

MIT SLOAN SCHOOL OF MANAGEMENT
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)



MODULE 4 UNIT 1
Interactive Video Set Video 1 Part 2 Transcript

Module 4 Unit 1 Interactive Video Set Video 1 Part 2 Transcript

DANIELA RUS: So, all autonomous systems today work in two phases. A first phase maps the environment; a second phase plans a path, once we have a map, and then executes the path from a starting location to a goal location. Now, the map is not what we usually think of when we look at Google Maps – the map is actually created in feature space, and as we follow along in the map, we track the features that we observe now against the features in the feature space.

Given a map, the vehicle can drive through the space, localizing in feature space, and then the vehicle can also follow roads, find drivable and not-drivable surfaces, and identify pedestrians and other vehicles that were not mapped at the time when the map was created.

Most autonomous systems are created using the following architecture: We start with a wheel-based vehicle, this is the robot platform. We extend it with sensing. The most commonly-used sensors today are laser scanners for distance and cameras for objects. And from these sensors, from the data extracted with these sensors, we are able to create maps and we are able to detect objects in the world.

With these sensors, we are also able to estimate the location of the vehicle on the road, which is very important for controlling the vehicle to go from point A to point B. And sometimes the sensors can be interfaced with the infrastructure.

Now, if we have a map, you have the ability to detect obstacles, and we have the ability to estimate our location in the map. The next thing we need is a planning module that allows us to figure out the trajectory the car needs to follow from point A to point B, and when an obstacle is encountered, the same system allows us to modify the trajectory to go around the obstacle.

So let's look at the details of each of these modules. Map-making is done in feature space, and for most vehicles it's done in the feature space extracted with a laser scanner. And here, in this image, you see some example laser scans that get aggregated over time, and that get processed for features that are being associated with specific locations.

The chart on the right is what a feature looks like for a particular defining point on the map. Each point in the map has different features associated with it. Given a map, a planning and control algorithm is able to compute a trajectory from the starting location to a desired goal location for the vehicle. That system is then also able to issue commands to the motors of the car to drive the vehicle along the desired trajectory, while constantly monitoring the location and the presence of other obstacles that are not included in the map. If an obstacle is found, this module is then able to issue a command that allows the car to go around the obstacle.

So the most important component that enables obstacle detection, localization, and map-making is the perception system of the car. In our system, perception is performed by looking for features in an urban environment, and the features consist of the curvature of

the curb, and the texture that we see on the buildings surrounding the road. And it turns out that, for urban settings, looking for these two types of features and aggregating them as a feature vector is actually sufficient to identify uniquely most locations in the world.

So here, we see a picture and we see example curb features and how these get mapped onto the feature space. We also see textures and different surfaces that comprise the buildings along the side of the road. By extracting the normals of these surfaces and combining the normals with the curvatures extracted from the curb, we can create these three-dimensional feature vectors that can then all be projected in a simple, two-dimensional feature vector to make the computation for localization and for creating maps, efficient and real-time.

It turns out that this approach performs very well. In fact, we have experimented with this approach to localization, mapping, and planning in urban settings, and we note that the performance of the system is actually very accurate. We get on the order of a few centimeters of accuracy when driving through a campus.

Here is an example of using the system, and what you see on the left is the on-board camera view of the autonomous car. What you see on the right is the set of features extracted from our laser scanner. You can kind of see how the features align, and you can see that even for situations where we don't have a lot of texture, for instance in the case of this road that is flanked by trees on both sides, the system has a rich enough set of features and it's performing quite well.

Now, given the map, given the capability to localize and the capability to detect obstacles, we're now ready for planning paths from start to goal, and this is done using a traditional waypoint planning and navigation system. The only modification is that we need to use something we call a dynamic virtual bumper.

We need to create a much bigger virtual bumper than what your car has, and this is so that if an obstacle shows up and makes contact with this virtual bumper, the car actually has enough time to stop or to respond to create an alternative trajectory that avoids the obstacle.

So here is the person. When the person is detected at the limit of the virtual bumper, the car is able to either stop, or to compute a path that will go around the person so there is no collision. We have implemented this detector, and here you see an example of running in autonomous mode in a system that is able to detect people in real time, and adjust the trajectory so as to avoid them. And this is the speed at which the perception and control technologies that are available today can operate in safe modes.

Another thing we can do with autonomous driving is to import traffic rules in the brains of our cars. So here's an example where an autonomous car is driving on a road with a solid lane marking, and the car encounters an obstacle. So, at this point, the question is what should the car do?

According to the solid line, the car should really stop, the car should not cross the line, but the set of traffic rules that are captured by the Vienna Convention actually enable cars to locally break rules. And identifying when it is possible to break a rule in order to ensure

safe and continuous passage is now possible to incorporate as part of the brain of an autonomous car.

So you have seen some autonomous golf carts. Here is what an autonomous drive-by-wire car looks like. It is our Toyota Prius that has been extended with some laser systems, a GPS system, some cameras, and motor encoders.

Now, there are two different ways of thinking about autonomous driving. One way is called series autonomy, and it is what companies like Google and Uber are pursuing.

The idea behind series autonomy is that you have a car and either the human is in control or the autonomous system is in control, but you can't have both human and autonomy operating at the same time. We have made extraordinarily exciting progress in series autonomy, but until we are able to deploy safe systems in all conditions, we will not be able to have series autonomy on public roads.

For this reason, we have begun to explore an alternative way to autonomy, which we call parallel autonomy. The idea behind parallel autonomy is that the human is in control, but there is also an autonomy system that is able to check the sensors of the car and is able to run in parallel with a human, kind of like a guardian angel, looking over the shoulders of the human.

If the parallel autonomy system notices, perhaps children running on a cross street that are outside of the field of view of the human, but the autonomy system can know about it because the autonomy system can network with sensors that are installed on that piece of the road, the autonomy system can correct the maneuvers of the human to ensure that driving is safe.

So the guardian angel, or parallel autonomy system, is sort of like an anti-lock brake system on steroids, where the autonomy system has access to information from all the sensors on the car to sensors on other cars, if the vehicles are networked, and to sensors installed in the infrastructure of the city, also through network communication. And this additional information allows the parallel autonomy system to know much richer information about the state of the road than the human can get with the naked eye, and for this reason, can truly act as a guardian angel and prevent accidents from happening.

So the autonomy system is implemented as a shared-control setup between human input, how the human wants to drive, and a series autonomy system. The shared controller, then, will issue low-level commands that will either allow the car to execute what the human wants to do, or will compute small modifications to what the human wants, but modifications that are safe in the current context.

So here is an example of parallel autonomy and how it can help. So, you can see the speed in bars, and the steering angle in arrows. And in blue, you see the decision of the autonomy system, and in red, you see the decision of the human. So for the race track, you see that the human driver wants to accelerate just before a steep curb, and that is not safe. However, the parallel autonomy system intervenes and makes the speed correction.

In the case of the T intersection, you see that the red car arrives at the T intersection and the driver wants to keep going, but because the autonomy system sees the cars coming from the left and from the right, the autonomy system prevents the human from entering the intersection until the intersection is safe. These new capabilities for improving human decision-making in autonomy will greatly improve transportation, will greatly improve the safety of transportation.

So here is an example of implementing a parallel autonomy system, and you see the human trying to turn the wheel too far to the left and too far to the right, but the car keeps steady forward because too far to the left there is woods, and too far to the right there are other obstacles.

In this video, you see another example of parallel autonomy with a human attempting to turn too far to the right, the human attempting to go too fast, and the human attempting to run over a pedestrian. In each of these cases, the parallel autonomy system takes over and makes small corrections that are consistent with the intentions of the driver, yet keep the car safe.

So there is extraordinary potential for parallel autonomy to improve transportation, even in the short run, when the cars would not be fully autonomous. And there will be a gradual change from where we are today, with anti-lock brakes that are helping us prevent skidding on roadways, to a point where situation awareness that is gathered through sensors mounted on the car will help human drivers even when the events detected by the car are not in the field of view of the human.