מכון ויצמן למדע

THE WEIZMANN INSTITUTE OF SCIENCE

FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Vision and Robotics Seminar

Room 1 ,Ziskind Building
on Thursday, Mar 28, 2019
at 12:15

Yizhak Ben-Shabat
Technion

Classification, Segmentation, and Normal Estimation of 3D Point Clouds using Deep Learning

Abstract:

Modern robotic and vision systems are often equipped with a direct 3D data acquisition device, e.g. a LiDAR or RGBD camera, which provides a rich 3D point cloud representation of the surroundings. Point clouds have been used successfully for localization and mapping tasks, but their use in semantic understanding has not been fully explored. Recent advances in deep learning methods for images along with the growing availability of 3D point cloud data have fostered the development of new 3D deep learning methods that use point clouds for semantic understanding. However, their unstructured and unordered nature make them an unnatural input to deep learning methods. In this work we propose solutions to three semantic understanding and geometric processing tasks: point cloud classification, segmentation, and normal estimation. We first propose a new global representation for point clouds called the 3D Modified Fisher Vector (3DmFV). The representation is structured and independent of order and sample size. As such, it can be used with 3DmFV-Net, a newly designed 3D CNN architecture for classification. The representation introduces a conceptual change for processing point clouds by using a global and structured spatial distribution. We demonstrate the classification performance on the ModelNet40 CAD dataset and the Sydney outdoor dataset obtained by LiDAR. We then extend the architecture to solve a part segmentation task by performing per point classification. The results here are demonstrated on the ShapeNet dataset. We use the proposed representation to solve a fundamental and practical geometric processing problem of normal estimation using a new 3D CNN (Nesti-Net). To that end, we propose a local multi-scale representation called Multi Scale Point Statistics (MuPS) and show that using structured spatial distributions is also as effective for local multi-scale analysis as for global analysis. We further show that multi-scale data integrates well with a Mixture of Experts (MoE) architecture. The MoE enables the use of semi-supervised scale prediction to determine the appropriate scale for a given local geometry. We evaluate our method on the PCPNet dataset. For all methods we achieved state-of-the-art performance without using an end-to-end learning approach.