Evaluating The Placement Of Roadside Cameras In Urban Intersections Towards Enhanced Collaborative Perception

Master's Thesis of Mohammed Othman

Mentoring:
Dr.-Ing. Mathias Pechinger
Mario Ilic, M.Sc.

Motivation

Current autonomous vehicle (AV) systems are constrained by the limitations of single-vehicle perception, as sensor ranges and fields of view are often hindered by occlusions—static or dynamic obstructions that restrict environmental awareness. The emerging integration of Cooperative Autonomous Vehicles (CAVs) and Vehicle-to-Everything (V2X) technologies offers a pathway to enhance situational awareness through collaborative perception (CP), which enables AVs to receive and share data from roadside cameras and other infrastructure. By addressing occlusion challenges, these technologies advance the goal of Vision Zero[1]—eliminating road fatalities—while paving the way for more reliable autonomous mobility systems in complex real-world environments.

Study Objectives

This study addresses research gaps by using synthetic data from a 3D virtual environment to assess roadside camera positioning, focusing on static and dynamic occlusions effects. Through data generated in Unity, the research aims to enhance road safety and VRUs detections (Figure 1).

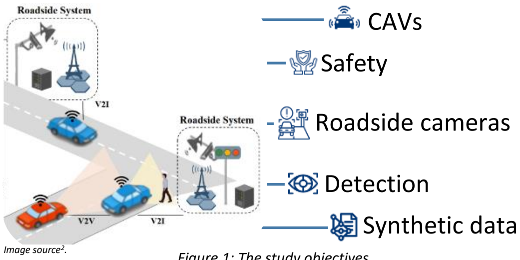

Figure 1: The study objectives.

Research Question

With the consideration of VRUs, how can synthetic data generated in a 3D environment be used to evaluate the placement of roadside cameras regarding the effects of occlusions?

Methodology

The study's evaluation process has two phases: a preliminary phase to assess the method's viability and an extensive phase to conduct a thorough evaluation with increased simulation replications (Figure 2). Two simulation-based methods are used to examine synthetic data from Unity for roadside camera placement analysis: Method 1: Record then Detect by YOLO, and Method 2: Dual-Camera.

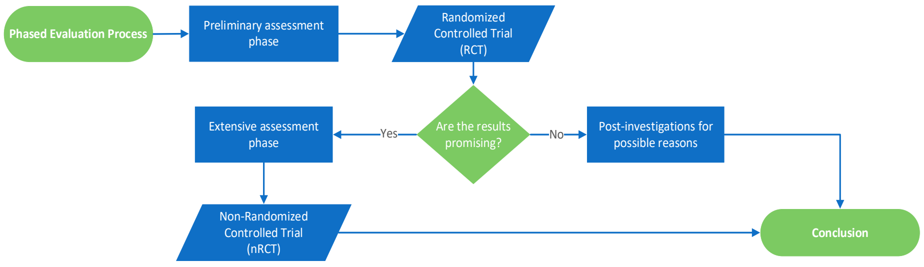

Figure 2: The evaluation approach of the study.

Method 1: Concept

YOLOv5n's pre-trained model is employed to process Unity-captured videos, leveraging its capabilities in accurately detecting and classifying vehicles in real-world environments.

Method 1: Results

The results showed that the pre-trained model largely failed to detect vehicles in the videos from both deployment positions, with an aggregated average of 74.09% of vehicles missing entirely. For the few detected vehicles, class assignment accuracy was also low, with an overall accuracy of just 14.29%. This underscores a key limitation of relying on a pre-trained model, which struggled to detect synthetically generated vehicles even without occlusion effects in the simulated scenarios.

Method 2: Concept

This method aims to isolate and analyse specific objects, such as vehicles and bicycles, that are impacted by dynamic and static occlusion. Using a 'Dual-Camera' setup—one camera capturing a normal roadside view and the other focusing solely on highlighted target objects—the approach enables frame-by-frame pixel analysis (Figure 3). Target objects are assigned a bright colour for clear differentiation and placed under a separate layer in Unity. This setup allows precise tracking and evaluation of roadside camera placement with respect to occlusion effects.

The total average occlusion percentage metric is developed to evaluate the effectiveness of each roadside camera placement per simulation run. It ranges between 0% and 100% indicating totally visible and totally non-visible targeted objects, respectively.

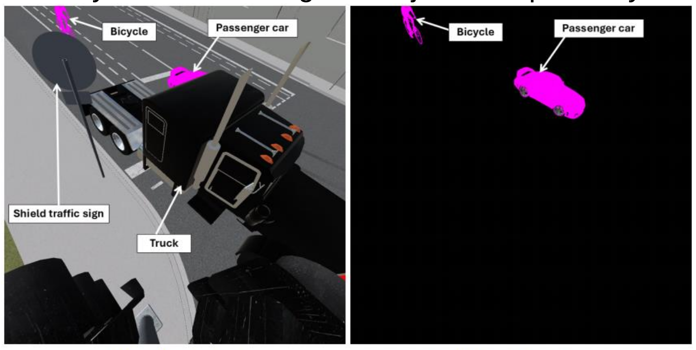

(a) Normal-view camera (b) Revealing camera

Figure 3: A snapshot showing typical views of the dual-camera.

Method 2: Results

In both preliminary and extensive phases (Figures 4 and 5) ,12 and 30 simulation runs were conducted per roadside camera position, respectively. Camera position (B) consistently recorded lower median occlusions of 2.45% and 4.02% for preliminary and extensive phases respectively, compared to position (A), which showed higher susceptibility to occlusions of 24.75% and 18.06%. The results confirm the dual-camera method's effectiveness in evaluating occlusion impacts and highlight the limitations of side-of-road camera placements.

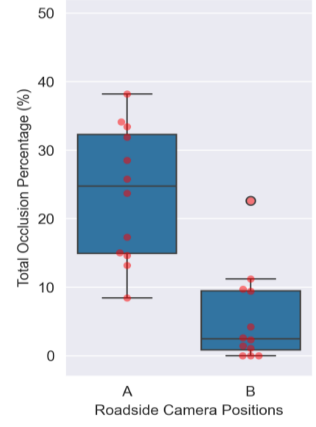

Figure 4: Preliminary phase results.

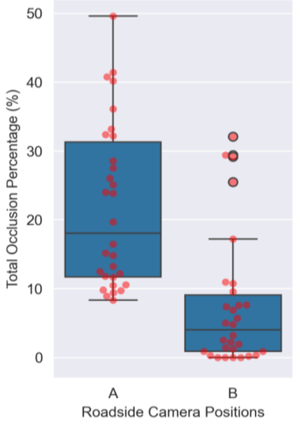

Figure 5: Extensive phase results.

Conclusion:

This study demonstrated the superiority of the dual-camera method in evaluating roadside camera placements, particularly regarding occlusion impacts in traffic monitoring. Using synthetic data from a 3D virtual environment with the Dual-Camera method effectively capturing both static and dynamic occlusions. Findings provide insights for optimizing camera placements to improve vehicle and VRU detection accuracy. The study also fills the research gaps by considering the 3D space, both occlusion types, and VRU.

References
[1] Stadtrat der Landeshauptstadt München. (2019, March). Beschluss zum Verkehrssicherheitskonzept "Vision Zero." https://muenchenunterwegs.de/visionzero
[2] Liu, Y., Wang, Z., Zhou, X., & Zheng, L. (2023). A Study of Using Synthetic Data for Effective Association Knowledge Learning. Machine Intelligence Research, 20(2), 194–206. https://doi.org/10.1007/s11633-022-1380-x