Vote3Deep: Fast Object Detection in 3D Point Clouds Using Efficient Convolutional Neural Networks

M. Engelcke, D. Rao, D. Z. Wang, C. H. Tong, I. Posner (ARXIV 2016)

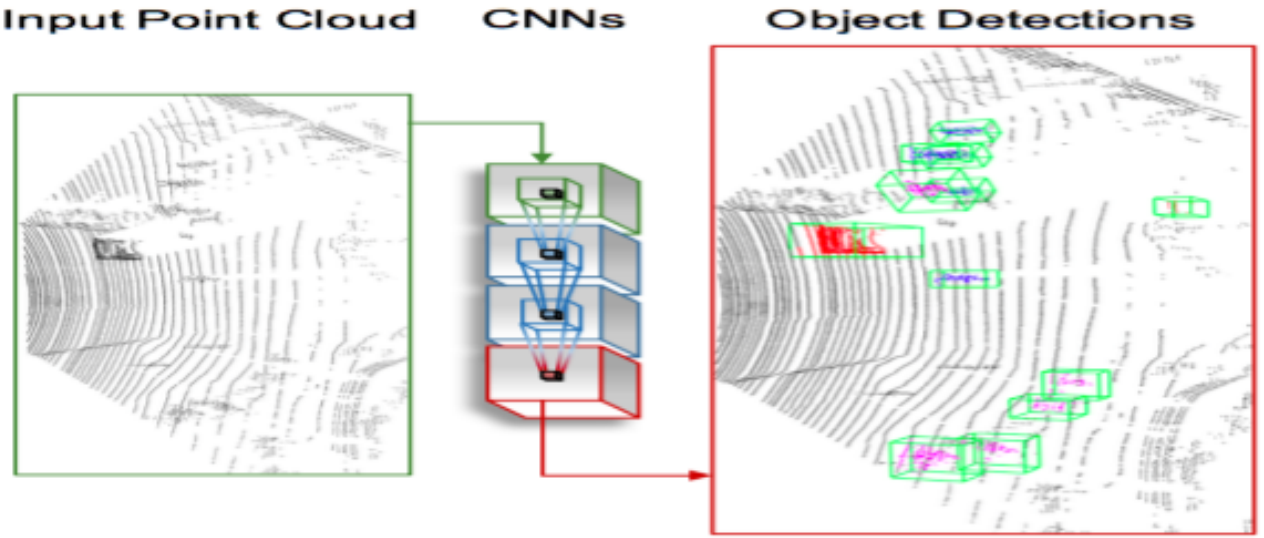

- Proposes a computationally efficient approach to detecting objects natively in 3D point clouds using convolutional neural networks
- Contributions:
  - Construction of efficient convolutional layers as basic building blocks for CNN-based point cloud processing by leveraging a voting mechanism exploiting the inherent sparsity in the input data
  - The use of rectified linear units and an L1 sparsity penalty to specifically encourage data sparsity in intermediate representations in order to exploit sparse convolution layers throughout the entire CNN stack
  - First work to propose sparse convolutional layers and L1 regularisation for efficient large-scale processing of 3D data
- Evaluates on KITTI object detection benchmark