Project Title: Evaluating Autonomous Vehicles' Safety Benefits in Mixed Autonomy Scenarios

Recipient/Grant (Contract) Number: Carnegie Mellon University, Grant #: 69A3552344811

Center Name: Safety21 National University Transportation Center for Promoting Safety

Research Priority: Promoting Safety

Principal Investigator(s): Carlee Joe-Wong, Carnegie Mellon University

PI Contact: cjoewong@andrew.cmu.edu

Project Partners: Southwestern Pennsylvania Commission

Research Project Funding: $201,765.00

Project Start and End Date: 07-01-2023 to 06-30-2024

Project 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 incidents themselves, especially when interacting with humans. 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. Our work has three parts: (i) estimating the effective incident rates of CAVs and how they are distributed across a city, leading to algorithms for prioritizing incident responses so as to reduce their overall impact on traffic flow and safety; (ii) incorporating CAVs' and human drivers' ability to react to human pedestrians, leading to algorithms for CAVs to reduce pedestrians' impact; and (iii) evaluating our models and analysis in our mixed-autonomy simulator. Towards modeling CAVs' effect on traffic incident rates, we will account for the fact that vehicle incident rates vary with the road congestion level and type, e.g., Pennsylvania data show that incidents are more common in heavy-traffic surface streets than sparsely populated highways. We will build on our prior Mobility21 work studying mixed-autonomy traffic patterns to account for changes in congestion levels across the road network due to vehicle incidents, e.g., if CAVs overall reduce the incident rate on highways, this might lead to better overall traffic flow and fewer subsequent incidents. Our results will enable prioritization of incident response so as to maximally reduce the resulting traffic congestion. We will then incorporate the effects of human pedestrians into our mixed-autonomy setting. Pedestrians can change safety dynamics as their actions may be more difficult to predict, especially for CAVs that may not be well-trained on pedestrian data. For example, CAVs can improve traffic flow by more closely following other vehicles; this is less feasible when human pedestrians are present. We therefore plan to incorporate these pedestrian "shocks" into our model of traffic flow and incident rates. We will use these results to propose new techniques for CAVs to predict and plan for pedestrian behaviors. We will use our existing mixed-autonomy simulator, developed with Mobility21â€™s support, to numerically evaluate our models and how the above safety effects vary for different amounts of CAVs. We will also leverage models and feedback from our deployment partner, the Southwestern Pennsylvania Commission (SPC), in our simulations. We will further measure how CAVs' effects are distributed around a city and implications for equity (see also "Outputs" below).
This project is synergistic with our concurrently submitted proposal entitled "Mitigating Cascading Failures for Safety in Transportation Networks in the era of Autonomous Vehicles," where the goal is to evaluate the safety impact of AVs from the perspective of their impact on cascading road failures and congestion. In contrast, the current project focuses on CAVs' safety impact in terms of the traffic incident rate in mixed-autonomy settings. As such, the two projects complement each other and can be combined at a total budget of $150,000 if preferred.

Outputs:

Our concrete deliverables will be (i) algorithms for emergency responders to prioritize incident response according to overall impact on congestion and safety in the road network; (ii) new algorithms that allow CAVs to plan for (potentially unexpected) human pedestrian behavior; (iii) quantitative estimates of how much CAVs will improve traffic flow for different types of roads, as a function of their prevalence on roads; (iv) quantitative estimates of how CAVs affect incident rates between vehicles and pedestrians, for different CAV penetration rates; (v) regulatory implications of these quantitative findings; and (vi) extensions of our existing mixed autonomy simulator to include better models of human pedestrian and driver behavior. We will work with our equity and deployment partner, the Southwestern Pennsylvania Commission (SPC), to ensure the usefulness of our project outputs, through regular meetings for feedback. We will also consider leveraging SPC's existing public participation processes to receive community feedback on our work; analysis of this feedback will be another output of this work. We expect that this feedback will allow us to account for metrics that are important for all communities affected by CAV deployments, helping to ensure that we address equity concerns associated with CAV deployments.

Outcomes/Impacts:

By quantifying CAVs' effects on traffic flow and accident rates with other vehicles and human pedestrians, we will enable regulatory agencies to anticipate the effects of gradual but increasing CAV deployment. Similar to our prior "Big Idea" project with Mobility21, on which this project builds, our work could inform the policies that agencies may wish to take in regulating CAVs (e.g., developing new metrics for how well they interact with human drivers and pedestrians). It will further provide quantitative validation for (or evidence against) claims that CAVs can benefit traffic and traffic safety as a whole. Finally, our simulator may aid such agencies and other researchers in studying CAV and human behavior in mixed-autonomy systems. Emergency responders may benefit from our work by prioritizing responses to incidents that have the greatest overall impact (in terms of traffic congestion and follow-on incidents) on the road network. CAV manufacturers may also be able to benefit from this work, as it will help them quantify the importance of developing AVs that interact "well" with humans and design learning algorithms that help CAVs to do so. These new learning techniques may enhance the benefits of CAVs and increase their potential for deployment.