WASP Project Course 2023

Enhancing NeRF-based reconstruction of automotive scenes

Background

In recent years, deep learning and computer vision have made significant strides, particularly in the area of 3D scene reconstruction. Neural Radiance Fields (NeRF) models have become state-of-the-art, providing excellent results in static scene reconstruction. However, their application in dynamic environments such as automotive scenarios presents a challenge due to the high level of dynamism and complexity inherent in these environments.

To address this challenge, this project aims to apply and enhance NeRF models for automotive scene reconstruction utilizing the Zenseact Open Dataset (ZOD). The integration of LiDAR data in ZOD can be beneficial in improving depth estimation and scene fidelity. Moreover, this project intends to develop strategies to deal with dynamic objects and improve 3D object annotations. The motivation behind obtaining a high-quality reconstruction is its potential applications. These include refining noisy annotations, generating novel views for data augmentation, creating rare training examples, and synthesizing views from new perspectives, such as generating synthetic side/back cameras from front camera and surround lidar data. Furthermore, we aim to explore the possibility of generating “free” semantic labels using a “language-field” consistent with CLIP features.

Constraints: Preferably Gothenburg to enable in-person collaboration

Participants

Industrial partner: Zenseact

Industrial supervisor: Christoffer Petersson, christoffer.petersson@zenseact.com

Academic supervisor: Lennart Svensson , Lennart.svensson@chalmers.se, Chalmers University of Technology

Coordinating WARA representative: -

Suggested WASP PhD students: Adam Tonderski, Georg Hess, William Ljungbergh, Adam Lilja

Challenges to investigate

- Understanding the limitations of existing SOTA NeRF models when applied to the automotive setting, particularly the ZOD.
- Determining the optimal way to incorporate LiDAR data to enhance depth estimation and scene fidelity.
- Exploring strategies to handle dynamic objects like pedestrians and cars, including their rigidity assumptions, parametrization for accurate reconstruction, and scene decomposition into independent parts.
- Examining the potential applications of NeRFs that are relevant to the automotive industry, such as refining noisy annotations, generating training data from novel views, manipulating the scene to create extremely rare but highly valuable training examples, generating synthetic side/back cameras from front camera and surround lidar data, and generating "free" semantic labels through a "language-field" consistent with CLIP features.

Resources

The project will utilize the Zenseact Open Dataset (ZOD) for the development and testing of the enhanced NeRF model. Open source software will be used throughout the project, such as PyTorch for general deep learning operations, and NerfStudio for nerfs in particular. As for computational resources, we plan on using personal computers, given the relative lightness of NeRFs, as well as the Berzelius cluster.

Deliverables

- A NeRF model capable of reconstructing automotive scenes with improved performance over current SOTA (on ZOD). This model includes LiDAR data and can handle dynamic objects in the scene. The quality of the model is verified with thorough evaluations and an ablation study.
- One or more of the following:
    - Decomposition of scene into separate subparts that can be manipulated (moved, added, removed)
    - NeRFs as noisy label refinement (e.g. refine 3d bounding boxes or semantic segmentation predictions)
    - Novel camera views in unseen directions (e.g. back-camera from front-camera + surround lidar)
    - Zero-shot labels from weak language supervision
    - …

References

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
https://arxiv.org/abs/2003.08934

Panoptic NeRF: 3D-to-2D Label Transfer for Panoptic Urban Scene Segmentation
https://arxiv.org/abs/2203.15224

UniSim: A Neural Closed-Loop Sensor Simulator
https://openaccess.thecvf.com/content/CVPR2023/papers/Yang_UniSim_A_Neural_Closed-Loop_Sensor_Simulator_CVPR_2023_paper.pdf

Nerflets: Local Radiance Fields for Efficient Structure-Aware 3D Scene Representation from 2D Supervision
https://arxiv.org/abs/2303.03361

Keywords

Deep Learning, Computer Vision, PyTorch, Weakly Supervised Learning, Neural Radiance Fields (NeRF), Image Reconstruction, Automotive Scene Reconstruction, LiDAR Reconstruction.