Proceedings of the 10th North American International Conference on Industrial Engineering and Operations Management, Orlando, Florida, USA, June 17-19, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA10.20250340
Published: June 17, 2025

AI-Driven Traffic Flow Optimization at Intersections Enhancing Efficiency and Safety

Kayle Peña Agront, Gaspar Chayer and Rawa Adla
Department of Electrical, Computer and Cybersecurity Engineering
Florida Polytechnic University
Lakeland, Florida, USA
{kpenaagront, gchayer, radla}@floridapoly.edu

Victoria Correa Andrade Doyon
Software Engineer
Microsoft
Atlanta, GA, USA
victoriadoyon@outlook.com

Abstract

Recent advancements have introduced solutions such as sensor-based monitoring and adaptive traffic signals, utilizing cameras, radars, and other sensors to collect and analyze real-time data for optimizing traffic flow. While these technologies offer promising improvements, they face challenges related to data accuracy, high infrastructure costs, and scalability, particularly in complex urban environments. By addressing the limitations of current traffic management solutions, this study investigates the integration of Artificial Intelligence (AI) as a key component in Vehicle-to-Everything (V2X) communication; a connected driving technology aimed at achieving smoother and safer traffic flow. This paper proposes a novel approach to traffic optimization, not only at signalized intersections but also in non-signalized scenarios. In such settings, vehicles would operate as independent agents, communicating through a shared cloud network to enhance coordination and efficiency.

Keywords
Traffic Optimization, Artificial Intelligence (AI), Vehicle-to-Everything (V2X), Intelligent Transportation Systems (ITS), Connected Vehicles

1. Introduction
As the number of cars on the road increases year after year, the need for greater efficiency and safety in traffic control mechanisms becomes exorbitantly prevalent. Traditional methods of traffic control struggle to adjust to the increasing number of vehicles congesting the roadways, resulting in bottlenecks, accidents, and inefficient usage of roads. Intelligent Transportation Systems (ITS) are traffic management technologies that use data and communication networks to enable a safe, coordinated, and smart traffic flow, which can be altered to work in tandem with new technologies capable of managing these changes. According to the National Highway Traffic Safety Administration, 94% of vehicle accidents are caused by human error (National Highway Traffic Safety Administration, 2015), which this study aims to minimize.

With the rise in traffic volume, implementing ITS in intersections is key to improving safety and efficiency. Intersections are points where two or more roads cross each other, and more specifically, this research targets signalized-intersections, which incorporate changing lights to signal drivers, although alterations can tailor it for non-signalized intersections as well. Intersections can accommodate vehicles, pedestrians, and cyclists, although they often differ in complexity. Different levels of complexity include roundabouts, multi-lane interchanges, or five-way intersections. Observing this intricacy solidifies the fact that interactions are a point of high traffic volume, where implementing ITS is ideal to reduce congestion and accidents.

In recent years, the integration of AI within ITS has been a major area of research. By leveraging machine learning, real-time data processing, and predictive analytics, AI enables smarter and more dynamic traffic management. Technologies such as V2X (vehicle-to-everything) and collaborative perception allow vehicles to share information from their sensors with other vehicles within connected vehicle networks, road infrastructure, and the cloud. Combining AI with these existing networks opens new possibilities for traffic management that show promising results in outperforming traditional signal methods. This paper explores how artificial intelligence and connected vehicle networks can be combined to optimize traffic flow at intersections. By creating a reinforcement learning model and training it in a simulation-based environment, the model was proven to be capable of reducing accidents at four-way intersections.

1.1 Objectives

The objective of this research is to achieve safer and more efficient intersections by leveraging AI and Intelligent Transportation Systems by developing a multiagent reinforcement learning model capable of deciding paths of aversion from accidents at four-way intersections and then validating the results utilizing a simulation-based environment. The developed simulation demonstrates how leveraging AI and connected vehicle networks in tangent can produce safer, smarter maneuvers compared to traditional signalized intersection methods.

2. Literature Review

While this paper aims to analyze the deployment of AI in intersections, the techniques that were implemented before these new developments pave the pathway for a streamlined deployment of AI into these existing solutions. These strategies fall under ITS, which, as introduced above, are traffic management technologies that use data and communication networks to enable a safe, coordinated, and smart traffic flow. Arterial ITS strategies that have been shown to improve vehicular throughput or reduce vehicular delay are adaptive signal control and traffic signal interconnection (Transportation Research Board, 2010). In areas where ITS has not been implemented, traffic platoons created by upstream signalized intersections can affect the operation of an access point by increasing the capacity of its traffic movements (Hummer and Ott 2015). Different methods of adaptive traffic control have been able to target common delays seen at intersections, like platoons, bottlenecks, etc.

In the past decade, digital technology introduced non-AI-based augmentation such as adaptive traffic control systems and networked infrastructure elements. These adaptive systems, although not an instance of machine learning, were designed to react to real-time traffic inputs via embedded sensors and dynamically alter signal plans (Skabardonis et al. 2003). Other instances of non-AI implementation are SPaT (Signal Phase and Timing) and MAP (Map Data) messages into the connected vehicle framework. SPaT messages allow infrastructure to inform vehicles of real-time signal status and upcoming phase changes, and MAP messages provide vehicle-mounted systems with rich digital maps of intersections, lane geometry, and allowed maneuvers (U.S. Department of Transportation, 2020). These V2X systems make vehicles and their users more informed in their decision-making, allowing for smoother traffic flow, safer turning maneuvers, and more efficient driving behavior.

Following these advancements, improvements were seen and scenarios that previously would have caused delays were now much more optimized, however, with the increase of road users in recent years, they now lack the flexibility and scalability to manage denser urban traffic flow with the vast variety of users all taking their actions and making decisions which the current traditional and rule-based systems were not designed to manage. The reliance on pre-programmed scenarios and static logical rules limits their ability to account for unpredictable human actions, sudden changes in traffic conditions, and the increasing heterogeneity of intersection geometries (Chen et al. 2024). Studying these limitations paved the way for the introduction of AI, with V2X having built a solid foundation that allows for AI implementation without causing a widespread overhaul of current traffic systems (Olayode et al. 2020). AI can bring data-driven flexibility and deep learning to the forefront of traffic management innovation and optimization, improving greatly upon the current methodology.

Artificial intelligence offers a promising alternative to enhance traffic control at intersections. Moving away from the traditional systems, AI finds efficient paths through drawbacks by allowing space for adaptive and intelligent real-time decision-making and presenting predictions as well. By utilizing real-time data from sensors, cameras, and vehicles, especially those with V2X capabilities, AI can optimize traffic flow, predict congestion patterns, and actively adjust signal timings to minimize delays and the probability of accidents (Miftah et al. 2025; Chen et al. 2024). The integration of AI in ITS is crucial to the transition from static, inefficient traffic control methods to dynamic, adaptive, and optimized strategies that improve both safety and efficiency at intersections (Olayode et al. 2020). AI systems in combination with V2X communication allow for connectivity that enhances situational awareness, enabling vehicles and traffic systems to collaboratively respond to changing conditions, reducing congestion and improving safety (Christopoulou et al. 2023).

Several different methods of AI implementation at intersections have shown an improvement in efficiency and safety for road users when tested and simulated. One of the implementations utilizes AI as the "controller" of traffic signal systems, making decisions based on dynamic traffic light adaptation to real-time vehicle flow, resulting in the avoidance of unnecessary stops and delays (Olayode et al. 2020). Collision warning systems use sensor fusion and computer vision to detect and classify objects in the vehicle's vicinity and give early warnings to drivers or take control with automated braking (Chen et al. 2024). Furthermore, AI-based systems that utilize V2X also maximize vehicle acceleration and braking patterns to minimize fuel usage and emissions (Miftah et al. 2025)

While this streamlined and optimized future is promising, the application of AI and connected vehicle technologies is not without its challenges. These technologies rely greatly on accurate, real-time sensors, cameras, and vehicle data. It also counts on the traffic lights working and does not account for weather adaptation or disruptions. Poor decision-making emerges when inconsistencies arise, such as noise or discrepancies in the data (Chen et al. 2024). The same occurs when faulty sensors, mislocated cameras, or unfavorable weather conditions such as rain or fog affect input quality. Failures such as these are capable of sending cars into congestion or creating flawed tracking, resulting in accidents and leaving room for human error once again. The high quantity of generated information requires quality and consistency of inputs from different system agents to ensure proper performance (Miftah et al. 2025). This dependency on live, ongoing, streams of car data—such as speed and location to provide rapid traffic-pattern adjustment—increases the required throughput time. These systems must process and react to this information almost in real-time, which demands tremendous computational power (Souri et al. 2024).

While these challenges persist, continued development and refinement of AI models and vehicle network communications offer promising results when considering the possibility of teaching these models to overcome the challenges through deep learning algorithms. Deployment of AI-based ITS is a crucial first step toward making intersections intelligent, adaptive, and connected with vehicles and infrastructure, especially in minimizing human error—responsible for 94% of traffic accidents (National Highway Traffic Safety Administration, 2015)—and improving roadway safety and efficiency.

3. Methods

This experiment was conducted utilizing a Proximal Policy Optimization (PPO) multiagent model, a form of reinforcement learning with a shared policy network. This algorithm was ideal for the simulated environment due to its scalability and ability to graphically represent the agents, with edges modeling agent interactions. The underlying model architecture consisted of a two-layer multilayer perceptron (MLP) with 256 units per layer and tanh activation functions. The actor and critic networks were implemented separately. Based on the actions taken, the model will receive a reward or a penalty. A reward entails positive reinforcement and encourages the model that its decisions are correct and accurate, while a penalty will punish the model and negatively reinforce that the actions taken were random and inaccurate. The model was trained using a learning rate of 5e-5, a discount factor (γ) of 0.99, and a minibatch size of 128. No reward shaping was applied. The reward function was designed to encourage safe navigation through the intersection by positively reinforcing forward movement and penalizing both collisions and off-track behavior. Penalties were applied symmetrically to agents involved in a collision to discourage unsafe interactions. The PPO model works well with a continuous action space, which the simulated environment utilizes, allowing the agents to take any number of actions and not be limited to a discrete set.

The dataset used to train the model consists of simulated accidents which take place in the CARLA Simulator software at four-way intersections with four cars (agents) navigating at once. This dataset was not developed as part of this study but was sourced from the DeepAccident benchmark (Wang et al. 2023), a publicly available dataset designed to evaluate accident prediction and traffic behavior in V2X-enabled environments. It includes a variety of simulated intersection crash scenarios with different driving patterns, numbers of vehicles, and types of collisions.

As stated above, the DeepAccident dataset utilizes the CARLA simulation environment (Dosovitskiy et al. 2017) to create high-fidelity four-way intersection scenarios. CARLA is an open-source urban driving simulator that supports dynamic traffic conditions, realistic physics, and various sensor configurations. These simulated environments include realistic vehicle behavior, environmental randomness, and high-stakes collision conditions, making it ideal for training reinforcement learning models to handle safety-centered tasks. Using CARLA and DeepAccident ensures that the data reflects a wide range of real-world driving complexities while maintaining repeatability and safety in training. Figure 1. Illustrates the agent-environment relationship, along with reward and state changes after actions are taken.

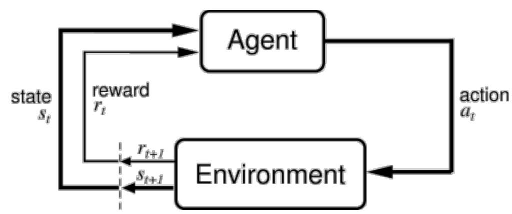

Figure 1. RL Agent and Environment Relationship

The PPO model and DeepAccident dataset were both chosen with the desire to use a continuous action space, however, for this preliminary experiment, options of the agents were limited. For future iterations of this experiment, there is capability to enforce a continuous action space, but for the present, agents were given four actions to take: left turn, right turn, forward, or remain in place. In order to prevent accidents, the model advises the agents on which action to take and is subsequently rewarded when the agents safely pass through the intersection. The results are then split into node-based output graphs displaying different averaged model behaviors, such as randomness of choices, number of rewards, and data loss.

4. Data Collection

After one hundred iterations—each iteration consists of running the entire dataset of one hundred and one CARLA crash simulations once, meaning the model has attempted to resolve over ten thousand crashes—observations and conclusions could be drawn by utilizing several output values that were collected and modeled as graphs. An *agent* refers to one of the four vehicles navigating through the intersection in the accident simulation data. Six graphs were generated, covering a variety of statistics. Average Reward: Overall success rate of agent, Policy Loss: Measures agent's behavior compared to optimal behavior, Entropy: Measures randomness in agent's actions, Value Function Loss: How well agent is estimating value of states, Total Loss: Combined loss from multiple sources, Episode Length: How long agents take to complete an episode.

5. Results and Discussion

While results are preliminary, the model excelled in terms of accuracy, quickly learning under the reward-penalty reinforcement learning system and improving upon itself through one hundred iterations. In the *Average Reward* graph (Figure 2.), it can be observed that when rewards begin to drop, the model corrects the behavior within very few iterations and had minor drops going forward. *Policy Loss* (Figure 4.) indicates how well the agent's actions match the optimal behavior. *Value Function Loss* (Figure 5.) reflects how well the agent is estimating the value of states, with a decreasing trend modeling an improvement in state prediction accuracy. The *Policy Loss* and *Value Function Loss* values can be seen averaged together under *Total Loss* (Figure 6.), which ideally wants to be seen decreasing over time. The graph remains steady through most iterations, with a significant drop at the tail-end to show that the model has improved its learning. *Entropy* (Figure 3.) measures the randomness of the actions taken, with the most desirable result being a downward trend. Our model relatively models this, with a few outliers spiking the data, but not enough to negate the fact that the model does show improvement in its randomness, not once returning to its highest value at the start. *Episode Length* (Figure 7.) depicts the time an agent takes to complete a task, a downward trend implying that it is becoming more efficient at completing its task. The model is not as efficient as it could be, however, for a preliminary model, the episode length does lower at a few points, leaving room for optimization.

5.1 Results

From this experiment, six output graphs were generated, modeling the results in different metrics, as stated in Results & Discussions. Through these results, the model can be seen efficiently detecting collisions (Average Reward, Figure 2.) more confidently (Entropy, Figure 3.) and accurately (Value Function Loss, Figure 5.) showing improvement overall in learning (Total loss, Figure 6.). The model's decision-making needs to be refined (Policy Loss, Figure 4.) as well as the time it takes to resolve a given accident and correctly direct all agents (Episode Length, Figure 7.).

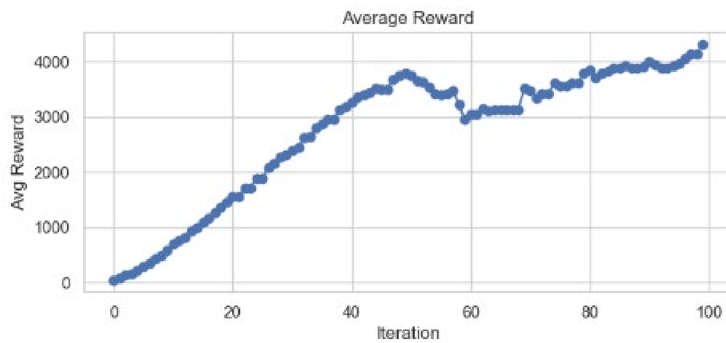

Figure 2. Average Reward Over 100 Iterations

This plot models the progression of the agent's average reward over the course of one hundred iterations. The upward trend demonstrates that the agent is effectively learning to improve its performance within the environment. A slight decline shown around iteration sixty may suggest a shift in dynamics, but the recovery and continued improvement indicate stability and learning progress.

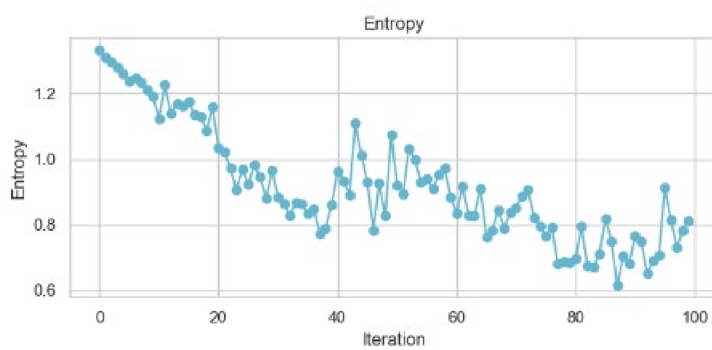

Figure 3. Entropy Over 100 Iterations

Entropy measures the level of randomness in the agent's action selection process. Initially, the entropy is relatively high, reflecting exploratory behavior, meaning the model is choosing actions at random to figure out the result of each. Over time, entropy decreases, signifying a transition toward more confident action choices. This decline is indicative of the agent adhering its learned policy and reducing exploratory behavior as it becomes more certain about which actions yield higher rewards.

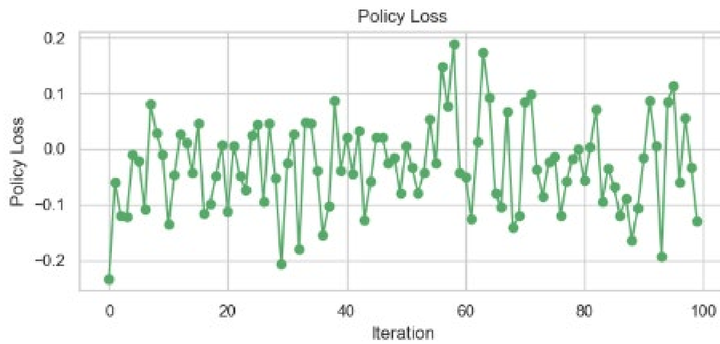

Figure 4. Policy Loss Over 100 Iterations

This figure displays the policy loss, which reflects the magnitude of the policy update at each iteration. Although the loss displays variability throughout the training process, it remains bounded, indicating consistent and stable policy updates. The fluctuations are common in reinforcement learning scenarios and may be attributed to changes in the environment or varying rewards.

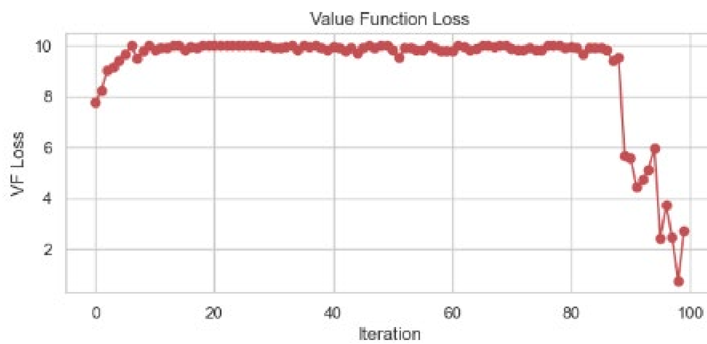

Figure 5. Value Function Loss Over 100 Iterations

The value function loss represents the error in the predicted state or state-action values. The graph shows a high initial loss, which quickly stabilizes and maintains a near-constant value for the majority of training. A significant drop in the loss toward the final iterations implies that the agent has improved its ability to estimate future rewards accurately, which is the best outcome for optimal policy performance.

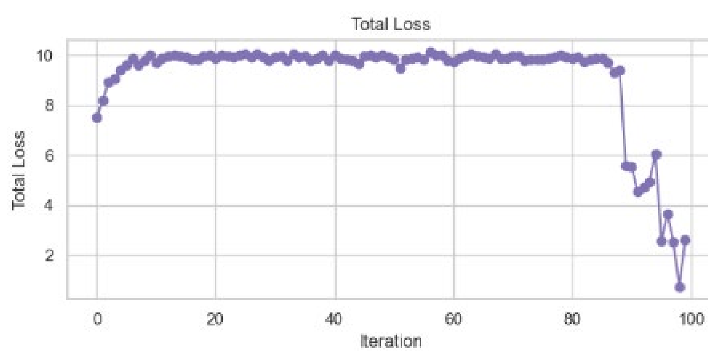

Figure 6. Total Loss Over 100 Iterations

The total loss, averaging the sum of the policy loss and value function loss follows a similar pattern to the value function loss. It remains relatively constant during the bulk of training, followed by a sharp decrease toward the final iterations. This trend indicates that the overall learning objective is being optimized and that the policy is nearing optimization.

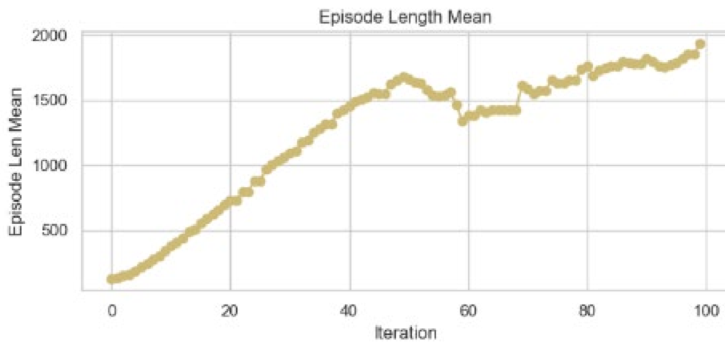

Figure 7. Episode Length Mean Over 100 Iterations

The average episode length increases steadily throughout training, indicating that the agent is lasting longer or performing more steps per episode. This trend matches with the increase in average reward and suggests that the agent is developing more efficient strategies for maintaining performance within the environment. The plateau and following recovery in episode length around iteration sixty align with similar patterns observed in the reward and loss metrics.

5.2 Proposed Improvements

While the model has shown relatively good yielding results during elementary testing, further refinement and strenuous testing will be conducted. Firstly, testing with greater amounts of iterations—such as five hundred, and then one thousand—as well as increasing the number of agents, and then altering the model to intake real-world collected data from the U.S. Department of Transportation to further train the model and familiarize it with realistic and unpredictable road-user behavior. Through iterative design, the model will learn to handle traffic at four-way intersections, with the ideal ability to accurately predict the agent behaviors and properly direct them to avoid collisions. As the focus of this research is to optimize intersection safety, later iterations will include different intersection styles, with a focus on signalized intersections as of now.

5.3 Validation

The results of this simulation are not only repeatable, but modifiable as well. The initial experiment was run with seventy iterations, outputting relatively good results, thus prompting the increase to one hundred iterations. The increase in iterations showed better results without any significant downfalls regarding computation time and policy loss. Overall, the experiment does not face validity concerns, as it can be replicated, altered, and adjusted as needed and still produce reliable results. Currently, iterations as high as five hundred are being tested.

6. Conclusion

Although testing is still in its early stages, the model has demonstrated strong initial performance and runs successfully. It quickly adapts under the reward-penalty structure of reinforcement learning, displaying noticeable improvements in accuracy and decision-making across iterations with few outliers. While the model is not yet fully optimized and further testing is required, the preliminary results are promising. With continued refinement and extended training, the system has the potential to operate more reliably and efficiently. Future iterations will also explore scalability, incorporating more agents, diverse datasets, and varied intersection layouts.

With that said, several limitations should be noted. The current simulation framework does not fully capture the unpredictability of real-world driving environments, including human driver behavior, sensor issues, and external conditions like weather or construction. Additionally, the model assumes idealized, low throughput communication among vehicles and infrastructure, which may not reflect the challenges of real-world deployment. The training process is also limited by the diversity of available simulation scenarios, potentially affecting the model's ability to navigate unseen intersections or traffic patterns. Addressing these limitations will be a key part of future development efforts.

As AI and connected vehicle technologies continue to advance, integrating reinforcement learning into intelligent transportation systems offers a promising pathway toward safer and more adaptive traffic control. By training these models in increasingly realistic environments and expanding their capabilities, this research contributes to the foundation for smarter urban mobility solutions. Ultimately, this model aims to support safer, more efficient traffic systems by minimizing collisions and enhancing the commuting experience for all road users.

References

Artificial Intelligence Research Area, Available: https://www.its.dot.gov/research-areas/Artificial-Intelligence/, Accessed on May 7, 2025.

Chen, S., Hu, X., Zhao, J., Wang, R., and Qiao, M., A review of decision-making and planning for autonomous vehicles in intersection environments, *World Electric Vehicle Journal,* vol. 15, no. 3, pp. 99, 2024.

Chen, Y., Zou, Y., Kong, X., and Wu, L., Investigating the impact of influential factors on crash types for autonomous vehicles at intersections, *Journal of Transportation Safety & Security*, vol. 16, no. 10, pp. 1089-1116, 2024.

Christopoulou, M., Barmpounakis, S., Koumaras, H., and Kaloxylos, A., Artificial intelligence and machine learning as key enablers for V2X communications: A comprehensive survey, *Vehicular Communications,* vol. 39, pp. 100569, 2023.

Dimitrakopoulos, G., and Demestichas, P., Intelligent Transportation Systems, *IEE Vehicular Technology Magazine,* vol. 5, no. 1, pp. 77-84, 2010.

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V., CARLA: An open urban driving simulator, *Proceedings of the 1st Annual Conference on Robot Learning*, pp. 1–16, 2017.

Gao, X., Zhang, X., Lu, Y., Huang, Y., Yang, L., Xiong, Y., and Liu, P., A Survey of collaborative perception in intelligent vehicles at intersections, *IEEE Transactions on Intelligent Vehicles,* vol. 14, no. 8, 2025.

Hummer, J. and Ott, S., Capacity and operational effects of midblock left-turn lanes, *NCHRP Report 745*, Transportation Research Board, Washington, D.C., 2015.

Miftah, M., Desrianti, D. I., Septiani, N., Fauzi, A. Y. and Williams, C., Big data analytics for smart cities: Optimizing urban traffic management using real-time data processing, *Journal of Computer Science and Technology Application,* vol. 2, no. 1, pp. 14–23, 2025.

National Highway Traffic Safety Administration, Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey, *Report No. DOT HS 812 115*, Traffic Safety Facts, Washington, D.C., February 2015.

Olayode, O.I., Tartibu, L.K. and Okwu, M.O., Application of artificial intelligence in traffic control system of non-autonomous vehicles at signalized road intersection, *Proceedings of the 30th CIRP Design Conference (CIRP Design 2020)*, Procedia CIRP, vol. 91, pp. 194–200, 2020.

Proven Safety Countermeasures, Available: https://highways.dot.gov/safety/proven-safety-countermeasures, Accessed on May 7, 2025.

Skabardonis, A., Varaiya, P. and Petty, K., Real-time measurement of traffic conditions using vehicle reidentification data, *Transportation Research Record*, vol. 1855, no. 1, pp. 82–88, 2003.

Transportation Research Board, *Highway Capacity Manual*, 5th Edition, National Research Council, Washington, D.C., 2010.

U.S. Department of Transportation, ITS Joint Program Office, Connected vehicle pilot deployment program: SPaT and MAP message implementation and use, Available: https://www.its.dot.gov, Accessed on May 7, 2025

V2X Deployment Research Area, Available: https://www.its.dot.gov/research-areas/V2X-Deployment/, Accessed on May 7, 2025.

Wang, T., Kim, S., Ji, W., Xie, E., Ge, C., Chen, J., Li, Z., and Luo, P., DeepAccident: A motion and accident prediction benchmark for V2X autonomous driving, *arXiv preprint arXiv:2304.01168*, 2023.

Xu, Y., Zhou, H., Chen, J., Qian, B., Zhuang, W., and Shen, S., A Survey of collaborative perception in intelligent vehicles at intersections, *IEEE Communications Standards Magazine,* vol. 4, no. 4, pp 18-25, 2020.

Yusuf, S., Khan, A., and Souissi, R., Vehicle-to-everything (V2X) in the autonomous vehicles domain – A technical review of communication, sensor, and AI technologies for road user safety, *Transportation Research Interdisciplinary Perspectives,* vol. 23, pp. 100980, 2024.

Biographies

Kayle Peña Agront is an Electrical Engineering graduate student at Florida Polytechnic University. She holds a Bachelor of Science in Computer Engineering and works as a Junior Quality Engineer at Draper's Advanced Packaging Facility in St. Petersburg, FL, where microelectronics are manufactured. She has experience in machine learning, demonstrated in academic projects such as a Tennis Match Prediction model as well as a project sponsored by the Polk County Sheriff's Office, where her team built three subsystems, each with its own model capable of reading license plates, detecting advancing objects, and detecting gunshots. In other projects, she has demonstrated an understanding of automotive technologies through designing an autonomous drivetrain and conducting a study on regenerative breaking in electric vehicles.

Gaspar Chayer, originally from Buenos Aires, Argentina, is pursuing a bachelor's degree in Computer Engineering at Florida Polytechnic University. His academic interests focus on AI and Intelligent Transportation. Gaspar is an intern at Microsoft with the Cloud Supply Chain Engineering Team where he has worked on projects enhancing global supplier collaboration

Victoria Correa Andrade Doyon is a Software Engineer at Microsoft with a Bachelor of Science in Computer Engineering from Florida Polytechnic University. Her professional background and academic training center around artificial intelligence, with a particular emphasis on applied machine learning, reinforcement learning, and multi-agent systems. She has worked extensively with natural language processing, anomaly detection, and cloud-based AI services, and has experience building and optimizing scalable machine learning pipelines. Her technical expertise includes the use of frameworks for multi-agent deep reinforcement learning and retrieval-augmented generation (RAG), as well as deploying AI models within enterprise environments. She is also proficient in vector embeddings, search indexing, and leveraging Azure's AI and data services for large-scale applications. Victoria's work reflects a strong understanding of both theoretical foundations and practical implementations of artificial intelligence. She is particularly interested in intelligent systems that operate in dynamic or uncertain environments, and continues to explore research directions related to AI safety, model interpretability, and collaborative decision-making under uncertainty. She contributes to research and development initiatives that aim to bridge academic advances with real-world applications, and her ongoing work demonstrates a commitment to innovation, precision, and technical excellence in the field of AI.

Rawa Adla, Ph.D. is an Associate Professor of Electrical and Computer Engineering at Florida Polytechnic University. She is a professional in Intelligent Transportation Systems (ITS), bringing together a strong blend of academic and industry experience in the field. Her research interests encompass ITS, autonomous vehicles, connected vehicle technologies (V2V, V2I, V2X), electric vehicles, and battery systems. Her work focuses on developing, analyzing, and testing innovative methodologies aimed at improving traffic safety, enhancing transportation system efficiency, and advancing the performance of electric vehicles and battery technologies.