Improving a Background Model for LiDAR 3D Point Clouds with Machine Learning

Gustav Nilsson and Seamus Doyle

You are sitting on a train, reading a book. You enter a tunnel. Suddenly, the train screeches to a halt and you feel a bump. Did we hit something... or someone?

A situation like this could have been caused by anything from people on the tracks to forgotten equipment. Whatever the cause, it might have been prevented if the problem was detected early enough. To do this however, a high precision sensor would be required, a sensor capable of seeing at far distances and preferably in darkness. Something like a LiDAR.

A LiDAR (Light Detection and Ranging) works by scanning its surroundings with a laser and collects points by measuring the time it takes for the laser to travel to an object and back. This produces a point cloud, a collection of points with x, y, z coordinates, see Figure 1. In our thesis, we worked with LiDAR point clouds for the purpose of detecting and tracking objects. More specifically, we worked on improving something called the *Background Model*.

Figure 1: *Point cloud of Grenden in Lund*

For detection and tracking to be possible, a background model is essential. The sheer size these point clouds alone, makes it impossible for conventional methods to work without it. The background model's main aim is to filter out all uninteresting points. We wondered whether it would also be possible to include knowledge of the environment into this model. In addition, could we include more in the background model than just this filter?

Thus, our entire thesis consisted of improving a very simple background model and this was done in mainly two ways. In the first, we used deep learning to gain information of the background. With this model, we were able to predict what every point in a point cloud of 45000 points represented. Such an example can be seen in Figure 2. This was also possible to do really fast, with speeds of up to 0.032 seconds per point cloud. This information was encoded into the simple background filter to make it more effective. Issues of wind in vegetation was successfully resolved in this way.

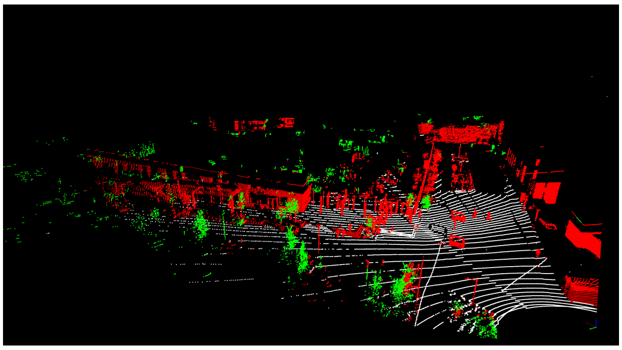

Figure 2: *Deep learning model's guess of what point cloud at Grenden in Lund contains. White is ground, green is vegetation and red is other.*

The second way in which we improved the background model, was to use machine learning models to find the ground. This could partly work as a filter for ground points, but also as a way of making the whole system more intelligent. Tracking and classification stages could become smarter by also knowing an object's height compared to the ground. An advantage of the model was its ability to predict height of ground in locations without any points. As long as these areas were not too sparse in points, the performance was quite good.

We trained the deep learning model on point clouds simulated using a game engine. On these point clouds the model achieved excellent results. On point clouds generated by real LiDAR recordings the model captured the environment correctly on locations that were relatively similar to the training. This was most cases. The ground model, on the other hand, did not have to be trained on any data, but could be fit directly to any point cloud.

Together, these models form the combined background model, a step towards a smarter and more robust LiDAR system so that you one day, never will have to experience a train accident like the one from before.