EVALUATION OF SAFETY AND MOBILITY AROUND LOW SPEED AUTOMATED VEHICLE THROUGH REAL WORLD DEPLOYMENT IN URBAN ROADWAY SYSTEM

Abhijit Sarkar
Virginia Tech Transportation Institute
Blacksburg, Virginia, USA

Vaibhav Sundharam
Virginia Tech Transportation Institute
Blacksburg, Virginia, USA

Kevin Grove
American Trucking Association
USA

Paper Number 23-0119

ABSTRACT

Low-speed automated vehicles (LSAVs) are a new type of road transportation option that can be deployed in densely populated areas to connect passengers to existing transit systems. These vehicles are designed to operate at low speeds (often in the < 25 mph range) in complex operational design domains and can be retrofitted to accommodate at-risk road users, thereby making transportation even more accessible. Also, as many LSAVs are electric vehicles, they also show great potential for benefiting the climate. Even though there are numerous benefits to deploying LSAVs, several hurdles must be overcome to achieve success. For example, it is unclear how an LSAV deployed in a regular lane with a 25 mph or higher speed limit may affect other traffic. Also, it is unclear how vulnerable road users like pedestrians and bicyclists behave around LSAVs, as human machine interfaces for such interactions are not yet properly developed. These variables pose multiple questions for the safety, mobility, and operation of LSAVs in unrestricted operational domains. For example, an LSAV's low speed may cause other vehicles to operate at a lower speed, causing more vehicles to queue behind it. This may also make some drivers frustrated and lead them to become involved in dangerous driving situations like overtaking and cut-ins. In this work, we studied real world deployment of an LSAV on the US roadway to understand driver behaviors via 360 degree camera views from cameras installed on the LSAV. We examined the problems encountered during the deployment of an EasyMile (EZ-10) LSAV. We specifically investigated events from a real-world deployment during which the EasyMile LSAV needed to stop. The EasyMile deployment studied in this work included cameras that captured the 360 degrees of roadway environment around the vehicle. We developed a scene perception algorithm using computer vision technology to track other roadway agents like cars, pedestrians, and bicyclists around the EasyMile LSAV. We used object detection and tracking algorithms to track the trajectories of each of the roadway agents. Then we used perspective geometry and camera specifications to find the relative distances and speeds of these agents with respect to the EasyMile. This helped us understand the configurations of the traffic around the LSAV and study other drivers' temporal behavior. For example, the collected data shows the approach of any vehicle towards the EasyMile. Finally, we used this information to study other vehicles' maneuvers and show how the information from the cameras can be used to study simple maneuvers of other vehicles such as cut-ins, lane changes, and following behavior.

INTRODUCTION

Recent efforts to design, test, and deploy AVs have resulted in the development of a variety of configurations. These configurations often differ in size, sensor capabilities, sensor integration, capacity, and other features depending on the operational design domain (ODD) being targeted. One particularly interesting category of AV that has emerged is the LSAV, as depicted in Figure 1. What makes this type of vehicle interesting is the relatively complex ODD in which it is designed to operate. LSAVs are often designed to operate around streets with a variety of other road users and with relatively complex traffic controls in place. Many of these vehicles are designed to hold multiple passengers and are deployed with the intent of supplementing or connecting existing transit systems. To accommodate this more complex ODD, LSAVs often operate at relatively low speeds below 25 mph.

LSAVs face several hurdles to deployment but also offer significant benefits if the hurdles can be overcome. The ability of LSAVs to operate on streets means they can connect riders to common destinations. The ability to operate on streets also means LSAVs could be deployed in densely populated locations where there is a larger pool of users. Finally, the ability of LSAVs to supplement transit can improve riders' connectivity, reliability, and access [1]. Importantly, LSAVs can also be designed to accommodate vulnerable road users whose mobility may be limited. This means LSAVs could provide benefits to an underserved segment of transit users.

Figure 1. Photo. A picture of VTTI's EZ-10 LSAV ready for operation.

The LSAV's autonomy can have great implications in certain operating scenarios. For example, LSAVs can be used for transportation of medical patients, older passengers, and other populations unable to drive (or who do not need to drive) [2-4]. Historically, LSAVs have been deployed on predefined short-range routes. In most cases, the vehicle's operating speed is below 20 mph. As a result, when deployed in a mixed-traffic scenario, the LSAV's average speed is often 10 to 15 mph below the average speed of the traffic surrounding it. This brings up questions and concerns about the effect of LSAVs in a mixed-traffic scenario and whether they pose any safety threat. In this study, we looked at multiple examples of other vehicles following the LSAV for an extended period. We specifically looked at vehicles behind the LSAV and vehicle in front of the LSAV. The helps us to evaluate the deployment of LSAV in real world through cameras mounted around the vehicle. The camera captures dynamics of the other roadway objects around the LSAV. We used advanced computer vision algorithm to track the trajectories and kinematics of each of the roadway objects that are present in the scene. we demonstrate how we can use the trajectories solely from the 2D camera data and retrieve 3D information about roadway agents, including their relative distance and speed in world coordinates. We further looked at the aggregated behavior of the traffic behind the LSAV and in front of LSAV to demonstrate the trends of the traffic.

EXPERIMENTATION AND DATA

Fairfax Relay Deployment

Shortly after exploring these issues at VTTI's facilities, there was an opportunity to explore them in a real-world test deployment. A public-private partnership that included Fairfax Country, VA, Dominion Energy, EDENS (Mosaic), The Virginia Department of Rail and Public Transportation, The Virginia Department of Transportation, VTTI, and George Mason University planned a pilot deployment of an LSAV in Northern Virginia for 2020. The goal of the pilot was to learn how the technology could be deployed safely and effectively and provide an opportunity to learn more about real-world interactions. The shuttle was purchased by Dominion Energy and operated by Fairfax Country with TransDev managing the vehicle, its maintenance, and its operators.

The vehicle was branded the Relay and configured to follow a pre-programmed path during operations. The Relay uses localization LIDAR, GPS, and internal sensors to follow the pre-programmed path. The Relay also uses 3D and 2D LIDARs to detect objects around the shuttle. The localization LIDAR has a range of approximately 492 feet, the 3D LIDAR has a range of approximately 262 feet, and the 2D LIDAR has a range of approximately 131 feet. The ranges of these LIDARs give the Relay the ability to detect changes in the environment or potential conflicts within the environment quickly and at a safe distance. VTTI was able to work with the partners to install data recording systems on the Relay. These systems recorded video 360 degrees around the vehicle in high definition. This video allowed VTTI to see the context of conflicts that emerged during operations. VTTI recorded the video on secure encrypted drives that were brought back to Virginia Tech for processing. Figure 2 shows an example of the video recording from the VTTI deployment. The Fairfax deployment had similar arrangements. In this paper we have mainly used camera data from outward looking videos.

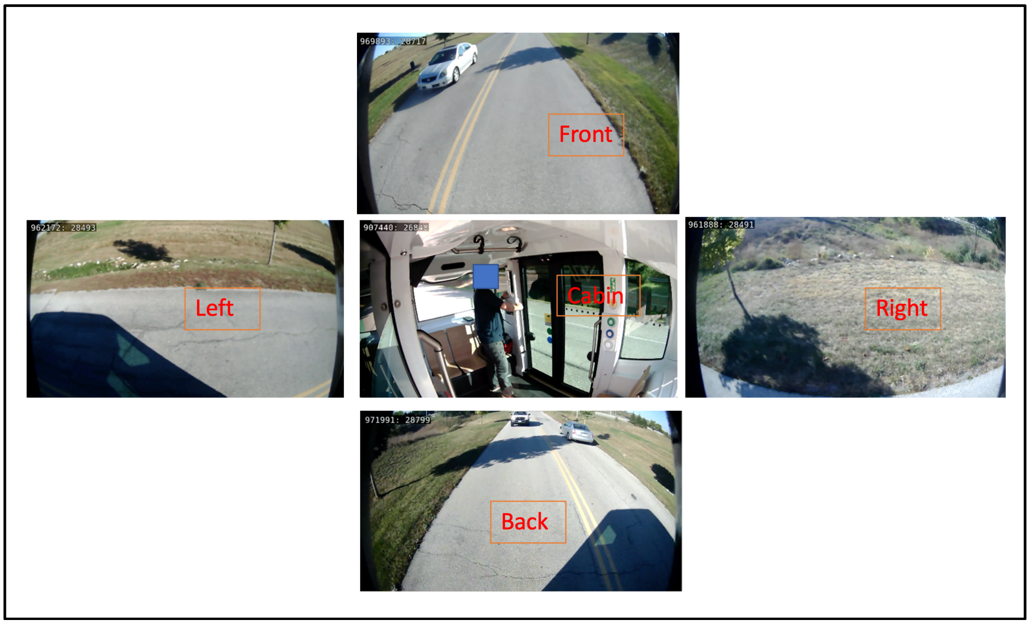

Figure 2. Video images. Example of scene cameras capturing 360-degree information around the LSAV.

COMPUTER VISION BASED CAMERA PROCESSING FOR SCENE PERCEPTION

Scene perception typically involves a detailed understanding of the scene using sensor data. In this case, we used monocular camera images to understand the construction of the roadway scene. For each roadway scene, there may be multiple roadway agents and geometric structures that play a key role. The roadway agents mainly involve the objects present in the scene, such as cars, trucks, pedestrians, and bicycles. Correspondingly, some other object classes also play a part in understanding the road, including road signs, traffic lights, work zone objects, etc. The geometric structural cues are mostly embedded in roadway boundaries, lane lines, crosswalks, etc. Sometimes, depth measurements also help provide an understanding of relative locations of objects and distance from the cameras. Each one of these objects and geometries can be perceived in pixel space by CV algorithms. In this paper, we discuss two specific aspects that are key to understanding the dynamics of other roadway agents around the EasyMile LSAV.

Object Detection and Tracking

Object detection refers to identifying an object in an image and locating that object in the pixel space. Over the last decade, deep learning methods have significantly enhanced the capability to detect objects in an image [5]. A typical object detection algorithm trains on a large set of annotated object classes and builds the capability to specifically localize the sets of pixels in an image to indicate a specific object class, such as a car, person, cat, dog, etc. Often the pixel space of an object is identified by drawing a BBox [6,7], as shown in Figure 3.

Object tracking, on the other hand, refers to the task where each object in a sequence of frames is tracked over time [8-10]. In a more formal definition of the task, each object in an image is assigned a unique object ID, and the object ID is maintained across the frames if the object is present in the FOV. This task mainly belongs to the standard CV task of "multi-object tracking," which also pertains to the subtasks of reidentification (i.e., identifying the same object from one frame to another) and motion prediction. Therefore, an effective algorithm should learn key features of the object from one frame and then should be able to transfer that knowledge into the next frame to reidentify the same object with high efficiency. As shown in Figure 3, each object is assigned a unique object id (e.g., 141, 147 etc.). Thus, through the object detection and tracking we can identify the class of the roadway object and track them over time.

(a) (b)

Figure 3. Example of detected objects in the EasyMile dataset.

Lane Detection

Lane lines define the boundary of a specific lane that a vehicle is supposed to follow unless it intends to change lanes. In the context of the project, lane line detection helps us understand the position of the ego vehicle (EasyMile) as well as the position of the other vehicles. We deployed an ultra-fast lane detection algorithm [11] to detect lane lines. It should be noted that the lane detection was performed on the original distorted images, and each lane line from the distorted image was subsequently undistorted and warped to the BEV. An example of detected lane lines on the EasyMile dataset is presented in Figure 4.

(a) (b)

Figure 4. Video images. Example of lane detection using ultra-fast lane detection algorithms on the (a) front and (b) rear video from the ego vehicle's camera.

Kinematics of Roadway Objects

Objects like cars, trucks, pedestrians, etc., populate a road scene and interact with each other. Their movements are codependent on other objects present on the road (other cars, pedestrians, etc.). To study the interaction and behavior of these objects around the ego vehicle, it is imperative to measure the distance and speed of such objects. A monocular camera does not preserve the depth. Therefore, we used knowledge of real-world measurements to calculate the distance of any pixel from the ego. To adequately understand the 360-degree behavior we need to study both longitudinal and lateral distance. Assuming a bird eye view of the scene, we need to create a pixel to real world 2D transformation map where we know how distance between consecutive pixels translate to real world coordinate (meter). Hence, we used two references. We used the lane width (12 ft) as the reference for lateral behavior. We used the length of right turn arrow as the reference for longitudinal distance. Once that is done, we use the bounding box from the object detection and tracking to find of distance of any roadway object. Then we used a 4th order Butterworth filter to remove noise. From the distance (both lateral and longitudinal), we used time derivative to compute speed in each direction.

TRAFFIC BEHAVIOR AROUND THE LSAV

In this section, we demonstrate how the information from the previous sections can be useful to understand the traffic behavior around the EasyMile. We specifically look at three scenarios and how each could result in potential safety hazards to the agents surrounding the AV, followed by a visualization of object behavior surrounding the LSAV to determine safety implications:

- A vehicle following the EasyMile AV
- A vehicle cut-in front of the EasyMile AV

Vehicle Following the LSAV

Figure 5 depicts a typical scene where a vehicle is following the AV. The location of the starting point of the detected vehicle along with its relative position with respect to the ego vehicle is plotted. As depicted in Figure 5(a), the vehicle with ID:2 is following the AV at an acceptable distance. However, after a few seconds, there is a rapid change in the velocity of the vehicle, which is evident from Figure 6 (red circle) followed by a 15-second halt. This sudden deceleration is due to the AV abruptly stopping in the middle of the road to give way to an oncoming distracted pedestrian, who is talking on the phone, as evident in Figure 5(e). A human driver would likely have analyzed the situation a priori, and hence such sudden braking would not have been required. Such a situation can lead to safety-critical conditions, possibly causing a crash if the driver following the AV is not maintaining sufficient distance. Also, any delay in the response time of the AV could have caused a serious roadside incident involving pedestrians. To better inform the road users about the deployment of LSAVs, adaptive warning signs can be installed onboard the AV that monitor the traffic right behind the AV and inform drivers to maintain a safe distance from the vehicle.

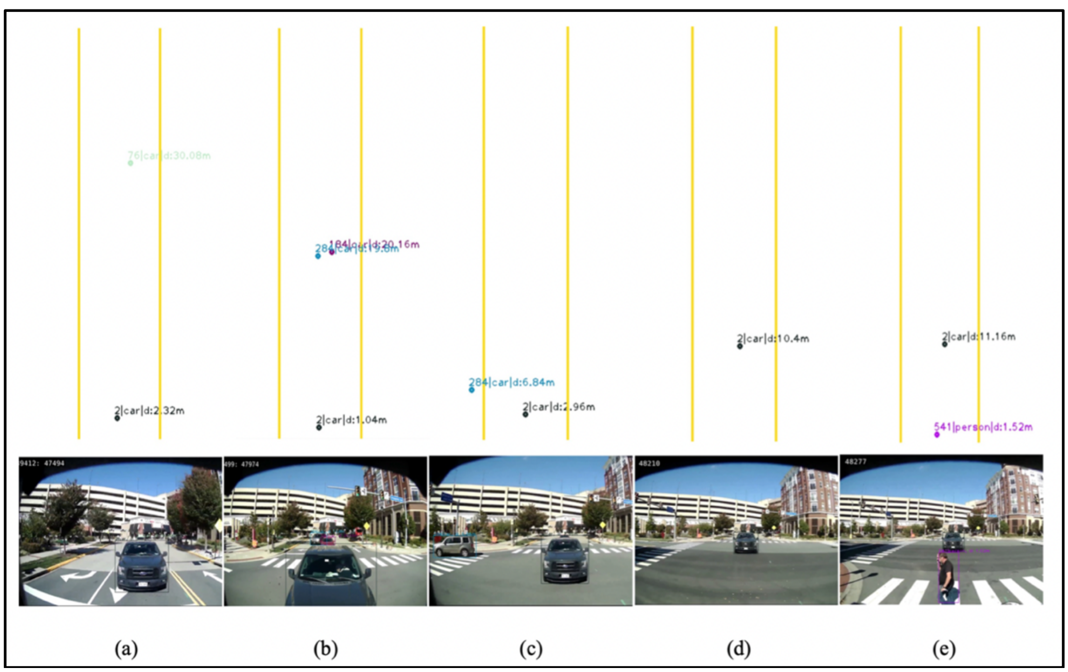

Figure 5. Diagrams with video images. Frames from a sample video taken at different time intervals depicting the trajectory of a vehicle following the AV.

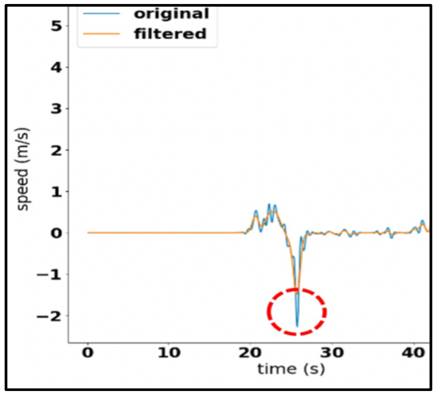

Figure 6. Rapid change in velocity of the vehicle (red circle) suggesting sudden deceleration.

Vehicle Cut-In Front of the LSAV

Being in the research and development phase, the EasyMile AV moves at speeds slower than recommended. Due to this, the drivers following the AV tended to get frustrated and try to overtake the AV, possibly under dangerous conditions. During our analysis of rear videos, we found multiple scenarios where the driver following the AV cut right in front of it to overtake it. Cut in refer to the behavior of another vehicle when then come in front of the another vehicle or between two vehicles [12]. In some situations, we saw drivers cutting into oncoming traffic to overtake the AV. One such case is presented in Figure 7, where the vehicle with ID:631 cuts in the front of the AV by crossing over to the opposite lane. A situation like this where individuals are involved in aggressive driving puts their lives and others' lives in jeopardy.

To mitigate such problems, the LSAV could be operated during non-peak hours when traffic is low. Also, during the initial stages of development, the LSAV can be operated on roads that have multiple lanes. This will give drivers additional room to maneuver their vehicles to overtake the AV, preventing safety-critical conditions. The data collected during this stage can be used to fine-tune the control systems onboard AVs to make them more robust for deployment in complex environments.

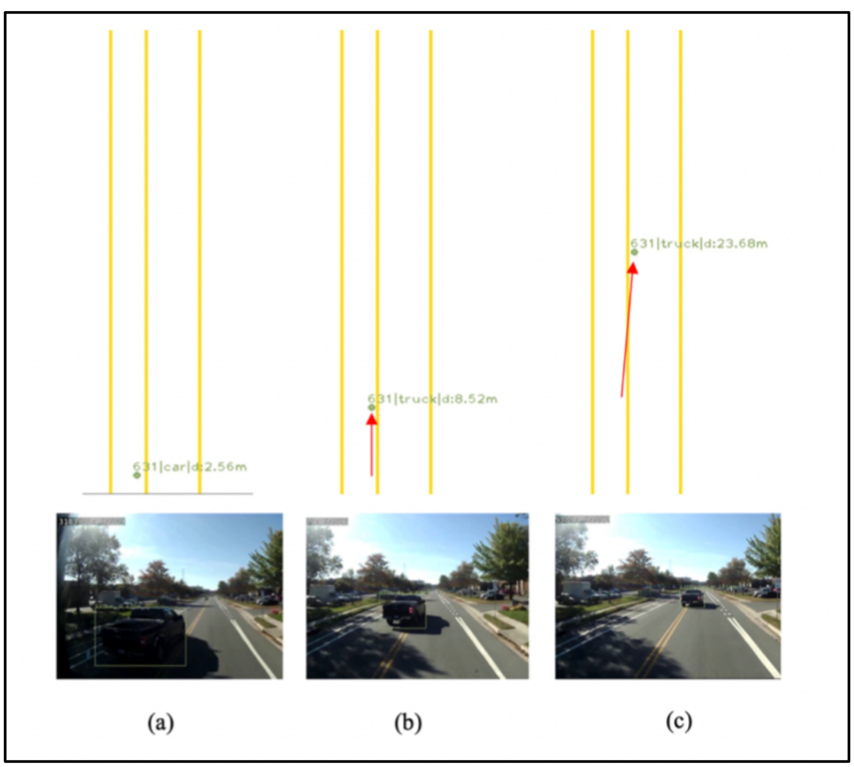

Figure 7. Diagrams with video images. Example of a vehicle cut-in front of the AV.

ANALYSIS OF TRAFFIC TRENDS

On the road, the LSAV is surrounded by numerous objects, and the way that these objects interact with the LSAV can help determine deployment safety implications. To analyze the interaction of LSAVs with various road agents, we plotted the traffic flow around the LSAV as a heatmap. To generate the holistic view of traffic density in the front and the rear of the LASV, we used 115 front camera and 131 rear camera videos. Furthermore, we aggregated 115 density arrays generated for front and 131 density arrays generated for the rear videos to get two 2D arrays (DF, DR) of shape (height, width, 1). We then normalized DF and DR independently, followed by Gaussian blurring with a 15x15 Gaussian kernel to get the final "flow" heatmap as depicted in Figure 8(a-b). Here, each pixel incorporates the density of traffic at that location for all front/rear videos. The distance resolution for both front and rear cameras is 36 meters. In each image, the location of the LSAV is shown as a brown dot marked "E" at location (0, 0). Followings are some of the observations from analyzing the front/rear traffic flow heatmap:

- From Figure 8(a), it is evident that there is a large concentration of traffic behind the LSAV. The vehicles stuck behind the LSAV would cut the lane to get ahead ("green arrow"). Additionally, the traffic building up behind the LSAV hindered the normal operation of the traffic by creating bottlenecks on the road.
- The previous observation is reinforced from the traffic distribution we achieved from the front camera. From Figure 8(b) we see that there is a large concentration of traffic ("red") right beside the LSAV as compared to any other area. This suggests a vehicle changing lanes from behind the LSAV ("green arrow") to overtake it due to the LSAV's slow operating speed.
- Additionally, from Figure 8(b), we can see that the traffic density is low in the lane where the LSAV is operating (i.e., vehicles are avoiding being in the same lane as the LSAV), increasing the traffic in other lanes and creating possible bottlenecks.

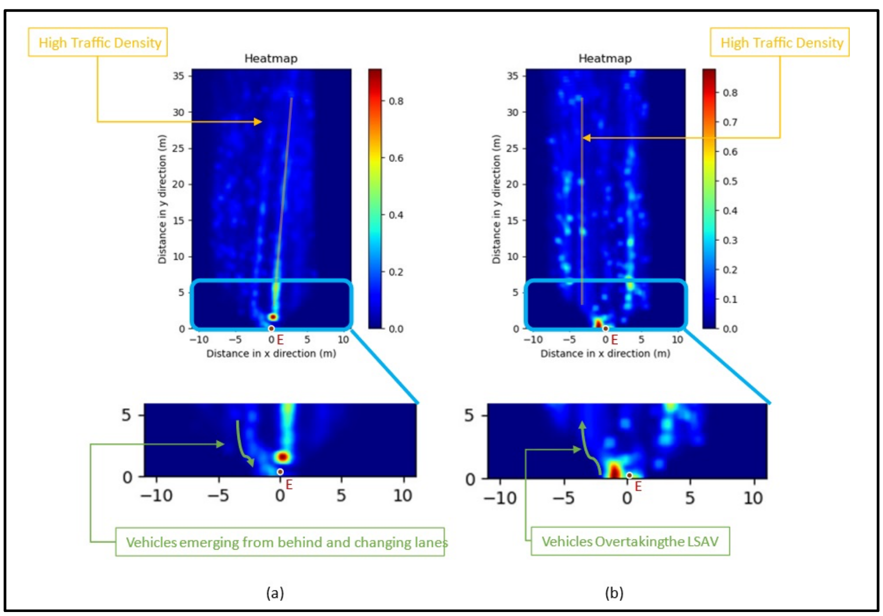

Figure 8. Traffic flow heatmap. (a) Rear camera and (b) front camera. "E" denotes the location of the LSAV.

CONCLUSION

LSAVs are expected to be deployed across the market very soon. In this work, we studied a real-world deployment of EasyMile LSAVs in Fairfax County. We developed and tested advanced CV algorithms to understand the interactions of LSAVs with various road agents. We specifically looked at vehicles in front and behind the LSAV. We investigated the behavior and movement of road agents such as cars and pedestrians around the LSAV. In the process, we developed a scene perception algorithm based on object detection and object tracking to measure kinematic behavior of these roadway agents with respect to the LSAV. We further demonstrated a process to summarize behavior of each of the agents over time and developed a composite understanding of the interaction between the two agents on the road.

Finally, we used these methods to demonstrate some key events such as following, cut-in front of vehicle, yield to pedestrian, etc.
We specifically looked at following behavior and Through experimentation, we quantitatively evaluated how (a) a long following vehicle queue may be generated behind the LSAV and (b) how vehicles try to change lanes and overtake the LSAV in front of the ego. In our experiment, we saw examples of both cases. The long queue suggests that LSAVs may reduce the overall traffic speed of the lane, hence creating traffic jams and reducing volumetric throughput. Also we have seen close following, occasional lane change, and overtaking by vehicles coming from behind; thus LSAV may create safety-critical situations [13] by encouraging such unnecessary driving maneuvers. This adds to LSAVs' operating principles and underlying algorithms, which define their perception and control methodologies. Our study shows that the LSAVs are often too overprotective and apply brakes during non-critical cut-in events. These conditions also raise the question of adaptability and acceptance of this new technology by the public and how other drivers should behave around an LSAV. Therefore, the entities deploying LSAVs need to consider these pros and cons before finalizing LSAVs' path and deployment plan. Similarly, LSAV developers should invest in developing technologies that would help these vehicles operate in mixed-traffic scenarios.

ACKNOWLEDGEMENT

The authors like to thank National Surface Transportation Safety Center for Excellence (NSTSCE) for funding this effort [14]. The team also like to thank Dr. Aditi Manke, Dr. Jon Antin, and Dr. Jon Hankey for their valuable feedback; Dr. Mike Mollenhauer, Jean Paul Talledo Vilela for their help in accessing the data.

REFERENCES

[1] Litman, T. (2021). *Autonomous vehicle implementation predictions implications for transport planning.* Victoria Transport Policy Institute.
[2] National Highway Traffic Safety Administration. (2020). *Automated vehicles for safety.* https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety
[3] Claypool, H., Bin-Nun, A., & Gerlach, J. (2017). *Self-driving cars: The impact on people with disabilities.* Ruderman Family Foundation.
[4] Coyner, K., Blackmer, S., Good, J., Lewis, P., & Grossman, A. (2021). *Low-speed automated vehicles (LSAVs) in public transportation* (TRCP Research Report 220). Transit Cooperative Research Program.
[5] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, *521*(7553), 436-444.
[6] Cai, Z., & Vasconcelos, N. (2019). Cascade R-CNN: High quality object detection and instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 43*(5), 1483–1498.
[7] Chen, K., Wang, J. Pang, J. Cao, Y. Xiong, Y. Li, X. Sun, S. Feng, W. Liu, Z. Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., … Lin, D. (2019). *MMDetection: Open MMLab detection toolbox and benchmark*. arXiv preprint arXiv:1906.07155.
[8] Papakis, I., Sarkar, A., & Karpatne, A. (2020.) GCNNMatch: Graph convolutional neural networks for multi-object tracking via sinkhorn normalization. arXiv preprint arXiv:2010.00067
[9] Bergmann, P., Meinhardt, T., & Leal-Taixe, L. (2019). Tracking without bells and whistles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 941-951).
[10] Papakis, I., Sarkar, A., & Karpatne, A. (2021, September). A Graph Convolutional Neural Network Based Approach for Traffic Monitoring Using Augmented Detections with Optical Flow. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)* (pp. 2980-2986). IEEE.
[11] Qin, Z., Wang, H., & Li, X. (2020). Ultra-fast structure-aware deep lane detection. arXiv:2004.11757.
[12] Sarkar, A., Engstrom, J., & Hanowski, R. (2022). Analysis of Car Cut-ins Between Trucks Based on Existing Naturalistic Driving Data. National Surface Transportation Safety Center for Excellence (NSTSCE)
[13] Sarkar, A., Krum, A., Hanowski, R., & Hickman, J. (2021, July). Responsibility Sensitive Safety Analysis of Truck Following in US Highway. In International Conference on Applied Human Factors and Ergonomics (pp. 119-126). Springer, Cham.
[14] Sarkar, A., Sundharam, V., Manke, A., & Grove, K. (2022). Camera-based Feature Identification for EasyMile Operation. National Surface Transportation Safety Center for Excellence (NSTSCE)