Vision-Based Algorithms for Obstacle Detection and Avoidance in Autonomous Vehicles

Ryan Diaz[1], *Yevgeniy Vorobeychik*[2]

1. University of Minnesota - Twin Cities, 2. Washington University in St. Louis

University of Minnesota

Washington University in St. Louis

INTRODUCTION

- Self-driving car technology has made much progress over the last few years, but development still needs to be done to ensure safety and robustness.
- This project aims to develop self-driving algorithms that can be deployed on a small-scale real-world autonomous vehicle environment.

RESEARCH QUESTION

How can we design vision-based algorithms to achieve safe and robust obstacle avoidance behaviors for autonomous vehicle agents?

METHODS

- **Environment:** We use the CARLA [1] simulation environment to train and test our vision-based algorithms.
- **Obstacle Detection:** Our object detection framework is YOLOv8 [2]. We first pretrained our model using a large annotated dataset collected in CARLA, which was then fine-tuned in the real-world environment *(Figure 1)*.
- **Obstacle Avoidance:** We utilize and build upon the "Roach" framework first established by Zhang et al. [3] *(Figure 2)*:
  1. An expert reinforcement learning policy is trained for route following and obeying basic traffic laws.
  2. We use the expert RL policy to collect demonstrations of the vehicle exhibiting obstacle avoidance behavior.
  3. The collected data is used to train a policy that takes in RGB image data as input through imitation learning.

DETECTION

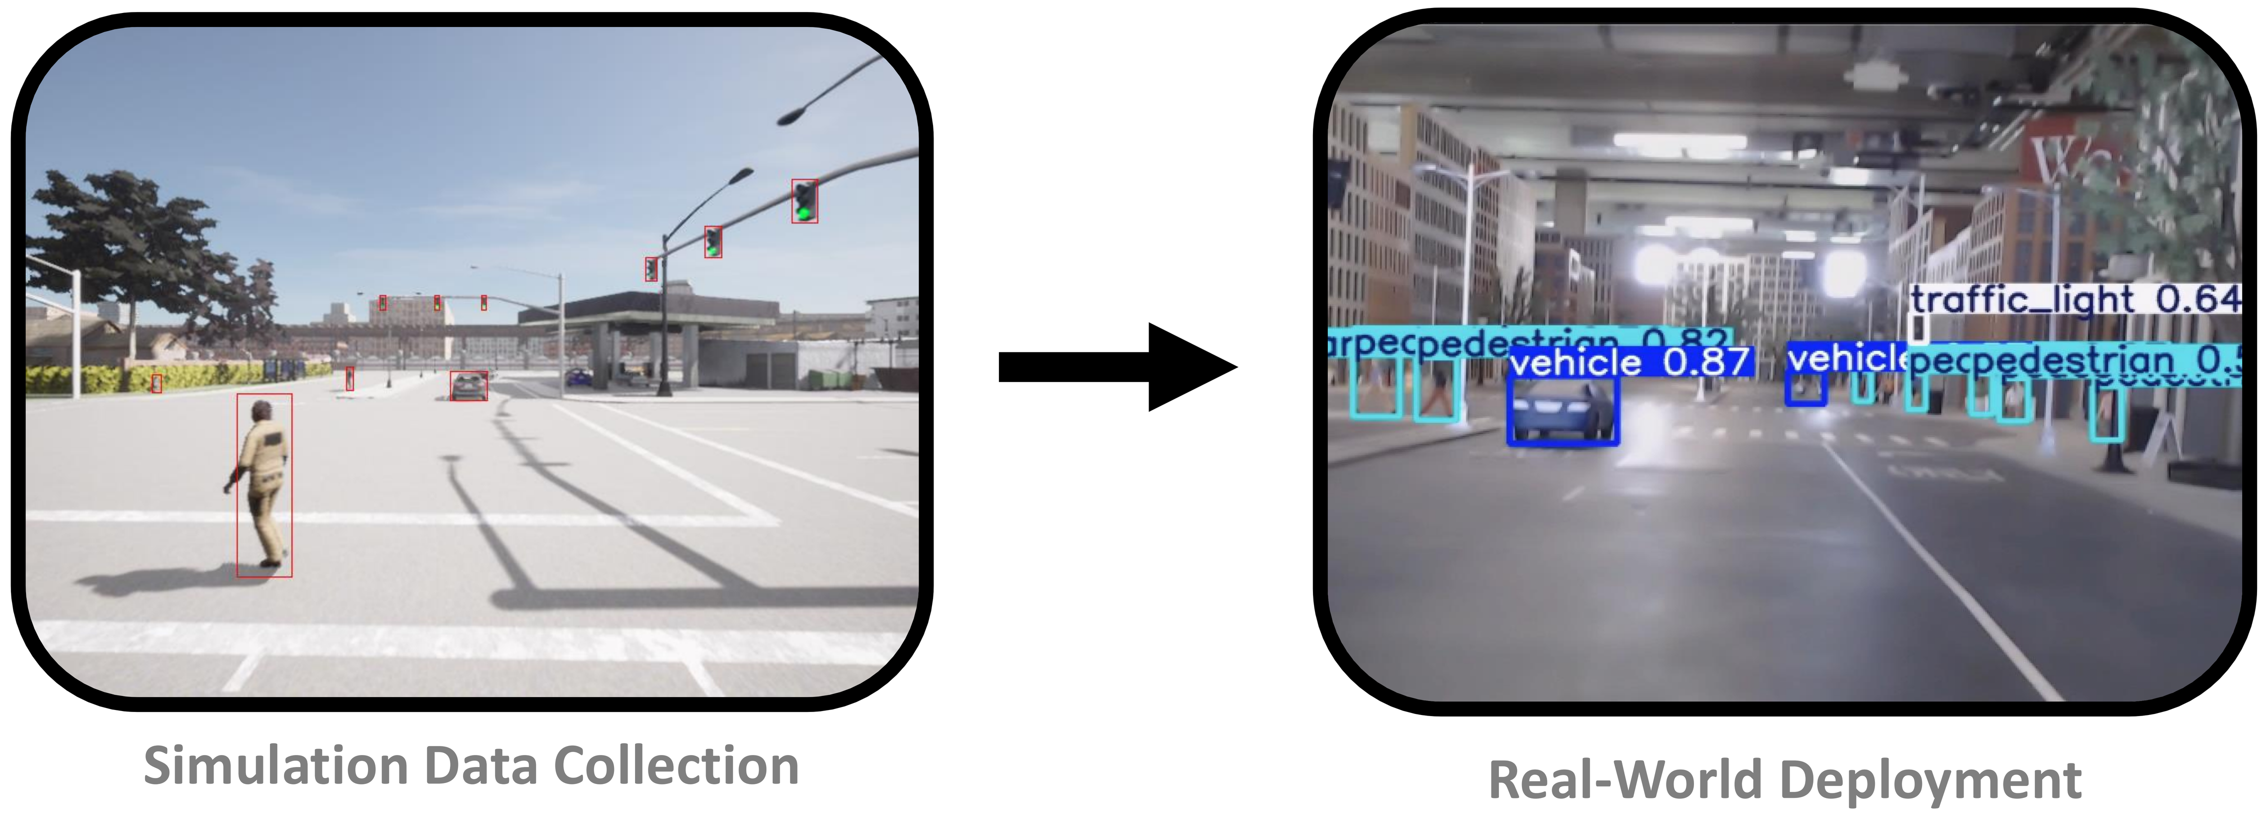

Figure 1: Pipeline for training our obstacle detection model.

AVOIDANCE

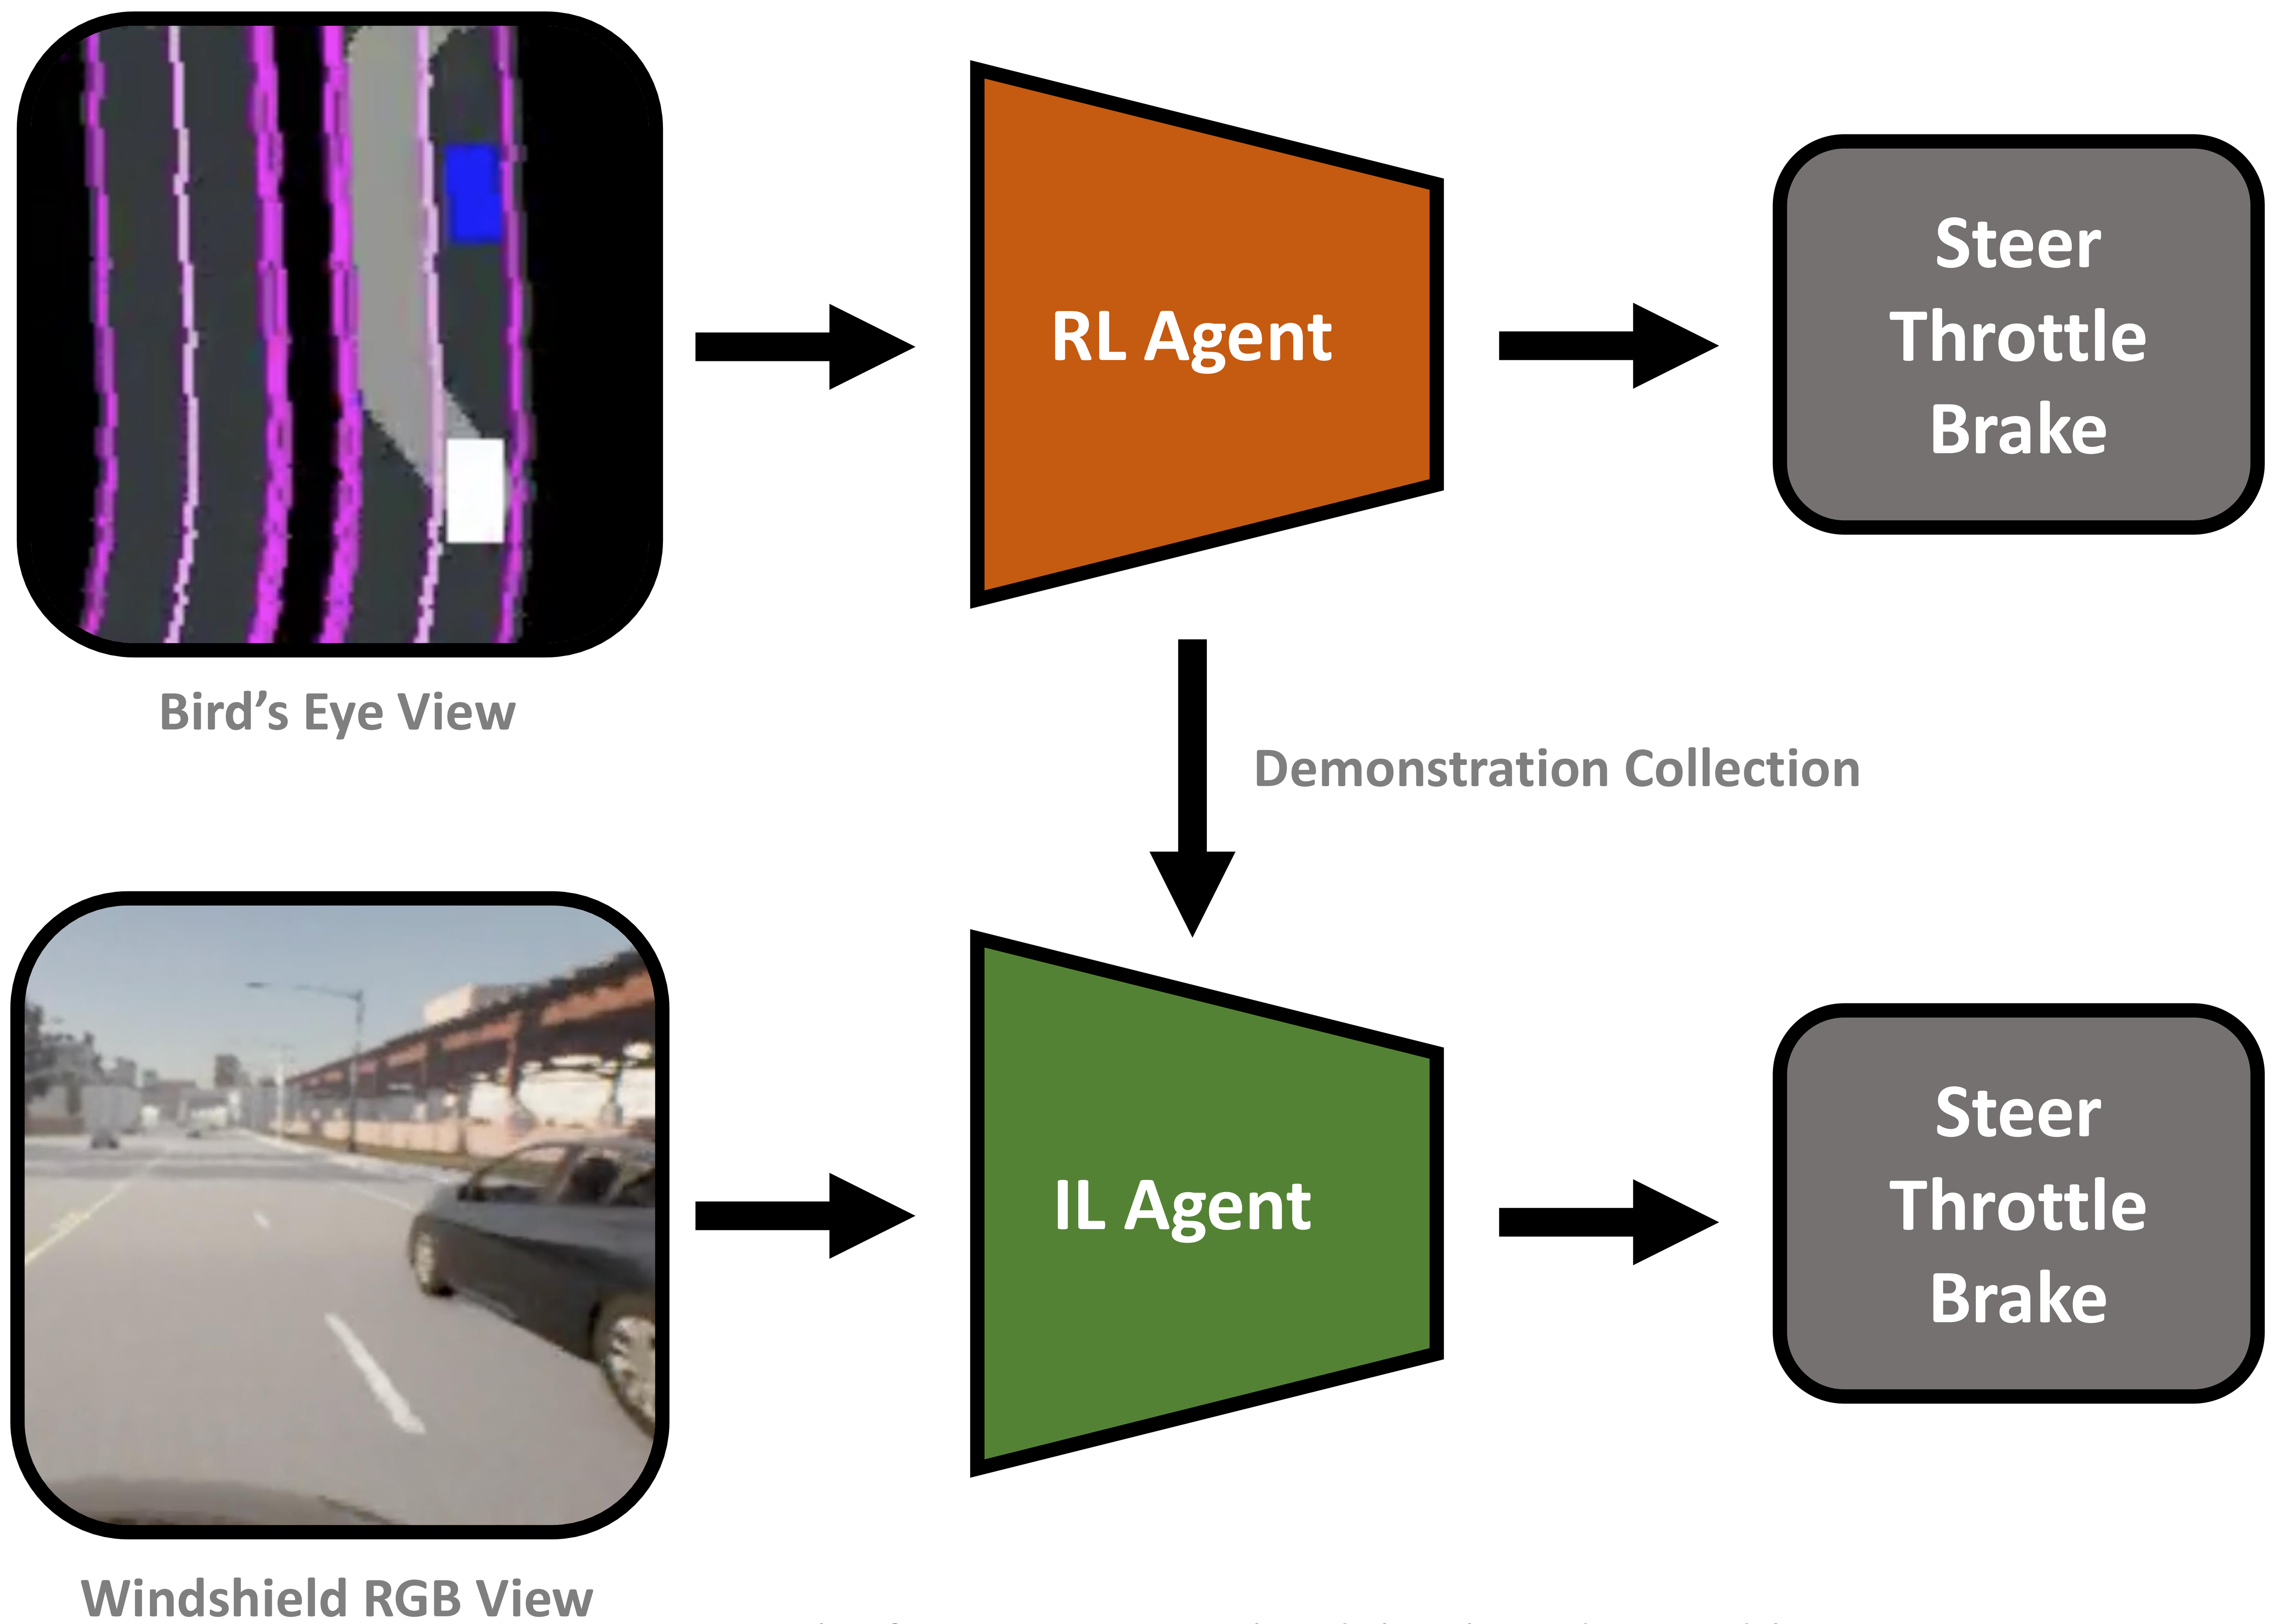

Figure 2: Pipeline for training our image-based obstacle avoidance model.

RESULTS AND DISCUSSION

- **Obstacle Detection Model:**
  - Pretrained YOLOv8 model achieved **0.742** mAP@50 on CARLA simulation data.
  - Fine-tuned YOLOv8 model achieved **0.886** mAP@50 on real-world data.
- **Obstacle Avoidance Model:**
  - Expert RL agent achieved a max cumulative reward value of **9944.709** during policy rollouts.
    - Qualitatively, it can sufficiently follow a route, change lanes, and stop at stop signs and stop lights at reasonable speeds.
  - IL agent is still a work in progress, but the data collection pipeline has been set up to train it.
- **Takeaways:**
  - Simulations can be an effective tool for collecting large datasets for real-world models.
  - Reward design and hyperparameter tuning for RL can be difficult and time consuming without expert domain knowledge.

FUTURE WORK

- **Real World Evaluation:** We hope to deploy our trained algorithm into the real-world setup to evaluate its sim-to-real transferability.
- **Multimodality:** We plan to build agents that incorporate additional input modes other than images, such as LIDAR.

REFERENCES

[1] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., vol. 78. PMLR, 13-15 Nov 2017, pp. 1-16.

[2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR),* 2016, pp. 779-788.

[3] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.