Robust Perception for Formula Student Driverless Racing

Gustaf Broström and David Carpenfelt
Supervisor: Björn Olofsson

Popular Science Summary of Master's Thesis
Broström, G., & Carpenfelt, D. (2021). *Robust Perception for Formula Student Driverless Racing.* TFRT-6151. ISSN 0280–5316.

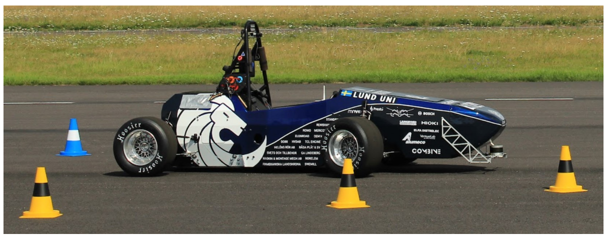

Figure 1: First self-driving car built by Lund Formula Student

Self-driving cars developed for racing require fast and accurate perception to stay on the track and not crash.
For a self-driving race car to follow a track made from blue and, yellow cones (see Figure 1), it first needs to detect the cones location and color. To accurately measure the location of an object, a LiDAR sensor is the perfect tool. A LiDAR uses laser rays and the known constant of the speed of light to range objects with millimeter precision. The color of the cones are of course most easily seen with a camera, but can also be calculated using the reflected light of the LiDAR.

The task of detecting cone-like objects using only the distance measurements requires some calculations. First of all the majority of laser rays of course hit the ground, and these have to be removed with some algorithm that can detect what measurements are actually from the ground. Since the ground is never perfectly flat, it can be seen as several interconnected flat surfaces and removed. After the ground has been removed, the remaining laser rays have either hit the cones or other objects in the surroundings. To distinguish between what is what, you first need to know if laser rays hit the same object, and this is done using a clustering algorithm.

When an object consisting of laser rays have been found, it needs to be determined whether or not it is a cone. Since the dimensions of a cone are known beforehand, it is easy to figure out if the object could possibly be a cone, by measuring how close the laser rays hit to a perfect model of a cone. When the object is known to be a cone, it is also possible to guess the color of the cone by just using the intensity of the reflected laser ray; dark colors reflect less than bright ones just as for sunlight. Yellow cones with black strips in the middle therefore have a distinct pattern from the blue ones with a white strip in the middle.

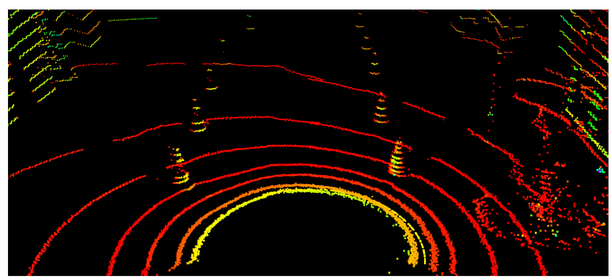

Figure 2: Laser rays from the LiDAR hitting the surroundings

To be sure that the cones are actually cones, we can have a look in the camera image at the locations where the LiDAR detected them. Since the placement of the LiDAR and camera relative to each other is known, it is easy to know where to look in the image produced by the camera. If we find cones here we know that the LiDAR was correct and the range is therefore easily extracted. The color is also easy to determine since the camera can see color. If the cone was outside of the camera's field of view, we can use the color estimate based on the reflected intensity of the laser.

The cone detection method with the LiDAR is fast enough to scan the environment 50 times per second. Together with the camera, the car saw 65% of cones each scan within 10 meters and detected cones that were not actually there less than 0.1% of the time. This Master Thesis was done by members of the Lund Formula Student 2021 team.