



Deep Neural Networks for Real-time 3D Object Detection

Using deep neural networks (DNNs) for 3D data instead of 2D images is challenging, because processing 3D data needs more computation time and memory than 2D data. Also usual 3D data formats, like point clouds, are not ideally suited for processing on GPUs.

There exist some DNN models for 3D data in robotics, such as VoxelNet, Vote3Deep and Frustum PointNets. The task of this master thesis is first to install, test and improve the Frustum PointNets algorithm. It should be compared to the other algorithms mentioned above using existing benchmarks. In this thesis, we focus on outdoor benchmarks such as KITTI and Sydney urban object data. These benchmarks are using LiDAR and RGB-D models.


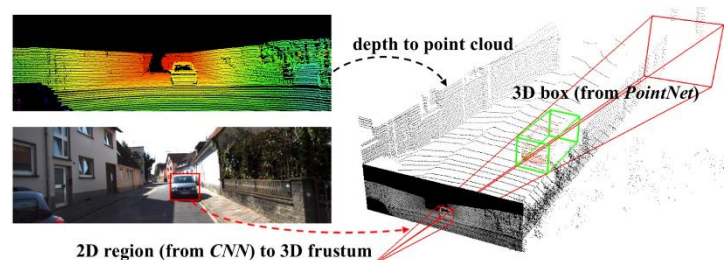
The optimized Frustum PointNets algorithm should then be implemented and tested on one of our outdoor robots, either on the Summit XL or on one of our self-built Quadrocopters.

Requirements:

- TensorFlow and ROS
- Python and C++
- Proficiency in English

References:

1. Frustum PointNets: 3D object detection, CVPR 2018
2. Vote3Deep: fast object detection in 3D point clouds, ICRA 2017
3. VoxelNet: end-to-end learning, arXiv 2017.



Kontakt

Ya Wang
Sand 1, Raum A305
Tel. (07071) 29-78987
francis.wang@uni-tuebingen.de