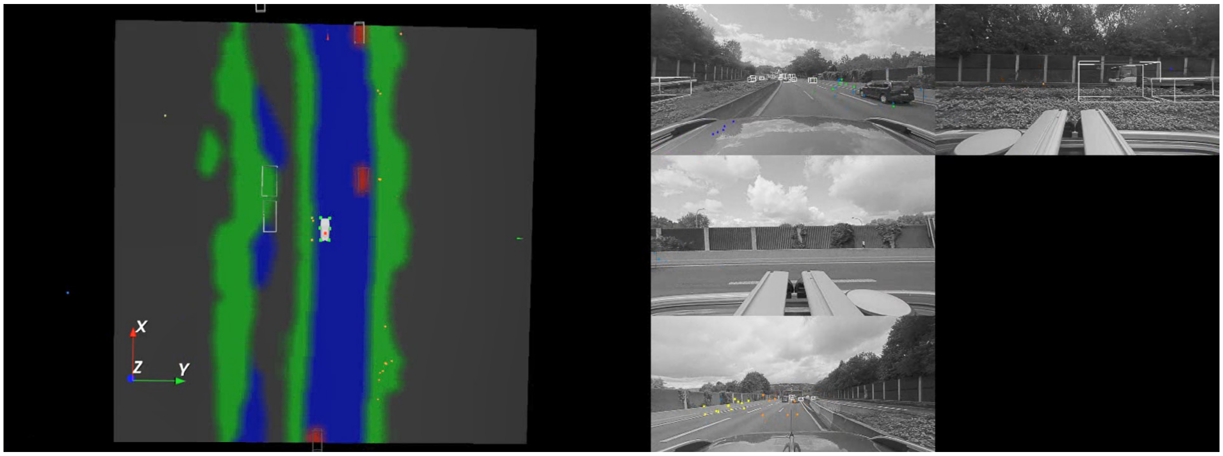

Blind spot

Sensor occlusion – a blind spot created by another object, like a large truck – is one of the biggest challenges of automated driving. It is less a problem of failing to detect occluded objects than it is the fact that today's systems are not fully aware of their blind spots. Human drivers have learned to account for unseen possibilities and guard against threats that may be hidden. Aptiv's perception approach creates this awareness and allows upstream functions to act defensively, as a human driver would.

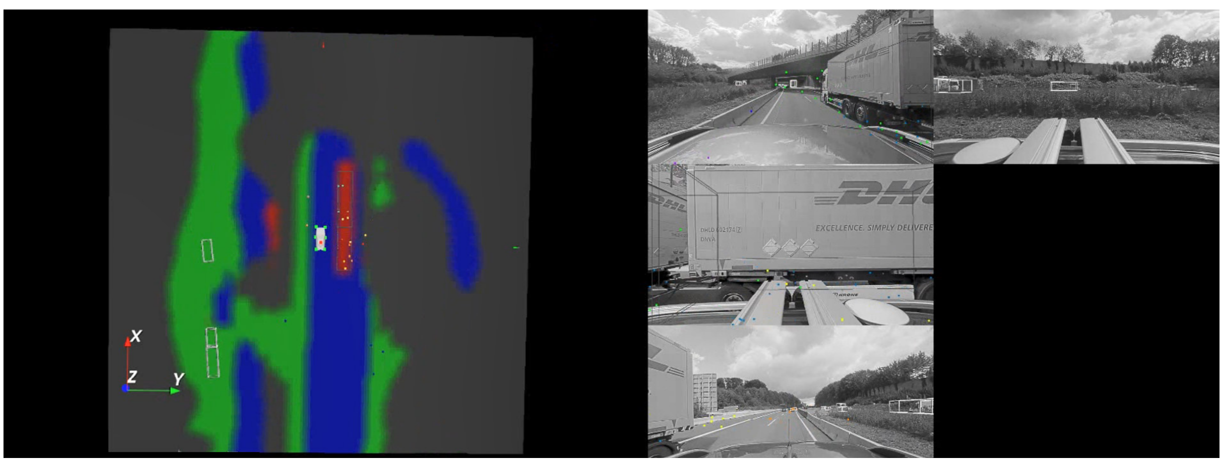

Stopped car in lane

Machine learning can help provide accurate object detection and tracking, including object boundaries and robust separation. With advanced processing methods, we can decrease position error and object-heading error by more than 50%, which means that the vehicle is better able to tell when another vehicle is stopped in its lane.

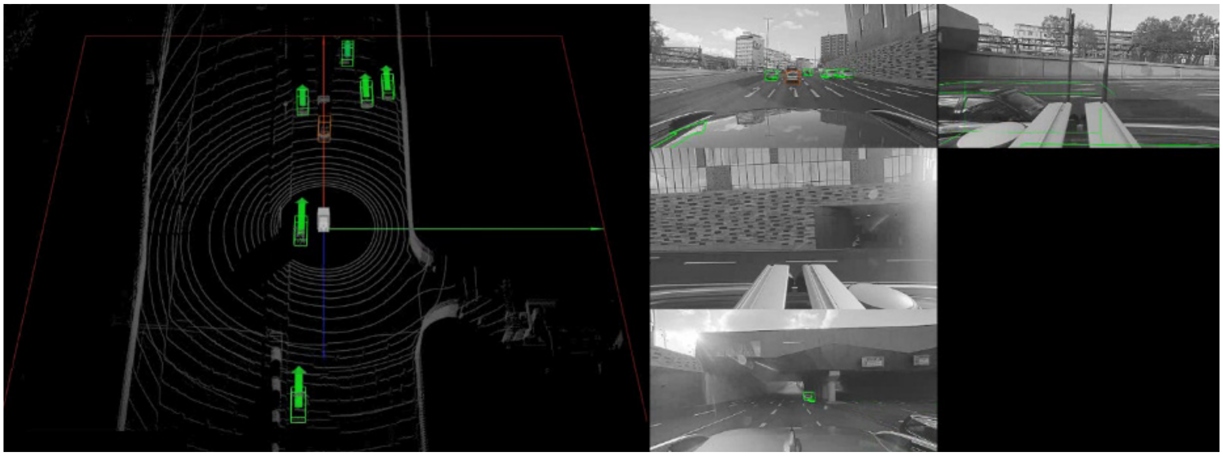

360-degree sensing

Aptiv's sensor fusion approach brings together inputs from various sensors around a vehicle. If the vehicle is equipped with enough sensors, this means it can have a 360-degree view of its environment, and that complete picture will help the vehicle make better decisions. Machine learning helps the system identify objects within that scope, classifying them as cars, trucks, motorcycles, bicycles, pedestrians and so forth. It can determine their heading. And it can help separate and identify stationary or slow-moving objects.

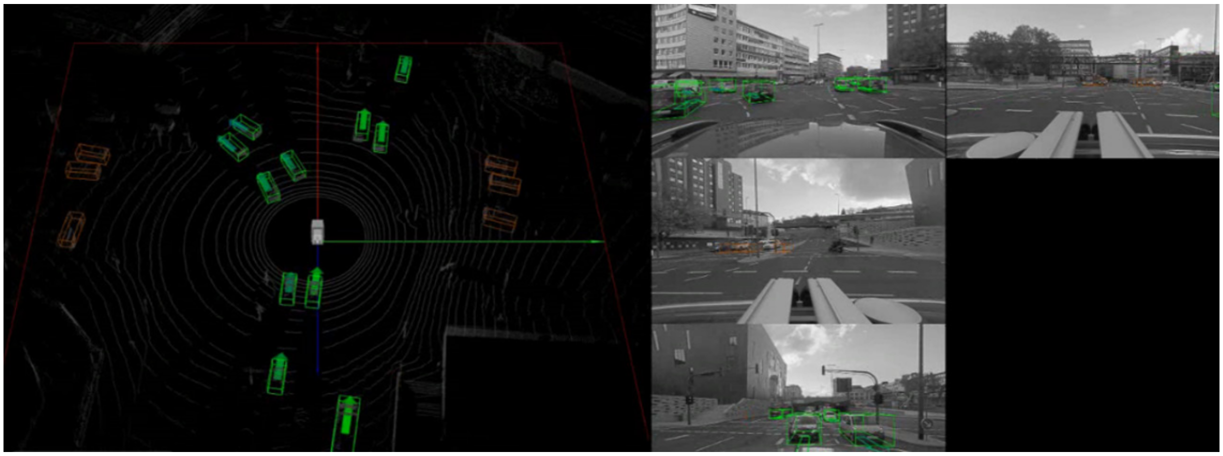

Tracking inside a tunnel

Machine learning can also help a vehicle understand when it is inside a tunnel. Tunnels have historically been a challenging environment for radar. The tunnel walls provide a reflective surface, which can result in a very high number of detections that can overwhelm a radar's capacity to process targets. Also, these reflections can come from high elevation angles, which can make stationary targets difficult to be identified as such. Further, tunnels will often have fans to help clear stagnant air, and the spinning blades of the fan could confuse a radar into thinking it is seeing a moving object. All of these issues can be mitigated by making adjustments to the radar processing when the vehicle is in a tunnel. By applying machine learning to radar data processing, the system is able to filter out noise from positive detections with much greater accuracy than classical methods have allowed. It can now better interpret radar returns in tunnels and other closed environments, classify targets such as fans, and effectively solve radar's tunnel challenge.

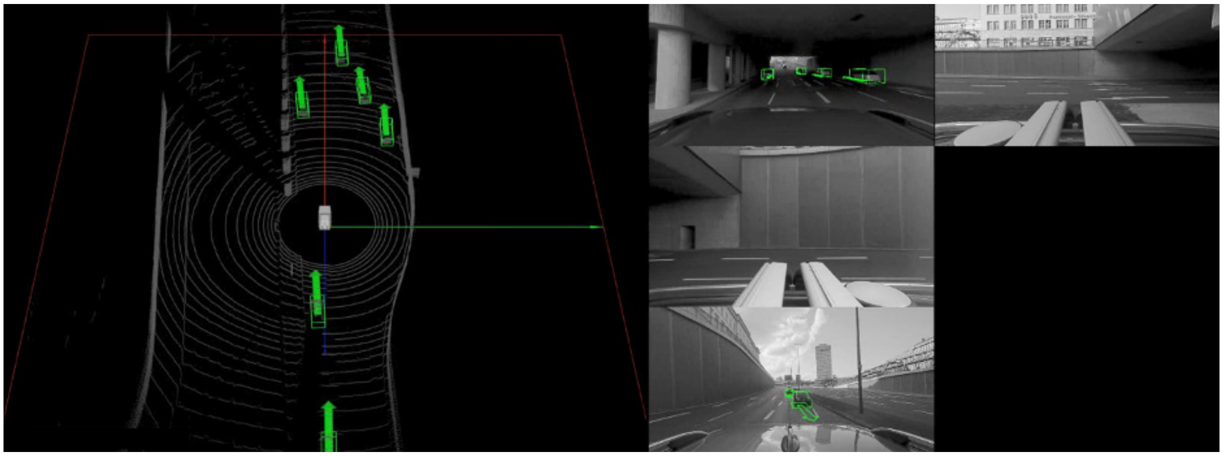

THE ROAD FROM HERE

As OEMs look to bring active-safety capabilities to their full range of vehicles, they will need sensors that are cost-effective and able to deliver data in challenging conditions, and the intelligence to get the most useful information from the data. They can achieve that through machine learning and a combination of sensors anchored by radar. Innovations such as Aptiv's RACam can package those sensors – in this case, radar and camera – into one compact unit.

Aptiv's Satellite Architecture centralizes the intelligence that receives data from those sensors, improving performance by keeping latency low and reducing sensor mass by up to 30%. OEMs can then develop differentiating features for various levels of automated driving on top of this robust base of sensing and perception technology, building from Level 1 automation to Level 2 and Level 2+.

Longer term, Aptiv's Smart Vehicle Architecture enables the overall vision by structuring the electrical and electronic architecture of a vehicle in a way that makes the most sense for its sensing and perception needs, creating a path to Level 3 and Level 4 automation. In the meantime, OEMs can take important steps today to help democratize active safety and ensure that everyone has access to these lifesaving technologies.

ABOUT THE AUTHOR

Rick Searcy
Advanced Radar Systems Manager

Rick manages the development of advanced radar systems for Aptiv, a position he has held since 2013. He has been involved in the development of every radar produced at the company since 1994.

Rick is located in Kokomo, Indiana. He earned his master's degree from the University of Michigan, where he studied applied electromagnetics and digital signal processing.

LEARN MORE AT APTIV.COM/ADVANCED-SAFETY →