

Influence of Autonomous and Partially Autonomous Vehicles on Minnesota Roads

Rajesh Rajamani, Principal Investigator

Mechanical Engineering
University of Minnesota

MAY 2023

Research Project
Final Report 2023-23



To request this document in an alternative format, such as braille or large print, call [651-366-4718](tel:651-366-4718) or [1-800-657-3774](tel:1-800-657-3774) (Greater Minnesota) or email your request to ADArequest.dot@state.mn.us. Please request at least one week in advance.

1. Report No. MN 2023-23	2.	3. Recipients Accession No.	
4. Title and Subtitle Influence of Autonomous and Partially Autonomous Vehicles on Minnesota Roads		5. Report Date May 2023	
		6.	
7. Author(s) Andre Espindola, Lee Alexander, and Rajesh Rajamani		8. Performing Organization Report No.	
9. Performing Organization Name and Address Laboratory for Innovations in Sensing, Estimation and Control (LISEC) Department of Mechanical Engineering University of Minnesota 111 Church St SE, Minneapolis, MN 55455		10. Project/Task/Work Unit No. CTS #2022002	
		11. Contract (C) or Grant (G) No. (c) 1036342 (wo) 10	
12. Sponsoring Organization Name and Address Minnesota Department of Transportation Research Services & Library 395 John Ireland Boulevard, MS 330 St. Paul, Minnesota 55155-1899		13. Type of Report and Period Covered Final Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes http://mdl.mndot.gov/			
16. Abstract (Limit: 250 words) <p>This project focuses on experimental tests of the performance characteristics of autonomous vehicles (AVs) on highways and local roads in Minnesota. The project provides detailed data characterizing AV performance, which in turn can be used to inform the transportation community on implications for infrastructure maintenance, winter road maintenance, work zone guidelines, safety, and traffic capacity. The experimental work presented here makes use of a new autonomous vehicle purchased by the Center for Transportation Studies at the University of Minnesota. The key aspects of the autonomous functions of the vehicle studied in this project include winter performance and implications for road maintenance, characterization of the driving performance of the AV and its likely influence on safety, traffic flow and fuel economy, and the ability of the AV to handle work zones and the implications on changes needed to the guidelines for work zones. The project documents the major challenges and obstacles ahead in the way of true autonomy on Minnesota roads, but also outlines further areas for research with which it will be possible to facilitate the improvement of the capabilities of autonomous vehicles in Minnesota in the future.</p>			
17. Document Analysis/Descriptors Autonomous vehicles, Driver support systems, Winter maintenance, Work zones, Safety		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Alexandria, Virginia 22312	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 56	22. Price

Influence of Autonomous and Partially Autonomous Vehicles on Minnesota Roads

FINAL REPORT

Prepared by:

¹Andre Espindola

¹Lee Alexander

¹Rajesh Rajamani, Ph.D.

¹Laboratory for Innovations in Sensing, Estimation and Control (LISEC)
Department of Mechanical Engineering
University of Minnesota

May 2023

Published by:

Minnesota Department of Transportation
Office of Research & Innovation
395 John Ireland Boulevard, MS 330
St. Paul, Minnesota 55155-1899

This report represents the results of research conducted by the authors and does not necessarily represent the views or policies of the Minnesota Department of Transportation or the University of Minnesota. This report does not contain a standard or specified technique.

The authors, the Minnesota Department of Transportation, the Local Road Research Board, or the University of Minnesota do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to this report.

ACKNOWLEDGMENTS

The authors would like to thank the Local Road Research Board and the Minnesota Department of Transportation (MnDOT) for funding this project. Special thanks to Vic Lund, St. Louis County, and Jackie Jiran, MnDOT, for their guidance and support throughout the project.

We very much appreciate the contributions of the members of the project's Technical Advisory Panel (TAP), specifically Cory Johnson (MnDOT), Tara Olds (MnDOT), Ben Worel (MnDOT), Jeff Tabery (MnDOT), Guy Kolnhofer (Dodge County), Bradley Estochen (Ramsey County), Mitch Bartelt (Washington County), and Ethan Peterson (MDOT). Members of the TAP were very interested in the project and had many good ideas and suggestions to share.

Finally, we would like to thank Gina Bass, Brian Davis, and Kyle Shelton from the Center for Transportation Studies at the University of Minnesota for allowing us use of the MnCAV vehicle for this project and for facilitating access to the vehicle software through permission from Dataspeed.

TABLE OF CONTENTS

CHAPTER 1: Introduction and Background.....	1
1.1 Background.....	1
1.1.1 Autonomous Vehicle Background	1
1.1.2 MnCAV Vehicle Overview.....	3
1.2 Project Description	4
CHAPTER 2: Initial MnCAV Vehicle Testing	5
2.1 Purpose.....	5
2.2 Methods.....	5
2.2.1 MnROAD Testing	5
2.2.2 Local Highway Testing	6
2.3 Key Findings	7
2.3.1 Lateral Control Results	7
2.3.2 Longitudinal Control Results	11
CHAPTER 3: Winter Testing	15
3.1 Purpose.....	15
3.2 Methods.....	15
3.3 Key Findings	16
CHAPTER 4: Vehicle Following Tests	21
4.1 Purpose	21
4.2 Methods.....	21
4.3 Key Findings	23
CHAPTER 5: Work Zone Testing	30
5.1 Purpose	30
5.2 Methods.....	30

5.2.1 Cone and Barrel Testing	30
5.2.2 Temporary Lane Marker Testing	31
5.2.3 Conflicting Lane Marker Testing.....	32
5.3 Key Findings	32
5.3.1 Cone and Barrel Testing	32
5.3.2 Temporary Lane Marker Testing	35
5.3.3 Conflicting Lane Marker Testing.....	36
CHAPTER 6: Discussion	38
6.1 Winter Performance and Winter Road Maintenance	38
6.1.1 Recommendations	39
6.2 Comparison with Human Drivers.....	39
6.2.1 Benefits.....	41
6.3 Implications of AV's for Construction Zone Guidelines	42
6.3.1 Recommendations	43
6.4 Suggestions on Areas for Future Research Support	43
CHAPTER 7: Conclusion	45
REFERENCES.....	46

LIST OF FIGURES

Figure 1. MnCAV vehicle with various sensors marked.	3
Figure 2. Satellite image of the MnROAD Low-Volume Road (LVR)	5
Figure 3. Deviation from lane center for a straightaway test on the LVR.	7
Figure 4. Demonstration of poor lane boundary estimation from the Mobileye camera.....	8
Figure 5. Deviation from lane center for a turnaround test on the LVR.....	9
Figure 6. Instance of MnCAV vehicle crossing inside lane boundary.	10

Figure 7. Deviation from lane center for traversal of a curve of a local highway.....	10
Figure 8. Performance results for longitudinal test with slower traffic moving at constant speed.	11
Figure 9. Performance results for longitudinal test with stop and go traffic.	12
Figure 10. Performance results for longitudinal test with stop and go traffic.	13
Figure 11. Wet snowfall road surface conditions.	16
Figure 12. Lane line detection quality for left lane (left) and right lane (right) for wet snowfall.	17
Figure 13. Lateral controller performance in wet snowfall.	17
Figure 14. Moderate snowfall road surface conditions.	18
Figure 15. Lane line detection quality for left lane (left) and right lane (right) for moderate snowfall.	18
Figure 16. Lateral controller performance in moderate snowfall.	19
Figure 17. Heavy snowfall road surface conditions.	19
Figure 18. Lane line detection quality for left lane (left) and right lane (right) for heavy snowfall.	20
Figure 19. Estimated and directly measured acceleration signal of the MnCAV vehicle.	23
Figure 20. Test 1, acceleration of MnCAV vehicle and preceding vehicle during light traffic.	23
Figure 21. Test 2, acceleration of MnCAV vehicle and preceding vehicle during light traffic.	24
Figure 22. Test 1, acceleration of MnCAV vehicle and preceding vehicle during medium traffic.	25
Figure 23. Test 2, acceleration of MnCAV vehicle and preceding vehicle during medium traffic.	25
Figure 24. Test 1, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.	26
Figure 25. Test 2, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.	27
Figure 26. Test 3, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.	27
Figure 27. Cone test setup on top of lane lines with 6ft spacing.	31
Figure 28. Temporary marker test setup.	31
Figure 29. Conflicting lane marker setup.	32
Figure 30. Lane line detection quality for baseline (left) and cones placed behind lane lines (right).	33
Figure 31. Cone lane line detection quality for baseline and 3ft, 6ft, and 10ft spacings.	34

Figure 32. Barrel lane line detection quality for baseline and 3ft, 6ft, and 10ft spacings.	35
Figure 33. Lane line detection quality for temporary lane markers.	36
Figure 34. Image showing detection of primary markers and no detection of conflicting marker.	37
Figure 35. Line Detection Quality with the Mobileye Camera System on MnRoad.	38
Figure 36. Camera view from inside the MnCAV vehicle while driving an hour after snowfall.	39
Figure 37. Example data segment showing accelerations of preceding vehicle and MnCAV vehicle under heavy traffic conditions.	41
Figure 38. Tests with traffic cones to evaluate what density of cones prevents the Mobileye lane boundary detection system from working.	42

LIST OF TABLES

Table 1. Definitions of autonomous driving levels	2
Table 2. Definitions of snow cover levels.....	15
Table 3. Definitions of traffic levels.	22
Table 4. Average acceleration and deceleration magnitudes for MnCAV vehicle and preceding vehicle.	28

EXECUTIVE SUMMARY

The goal of this study was to investigate the impact of autonomous and partially autonomous vehicles on Minnesota roads. In particular, the research team studied the readiness of current autonomous driving technology for various challenges encountered on Minnesota roads, such as winter driving and work zones. From this research, the team was able to develop recommendations for best practices in road maintenance as well as areas of potential future research focus. Additionally, the impact of adaptive cruise control (ACC) systems on traffic flow was studied, from which recommendations for further research were also made. To perform the tasks of this research project, the team used the MnCAV vehicle, an experimental autonomous vehicle designed by Dataspeed Inc. that was purchased by the Center for Transportation Studies at the University of Minnesota. The team performed initial testing when the vehicle was first received. This testing showed that the lateral (i.e. steering) controller performed well on straight roads but struggled on curved roads. In addition, it was found that the longitudinal (i.e. speed and spacing) controller performed well above 10 mph but was jerky and unstable below 10 mph. To resolve the longitudinal control issue before performing traffic flow tests, the team developed and improved the longitudinal controller, which resolved these initial issues. From the winter driving tests, it was determined that even a small amount of snow on the road surface can visually impede detection of the lane markings, making lateral control challenging. Thus, further research on how to perform lateral control on snow-covered roads will be necessary for states like Minnesota with significant winter seasons. Ideas in this regard were outlined in the report. From the traffic flow tests, it was found that the ACC system on the vehicle was capable of attenuating the propagation of acceleration waves through traffic, also known as string stability. This provided evidence that well-designed ACC systems can reduce the severity of backups if enough vehicles on the road use this technology. Further research will be required to determine the impact on traffic capacity for various percentages of vehicles on the road equipped with ACC. Finally, work-zone tests showed that many challenges were presented to autonomous driving systems in these scenarios. Humans navigate work zones by following cones and signs; however, most lane-centering systems rely entirely on lane markings to navigate and currently are not capable of detecting cones and signs. The team found that cones and barrels can block detection of lane markings when placed densely enough in a visually impeding position, and that conflicting lane markers create confusion for the system, which can only follow one set of lane markings at a time. Finally, it was found that temporary lane markings were not detected by the lane boundary estimation system, meaning these markings were not suitable for autonomous vehicles to use with current technology. Overall, this study showed that there were many obstacles in the way of true autonomy on Minnesota roads, but that with further research in the areas mentioned in this report, it will be possible to facilitate the improvement of the capabilities of autonomous vehicles in Minnesota in the future.

CHAPTER 1: INTRODUCTION AND BACKGROUND

1.1 BACKGROUND

1.1.1 Autonomous Vehicle Background

The concept of autonomous vehicles has been around since the 20th century, but only in recent years has this idea become much closer to reality through advancements in perception and artificial intelligence. While we are likely still many years from completely human-less driving with no geographical restrictions, partially autonomous features, such as adaptive cruise control (ACC), are now available on most new vehicles. Automatic lane centering is another feature that is becoming increasingly common on roads. Furthermore, companies such as Waymo and Cruise have begun to offer fully autonomous taxis (with no human driver) in small geographic regions within Phoenix and San Francisco (Waymo 2021), and Tesla has released beta software for fully autonomous driving (with a human safety driver) to thousands of its customers (Tesla). Given all these new features and stages within the autonomous driving industry, a method for classifying these different features has been developed. Table 1 provides definitions for the six levels in which an autonomous driving system (ADS) may be classified. These definitions are based on the dynamic driving tasks (DDTs) that the system may perform, which includes the collection of tasks required to maneuver the vehicle in a safe and effective manner such as lateral control, longitudinal control, navigation, etc. These DDTs are performed within an operational design domain (ODD), which encompasses the set of conditions under which the vehicle is designed to operate autonomously. This may include factors such as road types, weather conditions, geofence boundaries, and lighting conditions (SAE International, 2021).

Table 1. Definitions of autonomous driving levels

Level of Autonomy	L0	L1	L2	L3	L4	L5
Role of User	-Driver at all times -Performs entire DDT	- Performs all DDTs not performed by ADS	- Performs all DDTs not performed by ADS - Intervenes when required	- Does not perform any DDTs while ADS is engaged - Intervenes when required	- May request ADS disengages	- May request ADS disengages
Role of ADS	-Does not perform any DDT on sustained basis	- Performs some but not all DDTs (longitudinal OR lateral control)	- Performs both longitudinal and lateral control DDTs - Does not perform all DDTs	- Performs all DDTs within the ODD	- Performs all DDTs within the ODD - Automatically achieves minimal risk condition when failure is detected	- Performs all DDTs on all drivable roads - Automatically achieves minimal risk condition
What it Means	-Driver steers AND brakes/accelerates	- Driver steers OR brakes/accelerates - ADS performs remaining DDTs	- ADS steers AND brakes/accelerates - Driver must be ready to take over - Driver does whatever ADS can't do (ex-lane change)	- ADS performs all driving maneuvers within ODD - Driver must be ready to take over	- ADS performs all driving maneuvers within ODD - Driver becomes a passenger	- ADS does all driving maneuvers on all drivable roads - Driver is a passenger

Vehicles today that are enabled with only adaptive cruise control would be classified as L1, while vehicles that also have lane-centering features would be classified as L2. Waymo and Cruise vehicles would be classified as L4 since there is no human driver present when the vehicle is operating. There are currently no L5 ADS systems in existence.

Progress on autonomous driving technology has accelerated greatly in the past few years and will likely continue to grow rapidly with advancements in AI that will potentially lead to safe L5 systems being developed in the near future.

1.1.2 MnCAV Vehicle Overview

The Center for Transportation Studies at the University of Minnesota recently purchased an autonomous vehicle (AV) for research purposes, known as the MnCAV vehicle. The vehicle itself is a Chrysler Pacifica fitted with a vast sensor suite that allows the vehicle to perceive the surrounding environment. Information from these sensors is fed into control system software, which produces acceleration, braking, and steering commands to maneuver the vehicle. These commands are passed through a drive-by-wire (DBW) system, which is responsible for actuating the commands on the vehicle hardware itself. The system software for reading sensor data and producing control commands is implemented on the Robot Operating System (ROS), which is a distributed system that allows for modularity within the various subsystems. The sensor suite, initial software, and drive-by-wire system were all designed and implemented by Dataspeed Inc, based in Rochester Hills, MI.

The MnCAV vehicle features an array of sensors, including front and rear facing radar sensors, a front facing and two side facing Ouster lidar sensors, a Mobileye camera sensor, four external Flir RGB cameras, a thermal camera, an interior camera, a differential GPS sensor, and a V2V radio. Figure 1 shows an image of the MnCAV vehicle with the locations of the various sensors marked.




Figure 1. MnCAV vehicle with various sensors marked.

The sensors used most in this project were the front facing radar, the three lidars, and the Mobileye sensor. The radar sensor was used to measure the distance and relative velocity of the preceding vehicle in traffic, which was fed into the adaptive cruise control (ACC) system on the vehicle to control the vehicle speed. The lidar sensors were used for the same purpose when the preceding vehicle was

traveling below 5 mph. This was necessary because slow moving objects were filtered out of the radar detection to prevent stationary objects such as overhead signs from being misinterpreted as stopped vehicles by the system. The Mobileye sensor uses vision techniques to identify lane markings on the road. This sensor returns the estimated locations of lane boundaries relative to the vehicle frame, as well as a measurement of detection quality on an enumerated scale of 0-3 (0 meaning no detections, 1 meaning poor detections, 2 meaning good detections, and 3 meaning excellent detections). Information from this sensor is then fed into the lateral control system, which steers the vehicle into the center of its lane.

1.2 PROJECT DESCRIPTION

This project focused on performing experimental tests with the MnCAV vehicle on highways and local roads in Minnesota. The goal of the project was to provide detailed state-specific data characterizing AV performance, which will inform the transportation community on implications for infrastructure maintenance, winter road maintenance, work zone guidelines, road safety, and traffic capacity. The key aspects of the autonomous functions of the vehicle studied in this project were as follows:

1. Capabilities of the longitudinal and lateral controllers on the vehicle: How well do the controllers implemented on the MnCAV vehicle perform to control the speed and steering of the vehicle? Are the controllers smooth, comfortable, and safe, or does further work need to be done to improve the functioning of these controllers?
2. Winter performance and winter road maintenance: The ability of the lane-keeping-system (LKS) to function in Minnesota winter weather was studied, since all current LKS systems rely on cameras to find the lateral position of the vehicle with respect to the lane markers. The capabilities of the camera-based system to function in the presence of partially snow-covered and fully snow-covered environments was documented. This data will provide answers to the question: Will the level of winter road maintenance needed have to significantly go up to accommodate AVs?
3. Characterization of the driving performance of AVs and their likely influence on safety, traffic flow and fuel economy: Collection and analysis of data comparing AVs and human drivers in terms of lane keeping, time gaps, shock wave propagation, acceleration/ deceleration magnitudes and performance in low, moderate, and heavy traffic.
4. Handling of work zones: The ability of the AVs to handle work zones, which may have conflicting or multiple sets of lane markers in their vicinity, may require lane changes, may require other responses to work zone signs, and may even require transition to manual driving. The data from this project will help inform the transportation community about how guidelines on work zones may need to change and whether lane markers and signs in work zones would have to become more precise to support AVs.

CHAPTER 2: INITIAL MNCAV VEHICLE TESTING

2.1 PURPOSE

Upon first receiving the MnCAV vehicle from Dataspeed, initial testing was performed with the purpose of determining the capabilities and limitations of the MnCAV vehicle, gaining understanding of how to operate the system effectively, and providing baseline data describing vehicle performance in ideal conditions to compare with future testing in non-ideal conditions. This was important knowledge for the research team to obtain so that future testing plans could be developed in accordance with the capabilities of the system, and also so that potential issues with the system could be addressed and resolved. The focus of these tests centered on analyzing the performance of the longitudinal (i.e. speed) and lateral (i.e. steering) controllers of the system, as well as the perception capabilities of the sensors and estimation of the algorithms equipped on the vehicle.

2.2 METHODS

2.2.1 MnROAD Testing

The first set of initial testing was done on a closed track at the MnROAD facility in Albertville, MN. The closed track, low-volume road (LVR) at MnROAD consists of two long straightaway sections connected at each end by circular turnarounds, as pictured in Figure 2.



Figure 2. Satellite image of the MnROAD Low-Volume Road (LVR)

Two sets of tests were performed to assess the performance of both the lateral and longitudinal controllers on the vehicle. First, to assess the lateral controller, the MnCAV vehicle was set to fully autonomous mode and driven around the full LVR circuit with no lead vehicle. The set speed of the vehicle was controlled by the safety driver, and was set between 30 and 40 mph on the straightaway sections, and between 5 and 10 mph on the curved sections. Because there was no lead vehicle in front of the MnCAV vehicle, the speed control behaved exactly like a traditional cruise control. To analyze the performance of the lateral controller, measurements of deviation from the lane center were recorded from the Mobileye camera sensor, which was responsible for estimating the position of lane markings with respect to the body-fixed frame (i.e. from the perspective of the MnCAV vehicle). Because the Mobileye system gives us position measurements to both the left and right lane markings, the deviation from the lane center can be calculated according to Equation 1, where δ is deviation from the lane

center, d_l is the measured distance to the left lane marker, and d_r is the measured distance to the right lane marker:

$$\delta = \frac{d_l - d_r}{2} \quad (1)$$

By measuring the deviation from lane center on the straightaway sections of the circuit, the research team was able to assess the baseline performance of the lateral controller in ideal conditions. Optimal performance would show no deviation from lane center throughout the traversal of the straightaway. Additionally, the performance of the lateral controller during sharp turns was studied on the circular turnaround section of the circuit. The same quantitative deviation measurement was recorded in this section, as well as qualitative observations from the research team in the vehicle, which was presented in the form of video clips from the onboard cameras and 3D visualization of the vehicle environment.

A second set of tests were performed to study the performance of the longitudinal control of the vehicle. These tests were performed exclusively on the straightaway section of the LVR circuit, and featured the use of a lead vehicle, which was used to simulate various traffic scenarios the MnCAV vehicle may encounter in the real world. These tests focused on studying how the MnCAV vehicle would behave when approaching a slower moving vehicle, as well as how the ACC system would perform with cyclical deceleration and acceleration as experienced in stop and go traffic. To perform these tests, one member of the research team manually drove the lead vehicle to simulate one of the traffic scenarios listed above, while another member set the MnCAV vehicle to autonomous mode at a given set speed depending on the scenario being tested. Measurements of distance to the lead vehicle as well as velocity and acceleration of the MnCAV vehicle were taken using data obtained from the onboard radar, lidar, and IMU sensors and processed through the estimation algorithms provided by Dataspeed. From this data, the team was able to analyze the smoothness and safety of the longitudinal control system.

2.2.2 Local Highway Testing

In addition to the closed-circuit testing at MnROAD, initial tests were also conducted on local highways in the vicinity of MnROAD. These tests were performed to assess the baseline performance of the MnCAV vehicle in a real-world, non-contrived environment. Studies of both the lateral and longitudinal controllers were performed on these roads. The lateral controller was assessed by taking measurements of deviation from lane center on roads which featured slight curves, providing further insight to the performance of the lateral controller than what the straightaways and sharp curved sections of the LVR were able to provide. In addition, the longitudinal performance was studied through interactions of the system with real-world traffic. In particular, one interaction with a decelerating truck was analyzed using measurements of distance, velocity, and acceleration from the MnCAV vehicle. The same analysis techniques used at MnROAD were applied to the acquired data, which gave the team insight into the performance of both the lateral and longitudinal controllers in a low-traffic, real-world scenario.

2.3 KEY FINDINGS

2.3.1 Lateral Control Results

Results from initial lateral control testing are broken down into three sections: 1) performance on the straightaway section of the LVR, 2) performance on the curved turnaround section of the LVR, and 3) performance on local highways. First, results from the straightaway section of the LVR are presented. The MnCAV vehicle demonstrated excellent lateral control performance on the straightaway section of the LVR. Qualitative observations from the research team report that the steering behavior of the vehicle was smooth and stable throughout the duration of the testing, with the vehicle seemingly staying directly in the center of the lane as desired with no instability. These observations were consistent with the quantitative data the team acquired. Figure 3 shows the deviation from lane center of the MnCAV vehicle during a straightaway test. The blue points show the measured deviation from lane center, while the red points show the boundaries of the lane within which the vehicle must stay.




Figure 3. Deviation from lane center for a straightaway test on the LVR.

This result shows fast convergence of the vehicle to the lane center within 5 seconds from the beginning of the test, followed by stable, low noise lateral control for the duration of the test, with a maximum deviation from lane center of about 8cm, which is unperceivable to the passengers in the vehicle. This demonstrates that the MnCAV vehicle has excellent lateral control performance on straight roads given good estimation of the lane boundaries from the Mobileye system.

Tests on the curved sections of the LVR, however, showed very poor performance from the lateral control system of the MnCAV vehicle. Throughout the entirety of testing on the turnaround sections of

the LVR, the MnCAV vehicle struggled significantly to stay within the lane boundaries, which is due to a combination of various factors. First, the Mobileye system struggled in its estimation of the lane boundaries. Results showed significant instability in the lane boundary estimation, including dropped detections, incorrect detections, and rapid bouncing between various lane trajectories. For instance, Figure 4 shows a case where the true lane boundaries are curving to the left, but the Mobileye estimation shows the boundaries curving to the right.




Figure 4. Demonstration of poor lane boundary estimation from the Mobileye camera.

This results in the MnCAV vehicle following a trajectory curving to the right as shown in orange in the above figure, which is obviously not the desired trajectory when the road is actually curving to the left. Rapid bouncing between various lane trajectories and proposed vehicle trajectories was observed, making the lateral control very unstable. Another factor impacting the poor performance of the lateral controller was the sharpness of the curves which the MnCAV vehicle was navigating. The turnarounds on the LVR feature a relatively small radius of curvature, leading to much sharper turns than what would be observed in most real-world situations. This could certainly have an impact on the poor lane estimation that was observed, as the algorithms used for lane estimation within the Mobileye system may not be stable for turns below a certain radius of curvature. Finally, the lateral controller itself may have factored into the poor performance of the system. There is evidence based on real-world testing results, presented below, that the lateral controller does not track the desired trajectory effectively below a certain radius of curvature, which may also contribute to the poor lane centering seen in the LVR turnaround tests. Figure 5 shows the results of a test in the turnaround section, showing the performance of the lateral control system.




Figure 5. Deviation from lane center for a turnaround test on the LVR.

This result clearly shows significant error in the lane centering of the vehicle, crossing over the lane boundary multiple times and requiring the safety driver to take over to correct the vehicle, which is seen in the latter half of the test.

Lateral control testing on local highways also demonstrated some concerns with the lane centering system. Tests were performed on highways featuring mostly straight sections but also containing various slight curves. Results from the straight sections of highway matched what was observed at MnROAD, with the MnCAV vehicle showing excellent, stable performance with minimal deviation from the lane center. However, tests on the curved sections of highway showed the MnCAV vehicle crossing or nearly crossing the lane boundaries, indicating inadequate performance from the vehicle. Figure 6 shows an instance where the vehicle was crossing over the inside lane boundary while traversing a slight curve, which, in a situation involving multiple lanes of traffic, could bring the MnCAV vehicle uncomfortably close to colliding with another vehicle.




Figure 6. Instance of MnCAV vehicle crossing inside lane boundary.

Figure 7 shows quantitative results of lane center deviation for this test, which shows the vehicle crossing over the inside lane boundary.




Figure 7. Deviation from lane center for traversal of a curve of a local highway.

This instance of poor lane centering performance appears to be due to the lateral controller itself, rather than estimation of the lane boundaries as was seen during the turnaround tests at the LVR. The Mobileye system produced very stable and seemingly accurate estimates of the lane boundaries, and produced an acceptable proposed trajectory for the vehicle as shown in orange in Figure 6. However, the MnCAV vehicle did not follow the proposed trajectory well. Ideally, the center of the vehicle should be directly on top of the proposed trajectory. However, in this instance the vehicle is shifted to the right, which is why it is observed crossing the lane boundary. This indicates that the issue does not lie with the perception system but rather with the controller itself.

Overall, initial tests of the lateral control system showed excellent performance on straight sections of road but concerning behavior on curved sections of road. Particularly, on very sharp turns the perception system was found to be incapable of producing satisfactory estimates of the road boundaries, leading to poor desired vehicle trajectories, and on shallower turns the controller was found to be incapable of following the desired vehicle trajectory well, resulting in the vehicle crossing over the lane boundary. To improve the lateral control of the system, work must be done to improve both the perception system on sharp turns as well as the controller itself to better follow the desired trajectory.

2.3.2 Longitudinal Control Results

As described in the methodology section, tests of the longitudinal control system were also done at both the LVR and on local highways. Testing on the LVR featured simulations of various traffic scenarios using a manually driven lead vehicle. The first test was done with the lead vehicle traveling at a constant speed slower than the set speed of the MnCAV vehicle, and then speeding up once the MnCAV vehicle reached steady state. The goal was for the MnCAV vehicle to approach the lead vehicle from behind and be forced to slow down to match the lead vehicle speed, and then accelerate to match the acceleration of the lead vehicle. The MnCAV vehicle was set to 40 mph and the lead vehicle was traveling at 30 mph. Once steady state was reached, the lead vehicle accelerated up to about 40 mph. Figure 8 shows the performance results from this first test.




Figure 8. Performance results for longitudinal test with slower traffic moving at constant speed.

This test shows that the MnCAV vehicle successfully slowed down to match the speed of the lead vehicle, and then successfully accelerated to follow the lead vehicle. The following distance of the MnCAV vehicle never crossed below the set minimum distance, which demonstrates successful longitudinal control. The one negative observed in this test is that the acceleration signal of the MnCAV

vehicle appears to be fairly noisy in its tracking of the desired acceleration. The desired acceleration is the output from the high-level longitudinal controller (which converts distance and relative velocity measurements into acceleration commands), shown in the bottom chart in Figure 8. This indicates that the low-level controller, which converts from acceleration commands to pedal commands, may not be performing optimally to track the desired acceleration. This could potentially lead to jerky behavior in the vehicle even if the high-level desired acceleration signal is relatively smooth.

The second longitudinal test featured the MnCAV vehicle in a simulation of stop and go traffic. This situation required the MnCAV vehicle to slow down to a complete stop or near a complete stop, and then reaccelerate to follow the lead vehicle, repeating in a cyclical manner. Figure 9 shows the results from this test.

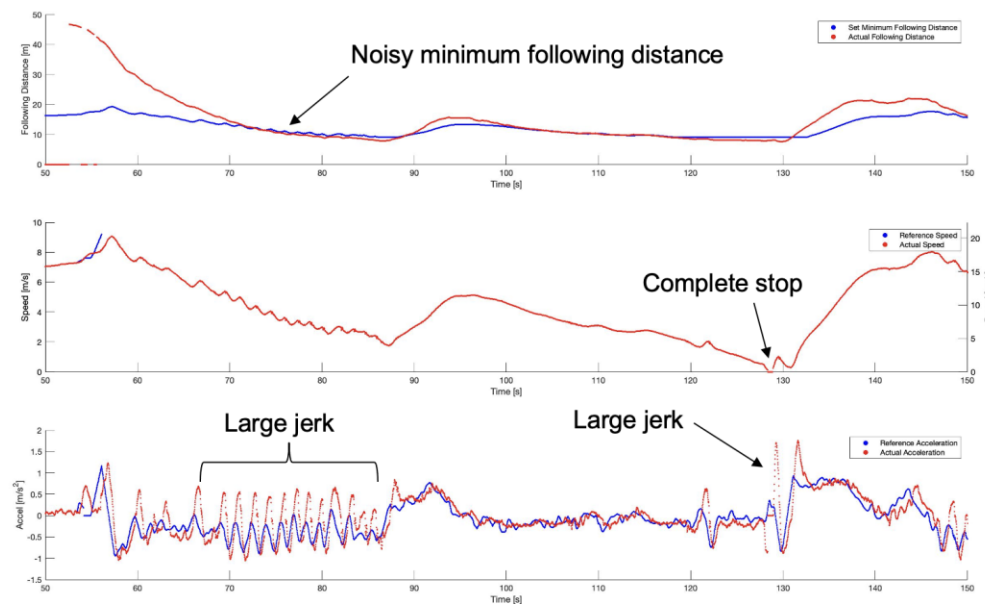


Figure 9. Performance results for longitudinal test with stop and go traffic.

The results from this test show much poorer performance of the longitudinal control system compared to the previous test. When approaching the lead vehicle, which is traveling initially below 5 mph, the MnCAV vehicle did successfully attempt to slow down to match the speed, however significant jerky behavior was observed in the MnCAV vehicle. This led to the MnCAV vehicle uncomfortably lurching continually towards the lead vehicle, which ended only once the lead vehicle accelerated, at which point the MnCAV vehicle successfully accelerated to follow. Again, when the lead vehicle slowed down and came to a complete stop, the MnCAV vehicle did successfully come to a complete stop, but only for a split second, and only after significant jerk and lurching behavior was observed. After a split second of being completely stopped, the MnCAV vehicle continued to lurch forward towards the lead vehicle. The lead vehicle then accelerated away and the MnCAV vehicle successfully followed the lead acceleration. Overall, the longitudinal controller was not smooth at all during this test, resulting in a very uncomfortable ride and nearly forcing a disengagement of the system. The source of the problem found

in this test was identified as poor estimation of the relative velocity at low speeds. Because the radar sensor on the vehicle cannot be used to detect objects below a certain threshold speed (to prevent stationary objects such as overhead passes from being identified as stopped vehicles), the system relies entirely on lidar segmentation to identify slow moving vehicles. The problem with this is that the lidar clustering algorithm implemented in the vehicle software is fairly noisy, meaning its detection of objects surrounding the vehicle is not very stable or precise. This noise then gets fed into the longitudinal controller, which results in unstable behavior as observed in this test. Significant work was invested by the research team to try to resolve this problem prior to working on the vehicle following task for this project. This work was successful in significantly improving the performance of the longitudinal controller, which will be discussed in Section 4.

Finally, the longitudinal controller was also tested in the real world on local highways. One instance of the vehicle approaching a turning truck is presented below. The MnCAV vehicle, initially traveling at 50 mph, approached the truck from behind, which was slowing down to below 15 mph to make a right turn. This forced the MnCAV vehicle to slow down to avoid colliding with the truck. The results from this test are presented in Figure 10.




Figure 10. Performance results for longitudinal test with stop and go traffic.

This test showed good performance of the MnCAV vehicle slowing down for the turning truck. While the actual following distance did fall below the set minimum following distance, it only very slightly did, and the controller did a good job following the set minimum following distance overall. The slowdown was also observed to be very smooth, with no jerky behavior occurring while slowing down. This shows that the longitudinal controller can be effective in a real world scenario when forced to decelerate for slower moving traffic.

Overall, the longitudinal controller was found to perform well at speeds above 10 mph, decelerating smoothly for slower moving vehicles and accelerating smoothly to follow the lead vehicle. The MnCAV vehicle was also shown to be capable of maintaining a safe and constant following distance behind a lead vehicle and to follow the set speed from the safety driver well. However, below 10 mph, the longitudinal controller was found to be extremely jerky and unstable, making the system unusable in heavy stop and go traffic. The source of this problem was found to be noisy relative velocity measurements from the lidar segmentation system. As stated earlier, the research team put in significant effort to improve the longitudinal control system in this scenario, which was proved to be successful and allowed the team to operate the vehicle in stop and go traffic during the vehicle following task of this project.

CHAPTER 3: WINTER TESTING

3.1 PURPOSE

One of the primary challenges facing the implementation of autonomous driving systems in snowy states such as Minnesota is finding a solution which allows for safe lateral control on snow covered roads. Because all autonomous driving systems operating today rely to an extent on detection of lane markings for lateral control, with a majority relying entirely on lane marking detection, snow cover on roads poses a significant threat to the functionality of autonomous driving systems (Zang, S. et al., 2019). The purpose of the winter driving task in this project was to analyze the capabilities of the MnCAV vehicle to drive autonomously in varying levels of winter driving conditions. The team was interested in better understanding the causes of failure of the system during winter driving conditions, from which conclusions could be drawn about best practices for road maintenance during snow events.

3.2 METHODS

To investigate the effects of snow cover on the performance of the lane marking detection system of the MnCAV vehicle, the research team drove the MnCAV vehicle on local freeways with various levels of snow cover. The team focused on gathering data in three different categories of snow cover, as defined in Table 2.

Table 2. Definitions of snow cover levels.

Snow Cover Category	Definition
Light	Snow cover is present on the road surface, but the lane markings are still clearly visible to the human eye
Medium	Snow is covering some lane markings, but some lane markings are still visible to the human eye
Heavy	Snow is covering all lane markings, no lane markings are visible to the human eye

When conditions allowed, the MnCAV vehicle was set to autonomous mode, but when lane lines were undetectable by the system, the vehicle was manually driven. The team gathered data from the Mobileye camera sensor, specifically measurements of lane line detection quality, which gave indication

as to how confidently the system was able to estimate the position of lane markings on an enumerated scale from 0 (no detections at all) to 3 (very confident detections). From this data, the team was able to draw conclusions about the performance of lateral controllers during winter driving and provide recommendations for next steps in addressing this challenge.

In addition to the study of lane marking detection, the research team also investigated the ability of the longitudinal controller to perform in various levels of active snowfall. The longitudinal controller relies on detection of preceding vehicles to adjust the vehicle speed accordingly, and thus it is important to know whether or not active snowfall will inhibit the detection of preceding vehicles, presenting a serious challenge to the functionality of the longitudinal control system. Preceding vehicles are detected primarily using the radar sensor equipped on the vehicle, and so the team acquired ranging data from the radar to examine how far the sensor could detect preceding vehicles in light, medium, and heavy active snowfall.

3.3 KEY FINDINGS

The first winter driving test took place in wet snowfall which melted on contact with the road. This created a wet driving surface that was slightly reflective, but not to an extent that would inhibit the clarity of lane lines. Figure 11 shows an image from the vehicle demonstrating the conditions of the road surface.



Figure 11. Wet snowfall road surface conditions.

To study the impact of the snowfall on lane marker detection, measurements of lane line detection quality were recorded from the Mobileye system. The measurements from this test are presented in Figure 12.

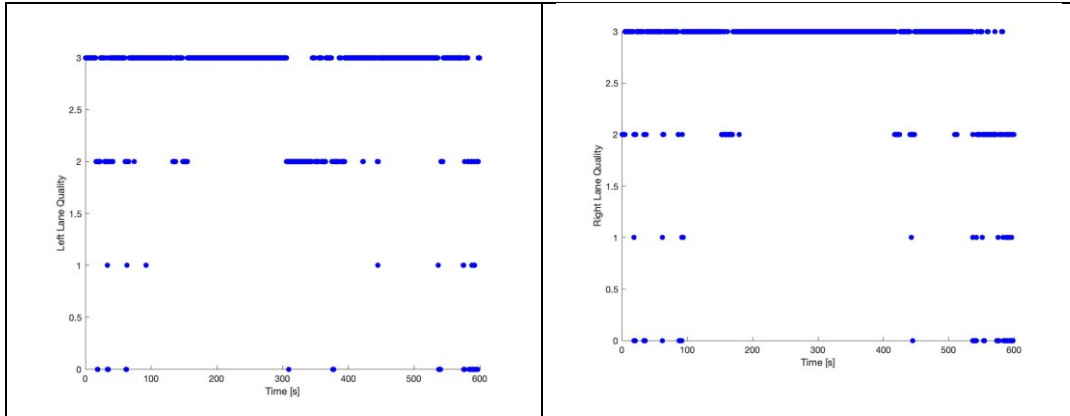


Figure 12. Lane line detection quality for left lane (left) and right lane (right) for wet snowfall.

The results show that for the vast majority of the test the Mobileye detected the lane markers with excellent quality, meaning the system is very confident in its estimation of the lane marker locations. This result matches well with baseline data acquired in ideal conditions at MnROAD, showing that the wet snowfall has no detrimental impact on the ability of the Mobileye system to detect the lane markings. This result is as expected, as from the human perspective there was no reduction in the clarity of the lane lines compared to ideal conditions. Figure 13 shows the performance of the lateral controller in this scenario, which demonstrates adequate performance of the lane centering system, staying within the lane boundaries. The slight deviations from the lane center are caused by problems with the lateral controller discussed in the previous section rather than issues with the lane boundary estimation.

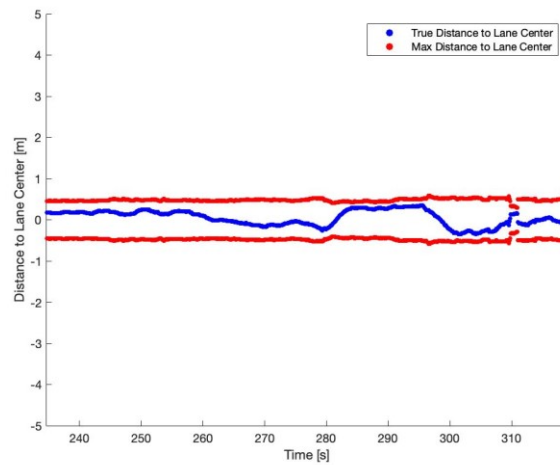


Figure 13. Lateral controller performance in wet snowfall.

The second winter driving test was done with a moderate amount of snow present on the road surface. In these conditions, the lane markers were somewhat visually impeded by the snowfall, but were for the most part still visible to the human eye. Figure 14 shows the road conditions from the perspective of the MnCAV vehicle.



Figure 14. Moderate snowfall road surface conditions.

Measurements of lane line detection quality for this test are presented in Figure 15.




Figure 15. Lane line detection quality for left lane (left) and right lane (right) for moderate snowfall.

These results show much more variation in the quality of lane detection compared to the wet snowfall test. The detection quality appears to jump almost constantly between 0 (no detections at all), 1 (poor detection), and 2 (good detection). While the system does have plenty of data points with good detection quality, because the detection was so inconsistent and would constantly jump to poor or no detections, the MnCAV vehicle was not capable of driving autonomously in these conditions and required manual driving for most of the test. Autonomous driving was tested for a short period, however, with the performance of the lateral controller shown in Figure 16.




Figure 16. Lateral controller performance in moderate snowfall.

This result clearly shows that the system was not able to control the vehicle adequately in these conditions, as the vehicle deviated from the lane boundary multiple times, forcing the safety driver to disengage the system and drive the vehicle manually. This means that the Mobileye system was able to detect the lane markers even in the presence of moderate snow but was simply not consistent enough to make autonomous driving feasible. Thus, we can conclude that even if the lane markers are only partially impeded, this has a significant impact on the ability to detect the markers and perform autonomous lateral control.

The third winter driving test featured heavy snowfall on the road which completely visually blocked the lane markers, even to the human eye. Figure 17 shows the road conditions from the MnCAV vehicle.



Figure 17. Heavy snowfall road surface conditions.

Measurements of lane line detection quality for this test are presented in Figure 18.




Figure 18. Lane line detection quality for left lane (left) and right lane (right) for heavy snowfall.

The results show that for the vast majority of the test, there were no detections of lane markers from the Mobileye system. Occasionally a spot was passed where lane markers were visible because there was no snow covering them, but for the vast majority of the test where no markers were visible, no detections were made from the Mobileye system. This is the expected result of the test, as the Mobileye system relies entirely on visual perception of the lane markers to estimate the lane boundaries. This indicates that if the lane markers are entirely blocked visually, they will be completely invisible to the detection system. This renders the lateral control system unusable since it relies entirely on lane line detection to function. Thus, for lateral control to be possible in moderate to severe winter driving conditions, a solution must be developed which allows for accurate estimation the lateral position of the vehicle on the road without the need for visual perception of the lane markings.

Overall, these results show that even partial visual impediment to the lane markings can result in significant reduction in detection quality, making lateral control very challenging in winter conditions. These tests indicate that a lateral control system which is entirely reliant on vision-based lane marker detection to function is not going to be sufficient for autonomous winter driving.

CHAPTER 4: VEHICLE FOLLOWING TESTS

4.1 PURPOSE

Traffic has a significantly detrimental impact on transportation in urban areas by increasing travel times, increasing emissions, and putting stress on drivers. Autonomous driving systems present an opportunity to significantly diminish traffic congestion by reducing the effects of human error and reaction time (Maurer, M. et al., 2016). The purpose of this portion of the study was to determine whether a simple ACC longitudinal control such as the one equipped on the MnCAV vehicle could actually have an impact on traffic flow given a sufficient number of vehicles equipped with this system, or if a more advanced longitudinal control algorithm would be required to eclipse the performance of human drivers. Previous studies have modeled the implications of autonomous vehicles on traffic flow through simulation (Muhammad, T. et al., 2020; Talebpour, A., Mahmassani, H., 2016), but this project sought to establish real world data using the MnCAV vehicle. The key parameter investigated by the research team was whether or not the longitudinal controller on the MnCAV vehicle was “string stable.” String stability refers to the ability of a longitudinal controller to reduce the amplitude of acceleration and deceleration waves propagating through the flow of traffic (Rajamani, 2011). Reduction in amplitude of such waves is associated with improved traffic flow, because small acceleration/deceleration magnitudes indicate that traffic is continuing to flow at cruising speed with no interruptions. If the ACC system was found to be string stable, this would provide evidence that opportunities exist to leverage ACC technology to improve the flow of traffic in urban areas.

4.2 METHODS

To perform these tests, the MnCAV vehicle was driven in a loop from I-94 to 35W south down to Hwy 62 and back again. Tests were performed during mid-afternoon such that the team encountered various light, medium, and heavy traffic scenarios throughout the drive. Traffic levels were defined using the average speed of the MnCAV vehicle according to Table 3:

Table 3. Definitions of traffic levels.

Traffic Level	Average speed range (mph)
Light	>40
Medium	20-40
Heavy	<20

The testing loop was driven once with the vehicle set to autonomous mode, and once with a manual driver to observe whether there were any significant differences between the two. Measurements of the MnCAV vehicle speed and acceleration as well as the preceding vehicle distance and relative velocity were recorded throughout the drive. In order to analyze the string stability of the MnCAV longitudinal controller, it was necessary to compare the acceleration of the MnCAV vehicle to the acceleration of the preceding vehicle. However, there is no sensor on the MnCAV vehicle which can directly measure the acceleration of the preceding vehicle. To address this problem, the research team estimated the acceleration of the preceding vehicle by differentiating the absolute velocity signal of the preceding vehicle and applying an aggressive low-pass filter to the resulting signal to reduce noise. The absolute velocity of the preceding vehicle was found according to Equation 2, where v_p is the absolute velocity of the preceding vehicle, v_{rel} is the velocity of the preceding vehicle relative to the MnCAV vehicle, and v_{ego} is the absolute velocity of the MnCAV vehicle:

$$v_p = v_{ego} + v_{rel} \quad (2)$$

The preceding vehicle acceleration could then be estimated from the absolute velocity. Because this result is a derived estimate and not a direct sensor measurement, it was necessary to validate that the resulting acceleration signal was a good estimate of the true acceleration signal. To do this, the same differentiation and filtering process was applied to the absolute velocity signal of the MnCAV vehicle and compared to the directly measured acceleration signal of the MnCAV vehicle. If this algorithm produced a good estimate of acceleration, the derived signal should match closely to the directly measured signal. Figure 19 shows a plot of the two signals, which demonstrates that the two signals match closely and therefore provides evidence that the estimation algorithm used is accurate.




Figure 19. Estimated and directly measured acceleration signal of the MnCAV vehicle.

4.3 KEY FINDINGS

The first set of testing was done in a light traffic scenario with the MnCAV vehicle traveling at cruising speed around 50 mph on the freeway. The acceleration of the MnCAV vehicle was measured, and the acceleration of the preceding vehicle was determined using the methods presented above. The results of the light traffic testing are presented in Figures 20 and 21, which show the acceleration of the MnCAV vehicle and preceding vehicle as a function of time. The average acceleration magnitude results are found in Table 4 below.

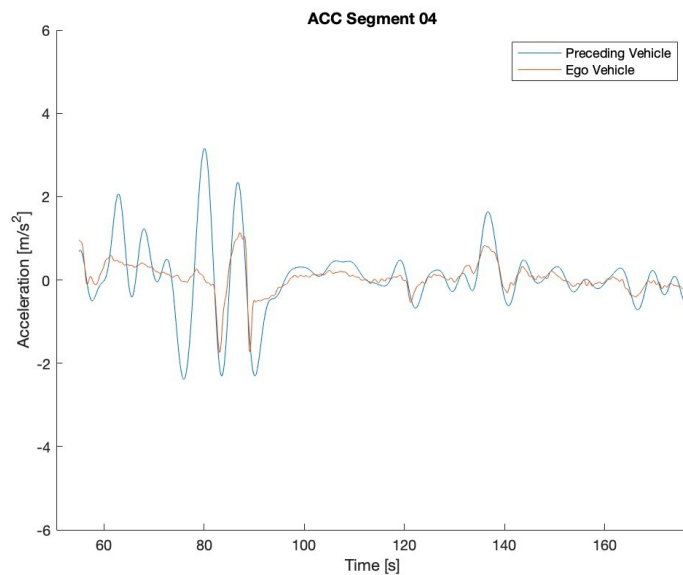


Figure 20. Test 1, acceleration of MnCAV vehicle and preceding vehicle during light traffic.




Figure 21. Test 2, acceleration of MnCAV vehicle and preceding vehicle during light traffic.

Both of these tests show relatively low frequencies in the acceleration signals, which is expected during a light traffic scenario, as the MnCAV vehicle is for the most part able to maintain a constant velocity. In both light traffic tests, the magnitude of acceleration of the MnCAV vehicle stays below that of the preceding vehicle (closer to 0 magnitude) for the duration of the test. This demonstrates that in light traffic, the MnCAV vehicle is accelerating and decelerating with less magnitude than the preceding vehicle, which means the wave of acceleration propagation through traffic is being attenuated by the MnCAV vehicle, improving traffic flow. This is because the ideal case for optimal traffic flow is to have no acceleration (i.e all vehicles are cruising at the speed limit with no interruption), and thus by attenuating the acceleration wave through traffic, the MnCAV vehicle brings the traffic flow closer to this goal.

The next set of tests were performed in medium traffic. In this case traffic was still flowing, but significantly below cruising speed. Figures 22 and 23 show the acceleration results for the medium traffic tests. Results for average acceleration magnitudes are also presented in Table 4 below.

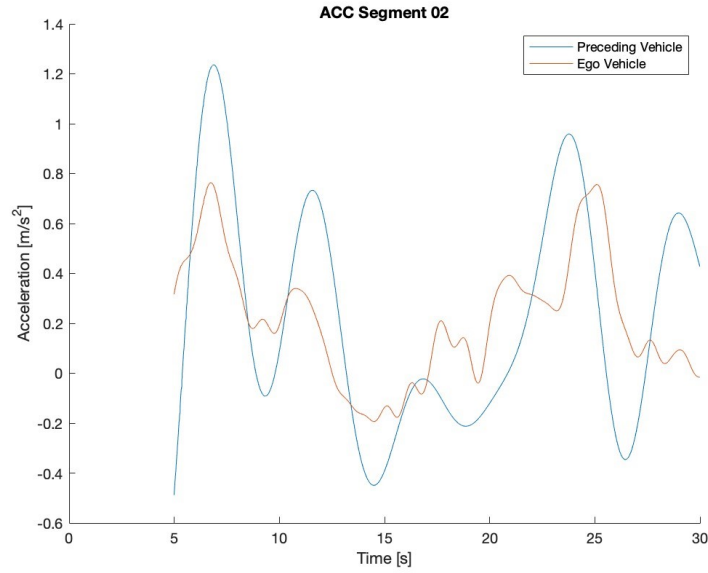


Figure 22. Test 1, acceleration of MnCAV vehicle and preceding vehicle during medium traffic.




Figure 23. Test 2, acceleration of MnCAV vehicle and preceding vehicle during medium traffic.

These tests demonstrate a higher frequency in the acceleration signals than what was observed in the light traffic tests, indicating the MnCAV vehicle is accelerating and decelerating with greater frequency, which is what would be expected in a heavier traffic situation. Once again, it is observed that the magnitude of acceleration of the ego vehicle is less than that of the preceding vehicle throughout the duration of the test, demonstrating that the MnCAV vehicle is attenuating the acceleration wave

through traffic even in a more congested scenario. This once again demonstrates the ability of ACC to improve the flow of traffic.

Finally, tests in heavy stop and go traffic were performed. In this case, traffic was moving at very slow speeds mostly under 10 mph with cyclical acceleration and deceleration at high frequency. The acceleration results for these tests are presented in Figures 24, 25, and 26.

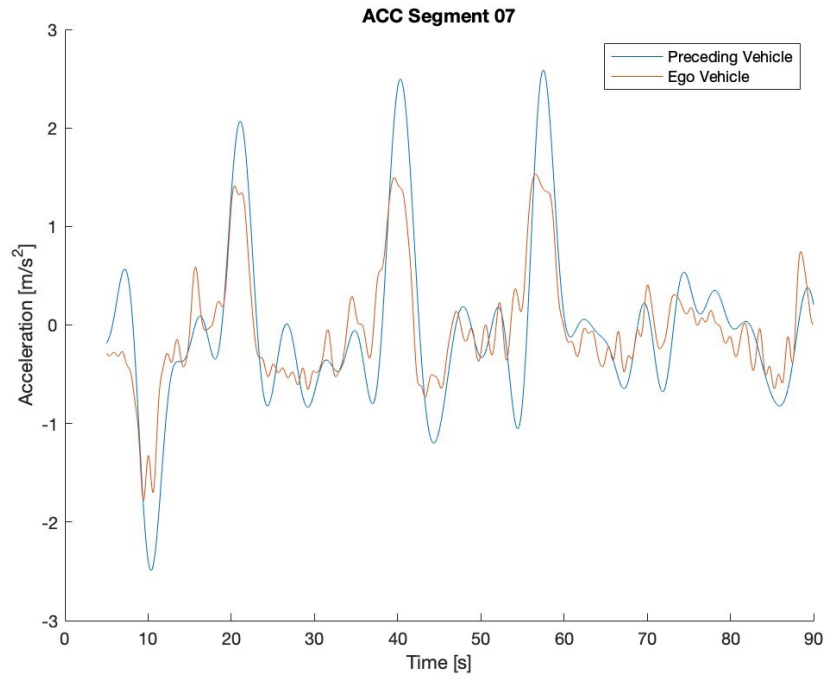


Figure 24. Test 1, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.




Figure 25. Test 2, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.

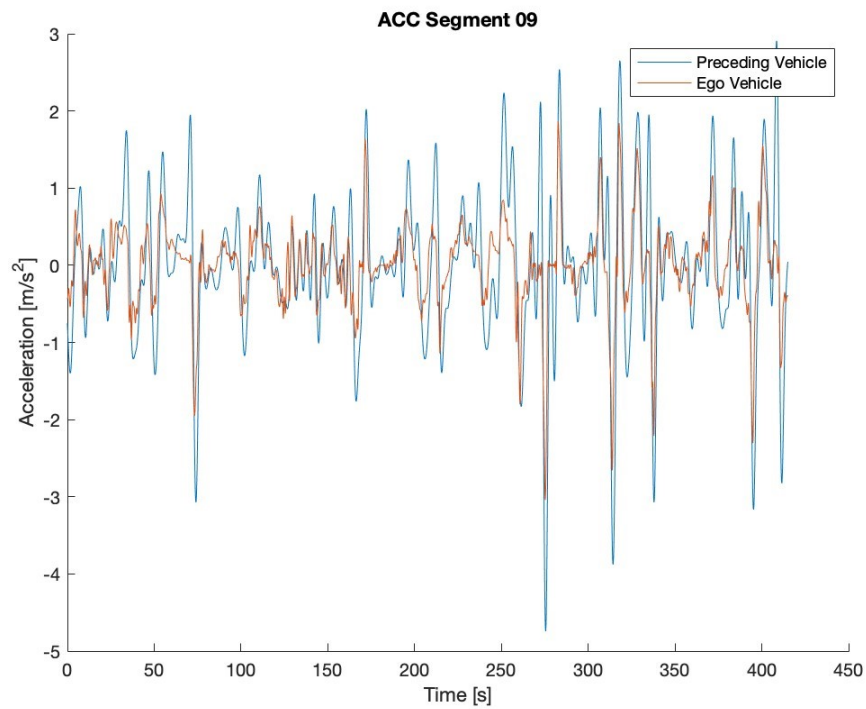


Figure 26. Test 3, acceleration of MnCAV vehicle and preceding vehicle during heavy traffic.

As expected, these tests show the highest acceleration frequency out of any of the three traffic scenarios tested. This indicates that the MnCAV vehicle was forced to rapidly accelerate and decelerate over and over in quick succession, which is the most challenging case for an ACC system to handle. Despite the challenging environment, the MnCAV vehicle still demonstrated a lower acceleration magnitude compared to the preceding vehicle throughout the duration of the tests. This indicates that even in heavy traffic, the MnCAV vehicle is able to attenuate the propagation of acceleration waves in traffic, helping to improve traffic flow, which is an excellent property of the ACC system.

Table 3 shows the average acceleration and deceleration magnitudes for the MnCAV vehicle and the preceding vehicle for each tested traffic scenario.

Table 4. Average acceleration and deceleration magnitudes for MnCAV vehicle and preceding vehicle.

Traffic Scenario	Average Acceleration [m/s ²]		Average Deceleration [m/s ²]	
	Ego Vehicle	Preceding Vehicle	Ego Vehicle	Preceding Vehicle
Light Traffic	0.1652	0.2812	-0.0452	-0.2084
Medium Traffic	0.2147	0.3397	-0.1190	-0.2082
Heavy Traffic	0.2337	0.3657	-0.2259	-0.3557

This table demonstrates the results that were discussed in the acceleration figures above, showing that the magnitude of acceleration and deceleration is lower for the MnCAV vehicle than the preceding vehicle for all traffic levels. Overall, this provides evidence that ACC systems are capable of attenuating acceleration waves in traffic, and are thus capable of being string stable. This means that while slowdowns on freeways will still occur even if all vehicles were using ACC (due to factors such as lane changes, merges, etc.), the severity of backups would be expected to decrease as more and more vehicles on the road start using ACC. It is still unclear what percentage of vehicles on the road must be using ACC to see significant improvements in traffic flow, though this is a topic that could be researched further in the future.

In conclusion, these tests provide clear evidence that ACC systems can attenuate acceleration waves in traffic and help improve traffic flow. As ACC systems become more and more common on vehicles on

the road, this provides opportunities to leverage the capabilities of these systems to help improve traffic flow on freeways. In the future, further research may investigate methods to best utilize the string stability properties of ACC systems to reduce traffic backups.

CHAPTER 5: WORK ZONE TESTING

5.1 PURPOSE

While snow poses a significant challenge to autonomous driving systems in the winter, work zones also pose a significant challenge to autonomous driving in the spring and summer seasons. Work zones often feature challenges such as blocked off lanes, pavement with poorly painted, unpainted, temporary, or conflicting lane markings, forced lane changes indicated by cones, barrels, and signage, as well as other scenarios that are out of the ordinary for typical driving, making work zones challenging for human drivers to navigate, but potentially downright unmanageable for many autonomous systems. In this task, the research team sought to better understand which components of a work zone are the most challenging for an autonomous system to interpret and handle, as well as determine some best practices for setting up work zones to minimize the difficulties encountered by autonomous systems in such scenarios. The team focused on the ability of the MnCAV vehicle to detect lane markings in various work zones scenarios, providing indication of the ability of the vehicle to perform lateral control in these situations.

5.2 METHODS

5.2.1 Cone and Barrel Testing

The first set of tests performed explored the impact of cone and barrel placement on lane marker detection from the Mobileye system. Tests were performed on the LVR at MnROAD. Two variables were studied by the research team: 1) the lateral placement of the cones and barrels (i.e., behind or on top of the lane markings), and 2) the longitudinal spacing between cones. The cones and barrels were placed on top of the lane markers with spacings of 3ft, 6ft, and 10ft. Once the cones or barrels were set up, the MnCAV vehicle was manually driven past the cones at a speed of about 5 mph. Measurements of lane marker detection quality were recorded for each test, which provided insight to the impact of cone placement on the lateral control system. Figure 27 shows an image of one test setup.



Figure 27. Cone test setup on top of lane lines with 6ft spacing.

5.2.2 Temporary Lane Marker Testing

To determine the impact of temporary lane markers on lane line detection, the research team created a testing zone by placing about 30 temporary raised pavement markings (TRPM's) in the center of an unmarked road, dividing the road into two lanes. The lane markings were spaced about 3ft apart. Once the testing zone was set up, the MnCAV vehicle was manually driven through it, alternating directions for each test. Measurements of lane line detection quality were recorded for each test. Figure 28 shows an image of the testing zone setup.



Figure 28. Temporary marker test setup.

5.2.3 Conflicting Lane Marker Testing

To investigate the impact of conflicting lane markers of the lane boundary estimation system, a method to produce a conflicting lane marker set needed to be developed. While ideally a conflicting set would have been painted over a large distance at MnROAD, this was found impossible for various reasons. Instead, the team decided to use freezer tape, the width of which is similar to that of lane markings, painted with lane marking paint, which was then placed on the road to create a conflicting marker. The marker was placed at an angle to the primary set so as to divert traffic to the opposite lane, as shown in Figure 29.



Figure 29. Conflicting lane marker setup.

Once the test was set up, the MnCAV vehicle was manually driven through the scene, following the conflicting marker into the opposite lane. The output of the Mobileye system was recorded to see which lane markers were detected by the system.

5.3 KEY FINDINGS

5.3.1 Cone and Barrel Testing

The first set of cone tests were performed with the cones placed just behind the center lane markings on the LVR. The purpose of these tests was to determine if the presence of cones near the lane markers had an impact on detection, even if there was no direct visual impediment to the markers. Results of lane line detection quality for one such test is presented in Figure 30.

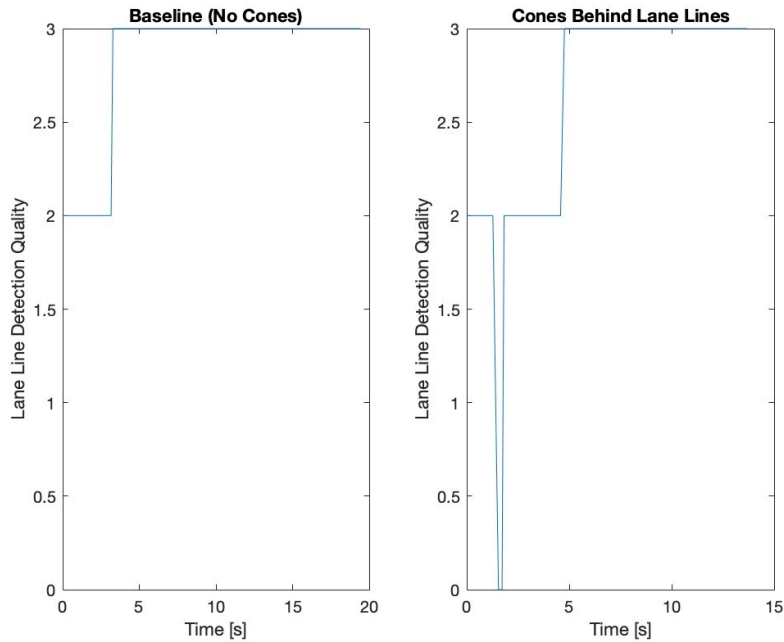


Figure 30. Lane line detection quality for baseline (left) and cones placed behind lane lines (right).

The results for the test with the cones placed behind the lane markings matches very closely to the baseline, with a detection quality of 3 for the majority of the test after jumping up from an initial quality of 2. There is one quick drop in detection for the cone test, but this was just an artifact of the test that was not connected to the presence of the cones. Overall, this test verifies that cones placed behind lane markers with no visual impediment to the markers have no impact of the quality of detection.

Next, cones were placed on top of the lane markers at varying longitudinal spacings. Results for baseline (no cones) and spacings of 3ft, 6ft, and 10ft are shown in Figure 31.

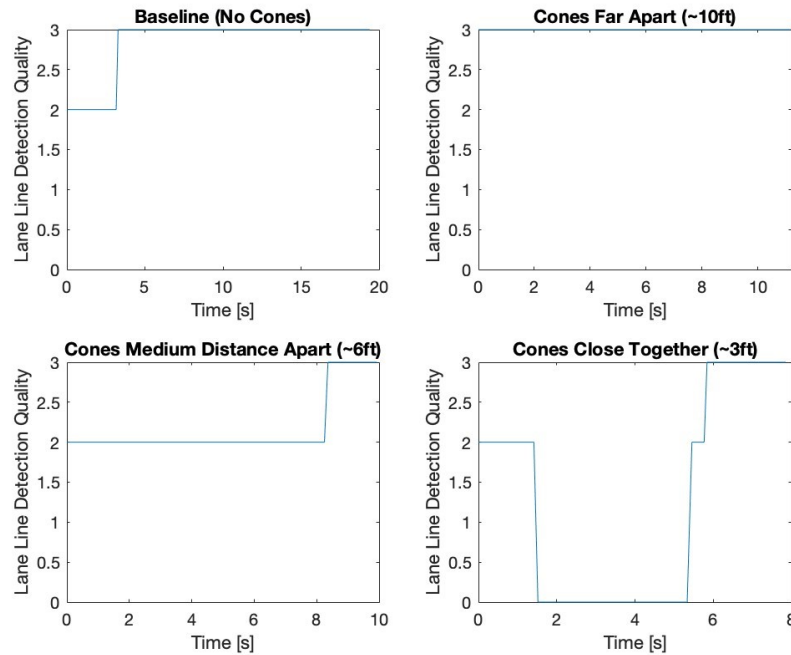


Figure 31. Cone lane line detection quality for baseline and 3ft, 6ft, and 10ft spacings.

When the cones are placed far apart (10ft spacing), there is no impact on the detection quality, as the test shows excellent detection quality throughout. However, as the cones are placed more and more densely together, it is observed that the detection quality begins to fall. For example, when the cones were placed 6ft apart, the detection quality dropped to 2 until the vehicle passed the cones, and when the cones were placed 3ft apart, the detection quality dropped to 0 (no detections) until the vehicle passed the cones. This demonstrates that if cones are placed too densely in a position that visually impedes the lane lines, this can result in a significant decline in detection quality, and can even completely block detection all together.

The same tests carried out with cones were also performed with barrels. Barrels have a much larger size than cones, so the research team was interested in investigating the impact of larger obstructions on lane detection quality. Figure 32 shows the measured lane line detection quality for baseline (no barrels) and barrels placed on top of the lane lines with longitudinal spacings of 3ft, 6ft, and 10 ft.

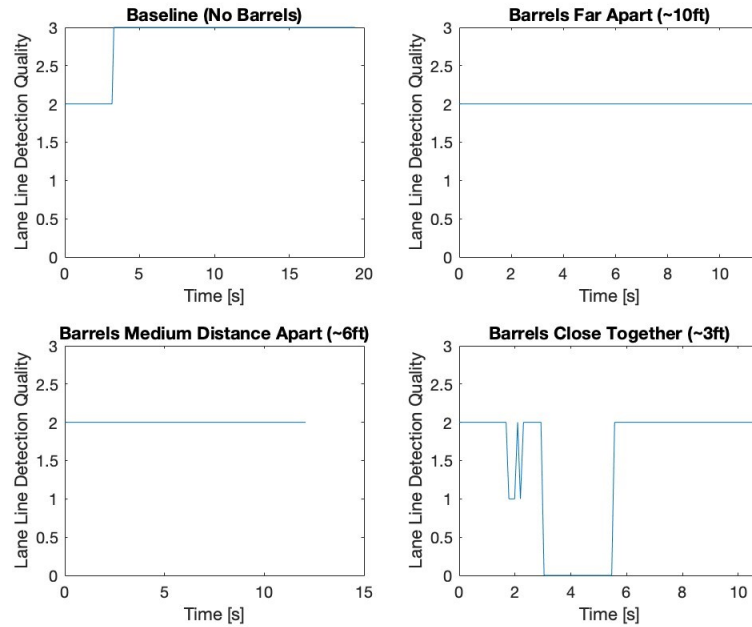


Figure 32. Barrel lane line detection quality for baseline and 3ft, 6ft, and 10ft spacings.

The barrel tests show similar results with what was observed for cones. With the barrels placed 10ft apart, a detection quality of 2 is observed, which is lower than the 3 that was observed with cones, indicating that the larger size of the barrels may have a slightly larger impact on detection at lower densities than cones. However, with the barrels placed 6ft and 3ft apart, the same results are observed as with the cones, with a detection quality of 2 in the presence of barrels for 6ft spacing, and a detection quality of 0 in the presence of barrels for 3ft spacing. This again demonstrates that as barrels are placed more and more densely while visually impeding the markers, it is possible to significantly reduce or completely block detection of the lane markers.

5.3.2 Temporary Lane Marker Testing

Figure 33 shows the results of lane line detection quality for a test with temporary lane markings.

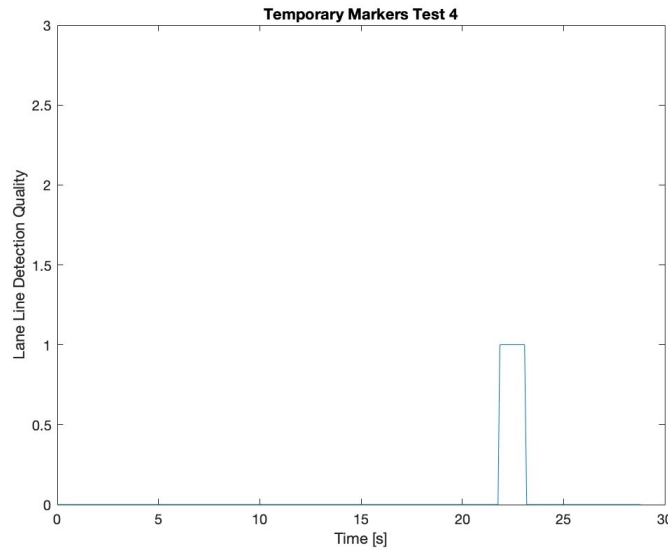


Figure 33. Lane line detection quality for temporary lane markers.

These results show that for almost the entire duration of the test, no detections were made of the temporary markers. For about 1 second near the end of the test, a slight detection with a quality of 1 was observed, but this was not a significant detection. This shows that while it is theoretically possible for the Mobileye system to detect temporary markers, the software is not yet designed to make such detections. Thus, this clearly demonstrates that the Mobileye system is not capable of producing lane boundary estimates with temporary markers, making lateral control impossible in such environments with the current control setup. The simplest solution to this challenge would be to produce software which is capable of detecting temporary markers; however, this cannot be implemented on the Mobileye system without an update from Mobileye itself.

5.3.3 Conflicting Lane Marker Testing

The tests with conflicting lane markers produced no detections of the conflicting markers. Rather, only the primary markers were detected and visualized by the Mobileye system. Figure 34 shows an image demonstrating the detection of the left, middle, and right lane boundaries, but no detection of the conflicting marker. The rightmost purple line appears to indicate a detection of the conflicting marker; however, this is simply how a case of no detection is visualized by the system, with the purple line parallel to the heading of the vehicle, and does not actually indicate a true detection of the conflicting marker.




Figure 34. Image showing detection of primary markers and no detection of conflicting marker.

This test indicates that in the presence of conflicting markers on the roadway, the Mobileye system is likely to follow whatever the primary, clearest set of markers is in the scene. While this testing setup was not perfect and could be improved by painting conflicting markers over a much larger distance, this does demonstrate the need for the Mobileye system to only follow one set of markers at a time. The Mobileye system does not have the cognitive ability to distinguish between different lane markers sets and pick the correct one using scene context. Rather, it will follow whatever the clearest and most prevalent set of markers is in the scene.

CHAPTER 6: DISCUSSION

6.1 WINTER PERFORMANCE AND WINTER ROAD MAINTENANCE

All Level-2 systems sold by car manufacturers in the market use cameras for detecting lane markers and finding the lateral distance of the vehicle from the center of the lane. The lateral distance from lane center is the key feedback variable used by the steering control system to do automatic lane centering (automatic lane-keeping).

The MnCAV vehicle used a commercially available Mobileye camera system for detecting lane markers/ boundaries and for finding the lateral distance of the vehicle from lane center. The Mobileye software returned a variable indicative of the quality of lane boundary detection (called Line Detection Quality, LDQ) and the distance to lane center. The LDQ varied from 0 to 3, with 0 being very poor or no lane boundary detection and 3 being very good detection. Under normal road conditions with clearly marked lane boundaries, the LDQ was typically 3 most of the time, as shown in Figure 35 below for automated driving on the low volume road at the MnRoad facility.




Figure 35. Line Detection Quality with the Mobileye Camera System on MnRoad.

The LDQ deteriorated with increasing presence of snow on the road. Figure 36 shows the situation on 35W north of Minneapolis on a day at a time after a heavy snowfall when the roads had already been cleared by snowplows.




Figure 36. Camera view from inside the MnCAV vehicle while driving an hour after snowfall.

As seen in Figure 36, although the highway is largely safe for driving and the locations of the lanes are visible to a human driver (from the tracks of previous vehicles), the lane boundaries themselves are covered with snow. The LDQ of the Mobileye system was very poor (0 most of the time) in this situation, as shown by the data recorded in Task 4. Snow is found to stay much longer on the lane boundaries than in the center of the lanes.

6.1.1 Recommendations

Given that all Level-2 AVs critically depend on cameras and visibility of lane boundaries, the research team recommends that MnDOT and other local transportation agencies can try to give AVs the best chance to function properly in winter by efficient removal of snow from lane lines. While current plowing operations may largely focus on clearing the center of lanes (as the snowplow moves inside a lane during plowing), particular attention to detail in clearing lane lines specifically could be done to allow current L2 technology to function during snow events.

As discussed in the report for Task 4, a number of other measurement and control methods (not relying on the visibility of lane markers) could be pursued to enable autonomous driving on roads with snow. These include GPS/ GNSS systems with RTK corrections, Lidar maps and more-human like driving systems based on AI. Of these, the research team feels that providing more consistent connectivity for corrections to GNSS systems could be done effectively by upgrading the MnCORS infrastructure in Minnesota. The current automotive industry (GM, Waymo, Tesla, Uber, etc) appears to be focusing on pursuing Level 4 driving in good conditions, while winter driving is left for later development years down the road. We may not see any improvements in this situation unless we (states that experience snowy winters) take the initiative. If the winter driving problem could be made less complex through smart infrastructure (such as upgrading the MnCORS network), this may accelerate the development of more robust systems for winter driving.

6.2 COMPARISON WITH HUMAN DRIVERS

In order to compare the performance of the AV with that of human drivers in terms of shock wave propagation and smoothness of driving, it was necessary to be able to do the tests under all traffic conditions (light, medium and heavy traffic conditions).

To begin with, the longitudinal controller which came preinstalled on the MnCAV vehicle (from the vendor Dataspeed) performed very poorly under speeds of 10 mph. Significant lurching behavior was observed at low speeds, which posed a danger of sometimes resulting in a collision without intervention from the safety driver. This meant that testing of the ACC system in heavy traffic was infeasible using the original Dataspeed controller on the vehicle. Even at higher speeds the smoothness of the ACC system was not as good as what could be expected on a consumer grade ACC system.

To resolve the challenges with the ACC at low speeds, we had to rewrite the entire longitudinal controller from scratch in order to achieve consumer grade performance, as well as good performance at low speeds. We were able to do this by rewriting the longitudinal control ROS node without impacting the rest of the vehicle controls. A significantly improved longitudinal controller was programmed onto the vehicle with the goal of resolving the lurching problems at low speeds and improving the overall smoothness and performance of the system. The controller developed included a steady state control law (a control law to keep the ego vehicle at a specified time gap behind the preceding vehicle), transitional maneuvers (control law to smoothly transition the ego from first detection of a preceding vehicle to steady state vehicle following and vice-versa), and a Stop mode (a strategy to control the ego vehicle to a complete stop despite significant noise in the relative velocity measurement). Adding these improvements provided us with significantly smoother longitudinal control all the way down to a stop, significantly decreasing the number of disengagements of the system in heavy traffic scenarios and allowing us to gather data in heavy traffic scenarios that is representative of a well-designed consumer level ACC system.

Once the controller had been adequately developed, tests were conducted under light, medium and heavy traffic conditions in which the longitudinal acceleration of the preceding vehicle was accurately estimated using radar measurements and signal processing, and the acceleration of the MnCAV vehicle was measured using an on-board accelerometer. The acceleration levels of the MnCAV vehicle were found to be consistently below those of the preceding vehicle under all traffic conditions. Figure 37 below shows an example segment of data under heavy traffic conditions where it can be seen that the ego (MnCAV) vehicle consistently has lower accelerations.




Figure 37. Example data segment showing accelerations of preceding vehicle and MnCAV vehicle under heavy traffic conditions.

Additional details of the tests and the data obtained are available in the Task 5 Report of this project. The smoother and lower accelerations of the MnCAV vehicle were judged to be due to:

- a) The larger time-gap of the MnCAV vehicle compared to the average time gaps of other drivers on the road.
- b) The smoother controller on the MnCAV vehicle due to which throttle acceleration and braking were more smooth while having lower time latencies.

It is not clear if the larger time gaps necessarily reduce the traffic flow rate because the MnCAV vehicle reduces shock wave propagation by attenuating disturbances from the preceding vehicle and thus is likely to help reduce traffic jams.

6.2.1 Benefits

- 1) The work done to improve the performance of the longitudinal controller has enabled the MnCAV vehicle to work smoothly under all traffic conditions, including heavy stop-and-go traffic. This will be beneficial for many future research projects in Minnesota, since many other research teams and projects will be able to use the vehicle and will not have any restrictions on the conditions under which the vehicle will be able to operate.

- 2) The work done to document the simultaneous acceleration histories of the preceding and MnCAV vehicles has demonstrated that the MnCAV vehicle now has consistently smoother and lower accelerations, and reduces shock wave propagation by attenuating deceleration disturbances from the preceding vehicle. These results imply that AVs in general have a significant potential to reduce fuel consumption, to provide safer driving (larger time gaps) and could also reduce traffic jams due to shock waves.
- 3) Further benefits in terms of higher traffic capacities could be obtained in the future if V2V and I2V/V2I wireless communication could be enabled so as to enable smaller time gaps without sacrificing smooth operation.

6.3 IMPLICATIONS OF AV'S FOR CONSTRUCTION ZONE GUIDELINES

How should transportation agencies prepare for AVs in terms of construction zone guidelines? The MnCAV vehicle uses the commercial Mobileye camera and software for detecting lane lines and finding lateral distance to lane center. The software is not set up to detect traffic cones or traffic barrels. Hence, the automatic steering control system in its current form will not respond to traffic cones directly. However, tests were conducted to see whether traffic cones/barrels would obstruct the visibility of lane lines for the Mobileye camera. Figure 38 below shows an example of tests conducted in which traffic cones were placed on a lane line at with longitudinal spacings of either 10 feet, 6 feet or 3 feet. Results showed that distances of 10 feet and 6 feet did not prevent the Mobileye camera from detecting the lane lines correctly and finding lateral distance to it. However, the distance of 3 feet between cones made the Mobileye system completely stop detecting the lane line.



Figure 38. Tests with traffic cones to evaluate what density of cones prevents the Mobileye lane boundary detection system from working.

In the case of conflicting lane markers where a second lane deviating from the primary lane was created, the Mobileye system rejected the presence of the deviating (branching off) lane and did not detect or return lateral distance to it at all. The current functioning of the Mobileye software is such that it only

detects the primary lane and is unlikely to pick up a branching-off new lane. Thus, creating a diversion for an autonomous vehicle due to the presence of a work zone is going to be difficult. Either the primary lane boundaries would have to be completely hidden (using traffic cones, for example), or else the vehicle would have to stop autonomous operation and the human driver would need to take over control.

Finally, the project also tested Temporary Raised Pavement Markers (TRPM's) using a standard longitudinal distance of 3 feet between markers. The Mobileye system was unable to detect a lane boundary from this set up, meaning the temporary lane markers were just inadequate for purposes of lane boundary creation.

6.3.1 Recommendations

Given that most (or all) current Level 2 AVs are unlikely to be able to drive safely through work zone related diversions, it is recommended that the human driver take over vehicle operation when a work zone on a highway is encountered. Figuring out the right path in the presence of conflicting lane markers is a significant challenge and so continued automatic steering control in such situations is not recommended. To enable human takeover, adequate early notice of an upcoming construction zone is needed, so that there is adequate time for a safe transition from autonomous driving to human driving.

Another finding of the project is that traffic cones spaced apart at 3 feet and placed on a lane boundary are adequate to completely prevent identification of the lane boundary by camera systems. This is an approach that can be taken up for hiding some lane boundaries and thus preventing the occurrence of conflicting lane boundaries when a work zone related diversion is needed. Very precisely creating lane boundaries and completely avoiding the presence of conflicting paths is a potential approach to possibly enabling autonomous steering in work zone related diversion situations.

Finally, the project found that temporary lane markers are inadequate for lane boundary creation and so their use is not recommended. Painted lane markers or equivalent wide tape markers on the road are needed.

6.4 SUGGESTIONS ON AREAS FOR FUTURE RESEARCH SUPPORT

There is already significant research, development and testing on autonomous vehicles being carried out by large companies such as Waymo (Google), Tesla, Uber, General Motors and others. However, all of these companies are largely focused on pursuing Level 4 autonomous driving in good weather and good road conditions, such as in California and Arizona. We may not see any improvements in this situation unless we (states that experience snowy winters) take the initiative. Hence, this research team suggests the following areas for future research support:

- A) Minnesota and other Midwest states have a significant winter season. As seen in this project, the presence of snow covering lane boundaries on the road will prevent automatic steering systems from being engaged. Further, snow on lane boundaries stays for significantly longer durations than snow in the interior of the lane. Hence, development of technologies enabling autonomous operation under such conditions needs to be encouraged.
- B) Low volume rural roads present special challenges for autonomous vehicles in all seasons. Such rural roads can be narrow, often will not have right-side lane markings, may not even have center lines, may not be plowed for snow removal, and can have trees and other objects close to the side of the road. Rural traffic intersections can have missing delineation and signage that are normally provided on higher volume roadways. All of these issues pose major challenges to autonomous driving, since AVs depend critically on such markings and signage, and do not have the sensing capability to search perpendicular roads at a traffic intersection for vehicle tracking. Supporting research related to autonomous operation in rural areas should therefore be encouraged.
- C) Due to significant autonomous vehicle research being pursued all over the world, many new technologies such as computer vision algorithms for object recognition and low-cost distance sensors for trajectory estimation are becoming available. The refinement and use of such new technologies to enable near-term applications for MnDOT and other transportation agencies should be supported. For example, the use of computer vision algorithms for automatic pedestrian counting and other similar applications could be enabled due to recent availability of well tested open-source object recognition software.

CHAPTER 7: CONCLUSION

The goal of this project was to perform experimental tests with the MnCAV vehicle on Minnesota roads and highways, with a focus on analyzing challenging scenarios autonomous vehicles may face such that recommendations for further research and best practices could be made to MnDOT. These scenarios included winter driving, vehicle following in traffic, and work zones. The project began with initial testing of the MnCAV vehicle such that a baseline understanding of its capabilities and limitations could be determined. The team found that the lateral control system performed well on straight roads but struggled on curves, and that the longitudinal control system performed well above 10 mph but was very jerky under 10 mph. The team addressed this by developing an improved longitudinal controller for the vehicle for use in the vehicle-following portion of the project. In the winter-driving portion of the project, the team found that even a small amount of snow on the road surface could impede visualization of the lane markings, making lateral control very challenging in such an environment. To address this issue, further research into snow removal on roads and alternative sensing techniques for lateral control are suggested. Through tests of vehicle following in traffic, the team found that the MnCAV vehicle demonstrated string stable behavior, attenuating traffic shockwave propagation and provided evidence that with enough vehicles on the road using ACC, reductions in backups could result. Finally, the team found that work zones provided significant challenges to autonomous driving systems. Cones or barrels placed densely enough on top of or in front of lane markers may result in a loss of detection of lane markings for the vehicle, which poses an issue to lateral control. In addition, conflicting lane markings can result in confusion for the system about which set to follow, which could result in unexpected behavior of the vehicle. Finally, it was found that temporary markers were not detectable by the Mobileye system, which means this method of lane marking in work zones was not suitable for current perception technology. Recommendations were made by the research team for further research and best practices in each of these areas. While the onboarding of autonomous systems onto Minnesota roads still faces many problems and challenges, through further research in the areas discovered through this project, this process could be greatly facilitated, bringing improvements to transportation in Minnesota in coming years.

REFERENCES

- Maurer, M., Gerdes, J., Lenz, B., & Winner, H. (2016). The effect of autonomous vehicles on traffic. In *Autonomes fahren* (pp. 317–334). Germany: Springer Berlin / Heidelberg.
- Muhammad, T., Kashmiri, F., Naeem, H., Qi, X., Chia-Chun, H., & Lu, H. (2020). Simulation study of autonomous vehicles' effect on traffic flow characteristics including autonomous buses. *Journal of Advanced Transportation*, 2020, 1–17.
- Rajamani, R. (2011). *Vehicle dynamics and control*. New York: Springer.
- SAE International. (2021). Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. Retrieved from https://saemobilus.sae.org/content/J3016_202104
- Talebpour, A., & Mahmassani, H. (2016). Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research. Part C, Emerging Technologies*, 71, 143–163.
- Tesla. (n.d.). Autopilot and full self-driving capability. Retrieved from https://www.tesla.com/en_eu/support/autopilot
- Waymo. (2021). Safety report: On the road to fully self-driving. Retrieved from <https://waymo.com/safety/>
- Zang, S., Ding, M., Smith, D., Tyler, P., Rakotoarivelo, T., & Kaafar, M. (2019). The impact of adverse weather conditions on autonomous vehicles: How rain, snow, fog, and hail affect the performance of a self-driving car. *IEEE Vehicular Technology Magazine*, 14(2), 103–111.