A USDOT NATIONAL
UNIVERSITY TRANSPORTATION CENTER

Evaluating Resilience in Mixed-Autonomy Transportation Systems

PI: Carlee Joe-Wong
OrcID: 0000-0003-0785-9291

Co-PIs: Osman Yagan and Sean Qian

FINAL RESEARCH REPORT

Contract # 69A3551747111

Problem Description

As autonomous vehicles near widespread deployments, we can envision a future where an increasing fraction of vehicles on the road will have some degree of autonomy. While many works have studied how autonomous vehicles might coordinate with each other to enhance transportation safety and efficiency, relatively few studies have examined interactions between human-driven and autonomous vehicles. One might expect that the safety and efficiency benefits of autonomous vehicles increase in proportion to their prevalence on roads, but the magnitude of this effect may be more complex; e.g., a threshold fraction of vehicles may need to be autonomous to have a measurable effect on safety or efficiency. Quantifying these benefits may aid regulators and policymakers in estimating the true benefits that autonomous vehicles can bring to transportation networks, particularly when the prevalence of autonomous vehicles can vary throughout a city's road network.

The adaptation of autonomous vehicles (AVs) can benefit the operation of the whole transportation network in terms of better utilization, lower travel time, and energy consumption. Many of these potential benefits stem from three core AV capabilities. First, the AVs can be programmed to make collaborative decisions that will improve the overall driving experience for all users, and look for ways to reduce traffic congestion even if it means the AVs might experience a longer travel time or a less efficient trip. Second, AVs are equipped with multiple sensors and use computers to automatically detect and respond to changes in the environment. Thus, the adaptation of AVs could enhance the safety of the transportation network and accommodate a denser traffic flow by allowing autonomous vehicles to follow each other more closely on roads without compromising safety. Third, since the computation of AVs takes less time than humans to respond, they generally have a lower reaction time than the human-driven vehicles (HVs). Moreover, since the AVs can share the information between them, they can respond to incidents a few blocks away that are not visible to human sight, further reducing the reaction time.

The objective of this project is to evaluate how autonomous vehicles can improve the resilience of transportation networks to exogenous disturbances like road construction, traffic accidents, and congestion. We plan to use our analysis to recommend policies and guidelines for regulators to specify the behavior of autonomous vehicles in a road network, so as to maximize resilience. We first consider the optimal strategy of the autonomous vehicle fleets and the equilibrium, or average steady-state, traffic conditions. Since the traffic flow in practice may vary over time, we then consider the resilience and robustness of the whole transportation system.

First, when the traffic flow is stable, vehicles should eventually reach an equilibrium state (e.g., level of congestion) on a road network. This equilibrium, however, may be suboptimal if the vehicles do not have full information of the city or only aim to optimize their own delays. We model the coexistence of HVs and AVs in routes chosen, finding that the system reaches a close-to-optimal equilibrium when only 20-50% of vehicles are programmed to make collaborative decisions that aim to optimize the overall delay of the whole transportation network.

Second, the adaptation of autonomous vehicles (AVs) and their co-existence with human-driven vehicles (HVs) form a multi-agent mixed-autonomy environment. This work aims to find the optimal operating policies for the fleet of AVs, considering HVs behave selfishly based on their own information.

Third, we find out how AVs could prevent cascading network congestion, which may result in large-scale congestion and failure initiated by a disaster or a particular event—reaching the optimal robustness of the whole transportation system.

Approach and Methodology

We describe our approach and methodology for each of the three major goals above.

Equilibrium analysis: To model the equilibrium of AV and HV traffic, we model HVs as acting selfishly to minimize its own travel delay regardless of the available resources and the capacity of the roads. Altruistic or collaborative vehicles, however, can take actions so as to maximize the overall social welfare of the system. Thus, we used a Stackelberg game formulation in which the AVs are the game “leaders,” making decisions on how to distribute themselves around congestion or blockages so as to maximize the collective system performance (e.g., minimizing overall travel time). These decisions account for the reactions of human users, who are modeled as the “followers” and make myopic, selfish re-distribution decisions. If too few vehicles are not selfish, their decisions will not effectively influence the equilibrium state of the traffic. After reaching a certain percentage of vehicles that do collaborate in aiming to optimize the overall delay of the whole transportation system, the equilibrium state starts to change, and the overall performance increases. When the percentage of collaborative AVs exceeds a certain level, the optimal equilibrium is reached, and increasing the ratio of AVs will no longer decrease total travel time. However, the AVs could maintain a shorter following distance from other vehicles. If we consider this, the overall travel time will keep decreasing slightly with the increasing ratios of the AVs.

AV and HV optimization: We next considered optimal AV and HV actions towards reaching an equilibrium. Following prior literature, the specific AV and HV follower decisions that we model are both vehicles' routing at intersections and the speed at which they travel. By turning onto different roads, vehicles can change their routes to their destinations, thus potentially rerouting around congestion. Since congestion can also be caused by too many vehicles on the road, controlling their speed will help to ensure sufficient flow of vehicles on each specific road to avoid congestion (and thus, another type of road blockage). As the resulting optimization problem is difficult to solve analytically, we utilize reinforcement learning to solve for them. Specifically, we use the COMA framework, in which AVs can collectively learn individual policies for their actions (the routes to take and how fast to travel). Our work uses a mesoscopic view, considering both network-level routing decisions and road level velocity decisions. We group the AVs in the same area as a small fleet, using multi-agent reinforcement learning (MARL) and the cell-transmission model (CTM) of the traffic to find the optimal policies for the AVs that maximize the social welfare —the average travel time of all vehicles. In each cell or region, vehicles can communicate locally with each other or to the cellular base station in that area. Our trained model can be stored at the base station, gathering all the information in the block, and determining the best actions.

We have built a simulator (see References for a publicly accessible link) to test a multi-agent reinforcement learning algorithm in which each vehicle is an "agent" who makes decisions after observing other vehicles' actions. We have simulated this initial setting in a grid network and found that AVs' intelligent decisions can alleviate failures. Moreover, even a limited number of AVs are able to make a significant difference in network failure, e.g., the number of surviving road segments (that are not blocked by congestion) is similar for 50% or 100% AVs, even though the former scenario includes many human-driven vehicles making selfish decisions.

Network resilience: We finally consider two networks that operate in parallel (for example, a road network and train network), where work can be offloaded or transferred between the two. Drivers, for example, may choose to take the train instead of driving if road conditions are bad. Prior work in this area has considered static offloading strategies, where a fixed percentage of users transfer from one network to another. Our work instead considers a dynamic strategy in which the proportion of users transferring may change depending on the current state of the network. We show that such a dynamic strategy outperforms static transfer strategies, in the sense that the network can survive larger initial failures when dynamic transfer strategies are followed.

We validated our models and results by presenting our preliminary findings to two focus groups of community stakeholders, including representatives from the Southwestern Pennsylvania Commission, PennDoT, Pittsburgh Port Authority, and Pittsburgh Department of Mobility and Infrastructure. The feedback from these focus groups was instrumental in revising our models of AV and HV actions and capabilities, as well as in suggesting the types of policies that might ensure that AVs behave in a collaborative manner.

Findings

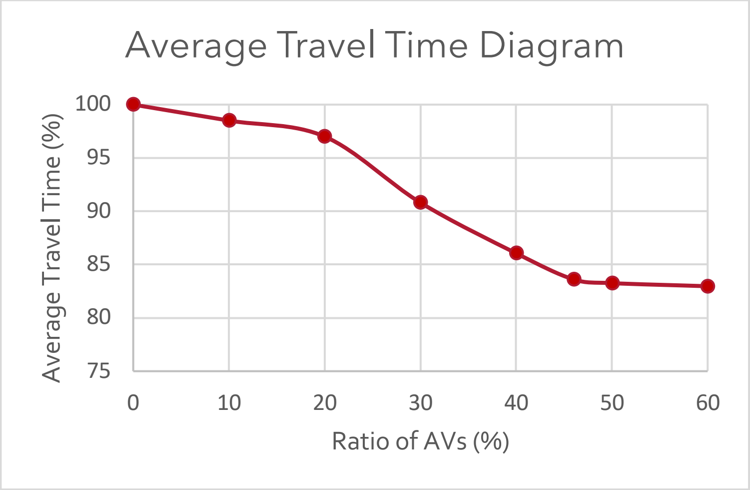

We can find the optimal equilibrium state of AV and HV traffic by solving a two-stage convex optimization problem. In the first stage, if we view each route as a bucket of water that fills up as more vehicles take that route, then we can treat selfish HVs as acting like water that flows into all buckets uniformly, eventually equalizing the delays on each route. This problem is analogous to a power allocation problem for the Gaussian noise channel in communication engineering. By applying specific methods, we could solve the optimal equilibrium state of parallel routes and figure out the portion of AVs that can reach this optimum by adopting water filling algorithms for the inner optimization problem for the HVs and then solving the AVs' optimization problem. We find, as shown in the figure above, that at around 50% AVs, we have achieved the maximum reduction in average travel time.

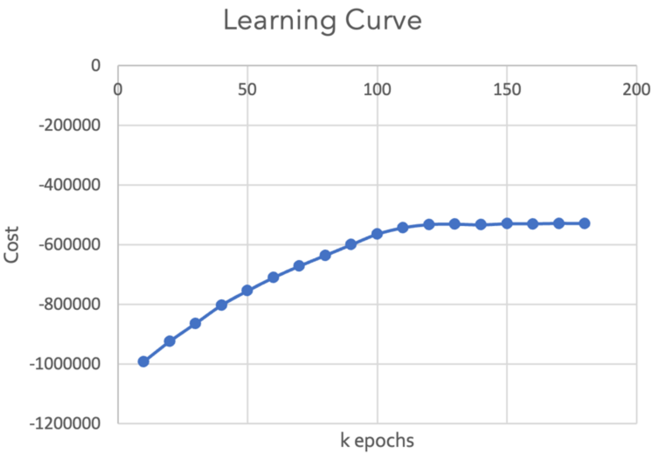

Our methods show that under different ratios of AVs, the model could lower the average travel time of the whole transportation system.

With the control of the velocity and routing of AVs, our model could learn the traffic pattern and react to the sudden traffic that may potentially fail the network and incentivize a better behavior of HVs. In short, reasonable control of the AVs could drastically change the performance of the whole transportation system even when the portion of AVs is much smaller than HVs. The figure at left shows the overall reward of the system, which increases as the MARL learns the best AV actions.

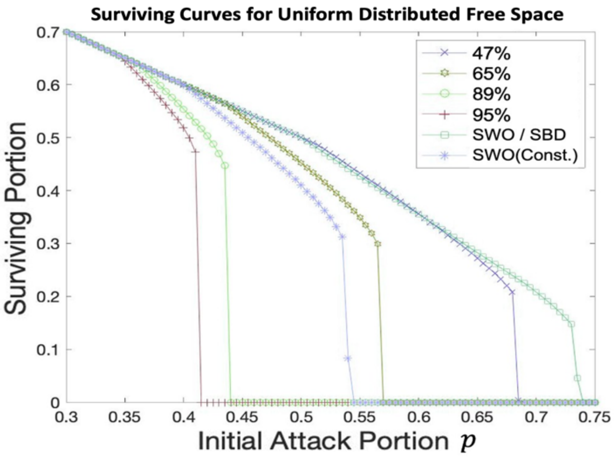

Our work on analyzing how users should transfer from one transportation network to another so as to avoid cascading failures can help network operators identify ways in which they can avert failures by diverting users to an alternative mode of transport. We have submitted these results to the IEEE Transactions on Network Science and Engineering and the IEEE Conference on Decision and Control. In this work, we introduce the dynamic coupling strategy with a step-wise optimization (SWO) framework between networks, which could determine the optimal coupling coefficients at each stage of the cascading failure process. We show that (under certain conditions) the SWO strategy reaches the global optimum regarding the final surviving number of areas. Thus, our step-wise dynamic optimization can lead to the optimal operation of the whole system. The figure at right shows the results of our SWO strategy compared to other fixed coupling strategies. Under the same initial failure size (the portion of the network that fails), the dynamic strategy could yield the largest surviving portion of the whole system.

We summarize the final policies suggested by our research (see also the section on Recommendations below) as follows:

1. Identify the optimal ratios (usually between 20%~50%) of collaborative vehicles and ensure that this proportion of vehicles follow the collaborative decisions.

2. Execute the dynamic micromanagement model in each area to determine the actions of micro AV fleets. Gather the observations and actions of each micro fleet every specific time interval and return to the centralized controller to judge the operations and update the model.

3. Compute traffic flow redistribution policies for AVs to avoid traffic jams and cascading failures on road networks. In multiple interdependent transportation systems, agencies may apply dynamic coupling strategies to optimize the robustness of the system against cascading failures of the transportation system.

Conclusions

AVs in the mixed-autonomy era could act as coordinators in the transportation system. With collaboration and information sharing between the AVs and the centralized server, we find that AV fleets can take the actions that incentivize the best response of the HVs. Even a partial adaptation of AVs not only increases the utilization of the transportation system, minimizing vehicles' delays, but also increases the resilience of the transportation network against failures and sudden traffic. Thus, drivers' on-road experience could be elevated to the next level in the new era of mixed-autonomy.

Opportunities for training and development have been provided to graduate students through mentorship of their research projects. Three Ph.D. students and two M.S. students were supported by this grant, and the PIs regularly meet with the students, give feedback on their work, and provide opportunities for the students to present the work to other researchers and stakeholders.

Recommendations

Our major recommendations take the form of three suggested policies, as listed below and elaborated in our policy brief.

1) Identify the optimal ratios of collaborative vehicles and ensure that this proportion of vehicles follow the collaborative decisions. In our simulations, usually, collaborative vehicles start to be effective after the ratio is between 20%~50%. We could reach such ratios with a combination of AVs and HVs following the collaborative decisions by compensating the expense of the vehicles or making a certain degree of obedience to the collaborative decisions a prerequisite for AV registration with local or state transportation agencies. This could be realized in practice by exemptions from tolls, or a dynamic tolling system with discounts for the collaborative vehicles. Agencies like PennDOT, for example, may impose vehicle-miles-travelled tolls on state highways that could be waived for collaborative vehicles. Other example incentives more suited to urban areas might include discounts at parking garages.

2) An AV fleet operator can apply our proposed method to train the model to manage AV fleets locally. The operator should (i) execute such a model in each area to determine the actions of micro fleets and (ii) gather the observations and actions of each micro fleet in each time interval and return them to the centralized controller to judge the operations. The update frequency can be determined by monitoring the traffic situation. If some new traffic pattern occurs, the performance of the whole network may decrease drastically; we can update the model more frequently to deal with such a situation. On the other hand, if the current trained models work well enough, we could lower the update frequency to ease the communication between local actors and the centralized critic. Such models can be adopted by the Incident response teams who need to direct traffic away from accidents.

3) For the road transportation system, AVs could identify the roads that are close to the capacity limits; or foresee the roads that may potentially receive lots of extra traffic due to the failures in adjacent areas. Based on such information, planning agencies could compute traffic flow redistribution policies for AVs to avoid traffic jams and cascading failures.

   Applying the proposed framework to multiple interdependent transportation systems can optimize the dynamic flow adjustment between different networks. Recommendations with transportation apps, e.g., the Port Authority mobile app, could achieve such traffic flow transferring between different network systems. Furthermore, ensuring the availability of public transportations and providing discounted parking areas, discounted tickets, etc. in these locations could also be a good policy to encourage passenger transfers.

Publications and Products

Lin, I., Joe-Wong, C., and Yagan, O. (2022, March). Dynamic Coupling Strategy for Interdependent Network Systems against Cascading Failures. *arXiv preprint 2203.01295*, Under major revision at the *IEEE Transactions on Network Science and Engineering*. Available online: https://arxiv.org/abs/2203.01295

Lin, I., Yagan, O., and Joe-Wong, C. (2022, July). Mixed-Autonomy Era of Transportation: Resilience and Autonomous Fleet Management. *Traffic21 Policy Brief*. URL pending.

Simulation code: https://github.com/delphi1618/Mixed-Autonomy-MARL

Project Participants

Chen Zhao
Carlee Joe-Wong (OrcID: 0000-0003-0785-9291)
I-Cheng Lin
Sean Qian (OrcID: 0000-0001-8716-8989)
Shaival Parikh
Osman Yagan (OrcID: 0000-0002-7057-2966)