

Sim-to-Real Domain Adaptation of Infrastructure Sensor Lidar Point Clouds using Generative Adversarial Networks

Amarin Kloeker, Yangyi Liu, Mischa Maier, Lutz Eckstein

The scarcity of publicly available datasets for infrastructure sensors has become a significant bottleneck in the development and deployment of advanced perception algorithms in the context of automated driving. In this paper we investigate the potential to overcome this challenge by leveraging the power of generative adversarial networks (GANs) to adapt lidar point clouds generated in simulation environments to the real-world domain. We propose a proof-of-concept architecture specifically designed to bridge the domain gap between simulated and real-world point cloud data, facilitating more accurate and robust perception algorithms for infrastructure sensors. Preliminary results on synthetic and real-world datasets show promising evidence that the proposed approach could improve the performance of object detection algorithms using the adapted point clouds. Our approach demonstrates the ability to significantly approximate simulated point clouds to the real-world domain, improving the average precision (AP) of a state-of-the-art object detection network from 20.37 % to 42.58 % solely based on low-feature point clouds from a simulation environment. This exploratory research contributes to the ongoing discussion on leveraging the potential of simulation environments for the development of perception systems in infrastructure sensor road side units and highlights the potential benefits of using GANs for domain adaptation in this context.

Full document at <https://doi.org/10.1109/ICECET58911.2023.10389244>