

**Top Level Newsletter: Connected Vehicle**  
**(Published by IEEE Orange County Section)**

**November 2021**  
**Volume 23.0**

Editor-in-Chief: Kay Das, IEEE Life Member,  
Systems Research Development and Innovation

Vol 23 comprises:

(1) Autonomous Vehicles Should Start, Go Slow

In this issue we depart from the norm and include in full just one article which hopefully gives the reader an overview of the technology and its progress, covering many aspects of previous material in this newsletter. It is time for a re-cap of the technology, several aspects of which have been highlighted in previous issues of the newsletter.

**General Notes**

This series of newsletters is intended to provide the IEEE member with a top level briefing of the many different subjects relevant to the research, development and innovation of the connected vehicle.

A note on the Connected Vehicle Newsletter development: Volume 18 was a top-level synopsis of the 74 most recent entries from the twelve previous volumes (vol 17 to vol 6) since March 2020. Volume 19 returned to the usual format. Volume 20 departed from the norm and included a complete article.

The objective is to provide a platform for fast learning and quick overview so that the reader may be guided to the next levels of detail and gain insight into correlations between the entries to enable growth of the technology. Intended audiences are those that desire a quick introduction to the subject and who may wish to take it further and deepen their knowledge. This includes those in industry, academia or government and the public at large. Descriptions will include a range of flavors from technical detail to broad industry and administrative issues. A (soft) limit of 300 to 600 words is usually set for each entry, but not rigorously exercised.

As descriptions are not exhaustive, hyperlinks are occasionally provided to give the reader a first means of delving into the next level of detail. The reader is encouraged to develop a first level understanding of the topic in view. The emphasis is on brief, clear and contained text. There will be no diagrams in order to keep the publication concise and podcast-friendly. Related topics in the case of Connected Vehicle technology, such as 5G cellular and the Internet of Things will be included. The terms Connected Vehicle

and Automated Driving will be used inter-changeably. Articles from other published sources than IEEE that add to the information value will occasionally be included.

This newsletter forms part of the regional Advanced Technology Initiative (ATI) of which connected vehicles form a constituent part. Technical articles solely from IEEE journals/magazines are referred to by their Digital Object Identifier (DOI) or corresponding https link. The link for each article is provided. Those readers who wish to delve further to the complete paper and have access to IEEE Explore ([www.ieeexplore.ieee.org](http://www.ieeexplore.ieee.org)) may download complete articles of interest. Those who subscribe to the relevant IEEE society and receive the journal may already have physical or electronic copies. In case of difficulty please contact the editor at [kaydas@mac.com](mailto:kaydas@mac.com). The objective is to provide *top level guidance* on the subject of interest. As this is a collection of summaries of already published articles and serves to further widen audiences for the benefit of each publication, no copyright issues are foreseen.

Readers are encouraged to develop their own onward sources of information, discover and draw inferences, join the dots, and further develop the technology. Entries in the newsletter are normally either editorials or summaries or abstracts of articles. Where a deepening of knowledge is desired, reading the full article is recommended.

### **1. Autonomous Vehicles Should Start, Go Slow, by Shaoshan Liu et al**

Published: IEEE Spectrum Magazine (March 2020, pp 36-49)

MANY YOUNG URBANITES don't want to own a car, and unlike earlier generations, they don't have to rely on mass transit. Instead they treat mobility as a service: When they need to travel significant distances, say, more than 5 miles (8 kilometers), they use their phones to summon an [Uber](#) (or a car from a similar ride-sharing company). If they have less than a mile or so to go, they either walk or use various "micromobility" services, such as the increasingly ubiquitous [Lime](#) and [Bird](#) scooters or, in some cities, bike sharing.

The problem is that today's mobility-as-a-service ecosystem often doesn't do a good job covering intermediate distances, say a few miles. Hiring an Uber or [Lyft](#) for such short trips proves frustratingly expensive, and riding a scooter or bike more than a mile or so can be taxing to many people. So getting yourself to a destination that is from 1 to 5 miles away can be a challenge. Yet such trips account for [about half](#) of the total passenger miles traveled.

Many of these intermediate-distance trips take place in environments with limited traffic, such as university campuses and industrial parks, where it is now both economically reasonable and technologically possible to deploy small, low-speed autonomous vehicles powered by electricity. We've been involved with a startup that intends to make this form of transportation popular. The company, [PerceptIn](#), has autonomous vehicles operating at tourist sites in Nara and Fukuoka, Japan; at

an industrial park in Shenzhen, China; and is just now arranging for its vehicles to shuttle people around [Fishers](#), Ind., the location of the company's headquarters.

Because these diminutive autonomous vehicles never exceed 20 miles (32 kilometers) per hour and don't mix with high-speed traffic, they don't engender the same kind of safety concerns that arise with autonomous cars that travel on regular roads and highways. While autonomous driving is a complicated endeavor, the real challenge for PerceptIn was not about making a vehicle that can drive itself in such environments—the technology to do that is now well established—but rather about keeping costs down. Given how expensive autonomous cars still are in the quantities that they are currently being produced—an experimental model can cost you in the neighborhood of US \$300,000—you might not think it possible to sell a self-driving vehicle of any kind for much less. Our experience over the past few years shows that, in fact, it is possible today to produce a self-driving passenger vehicle much more economically:

PerceptIn's vehicles currently sell for about \$70,000, and the price will surely drop in the future. Here's how we and our colleagues at PerceptIn brought the cost of autonomous driving down to earth. Let's start by explaining why autonomous cars are normally so expensive. In a nutshell, it's because the sensors and computers they carry are very pricey. The suite of sensors required for autonomous driving normally includes a high-end satellite-navigation receiver, lidar (light detection and ranging), one or more video cameras, radar, and sonar. The vehicle also requires at least one very powerful computer. The satellite-navigation receivers used in this context aren't the same as the one found in your phone. The kind built into autonomous vehicles have what is called [real-time kinematic](#) capabilities for high-precision position fixes—down to 10 centimeters. These devices typically cost about \$4,000. Even so, such satellite-navigation receivers can't be entirely relied on to tell the vehicle where it is. The fixes it gets could be off in situations where the satellite signals bounce off nearby buildings, introducing noise and delays. In any case, satellite navigation requires an unobstructed view of the sky. In closed environments, such as tunnels, that just doesn't work.

Fortunately, autonomous vehicles have other ways to figure out where they are. In particular, they can use [lidar](#), which determines distances to things by bouncing a laser beam off them and measuring how long it takes for the light to reflect back. A typical [lidar unit for autonomous vehicles](#) covers a range of 150 meters and samples more than 1 million spatial points per second. Such lidar scans can be used to identify different shapes in the local environment. The vehicle's computer then compares the observed shapes with the shapes recorded in a high-definition digital map of the area, allowing it to track the exact position of the vehicle at all times. Lidar can also be used to identify and avoid transient obstacles, such as pedestrians and other cars. Lidar technology suffers from two problems. First, these units are extremely expensive: A high-end lidar for autonomous driving can easily cost more than \$80,000, although costs are dropping, and for low-speed applications a suitable unit can be purchased for about

\$4,000. Also, lidar, being an optical device, can fail to provide reasonable measurements in bad weather, such as heavy rain or fog.

The same is true for the cameras found on these vehicles, which are mostly used to recognize and track different objects, such as the boundaries of driving lanes, traffic lights, and pedestrians. Usually, multiple cameras are mounted around the vehicle. These cameras typically run at 60 frames per second, and the multiple cameras used can generate more than 1 gigabyte of raw data each second. Processing this vast amount of information, of course, places very large computational demands on the vehicle's computer. On the plus side, cameras aren't very expensive.

The radar and sonar systems found in autonomous vehicles are used for obstacle avoidance. The data sets they generate show the distance from the nearest object in the vehicle's path. The major advantage of these systems is that they work in all weather conditions. Sonar usually covers a range of up to 10 meters, whereas radar typically has a range of up to 200 meters. Like cameras, these sensors are relatively inexpensive, often costing less than \$1,000 each.

The many measurements such sensors supply are fed into the vehicle's computers, which have to integrate all this information to produce an understanding of the environment. Artificial neural networks and [deep learning](#), an approach that has grown rapidly in recent years, play a large role here. With these techniques, the computer can keep track of other vehicles moving nearby, as well as of pedestrians crossing the road, ensuring the autonomous vehicle doesn't collide with anything or anyone.

Of course, the computers that direct autonomous vehicles have to do a lot more than just avoid hitting something. They have to make a vast number of decisions about where to steer and how fast to go. For that, the vehicle's computers generate predictions about the upcoming movement of nearby vehicles before deciding on an action plan based on those predictions and on where the occupant needs to go. Lastly, an autonomous vehicle needs a good map. Traditional digital maps are usually generated from satellite imagery and have meter-level accuracy. Although that's more than sufficient for human drivers, autonomous vehicles demand higher accuracy for lane-level information. Therefore, special high-definition maps are needed. Just like traditional digital maps, these HD maps contain many layers of information. The bottom layer is a map with grid cells that are about 5 by 5 cm; it's generated from raw lidar data collected using special cars. This grid records elevation and reflection information about the objects in the environment.

On top of that base grid, there are several layers of additional information. For instance, lane information is added to the grid map to allow autonomous vehicles to determine whether they are in the correct lane. On top of the lane information, traffic-sign labels are added to notify the autonomous vehicles of the local

speed limit, whether they are approaching traffic lights, and so forth. This helps in cases where cameras on the vehicle are unable to read the signs.

Traditional digital maps are updated every 6 to 12 months. To make sure the maps that autonomous vehicles use contain up-to-date information, HD maps should be refreshed weekly. As a result, generating and maintaining HD maps can cost millions of dollars per year for a midsize city. All that data on those HD maps has to be stored on board the vehicle in solid-state memory for ready access, adding to the cost of the computing hardware, which needs to be quite powerful. To give you a sense, an early computing system that [Baidu](#) employed for autonomous driving used an Intel Xeon E5 processor and four to eight Nvidia K80 GPU accelerators. The system was capable of delivering 64.5 trillion floating-point operations per second, but it consumed around 3,000 watts and generated an enormous amount of heat. And it cost about \$30,000.

Given that the sensors and computers alone can easily cost more than \$100,000, it's not hard to understand why autonomous vehicles are so expensive, at least today. Sure, the price will come down as the total number manufactured increases. But it's still unclear how the costs of creating and maintaining HD maps will be passed along. In any case, it will take time for better technology to address all the obvious safety concerns that come with autonomous driving on normal roads and highways.

We and our colleagues at PerceptIn have been trying to address these challenges by focusing on small, slow-speed vehicles that operate in limited areas and don't have to mix with high-speed traffic—university campuses and industrial parks, for example.

The main tactic we've used to reduce costs is to do away with lidar entirely and instead use more affordable sensors: cameras, inertial measurement units, satellite positioning receivers, wheel encoders, radars, and sonars. The data that each of these sensors provides can then be combined through a process called sensor fusion. With a balance of drawbacks and advantages, these sensors tend to complement one another. When one fails or malfunctions, others can take over to ensure that the system remains reliable. With this sensor-fusion approach, sensor costs could drop to something like \$2,000. Because our vehicle runs at a low speed, it takes at the very most 7 meters to stop, making it much safer than a normal car, which can take tens of meters to stop. And with the low speed, the computing systems have less severe latency requirements than those used in high-speed autonomous vehicles.

PerceptIn's vehicles use satellite positioning for initial localization. While not as accurate as the systems found on highway-capable autonomous cars, these satellite-navigation receivers still provide submeter accuracy. Using a combination of camera images and data from inertial measurement units (in a technique called visual inertial odometry), the vehicle's computer further improves the accuracy, fixing position down to the decimeter level.

For imaging, PerceptIn has integrated four cameras into one hardware module. One pair faces the front of the vehicle, and another pair faces the rear. Each pair of cameras provides binocular vision, allowing it to capture the kind of spatial information normally given by lidar. What's more, the four cameras together can capture a 360-degree view of the environment, with enough overlapping spatial regions between frames to ensure that visual odometry works in any direction. Even if visual odometry were to fail and satellite-positioning signals were to drop out, all wouldn't be lost. The vehicle could still work out position updates using rotary encoders attached to its wheels—following a general strategy that sailors used for centuries, called dead reckoning. Data sets from all these sensors are combined to give the vehicle an overall understanding of its environment. Based on this understanding, the vehicle's computer can make the decisions it requires to ensure a smooth and safe trip.

The vehicle also has an anti-collision system that operates independently of its main computer, providing a last line of defense. This uses a combination of millimeter-wave radars and sonars to sense when the vehicle is within 5 meters of objects, in which case it's immediately stopped. Relying on less expensive sensors is just one strategy that PerceptIn has pursued to reduce costs. Another has been to push computing to the sensors to reduce the demands on the vehicle's main computer, a normal PC with a total cost less than \$1,500 and a peak system power of just 400 W.

PerceptIn's camera module, for example, can generate 400 megabytes of image information per second. If all this data were transferred to the main computer for processing, that computer would have to be extremely complex, which would have significant consequences in terms of reliability, power, and cost. PerceptIn instead has each sensor module perform as much computing as possible. This reduces the burden on the main computer and simplifies its design. More specifically, a GPU is embedded into the camera module to extract features from the raw images. Then, only the extracted features are sent to the main computer, reducing the data-transfer rate a thousand-fold.

Another way to limit costs involves the creation and maintenance of the HD maps. Rather than using vehicles outfitted with lidar units to provide map data, PerceptIn enhances existing digital maps with visual information to achieve decimeter-level accuracy. The resultant high-precision visual maps, like the lidar-based HD maps they replace, consist of multiple layers. The bottom layer can be any existing digital map, such as one from the [OpenStreetMap](#) project. This bottom layer has a resolution of about 1 meter. The second layer records the visual features of the road surfaces to improve mapping resolution to the decimeter level. The third layer, also saved at decimeter resolution, records the visual features of other parts of the environment—such as signs, buildings, trees, fences, and light poles. The fourth layer is the semantic layer, which contains lane markings, traffic sign labels, and so forth.

While there's been much progress over the past decade, it will probably be another decade or more before fully autonomous cars start taking to most roads and highways. In the meantime, a practical approach is to use low-speed autonomous vehicles in restricted settings. Several companies, including [Navya](#), [EasyMile](#), and [May Mobility](#), along with PerceptIn, have been pursuing this strategy intently and are making good progress.

Eventually, as the relevant technology advances, the types of vehicles and deployments can expand, ultimately to include vehicles that can equal or surpass the performance of an expert human driver. PerceptIn has shown that it's possible to build small, low-speed autonomous vehicles for much less than it costs to make a highway-capable autonomous car. When the vehicles are produced in large quantities, we expect the manufacturing costs to be less than \$10,000. Not far in the future, it might be possible for such clean-energy autonomous shuttles to be carrying passengers in city centers, such as Manhattan's central business district, where the average speed of traffic now is only 7 miles per hour. Such a fleet would significantly reduce the cost to riders, improve traffic conditions, enhance safety, and improve air quality to boot. Tackling autonomous driving on the world's highways can come later. (2477 words)

**DOI: none**