

Building Foundational Models for 3D Point Cloud Segmentation

Motivation

Three-dimensional mapping of the environment (e.g., for infrastructure or vegetation monitoring) is becoming increasingly common in the geospatial domain, generating large volumes of 3D point cloud data. In order to extract relevant information from 3D point clouds, the raw data must be enriched with semantic information. A fundamental step in this process is semantic segmentation, where each point is assigned to a semantic class. Based on PointNet as one of the first deep learning (DL) architectures for 3D point cloud processing, several DL architectures for 3D point cloud segmentation have been developed (e.g. KP-FCNN, PointTransformer). However, the limited availability of annotated data is a major issue in training such DL models for domain-specific segmentation tasks. To address this issue, methods for pre-training DL models have recently become a focus of research. Pre-training can be done either in a self-supervised way or by supervised training on large, automatically generated training datasets. The goal of these methods is to create generic models (foundational models) that can be applied to new datasets either without further training (zero-shot learning, promptable models) or after fine-tuning on a small amount of annotated data (few-shot learning). In particular, in the image domain, a promptable segmentation model called SegmentAnything has been introduced, which can solve a wide range of image segmentation tasks based on user-defined prompts such as bounding boxes or point queries.

Description

The goal of this project is to investigate how pre-training methods and promptable architectures can contribute to the creation of foundational models for 3D point cloud segmentation. For this purpose, the state-of-the-art will be surveyed and a suitable approach will be developed, implemented and evaluated. Depending on your specific interests, the project can cover the following topics:

- (1) Automatic generation of large annotated point cloud datasets from virtual laser scanning.
- (2) Implementation of state-of-the-art DL architectures for 3D point cloud segmentation and their extension towards a promptable design.
- (3) Development of self-supervised pre-training methods for 3D point cloud segmentation models.
- (4) Integration of the point cloud segmentation models into a point cloud labeling tool to assist humans in point cloud labeling.

Tech Stack

Python, PyTorch, (C++, Blender)

What you should bring


Participation in this project requires a willingness to familiarize yourself with current research, a basic understanding of deep learning (not necessarily in the point cloud domain) and basic Python programming skills.

Classification

The master project is directly related to the ongoing research and software development projects of the HPI Computer Graphics Systems group. It forms a basis for a subsequent master thesis or a future doctoral thesis. In the context of the master's project, there are generally opportunities to work as a student assistant at HPI or as a software developer at research partners.

Contact

Prof. Dr. Jürgen Döllner (doellner@hpi.de), Josafat-Mattias Burmeister (josafat-mattias.burm@hpi.de)



Human-annotated 3D point cloud (left), point cloud segmentation result from KP-FCNN (middle) and image segmentation result from SegmentAnything (right).