3D Laser Based Obstacle Detection for Autonomous Driving

Dirk Hähnel Sebastian Thrun

Stanford Artificial Intelligence Lab, Stanford University

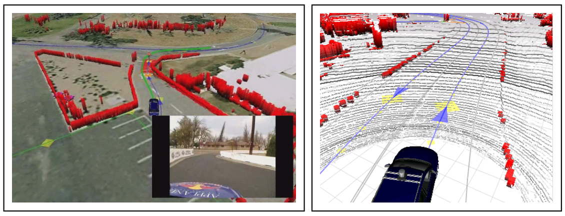

Figure 1: *Obstacle detection: Here the vehicle detects vegetation, a road sign, and also a small curb.*

In urban environments, the vehicle encounters a wide variety of static and moving obstacles. Obstacles as small as a curb may trip a fast-moving vehicle, so detecting small objects is of great importance. Overhangs and trees may look large obstacles at a distance, but traveling underneath is often possible. Thus, obstacle detection must consider the 3D geometry of the world. Figure 1 depicts a typical output of the obstacle detection routine in an urban environment. Each red object corresponds to an obstacle. Towards the bottom right, a camera image is shown for reference.

Our robot's primary sensor for obstacle detection is the Velodyne laser. A simple algorithm for detecting obstacles in Velodyne scans would be to find points with similar x-y coordinates whose vertical displacement exceeds a given threshold. Indeed, this algorithm can be used to detect large obstacles such as pedestrians, signposts, and cars. However, range and calibration error are high enough with this sensor that the displacement threshold cannot be set low enough in practice to detect curb-sized objects without substantial numbers of false positives.

An alternative to comparing vertical displacements is to compare the range returned by two adjacent lasers, where "adjacency" is measured in terms of the pointing angle of the sensor. Each of the 64 lasers has a fixed pitch angle, and thus would sweep out a circle of a fixed radius on a flat ground plane as the sensor rotates. Sloped terrain locally compresses these rings, causing the distance between adjacent rings to be smaller than the inter-ring distance on flat terrain. In the extreme case, a vertical obstacle causes adjacent beams to return nearly equal ranges. Because the individual beams strike the ground at such shallow angles, the distance between rings is a much more sensitive measurement of terrain slope than vertical displacement. By finding points that generate inter-ring distances that differ from the expected distance by more than a given threshold, obstacles that are not apparent to the vertical thresholding algorithm can be reliably detected in this fashion.

In addition to terrain slope, rolling and pitching of the vehicle will cause the rings traced out by the individual lasers to compress and expand. If this is not taken into account, rolling to the left can cause otherwise flat terrain to the left of the vehicle to be detected incorrectly as an obstacle. This problem can be remedied by making the expected distance to the next ring a function of range, rather than the index of the particular laser. Thus as the vehicle rolls to the left, the expected range difference for a specific beam decreases as the ring moves closer to the vehicle. Implemented in this way, small obstacles can be reliably detected even as the sensors rolls and pitches.

Two more issues must be addressed when doing obstacle detection in urban terrain. First, trees and other objects frequently overhang safe driving surfaces and should not be detected as obstacles. Overhanging objects are filtered out by comparing their height with a simple ground model. Points that fall in a particular x-y grid cell that exceed the height of the lowest detected point in the same cell by more than a given threshold (the height of the vehicle plus a safety buffer), are ignored as overhanging obstacles.

Second, the Velodyne sensor possesses a "blind spot" behind the vehicle. This is the result of the sensor's geometry and mounting location. Further, it also cannot detect small obstacles such as curbs in the immediate vicinity of the robot due to self-occlusion. Here additional IBEO and SICK LDLRS sensors are used to supplement the Velodyne data. Because both of these sensors are essentially 2-D, ground readings cannot be distinguished from vertical obstacles, and hence obstacles can only be found at very short range (where ground measurements are unlikely). Whenever either of these sensors detects an object within a 10 meter range, the measurement is flagged as an obstacle. This combination between short-range sensing in 2-D and longer range sensing using the 3-D sensor provides high reliability.