Safety21
INNOVATING SAFETY FOR ALL
The National University Transportation Center for Promoting Safety
Carnegie Mellon University
Penn University of Pennsylvania
The Ohio State University
CCAC
Community College of Philadelphia
The University of Texas Rio Grande Valley
Morgan State University

Evaluating Autonomous Vehicles' Safety

Carlee Joe-Wong, https://orcid.org/0000-0003-0785-9291
Osman Yağan, https://orcid.org/0000-0002-7057-2966
I-Cheng Lin, https://orcid.org/0000-0002-5306-3262

FINAL REPORT

August 1, 2024

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under [grant number 69A3552344811] from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
N/A		

4. Title and Subtitle	5. Report Date
Evaluating Autonomous Vehicles' Safety Benefits in Mixed Autonomy Scenarios	August 1, 2024
	6. Performing Organization Code N/A

7. Author(s)

Carlee Joe-Wong, https://orcid.org/0000-0003-0785-9291

Osman Yağan, https://orcid.org/0000-0002-7057-2966

I-Cheng Lin, https://orcid.org/0000-0002-5306-3262

8. Performing Organization Report No.
N/A

9. Performing Organization Name and Address	10. Work Unit No.
Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213	
	11. Contract or Grant No. Federal Grant No. 69A3552344811

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Safety21 University Transportation Center Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213	Final Report (July 1, 2023-June 30, 2024)
	14. Sponsoring Agency Code USDOT

15. Supplementary Notes
Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

16. Abstract

Connected autonomous vehicles (CAVs) are gradually advancing towards widespread deployments. CAVs promise to improve transportation safety by operating more efficiently and avoiding incidents like crashes due to human driver error. However, they may cause crashes or other safety incidents themselves, especially when interacting with humans. Our work has three parts: (i) estimating the effective incident rates of CAVs and how they are distributed across a city; (ii) incorporating CAVs' and human drivers' ability to react to human pedestrians; and (iii) evaluating our models and analysis in our mixed-autonomy simulator for a variety of road topologies. The main conclusion of our project is that safety dynamics with CAVs are complex and difficult to predict, requiring sophisticated simulators that are flexible enough to model a range of CAV and human driver behavior. Even complex reinforcement learning models, which can theoretically capture different vehicle objectives and complex decision-making, can struggle to accurately capture vehicle behavior and traffic dynamics, due to the complexity of training such models.

17. Key Words	18. Distribution Statement
Autonomous vehicles, Mixed-autonomy, Safety	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	6	

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Problem Description

Connected autonomous vehicles (CAVs) are gradually advancing towards widespread deployments. CAVs promise to improve transportation safety by operating more efficiently and avoiding incidents like crashes due to human driver error. However, they may cause crashes or other safety incidents themselves, especially when interacting with humans. For example, Cruise famously was forced to pause its testing of driverless vehicles when one such vehicle appeared to severely injure a pedestrian in San Francisco. On a systemic level, CAVs may encourage more risky behavior, as CAVs may not be able to accurately predict human actions, or humans may overestimate CAVs' capabilities of avoiding safety incidents, leading to their taking riskier actions.

As CAVs advance towards larger-scale deployments, we can expect that they will share the road with human-driven vehicles and human bicyclists and pedestrians. Thus, from a safety perspective, it is important to understand how CAVs will interact with humans and in particular whether their presence will have a significant effect on safety. These effects can be both direct and indirect: for example, if CAVs increase the density of vehicles on roads, then these vehicles (both CAVs and human-driven vehicles) may be more likely to get into accidents themselves, due to the increased density of traffic. Conversely, if CAVs can facilitate smoother flow of traffic, thus reducing the overall likelihood of traffic accidents, then they may have indirect positive benefits on the overall rate of vehicular accidents and thus safety. The goal of this project is to evaluate the potential safety benefits of CAVs in mixed-autonomy settings, in which CAVs and human vehicles share the road.

Approach and Methodology

Our work has three parts: (i) estimating the effective incident rates of CAVs and how they are distributed across a city; (ii) incorporating CAVs' and human drivers' ability to react to human pedestrians; and (iii) evaluating our models and analysis in our mixed-autonomy simulator for a variety of road topologies.

Our approach is fundamentally simulation-based, as it is difficult to accurately model the dynamics of human and CAV interactions at the scale of a city with hundreds or thousands of vehicles operational at any given time. We draw from existing models of vehicle flow, which typically use microscopic models for individual vehicle behavior and macroscopic models to model a network of interacting vehicles. Neither are a good fit for our work, as we would like to model how individual vehicle decisions (whether made by human drivers or CAVs) affect the overall dynamics of vehicular traffic and thus vehicular safety. We therefore take a hybrid, hierarchical approach that divides the road network into smaller "cells." Figure 1 illustrates this cell-based approach. Vehicles in the same cell are in close physical proximity and thus may directly interact with each other. Since there are relatively few vehicles in each cell, we can use microscopic models for their interactions. These interactions can then be summarized with cell-level statistics (e.g., average speed, number of vehicles) that determine how each cell is affected by its neighboring cells and thus the overall dynamics of the road network.

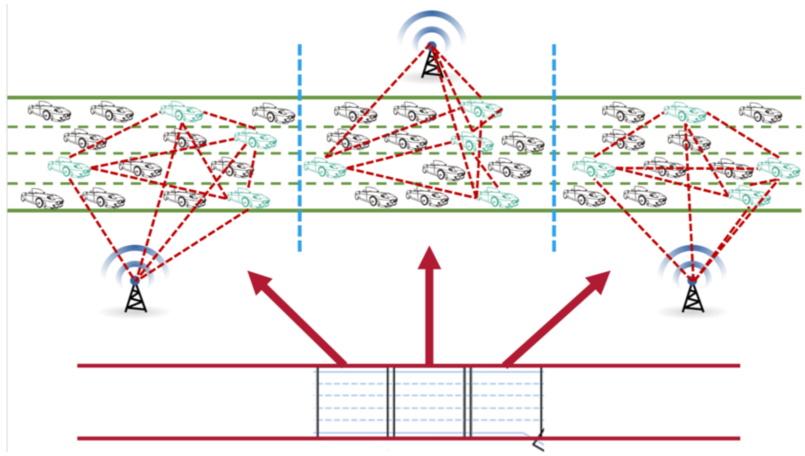

Figure 1: Illustration of our cell-driven model that we use to simulate the actions of connected autonomous vehicles and human-driven vehicles sharing a road.

While our cell-level model is more tractable than using microscopic models of each individual vehicle and extrapolating to the overall road network, it is still difficult to directly track each vehicle's actions in this model. Thus, instead of seeking closed-form solutions, we utilize reinforcement learning simulations to model the states, actions, and rewards experienced by individual vehicles and cell-level dynamics. Using simulations allows us to model a much larger-scale network, with many distinct cells along different roads. Reinforcement learning provides a flexible decision framework that allows us to define different notions of "safety" by changing the reward function, e.g., we can compare outcomes if vehicles aim to optimize their own safety or overall travel times compared to whether they selfishly aim to minimize their individual travel times. We can further model human pedestrians by incorporating their presence into the states of individual vehicles, which then affects those vehicles' actions.

Findings

To estimate the effective incident rates of CAVs given how they are distributed around a city, we built on the mixed autonomy simulator that we created in a previous Mobility21 project. We aimed to incorporate safety concerns into the simulator, and in particular to evaluate how CAVs could change their behavior in mixed-autonomy settings with human-driven vehicles.

Towards our first objective of estimating effective incident rates of CAVs, we found that there is little reliable data on CAV-involved traffic accidents, due to their limited deployment. Thus, we decided to make our simulator flexible to different vehicle incident statistics, allowing them to be input into the simulation.

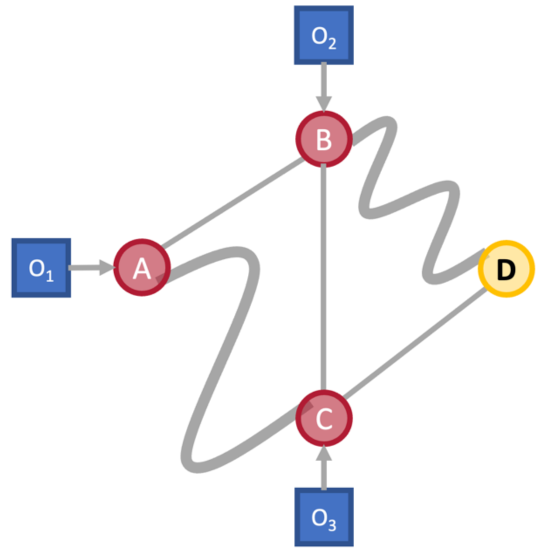

Figure 2: Illustration of a network satisfying Braess's paradox, in which lower traffic volumes can increase travel time from the origins (A, B, C) to destination D.

Towards our second and third objectives, our original mixed-autonomy simulator assumed that CAVs chose their movements around a city so as to minimize congestion and that human-driven vehicles took fixed routes with random perturbations. We extended the simulator to incorporate arbitrary CAV objectives, which will enable us to include different forms of safety. We have also added the ability to discretize CAV and human-driven vehicles' actions, which will allow us to have CAVs more easily learn how to optimize more complex objectives, including safety concerns, in environments that include complex safety-relevant characteristics like vehicle incident frequency.

We used this simulator to evaluate whether our simulated vehicular behavior matches known behavior that has been observed in practice. For example, Figure 2 illustrates Braess's paradox, in which decreasing traffic volume leads to increased travel time for vehicles, given that each vehicle acts so as to myopically minimize its individual travel time. We find that our simulator is able to reproduce this behavior, but that introducing CAVs whose goal is to minimize collective travel times for all vehicles, can resolve the paradox as long as a sufficient proportion of the vehicles are CAVs, as one would intuitively expect. Thus, our simulation results are consistent with prior findings, lending credence to its ability to simulate CAV actions.

Conclusions

The main conclusion of our project is that safety dynamics with CAVs are complex and difficult to predict, requiring sophisticated simulators that are flexible enough to model a range of CAV and human driver behavior. Even complex reinforcement learning models, which can theoretically capture different vehicle objectives and complex decision-making, can struggle to accurately capture vehicle behavior and traffic dynamics, due to the complexity of training such models. Thus, having more data to train these models is essential to developing realistic, safety-aware mixed-autonomy simulators. Finding such data in practice, however, is challenging, due to the low current CAV penetration rate. The dynamics of vehicle traffic can also depend heavily on specific road topologies, making it difficult to generalize findings on safety from one environment to another. Future work may include characterizing the types of road networks that have similar safety and vehicular traffic characteristics in mixed-autonomy settings.

Project Team

PI: Carlee Joe-Wong, https://orcid.org/0000-0003-0785-9291
Co-PI: Osman Yağan, https://orcid.org/0000-0002-7057-2966
PhD Student: I-Cheng Lin, https://orcid.org/0000-0002-5306-3262