12-2023

Advancing Winter Weather ADAS: Tire Track Identification and Road Snow Coverage Estimation Using Deep Learning and Sensor Integration

Parth Kadav
Western Michigan University

ADVANCING WINTER WEATHER ADAS: TIRE TRACK IDENTIFICATION AND ROAD SNOW COVERAGE ESTIMATION USING DEEP LEARNING AND SENSOR INTEGRATION

Parth Kadav, M.S.E.

Western Michigan University, 2023

Modern vehicles have undergone a transformation with the widespread integration of Advanced Driver Assistance Systems (ADAS) technology becoming the new standard and are set to be mandated by the by the National Highway Traffic Safety Administration (NHTSA) for all passenger vehicles and light trucks. ADAS features have proven to prevent or mitigate crashes by either alerting or assisting the driver. ADAS typically utilizes a forward-facing camera, which comes standard in modern vehicles to provide limited automation features such as Lane Keeping Assist (LKA), and Lane Centering Assist (LCA) to improve driver safety. These systems rely on the assumption that vehicle surroundings, and lane markings are clear and visible, but when a vehicle operates in adverse weather conditions like heavy snow, these systems fail. In this study we address this research gap using two novel studies. In the first study, we outlined a novel way to determine the safe driving region in lanes covered with snow, using unique features like tire tracks. It is anticipated that these research findings can inform new ways to improve the drivable region detection in regions of snow-occluded lanes and expand the operational design domain (ODD) of ADAS. In the second study, a data processing pipeline is detailed that estimates the amount of road snow coverage using an on-vehicle camera and infrastructure weather sensor data inputs. These studies together take a step towards broadening the ODD of ADAS, enhancing their performance and safety in inclement weather.

ADVANCING WINTER WEATHER ADAS: TIRE TRACK IDENTIFICATION AND ROAD SNOW COVERAGE ESTIMATION USING DEEP LEARNING AND SENSOR INTEGRATION

by

Parth Kadav

A thesis submitted to the Graduate College
in partial fulfillment of the requirements
for the degree of Master of Science in Engineering
Mechanical Engineering
Western Michigan University
December 2023

Thesis Committee:

Zachary D. Asher, Ph.D., Chair
Richard T. Meyer, Ph.D.
Guan Yue Hong, Ph.D.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my heartfelt gratitude for the outstanding guidance and mentorship provided by my advisor Dr. Zachary D. Asher. I want to express my deep appreciation for the enjoyable and fulfilling experience I've had working with you. Your expertise, assurance, passion, and trust have been invaluable, and I'm truly thankful for the chance to collaborate with you and gain from your guidance and resources.

I thank my committee members, Dr. Richard T Meyer, and Dr. Guan Yue Hong for their support and insightful feedback towards my research and studies. I would like to thank our external collaborators at the National Center for Atmospheric Research (NCAR), Dr. Curtis Walker and Dr. Amanda Siems-Anderson. I would like to thank the Michigan Translational Research And Commercialization (MTRAC), and the National Science Foundation (NSF) for funding my research and studies.

I would like to thank my fellow lab mates and student collaborators Dr. Nick Goberville, Dr. Johan Fañas Rohas, Dr. Farhang Motallebi-Araghi, Dr. Nicholas Brown, Kyle Carow, Kyle R Prins, Sachin Sharma, and Anika Tabassum, for their contributions to my research and publications.

Finally, I want to express my gratitude to my parents, my sister Dr. Dnyanada Kadav, my girlfriend, my best friend and brother Ashwamedh, and all my friends for their support and assistance in helping me complete this process and reach my goals.

TABLE OF CONTENTS

Table of Contents—Continued

LIST OF TABLES

LIST OF FIGURES

List of Figures—Continued

LIST OF ABBREVIATIONS

NHTSA	National Highway Traffic Safety Administration
ADAS	Advanced Driver Assistance Systems
ODD	Operational Design Domain
LKA	Lane Keeping Assist
LCA	Lane Centering Assist
ACC	Adaptive Cruise Control
FCW	Forward Collision Warning
AEB	Automated Emergency Braking
BSD	Blind Spot Detection
TSR	Traffic Sign Recognition
ABS	Anti-lock Braking System
ESC	Electronic Stability Control
IoU	Intersection over Union
mIoU	Mean Intersection over Union
CVAT	Computer Vision Annotation tool
ML	Machine Learning
DL	Deep Learning
YOLO	You Only Look Once
CNN	Convolutional Neural Network
CADC	Canadian Adverse Driving Conditions

List Of Abbreviations—Continued

ASOS Automated Surface Observing System

FPS Frames Per Second

BCE Binary Cross Entropy

RGB Red Green Blue

ROI Region of Interest

CHAPTER I

REVIEW OF RESEARCH GAPS TO ADVANCED DRIVER ASSISTANCE SYSTEMS IN INCLEMENT WEATHER CONDITIONS

Vehicles today reflect a century of collaborative engineering progress across areas such as performance, range, fuel efficiency, emissions, comfort, and safety. Looking ahead, the vehicles of the future face even higher standards in terms of safety, efficiency, and performance. This has paved the way for the integration of Advanced Driver Assistance Systems (ADAS) in the realm of automotive technology. ADAS has become an integral part, enhancing both the safety and convenience of drivers. ADAS use sensors such as cameras, and radar to enhance driving safety that include Adaptive Cruise Control (ACC), Lane Keeping Assist (LKA), Automatic Emergency Braking (AEB), Blind Spot Detection (BSD), and Traffic Sign Recognition (TSR). These systems aim to prevent accidents and improve overall road safety. However, as literature suggests, they need visible lane markings and road lane lines to function and perform as intended. ADAS has shown poor performance in inclement weather conditions such as snow, rain and fog which restricts their Operational Design Domain (ODD). Given the wide range of inclement weather conditions, I chose to focus on snowy weather conditions, aiming for a more specific focus within such challenging weather. There is a noticeable gap in research concerning the detection of drivable region in snow-covered lanes, a key factor in expanding the capabilities of ADAS. We have broken down our research into two chapters which collectively aim to tackle the research problem.

ADAS is currently limited to near-optimal scenarios, where clear road visibility is ensured [10] [16]. In situations where lane markings are obscured, such as during snowfall, the perception subsystem responsible for automated driving lacks crucial inputs, resulting in significantly reduced to potentially failure. In order to expand the ODD of ADAS in regions of snow-occluded lanes, we can first focus on leveraging existing on-vehicle sensors. Camera sensors have been used for various ADAS technologies and come standard in a majority of modern vehicles. By improving ADAS in snow weather conditions using the on-vehicle camera sensor we can target at improving safety and mitigating fatal crashes. ADAS is a crucial requirement in such conditions as the driver needs to be aware of their position in order to safely navigate the lane. A few studies have been conducted to address this research gap, study [20] shows drivable region detection in snow conditions using an array of sensors increasing the computational load adding inefficiencies. Additionally, literature in Chapter II and III suggests that previous studies have used an open-source dataset which is not necessarily catered for drivable region detection. There lacks a means to detect the drivable region using a single on-vehicle camera sensor and using a purpose-built dataset. The first research question thus is "Can we create and implement a method of detecting the drivable region in snow-occluded lanes using a single on-vehicle camera sensor?" To address this, the first study details an end-to-end pipeline that identifies the drivable region when given an image of a snowy road from the on-vehicle camera sensor. enabling drivable region detection even in snow-occluded lanes. This is an important step towards safer driving in adverse weather. This work is found in Chapter 2.

Automobile safety has long been a cornerstone of vehicle engineering, giving rise to crucial features like Anti-lock Braking System (ABS), Electronic Stability Control (ESC), Traction

control, and ADAS. Accurate environmental information is crucial for these technologies to function optimally, enabling them to adapt to real-time conditions with precise feedback from the environment. Sensor performs varies with changing weather conditions. There have been studies that have analyzed performance of Light Detection and Ranging (LiDAR), radar and camera sensors in inclement weather conditions. To enable ADAS performance in adverse weather conditions with our focus on snowy weather conditions, it is essential to first establish essential information about the vehicle's surrounding primarily focusing on road conditions. I reviewed studies that have addressed this research gap but they all use a biased, and small dataset [52,53]. As we proposed a new method of identifying the drivable region in regions of snow occluded lanes using camera inputs in Chapter 2, it is imperative to initially assess the local road condition to implement appropriate systems. While widely utilized, existing infrastructure weather stations often offer limited insight into ground-level conditions near the vehicle. The second research question thus is "Can we utilize on-vehicle camera sensor and infrastructure weather sensor data inputs to estimate the amount of snow coverage on the road local to the vehicle?". The second study addresses this research question by outlining a novel methodology for estimating the amount of snow on the road. This involves using a custom collected image dataset from an on-vehicle camera sensor and leveraging open-source, publicly available infrastructure weather sensor features using purpose-built machine learning models and feature engineering of inputs. This will pave the way for safe vehicle automation in regions of lane occlusion and expanding the ODD of ADAS. This work is found in Chapter 3.

To summarize the research questions of this thesis, Chapter 2 addresses the question "Can we create and implement a method of detecting the drivable region in snow-occluded lanes using

a single on-vehicle camera sensor?" by explaining a detailed end-to-end pipeline that outputs the drivable region in regions of snow-occluded lanes, and Chapter 3 addresses the question "Can we utilize on-vehicle camera sensor and infrastructure weather sensor data inputs to estimate the amount of snow coverage on the road local to the vehicle?" by using on-vehicle camera sensor and infrastructure weather sensor data inputs to predict the road condition specifically for snowy weather scenarios.

CHAPTER II

TIRE TRACK IDENTIFICATION: APPLICATION OF U-NET DEEP LEARNING MODEL FOR DRIVABLE REGION DETECTION IN SNOW OCCLUDED CONDITIONS

This study investigates a novel method to identify visual features in regions of snow occluded lanes for estimating the drivable region using deep learning and computer vision. This research was a follow-up study to [1]. I was able to expand the scope of the original research and publish the research presented in this chapter at ITS World Congress 2022 [2]. Please note that much of this chapter is presented verbatim from the publication. Parth Kadav led the development of the methodology and authorship of this publication, with the co-authors providing suggestions and guidance. The citation for this work is as shown below.

Kadav, P., Goberville, N., Motallebiaraghi, F., Fong, A., and Asher, Z.D., "Tire track identification: Application of U-net deep learning model for drivable region detection in snow occluded conditions," *Intelligent Transportation Systems World Congress.*

Abstract

Advanced Driver Assistance Systems (ADAS) typically utilize cameras to provide limited automation features to improve driver safety. ADAS utilizes computer vision (CV) to extract vehicle surrounding information. However, when the vehicle is operating in bad weather (e.g., obstructed lane lines), ADAS products fail. We have developed a new technique to detect tire tracks which was evaluated in conditions of variable snow coverage and lane line occlusion.

Previously we focused on using basic machine learning (ML). We expanded this to a convolutional neural network (CNN). A custom dataset was collected using an instrumented automated research vehicle. The CNN model had an intersection over union (IoU) score of 89% in detecting tire tracks and outperformed the traditional ML model on key metrics (precision, recall, and more). Overall we have demonstrated that this method works as an end-to-end pipeline to detect tire tracks and expand the operational design domain of ADAS.

2.1 Introduction

Advanced Driver Assistance Systems (ADAS) such as Forward Collision Warning (FCW), Automatic Emergency Braking (AEB), Lane Departure Warning (LDW), Lane-keeping Assistance (LKA), blind-spot warning assistance, and many more have the potential to prevent or mitigate approximately 40% of all passenger vehicle crashes [3]. Because human error causes the majority of road accidents, ADAS was created to automate and improve aspects of the driving experience in order to increase safety and safe driving practices. If the vehicle crosses the lane and no turn signals or necessary steering movements are detected, lane-keeping systems detect reflective lane markings in front of the car and inform the driver via various sorts of audio, tactile, and/or visual cues [4]. From the 1853 driver injury crashes studied in [5,6], it was discovered that LDW/LKA systems were able to reduce head-on and single-vehicle crashes on roads with higher speed limits (45-75mph) and visible lane markings by 53%. The greatest benefit of such systems, according to [7], is at lower operating speeds (5-20 mph), where between 11 and 23% of drift-out-of-lane incidents and 13 to 22% of seriously to fatally wounded drivers could have been avoided if the technology was used. FCW and AEB alone cut front-to-rear collisions by nearly half [8]. By

2023, the market for ADAS is expected to be worth more than $30 billion [9] where ADAS will not only be confined to safety but will also help increase vehicle efficiency [10–15].

Despite these successes of ADAS technology, there is a glaring unresolved problem: inclement weather. During 2007–16, weather-related vehicular crashes accounted for 21% (1,235,145) of all reported crashes annually resulting in 16% (5,376) of crash fatalities and 19% (418,005) of crash injuries throughout the United States [16]. Fundamentally, adverse weather conditions can cause impairment to situational awareness and inhibitions to vehicular maneuverability which can occur in a variety of ways depending on the type of adverse weather [16].

Developing techniques for the operation of ADAS in inclement weather is a current research challenge. Because there are significant ramifications for safety as outlined above, the initial goal is to recognize and classify road lanes during inclement weather in order to aid in the location of both the ego vehicle and other vehicles [17]. The challenge is that inclement weather such as heavy rain, snow, or fog lowers the maximum range and signal quality for ADAS sensors such as cameras and it occludes the high contrast lane markers [17]. This is a well-documented problem and has been demonstrated in cameras and lidars in particular [18]. A specific instance of this issue can be found in [5], where it is stated that LDW/LKA was only able to reduce head-on and single-vehicle crashes on roads with operating speeds of 45-75 mph by 53% if the roads had visible road markings and specifically "the road surface was not covered by ice or snow." New sensor technologies are getting better in these performance areas but are still far from addressing

the issue of reliable ADAS operation in inclement weather [10]. For now, to achieve a feasible research scope for this paper, we will focus on only snowy weather.

There are just a few key studies that address the issue of reliable ADAS operation in snowy weather. The first study developed a custom snowy weather dataset and determined the driveable region through semantic segmentation [19]. When evaluated on a non-snow dataset, the model had a mean Intersection over Union (mIoU) of 80%, when trained on a snowy dataset it dropped to 19% and when both models were combined, it provided a mIoU of 83.3%. However, the model still must be improved and made more robust because it considers the entire road rather than just the Region of Interest (ROI), which can be computationally expensive. The second study employed a CNN model with a specified architecture and used sensor fusion between the camera, lidar, and radar [20]. The results showed that there was an increase in driveable region detection (mIoU of 81.35%) and non-driveable region detection (mIoU of 93.85%) after fusing the information from various sensors and testing it on the dataset. This is an improvement, but it comes with drawbacks, the most significant of which is that the method requires additional sensors, which increases the cost and computational power required. Additionally, this method, like the first study, examines the entire driveable region, rather than just a ROI [20]. In the third relevant study, a method to improve the detection in adverse weather conditions using "You Only Look Once" (YOLO) was developed by merging it with a CNN and the Federated Learning (FL) framework [21]. This was tested on the Canadian Adverse Driving Conditions (CADC) dataset. The method resulted in the average test accuracy of the model, gossip, and centralized approaches which are the three different methods they use in their study to be 90.4−95.2%, 82.4−88.1%, and 71.4−76.16%, respectively.

The FL method, which utilizes an edge server, is the foundation for this model. After training a global YOLO CNN model on a publicly available dataset, the edge server sends the initial parameters to the AVs. These parameters are then used by the AVs to locally train the model on their own dataset. The number of AVs collecting data, the connection between the edge server and each vehicle, and the computational power in each vehicle all contribute to the FL method's training time. Furthermore, the vehicle chassis has been equipped with eight cameras, increasing the cost [21]. All the above studies provide methods for improving the detection of objects and regions in the entire driveable space and not necessarily the lane information, these studies are both computationally and monetarily costly and rely on multiple sensors. None of these studies demonstrates high accuracy driveable region detection for snow-covered roads using a single camera sensor that is implementable in modern ADAS products.

To address this research gap, we are utilizing a computationally light, cost-effective, and high-accuracy method of extracting driveable region information using a single camera which is a ubiquitous automotive sensor [18,19]. ML techniques such as CNN have established themselves as a dominating methodology in modern computer vision algorithms and applications, as well as in segmentation research. Based on our previous study for detecting tire tracks in snowy weather conditions [1], the ML model required a lot of image pre-processing and feature engineering, which is addressed in this study by using a CNN. In this study, both supervised ML semantic segmentation models and CNNs were developed. These methods were then compared for detecting tire tracks in the snow. The paper addresses the following novel topics:

1) Custom data acquisition method for tire track data collection and labeling

2) Snow tire track image preprocessing and feature extraction
3) Tire track identification CNN architecture
4) CNN and ML model performance comparison for snow tire track identification

2.2 Methodology

In this section, we will first discuss the methods we used to collect and prepare the data. The data that has been processed is then used to develop models.

2.2.1 Data Collection

The route we chose consisted of two-lane arterial roads in Kalamazoo that had the road characteristics we were looking for. This drive cycle replicated roads that are rarely cleaned after snowfall and are maintained much less frequently than highways and other multi-lane roads. We collected the data during the 2020 winter season. The lanes had snow occlusion with distinct tire track patterns, with the tire tracks visible to show the tarmac below and the lane line markings covered in snow. Data was collected using our Energy Efficient and Autonomous Vehicles (EEAV) Lab's instrumented automation development platform shown in Fig. 1.1 This development platform is built upon a drive-by-wire capable 2019 Kia Niro and the relevant sensor for this study is a forward-facing ZED 2 RGB stereo camera made by Stereolabs. The ZED 2 has a 120-degree field of view wide-angle lens that captures images and videos using stereo vision, although only one of the lenses was used for this study. The camera was set to record video at a frame rate of 29 frames per second with a resolution of *1280 x 720* pixels.

The ZED 2 was connected to the in-vehicle computer and data was collected as *.mp4 files over arterial roads with visible tire tracks and occluded lane lines. From these video files, a total of 1,500 individual frames were extracted for ML training. Figure. 1.1 shows an overview of this data collection process. The *1,500* frames of images were divided into three batches, each with *500* images. Different parameters such as exposure, resolution, and occlusion were assessed in the images. Clear tire tracks with distinct tarmac and snow boundaries were chosen from the images.

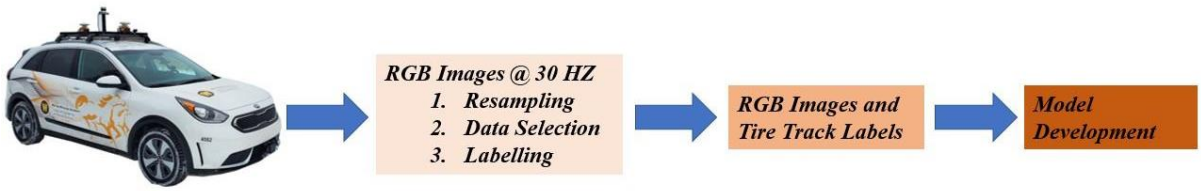

Figure 1. Flow diagram for the data collection, resampling of the data, extracting 1500 RGB Images and corresponding Tire Track Labels, and labeling of data

2.2.2 Data Preparation

The images that were previously segregated into different batches of frames are then used for labeling. Every frame's tire tracks were labeled by hand using an open-source, online tool known as the Computer Vision Annotation Tool (CVAT). Images were uploaded in respective batches and the labeled dataset of each batch was exported with their corresponding raw images using the format: CVAT for images 1.1. This process was again repeated for all the batches.

Each exported dataset contained the raw images and an Extensive Markup Language (XML) file which contained the attributes for the labels, such as the position of the tire-track with their corresponding pixel location on the image, image file name, and their assigned tags (tire-

track, road, road-edge boundary). This process can be updated and more labels can be added according to the use case. The exported labels were then further assessed for post-processing and training the ML and CNN model. The overall data preparation pipeline is described in the next section

2.2.3 Model Development Pipeline

To develop the ML model we must preprocess the data and then perform feature extraction. The process of converting raw data into numerical features that the model can process while preserving information from the original data set is referred to as feature extraction. This is done because it produces significantly better results than applying machine learning to the raw dataset directly.

To improve feature detection and reduce the computational cost, images were masked with a ROI that includes just the road surface and not the entire frame. As stated in [19,20], it is seen that different methods are used to detect road surfaces with high accuracy with an array of sensors. We implemented these road surface detections by using a static ROI in which the pixels inside the ROI are the road surface and every other pixel outside the ROI is considered to be the background. Figure 1.2 shows the process to extract the masked images for the ROI.

The raw images were first resized to the desired shape from their original size of *1280 x 720*. In our case, we chose the images to be of shape *256 x 256*. The road ROI mask was obtained from the raw image to reduce the number of pixels used for training and reduce the computational cost.

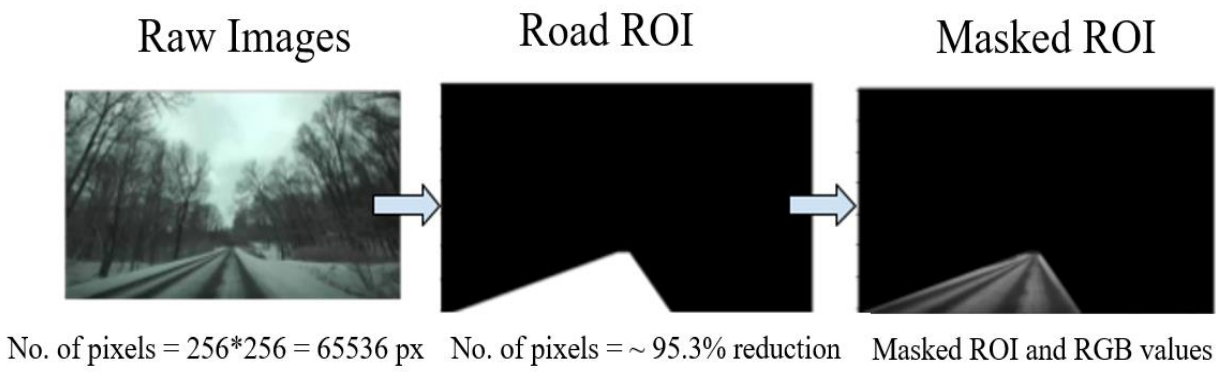

Figure 2. The feature extraction procedure, which begins by extracting only the frames within the ROI and then extracts the features from those pixels.

The Road ROI only consists of *3099* pixels which are only *~ 5%* of the total pixels in the raw image. The ROI mask was then fused with the raw image to obtain all the pixels within the ROI. This will in turn be the input to the model. The different features extracted from the masked images include the red, green, blue, grayscale pixel values, and the pixel X, Y locations as done in the previous study [1].

The different feature vectors shown in Table 1.1 are grouped into different sets and are individually selected to be the final input to the model. The results from these will show the features that contribute the most to the model and yield the highest performance. The model was split into a *55 - 45%* train test split. The entire model was trained using a single input array X having the shape = ((m*p), n) where m is the total number of images, p is the number of pixels in the ROI of each image (*3099* pixels for the *256 x 256* sized images), and n is the number of feature vectors in the array. An overview of this process is shown in Figure 1.3.

Table 1. Feature set properties

Feature set	***Included Feature Vector***	***Train Array Shape (m = 1200)***	***Test Array Shape (m =300)***
0	Gray	(3718800,1)	(929700,1)
1	Gray, X loc, Y loc	(3718800,3)	(929700,1)
2	Red, Green, Blue	(3718800,3)	(929700,3)
3	Red, Green, Blue, X loc, Y loc	(3718800,5)	(929700,5)

2.2.4 Machine Learning Implementation and Evaluation

As seen in our previous study [1], we trained various ML models from the input features and their respective labels. The input feature array X and label vector y were extracted from the image preprocessing and feature extraction block and then fed as inputs to the ML model. Six different models were evaluated to determine the feature set/model combination for the highest performance metrics. Models that were evaluated include K - Nearest Neighbor (KNN), Naive-Bayes, Decision Trees (Dtrees), Random Forest, Linear Regression, and Logistic Regression. These models were chosen for their characteristics and capabilities in commuting binary classification [22–24].

The outputs from the predicted model y_{pred} were compared with ground truth for evaluation. The metrics used for evaluation were the intersection over union (IoU), mIoU, pixel prediction accuracy, precision, recall, F1 score, and frame per second (FPS). These metrics were evaluated based on the ability to draw strong conclusions from the model's performance [22]. Table 1.2 shows the equations demonstrating these calculations as well as the four corners of a confusion matrix, which define the true positives, true negatives, false positives, and false negatives.

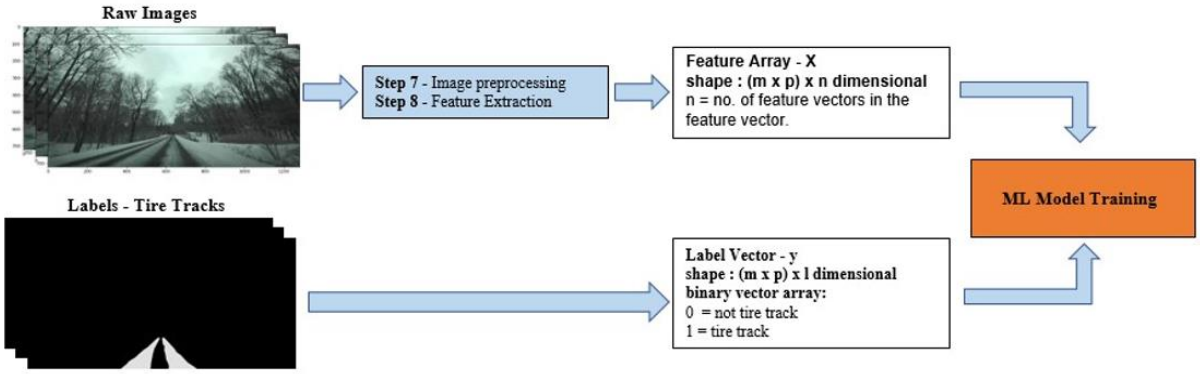

Figure 3. A flow diagram for training the ML model. The features recovered from the raw photos are stored in the input feature array X, and the label vector y contains the pixel status as either tire track (1) or non-tire track (0).

- True Positive (TP): no. of pixels classified correctly as in a tire track
- False Positive (FP): no. of pixels classified incorrectly as in a tire track
- True Negative (TN): no. of pixels classified correctly as not in a tire track
- False Negative (FN): no. of pixels classified incorrectly as not in a tire track

Following the creation of the ML models, we discovered that this method, in our instance, necessitates a significant amount of feature engineering or image pre-processing. The raw images

are cropped and turned to grayscale. Similarly, the segmentation masks are cropped to generate the ROI mask, and the X and Y pixel locations from the segmentation masks are saved to feed into the model, as explained in our image pre-processing and feature extraction sections. Furthermore, the ROI is static, which means it is fixed for each image and does not account for changing road curvature. overall this process necessitates a substantial level of effort, which CNN will address.

Table 2. Equations for metrics used for ML models

Accuracy	$= \frac{total\ correct\ predictions}{all\ predictions} = \frac{TP + TN}{TP + TN + FP + FN}$	(1.1)
IoU (Jaccard Index)	$= \frac{\lvert A \cap B \rvert}{\lvert A \cup B \rvert} = \frac{\lvert A \cap B \rvert}{\lvert A \rvert + \lvert B \rvert - \lvert A \cap B \rvert}$	(1.2)
mIoU	$= 1/n * \sum_{i=1}^{n} \frac{intersection}{union} = 1/n * \sum_{i=1}^{n} \frac{TP_i}{TP_i + FP_i + FN_i}$, where n = # of classes	(1.3)
Precision	$= \frac{TP}{TP + FP}$	(1.4)
Recall	$= \frac{TP}{TP + FN}$	(1.5)
F1 Score	$= 2 * \frac{precision * recall}{precision + recall}$	(1.6)

2.2.5 Convolutional Neural Network Implementation and Evaluation

Deep learning has been shown to perform significantly better on a wide range of tasks, including image recognition, natural language processing, and speech recognition. Deep networks, when compared to traditional ML algorithms, scale effectively with data, do not require feature

engineering, are adaptable and transferable, and perform better on larger datasets with unbalanced classes [25].

CNNs are a type of deep neural network whose architecture is designed to automatically conduct feature extraction thus eliminating this step [26]. CNN's create feature maps by performing convolutions to the input layers, which are then passed to the next layer. In contrast to basic ML techniques, CNNs can extract useful features from raw data, eliminating the need for manual image processing [27,28].

As previously stated, our ML model required feature engineering and did not function as an end-to-end pipeline for tire track identification. To make this process easier and to improve the overall accuracy we have implemented a CNN.

2.2.6 Architecture

The U-net architecture has demonstrated excellent performance in computer vision segmentation [29]. CNN's basic premise is to learn an image's feature mapping and use it to create more sophisticated feature maps. This works well in classification problems since the image is turned into a vector, which is then classified. In image segmentation, however, we must not only transform a feature map into a vector but also reconstruct an image from this vector [29]. U-net architecture was developed specifically for this problem and was first introduced in a medical application [29]. Its structure is depicted in Figure 1.4.

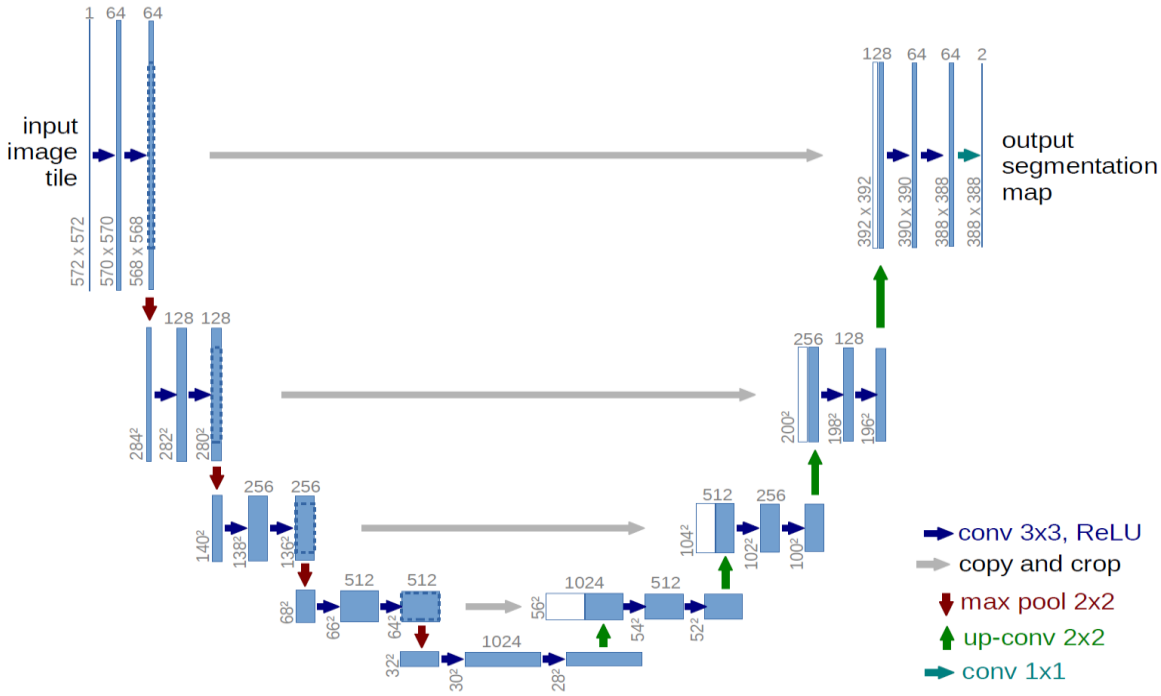

Figure 4. U-network architecture (example for 32x32 pixels in the lowest resolution) [30]. A multi-channel feature map is represented by each blue box. The number of channels is indicated on the box's top. The x-y size is indicated at the box's lower-left edge.

The U-net architecture learns the image's feature maps while converting it to a vector, and the same mapping is used to convert it back to an image. The left side of the U-net architecture is known as the contracting path, while the right side is known as the expansive path. The number of feature channels/filters doubles after each downsampling block to learn more complicated structures from the previous layer's output, while the image size decreases. This path consists of numerous contraction blocks. Each block takes an input and applies it to a 3×3 convolutional layer with a rectified linear unit (ReLU) activation function. The padding is set to 'same' which is followed by a 2×2 max-pooling layer for downsampling. We start off with 32 feature channels and double them with every contraction block until we reach 512 feature channels, which is when

we move onto the expansive path. Each block in the expansive path (shown on the right side of the image) is composed of two 3×3 convolution layers and one 2×2 up-sampling or up-convolution layer with a ReLU activation function and padding set to 'same'. The input is appended by the feature maps of the matching contraction layer with each block in the up-convolution, which is known as concatenating and is indicated by the gray arrow between the two layers. The number of feature channels is halved with each block in this layer. A 1×1 convolution layer is applied in the final layer, with the number of feature maps equaling the number of required classes/segments. In addition, in both the expansive and contraction paths, we add a dropout layer between each convolution layer. This reduces model overfitting by randomly shutting down the necessary number of neurons in that layer [31,32].

2.2.7 Metrics

As mentioned in the ML section, the different metrics are shown which are used to evaluate the model's performance. From equation (1.1) in Table 1.1, the accuracy shows the fraction of predictions our model got right. But accuracy alone doesn’t tell the complete story when working with a class-imbalanced dataset [33] In our dataset, there is a great amount of imbalance between the tire tracks and the background, which is why accuracy is not a good metric for evaluation. This means that the inaccuracy of minority classes is overshadowed by the accuracy of the majority classes when compared to pixel-wise accuracy. IoU, which is also known as Jaccard Index is substantially more suggestive of success for segmentation tasks, especially when the input data is significantly sparse. When training labels contain 80-90% background and only a tiny fraction of positive labels, a basic measure like accuracy can score up to 80-90% by categorizing everything

as background. Because IoU is unconcerned about true negatives, even with extremely sparse data, this naive solution will never arise. IoU computes the overlapping region for the true and anticipated labels by comparing the similarity of finite sample sets *A*, *B* as the IoU [34]. As stated in equation (1.7)

$$Jaccard\ Index\ (IoU) = \frac{|T \cap P|\ (Area\ of\ Overlap)}{|T \cup P|\quad (Area\ of\ Union)} \tag{1.7}$$

T stands for the true label image and P stands for the prediction of the output image. This is used as a metric, providing us with a more accurate means of measuring IoU in our model's segmentation region.

2.2.8 Loss Function

We use two loss functions in our model. Loss functions are used to reduce loss and the number of incorrect predictions made. The loss function Binary Cross-Entropy (BCE) is used in binary classification [35] The BCE function is shown in equation (1.8)

$$BCE\ = -t_1 log(s_1)\ -\ (1 - t_1) log(1 - s_1) \tag{1.8}$$

where t_1 denotes the label/segmentation mask and s_1 denotes the label's predicted probability across all images. We use BCE because our model needs to predict the segmentation mask of the tire track.

The Jaccard Loss, which is equal to the negative $Jaccard\ Index$ from equation (1.7), is the second loss function used. A higher IoU value indicates that there is more overlap between the

true label and the predicted label, but the loss function is concerned with minimizing IoU, which is why we use a negative Jaccard Index as the loss function to reduce loss.

2.2.9 Convolutional Neural Network Model Training

The model was trained using the input images and their associated segmentation masks. We used google colab pro's cloud GPU to train our model. The ML model's input feature vector array was used with feature set 2 (RGB images). The shape of the training array is $(m \times n \times p \times l) = (1300,256,256,3)$ where m is the number of images in the training set, n is the image height, p is the image width and l is the number of channels in the image. In our case, we resize the images to the desired size in feature extraction (2.6b) and use feature set 2, which uses the image's RGB values. We can use the raw RGB images without any pre-processing because no image pre-processing is required.

We consider stochastic gradient descent (SGD) and Adaptive moment estimation (Adam) for our optimizers. Optimizers update the model in response to the loss function's output, attempting to minimize the loss function's output. SGD begins with a random initial value and continues to take steps with a learning rate to converge to the minima. SGDs are simple to implement and fast for problems with a large number of training examples but have a disadvantage in that they necessitate extensive parameter tuning [36] Unlike stochastic gradient descent, Adam is computationally efficient and is better suited to problems with very noisy/or sparse gradients because it computes adaptive learning rates [37] For image segmentation, Adam is thought to be a very powerful loss function [38], which is why we chose Adam as our optimizer. BCE and

Jaccard loss are two different loss functions that we use which is covered in section 8.c. The batch size is set to 16 and the model is run for 25 epochs with an early callback to save the model at the best epoch for the validation loss. For testing, training, and validation, the predicted images are thresholded, so anything above 50% is saved as a correct prediction. There are 7,760,097 trainable parameters in total.

2.2.10 Convolutional Neural Network Model Evaluation

In contrast to our ML models, the model's predicted output was an image. The predicted segmentation masks were then assessed using a variety of metrics. We test the model for IoU, precision, recall, and F1 score, as mentioned in the metrics section. Equations (1-6) show how the confusion matrix is used to perform these calculations. Figure 1.5 shows the outputs from CNN.

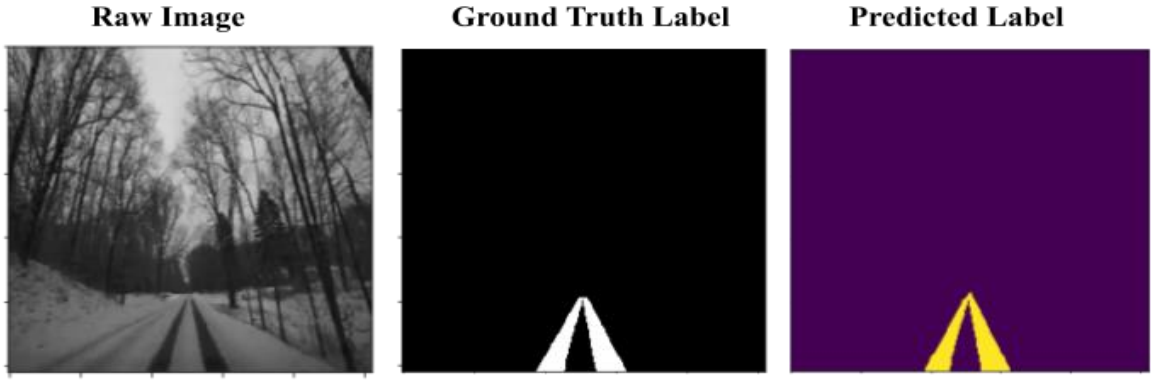

Figure 5. CNN output (The raw image is on the left, the labeled segmentation mask is in the middle, and the predicted segmentation mask from the CNN is on the right)

2.3. Results

When we run the model with the loss function set to BCE and Adam as the optimizer, we see that the model's accuracy increases to ~98%. However, as discussed in the metrics section,

accuracy is not a good metric for datasets with a lot of class imbalance, which is why it produces such high values. Therefore we must also test the IoU.

Figure 1.6a shows that the model with Jaccard loss function has an IoU score of 93% and a validation IoU of 88%. Figure 1.6b shows the IoU of the model with the loss function set to BCE is 89 %, and the validation IoU is 84%. This means that, when compared to BCE, the Jaccard loss function does a better job of finding the intersection/overlapping region for the segmentation masks between the true and predicted. Even though this is true, BCE is still regarded as a good performer because it is only 3% less accurate. The two models have an average frame rate of nearly 350 FPS.

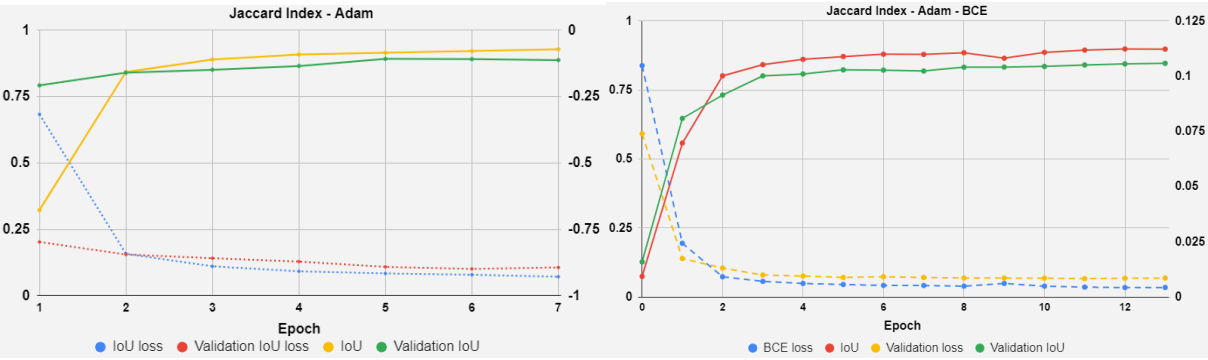

Figure 6.(a) Jaccard loss function, Jaccard Index (IoU) as the metric (b) BCE loss function, Jaccard Index (IoU) as the metric

The results of the best CNN and ML models are summarized in Table II. Dtress with feature set 1 was found to be the model with the best performance in our prior study. We compare the metrics for that model to our CNN model with feature set 2 since we don't have to perform any preprocessing in our case.

Table 3. CNN and ML metrics

Model	***Feature set***	***Accuracy***	***Precision***	***Recall***	***F1 Score***	***FPS***
CNN	2	0.98	0.96	0.95	0.96	350.32
Dtress	1	0.90	0.905	0.911	0.908	1084.1

We observe that the CNN model performs better than the ML model without any image preprocessing on metrics like accuracy, precision, recall, and F1 score, shown in Figure 1.7.

Limitations of this study include comparing metrics such as mIoU with the previous ML models. The ML model with Dtress and feature set 1 obtains a mIoU of 83 %, whereas the CNN achieves a mIoU of 65 %. This could imply that the ML model is more accurate at predicting tire tracks, but it is not the whole story. A static ROI for the ML model was employed, which means that the ML model only receives a portion of the raw image and the segmentation masks. The mIoU calculates the IoU for each class before averaging the results across all of them. Because we just feed a section of the image into the ML model rather than the complete image, it performs better at detecting these tire tracks only in that precise region, which implies the model will not do well if the road geometry shifts or if the model is tested on the entire image.

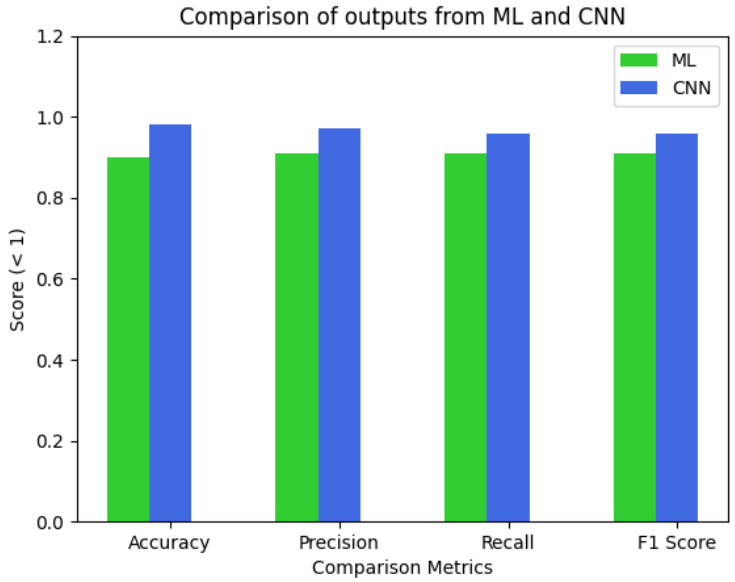

Figure 7. Precision, Accuracy, Recall and F1 score metric comparison between CNN and Dtress.

The CNN, on the other hand, does not require a ROI but instead takes in the full image as input, lowering the mIoU because it is no longer simply looking at the ROI but the complete image. Another explanation for CNN's lower mIoU is the significant class imbalance (more background pixels and fewer tire track pixels), as well as the fact that deep neural networks require more training data than ML models which means to improve the mIoU we will need to train the model on larger datasets. Another way to attain a higher mIoU would be to crop the ROI for images and segmentation masks in the same way as our ML models, and then use that as the input to the CNN. However, this would necessitate preprocessing and feature engineering, which is one of the drawbacks from the ML models addressed in this paper

2.4 Conclusion

This study addresses the research gap of driveable region detection for snow-covered roads using a single camera sensor that is implementable in modern ADAS products. We proposed a new method for extracting the drivable region for snowy road conditions when the lane lines are occluded by instead focusing on identifying tire tracks. First data was collected on our instrumented vehicle and then the data was processed by extracting the frames from the videos, segmenting them into batches, and labeling them with CVAT.We have showcased how this information was used in the model development process. Using just the raw image and no image pre-processing or feature extraction, we evaluated a U-net-based CNN for IoU, Accuracy, Feature set, Recall, F1 score, and FPS. The IoU score for the model with the Jaccard loss function was 93%. The model had an accuracy of 98%, a 95% recall, a 96% precision, and a 96% F1 score. Furthermore, we found a significant improvement in these metrics when compared to the ML model from the previous study. By feeding in the raw image and obtaining the predicted tire tracks, this method offers a full end-to-end solution for detecting drivable regions in snowy road conditions.

Overall this study demonstrates that drivable region detection in inclement weather is feasible using current technology in a single camera. The results can be improved by improving image processing and tuning the CNN. Beyond this study, there are many other research gaps in inclement weather automation that need to be addressed to combat the significant loss of life that comes from these scenarios.

Our research question for this study was “Can we create and implement a method of detecting drivable region in snow-occluded lanes using a single on-vehicle camera sensor?”. In this chapter we developed a novel way of identifying the drivable region by employing a purpose-built end-to-end pipeline using our custom dataset and deep learning to address the question. The work established in this study will serve as a tool to further expand the ODD of ADAS in snow conditions.

Future work includes conducting an expansive data collection process and creating a larger dataset with various scenarios such as different lighting conditions, road conditions, sun angles, intersections, roundabouts, and other corner cases. We have addressed one section of this future work in our 1000+ mile data collection task for our National Science Foundation (NSF) Partner For Innovation (PFI) project during the winter of 2022-2023 in Kalamazoo MI.

APPENDIX

This study was further expanded and published in a book chapter. The citation for this work is as shown below.

Kadav, P., Sharma, S., Araghi, F.M., and Asher, Z.D., “Development of Computer Vision Models for Drivable Region Detection in Snow Occluded Lane Lines,” in: Kukkala, V. K. and Pasricha, S., eds., *Machine Learning and Optimization Techniques for Automotive Cyber-Physical Systems*, Springer International Publishing, Cham, ISBN 9783031280160: 591–623, 2023

CHAPTER III

ROAD SNOW COVERAGE ESTIMATION USING CAMERA AND WEATHER INFRASTRUCUTRE SENSOR INPUTS

This chapter consists of work developed in a project sponsored by the Natioanl Science Foundation - Partnership for Innovation (NSF - PFI) program and presented work at the SAE World Congress Experience conference 2023. Please note that much of this chapter is presented verbatim from the publication. Parth Kadav presented and led the development of the methodology and authorship of this publication, with the co-authors providing suggestions and guidance. The citation for thie work is as shown below.

Kadav, P., Goberville, N., Prins, K., Siems-Anderson, A. et al., "Road Snow Coverage Estimation Using Camera and Weather Infrastructure Sensor Inputs," SAE Technical Paper 2023-01-0057, 2023, https://doi.org/10.4271/2023-01-0057.

Abstract

Modern vehicles use automated driving assistance systems (ADAS) products to automate certain aspects of driving, which improves operational safety. In the U.S. in 2020, 38,824 fatalities occurred due to automotive accidents, and typically about 25% of these are associated with inclement weather. ADAS features have been shown to reduce potential collisions by up to 21%, thus reducing overall accidents. But ADAS typically utilize camera sensors that rely on lane visibility and the absence of obstructions in order to function, rendering them ineffective in

inclement weather. To address this research gap, we propose a new technique to estimate snow coverage so that existing and new ADAS features can be used during inclement weather. In this study, we use a single camera sensor and historical weather data to estimate snow coverage on the road. Camera data was collected over 6 miles of arterial roadways in Kalamazoo, MI. Additionally, infrastructure-based weather sensor visibility data from an Automated Surface Observing System (ASOS) station was collected. Supervised Machine Learning (ML) models were developed to determine the categories of snow coverage using different features from the images and ASOS data. The output from the best-performing model resulted in an accuracy of 98.8% for categorizing the instances as either none, standard, or heavy snow coverage. These categories are essential for the future development of ADAS products designed to detect drivable regions in varying degrees of snow coverage such as clear weather (the none condition) and our ongoing work in tire track detection (the standard category). Overall this research demonstrates that purpose-built computer vision algorithms are capable of enabling ADAS to function in inclement weather, widening their operational design domain (ODD) and thus lowering the annual weather-related fatalities.

3.1 Introduction

According to the Fatality Analysis Reporting System (FARS) encyclopedia by the National Highway Traffic Safety Administration (NHTSA), there were nearly 103,172 fatal crashes from the year 2018-2020 in the United States [39]. Out of these fatal crashes, nearly 10% were related to inclement weather such as snow, ice, sleet, and rain. Similarly, during 2007-2016, weather-related vehicular crashes accounted for nearly 21% of all reported crashes annually resulting in 16% of crash fatalities and 19% of crash injuries throughout the United States [40]. It is really

crucial to understand how different weather conditions can affect the transportation network. Fundamentally, adverse weather conditions can cause 1) Impairment of situational awareness and 2) Inhibitions to vehicular maneuverability [41]. Due to poor visibility caused by heavy rain, blowing dust or snow, or dense fog, multi-vehicle collisions can occur when drivers lose awareness of their position, location, and speed in relation to other cars. Automated vehicles can open the way for dependable and safe driving in any weather [3,8,10,42].

Nearly 94% to 96% of all auto accidents are caused due to human errors (speeding, aggressive/reckless driving, distracted driving, chemical impairment, and drowsy driving), which are preventable according to a study conducted by NHTSA in 2016 [43]. ADAS systems were created to automate driving tasks, improve aspects of the driving experience, and increase safety and safe driving practices [44]. About 40% of all accidents in passenger vehicles can be prevented or significantly reduced with the use of ADAS features including Forward Collision Warning (FCW), Automated Emergency Braking (AEB), Lane Departure Warning (LDW), Lane Keeping Assistance (LKA), blind spot warning assistance, and many more. [1–3,13]. Furthermore, ADAS features such as FCW and AEB alone reduce front-to-rear crashes by nearly 50% [8]. From the 1,853 driver injury crashes studied in [5,6], it was discovered that LDW and LKA systems were able to reduce head-on and single-vehicle crashes on roads at higher speed limits (45-75 mph) and visible lane markings by nearly 53%. Based on the statistics, ADAS features such as LDW, LKA, AEB, and FCW significantly cut down on collisions caused by human and external variables[7].

One of the ways that ADAS improves safety is to provide vital information about the vehicle and its surroundings by classifying road lanes [17,45]. Lane recognition is the foundation

of many driving assistance systems such as LKA, LDW, and Lane Centering Assist (LCA), specifically identifying lane markings. During snowy conditions, lane markings can get obscured or hidden which can render driving assistance systems ineffective. In reality, snow accumulation on highways frequently leads drivers to disregard lane positions and drive on different regions of the road as necessary, in other words, forming informal auxiliary traffic lanes [41]. The poor performance of driver assistance systems in adverse weather conditions, such as rain, snow, fog, and hail, is among the most crucial challenges in vehicle automation. Unfortunately, just like a human's vision, the sensors used by driving assistance systems can be negatively affected by inclement weather. Rainy and foggy conditions cause significant degradation in the performance of Camera, Radar, and LiDAR [46,47]. The LiDAR will misdetect objects under rainy and snowy conditions due to rain droplets, snow particles, and ice [48]. Similarly, Radar, which is used for many driver assistance systems such as adaptive cruise control (ACC) and AEB, has an issue with signal attenuation in the rain [46,49,50]. On-board vehicle cameras are essential in providing both the systems and the driver with crucial information. Cameras come standard in all vehicles with level 1 and level 2 autonomy [51]. Various sensors operate differently in various weather conditions, according to the literature review conducted in this section. To enable ADAS performance in inclement weather conditions and actively toggle between sensors based on environmental conditions, a method to determine the category of road conditions in inclement weather needs to be established so that purpose-built perception techniques can be deployed.

There are few studies in the literature that address the issue of estimating road weather conditions for inclement weather. One such study conducted in 2011 introduced a method of estimating road weather using a ML model trained with camera images and Road Weather Information Systems (RWIS) data [52,53]. The results from this study indicate that the model was capable of achieving a 91% accuracy on the test set for classifying the road conditions into five different categories (dry, ice, snow, track and wet). This study utilized Principal Component Analysis (PCA) to determine which inputs contributed the greatest to model performance. The model used limited training data and had a biased dataset gathered from static images at intersections. Another study conducted by Qian proposes a system that categorizes road conditions using static images using a camera [54]. This study obtained an accuracy of 68% on classifying the road conditions into dry, wet, and snow. However, this study only uses a dataset of 100 images with a 50-50 train test split. Having such a small dataset, specifically a small training set can lead to poor performance and generalization. The methods and results of these studies provide ways to estimate the weather conditions mainly for object-dependent ADAS purposes and do not talk about lane-dependent features such as lane lines, road type, and amount of snow coverage in the lane which are independent of any objects in the environment. Additionally, they only employ camera data using a small dataset as the input, and no additional input is provided to the models. Therefore, a more rigorous study of snow coverage estimation using a multi-input model is needed to move this research forward. It is crucial for estimating the road snow coverage in order to expand the ODD of ADAS and use algorithms that detect the drivable region in snow-occluded lane lines as done in our previous studies [2,55].

To address the need for real-time estimation of road snow coverage, the proposed method uses Machine Learning (ML) models that use camera data and infrastructure weather sensor data as inputs to predict road snow coverage. We recorded and labeled each image in different categories based on the subjective level of snow coverage on the road. The three different snow coverage categories were none, standard and heavy. The images were processed using feature engineering, and different image features were obtained. The inputs to the ML models were the image-level features and ASOS infrastructure weather sensor features. We tested the performance of the different models on key metrics such as accuracy, precision, recall, and F1 score. The goal of this work is to provide a robust snow coverage estimation method for ADAS perception systems using a single-camera sensor and infrastructure-based weather sensor data. The methods discussed in the next section talk about the details of the different feature sets, ML methods, and the overall performance of the various models in classifying road snow coverage.

3.2 Methodology

In this section, we will first examine the drive cycle that was selected, the vehicle platform, and the equipment used, followed by a discussion of the methods to collect and prepare the data. Following that, several ML models will be developed and assessed.

3.2.2 Drive cycle

The drive cycle consisted of the two-lane arterial roads in Kalamazoo, MI. The route was selected based on having low traffic volume, two lanes, clear visible lane lines, and occluded lane lines. Arterial roads receive snow level variation as they are plowed irregularly and have a low

amount of traffic which results in varying amounts of snow coverage. The route consisted of 5 different road sections, which were A, B, C, D, and E each one mile in length with different cardinal directions. To add variation to the dataset, the data was collected on different days with changing snow precipitation forecast through the 5 different road segments during the winter of 2020-2021. Figure 2.1 shows the different road segment.

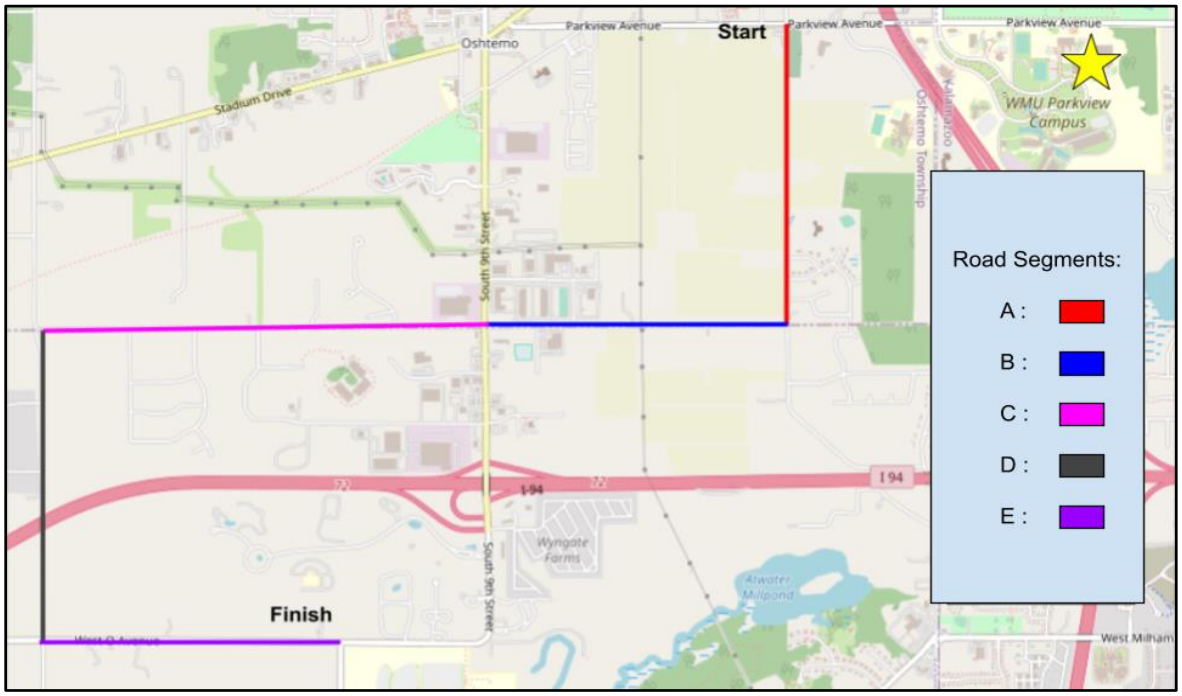

Road segments used for data collection during the winter of 2020-2021 in Kalamazoo, MI.

3.2.3 Vehicle Platform and Sensors

The Energy Efficient Autonomous Vehicles (EEAV) research vehicle platform, shown in Figure 2.2a, was used to collect data. This is a 2019 Kia Niro and includes a forward-facing RGB camera, Polysync Drivekit, Neousys in-vehicle computer, vehicle Controller Area Network (CAN)

bus interface, and a Mobileye camera. We used the forward-facing ZED 2 RGB camera from Stereolabs [56]. The ZED 2 is a widely available machine vision camera, which is available with a Software Development Kit (SDK) that provides greater functionality for our instrumented research vehicle. The ZED 2 provided us with the raw RGB images used to build the dataset. The images were captured at a resolution of 1280 x 720 and at a frame rate of 30 frames per second

Figure 8. (a) Kia Niro Instrumented Research Vehicle, (b) ZED 2 stereo camera.

3.2.4 Infrastructure Weather Sensor

This study used historical weather data collected by the Automated Surface Observation System or ASOS station located at the Kalamazoo Battle Creek International Airport. ASOS is considered a “gold standard” observation, used widely in the atmospheric sciences [57]. Figure 2.3 depicts an ASOS station that was deployed at airports around the United States to enhance the nation's weather services. The intention was to provide reliable and useful automated weather observations in a cost-effective manner [58]. The ASOS dataset used contains weather data observations for the corresponding days of collected drive cycles. This data is published in one-minute intervals for parameters such as visibility, temperature, wind characterization,

precipitation, and atmospheric pressure. While ASOS stations are capable of observing falling precipitation, there are a number of issues that can lead to erroneous precipitation reports. These include the inability to recognize precipitation type for frozen or mixed precipitation events [59] and undercatch of snowfall amount or intensity in strong winds [60]. However, in the U.S., snowfall intensity is measured not by accumulation but by visibility, with light snow categorized as >1 km visibility, moderate between 0.5 and 1 km visibility, and heavy snow less than 0.5 km visibility [61]. Due to the more reliable automated visibility observations, for this study, we focused on the visibility coefficient.

Figure 9. ASOS weather station [62].

3.2.5 Data Pipeline

Figure 2.4 shows the overall model development pipeline. This pipeline shows the different steps taken to achieve model results.

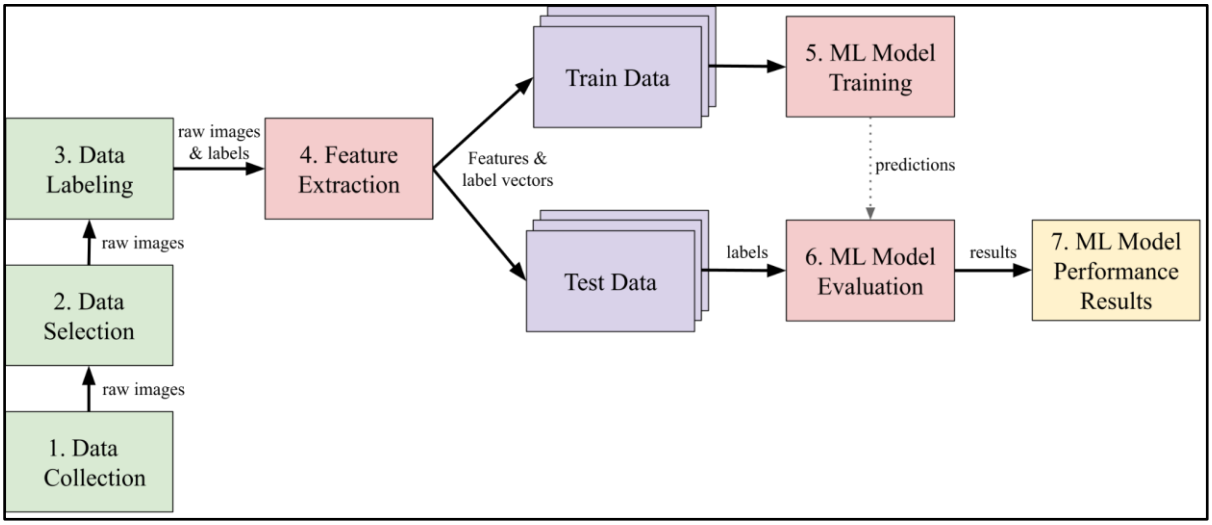

Figure 10. Overall model development pipeline.

3.2.6 Data selection and filtering

We collected ~ 100,000 RGB images. The images were resampled from 30 fps to match the ASOS dataset. As the ASOS data was sampled every minute (0.167 Hz), we had to map the images with ASOS data based on the timesteps. Further quality control was taken into account and these images were assessed for poor quality such as over-exposed images from sun glare, windshield wiper obstruction, image noise, distortion, etc. When finished the final dataset had a total of 20,883 images spanning across the five road sections on different days.

3.2.7 Labeling

A subjective method was used to place data from each road segment into three categories: none, standard, or heavy. Each of the road segments were assigned into one of these categories based on how much snow was covering the surface of the road during the entire road segment video. Figure 2.5 shows the three different snow conditions. Figure 2.5a shows the none condition, Figure 2.5b shows the standard condition, and Figure 2.5c shows the heavy condition. We labeled all the images in the dataset based on the subjective snow condition

Figure 11. (a) None condition, (b) standard condition, and (c) heavy condition.

3.2.8 Feature extraction

To build and train the ML models, we first needed to preprocess the data and then extract features. Feature extraction transforms raw data into numerical features the model can process while retaining original data. This works better than applying ML to the raw dataset [63]. To start the process of feature engineering, the raw image was first down sampled to *256 x 256* from its original dimensions of *720 x 1280*. Resizing results in reduced computational load while training models. To further improve feature detection and reduce computational complexity, images were

masked with a static Region of Interest (ROI) that only included the road surface. The Road ROI mask was then fused with the raw image to output the Masked ROI. The masked ROI contains less than 10% of the total pixels when compared to the raw image. Similar to our previous study, we decided to create different feature sets, each containing various image features, which will help in identifying features that perform better compared to others [1].

Images contain pixel-level color channel values which are contained in 3 dimensional arrays which contain the red, green and blue values for each pixel (RGB). For this study we decided to use the RGB mean and standard deviation values as the image-level features. The RGB values change as the level of snow coverage changes in the image, with a lower road snow coverage, we have lower overall RGB intensities in the image and as the snow coverage increases the RGB intensities increase. These features strongly correlate with the changing snow coverage on the road.

Table 2.1 shows the different feature sets that were created for this study. We organized these features into sets where each set has its corresponding feature vector. For example, feature set 0 has three feature vectors which are the mean values for the red, green and blue color channels in the masked ROI image, feature set 2 has six feature vectors which are the mean (Equation 2.1) and standard deviation (std. dev) (Equation 2.2)

$$\underline{x} = \frac{1}{N}\sum_{i=1}^{N} x_i \tag{2.1}$$

$$\sigma = \sqrt{\frac{\Sigma(x_i - \mu)^2}{N}} \tag{2.2}$$

values for red, green and blue color channels respectively. Feature set 3, 4 and 5 include the ASOS visibility coefficient input along with the image-level features. Each feature set has its own feature array X, the shape of feature array $X = (m\ x\ n)$ dimensions where m = number of images in the array and n = number of features. The feature array X is the input. Similarly label vector $y = (m\ x\ 1)$ dimensions, where m is the number of images in array corresponding to the feature array X representing the subjective snow coverage, as mentioned in the labeling location section ($none: 0, standard: 1, heavy: 2$). Each element in the label vector maps the label to its corresponding input from the feature array X. The dataset was split into 70 - 30% for training and testing.

Table 4. Included feature sets used in model development along with their array shapes

Feature Set	**Included Feature Vector**	**Train Array Shape (m = 14,618)**	**Test Array Shape (m = 6265)**
0 *(Img-level)*	R,G,B (mean)	(14,618, 3)	(6,265, 3)
1 *(Img-level)*	R,G,B (std. dev)	(14,618, 3)	(6,265, 3)
2 *(Img-level)*	R,G,B (mean), R,G,B (std. dev)	(14,618, 6)	(6,265, 6)
3 *(Img-level + ASOS)*	R,G,B (mean), visibility coefficient	(14,618, 4)	(6,265, 4)
4 *(Img-level + ASOS)*	R,G,B (std dev), visibility coefficient	(14,618, 4)	(6,265, 4)
5 *(Img-level + ASOS)*	R,G,B (mean), R,G,B (std. dev), visibility coefficient	(14,618, 7)	(6,265, 7)

3.2.9 Machine Learning Techniques

We evaluated different types of ML algorithms to test which models perform better in combination with the different types of feature sets. The six different ML models that were evaluated were: Decision Trees (dtrees), Random Forests (rforest), K-Nearest Neighbors (KNN), Logistic Regression, Support Vector Machines (SVM), and Naive-Bayes (naive). These models were selected based on their capabilities and demonstrated performance for computing classification tasks for computer vision applications. [22,24,64].

Let us look at an overview of all the models used in this study and their computational capabilities. Dtrees and rforest work by making a series of logical decisions mapped as nodes on a tree. This offers insight into relevant features. Training these models is computationally heavy. Both decision trees and random forest work well with less number of features. Logistic Regression works by fitting a logistic curve to the data and works well on datasets in which there is minimal overlap on the classes. Naïve Bayes offers a relatively simple model and performs well on datasets with less features that are independent of each other. Support Vector machines work by mapping the data points onto a space with more than two dimensions and then finding a hyperplane that groups them. K Nearest Neighbors is a simple algorithm that performs well in classification tasks. With our dataset k neighbors are used to label new data based on proximity to neighboring data-point. KNN works well with large, noisy datasets. [65,66]. The work in this paper was performed in Python using models provided by the open-sourced ”scikit-learn” python package [67]

3.2.10 Evaluation Metrics

The predicted outputs of the model y_{pred} were compared with the ground truth labels y and then evaluated for various metrics. The metrics used for evaluation were prediction accuracy, precision, recall, F1 score, and average model compute time. Equations 2.3 to 2.6 show how these metrics are calculated using the four corners of the confusion matrix: true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN). Accuracy is the fraction of predictions the model got right which means the number of images were correctly classified as none, standard or heavy snow based on their condition. Precision measures the quality of a model's positive prediction. Recall displays the proportion of accurate positive predictions made among all possible positive predictions. Precision and recall together make up the F1 score.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (2.3)$$

$$Precision = \frac{TP}{TP + FP} \quad (2.4)$$

$$Recall = \frac{TP}{TP + FN} \quad (2.5)$$

$$F1\ Score = 2 * \frac{precision * recall}{precision + recall} \quad (2.6)$$

3.3 Results

The results of this research include an overview of the analyses conducted for image level features and ASOS weather data features as well as the results from the ML training for estimating

the snow coverage using different features as inputs. Results were obtained for a total of 35 different ML models. When using only image-level features, the order of the best-performing ML models was: SVM, Naive-Bayes, Logistic Regression, KNN, Random Forests, and Decision Trees. When we use variation in the feature sets as input to the model, such as feature set 5 which includes all image-level features and the snow visibility, we obtained the best-performing model. The results indicate that using image-level features along with the visibility coefficient from the ASOS dataset improves the performance of the model in key metrics such as accuracy, F1 score, and precision by a significant margin irrespective of the model used. To look at one such example, figure 2.6 highlights the most important feature in feature set 5 for dtrees. The two most important features for this model and feature set combination are the blue mean value from the image-level feature and the visibility coefficient from the ASOS dataset. This implies that both image-level features and infrastructure weather sensor data input play an important role in enhancing the models performance which is consistent with the results from other models as well.

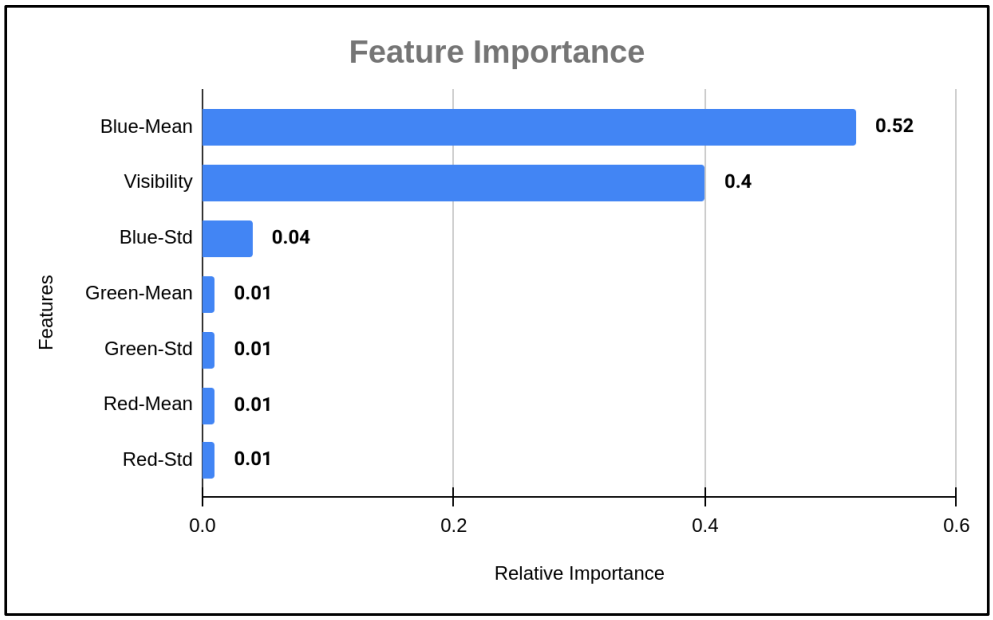

Figure 12. Feature importance for Dtrees with the feature set 5.

To further illustrate the importance of adding the weather sensor data as an input, we obtained the confusion matrix for dtrees with all image-level features (feature set 2) in Figure 2.7b, and all image-level and infrastructure weather sensor data (feature set 5) in Figure 2.7a. The vertical axis shows the true labels and the horizontal axis shows the predicted classes. The diagonal shows the classifications for each of the snow coverage conditions as the first element in the diagonal shows the True Positives for class 0 (none), class 1 (standard), and class 2 (heavy). The confusion matrix heatmap shows that feature set 5 outputs more TP's for each class than feature set 2.

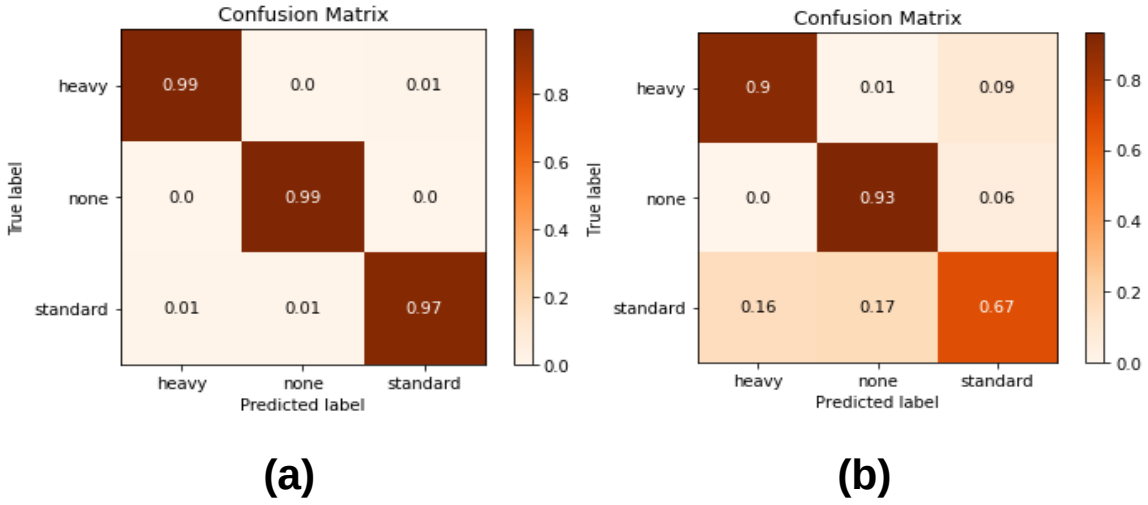

Figure 13. Confusion matrix heat map for (a) Dtrees with feature set 5, and (b) Dtrees with feature set 2.

Figure 2.8 shows the comparison between the 6 different models for feature set 2 and feature set 5. As seen in Figure 2.7 all of the models perform at least ~67% better with both image-level and weather data features (feature set 5) when compared to only image-level features (feature set 2). The best performing model for feature set 2 was svm which tied with logistic regression, and naive. Dtrees with feature set 2 performed poorly when compared to the other models. Contrastingly, adding the snow visibility input from ASOS improved the model performance significantly for all models which is shown by the blue bars. The best performing model for feature set 5 is dtrees.

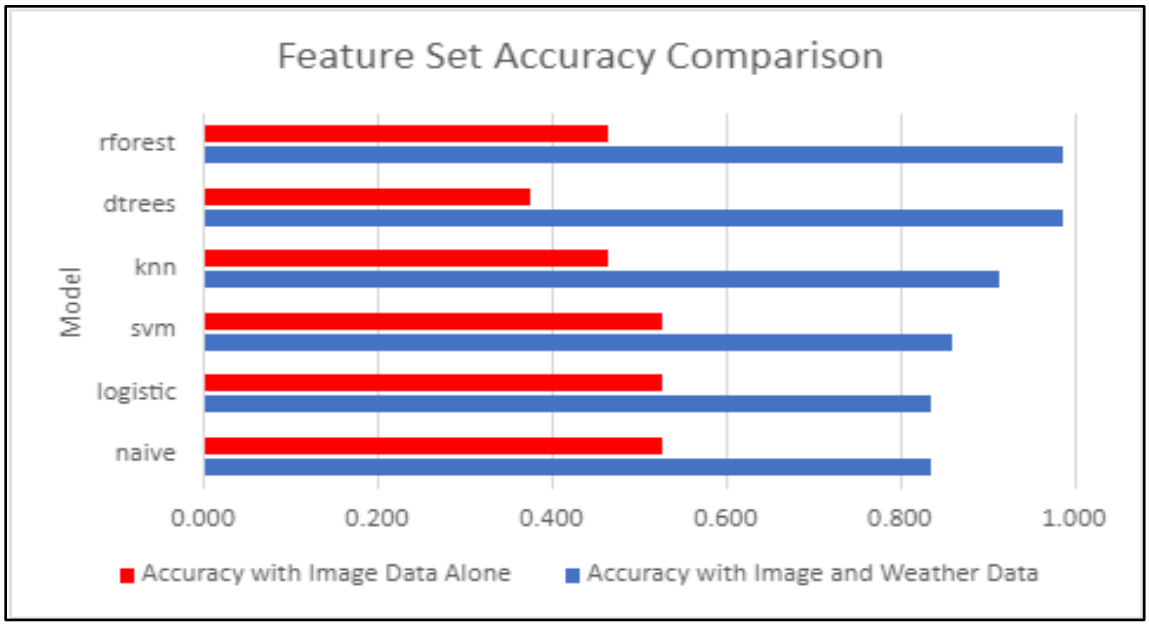

Figure 14. Feature Set Accuracy Comparison between feature set 2 (Image data alone) and feature set 5 (Image and Weather Data).

Figure 2.9 shows Accuracy, and F1 Score by the models using all features (feature set 5). Dtrees achieved an average compute time of 9.51 seconds and rforest achieved an average compute time of 0.09 seconds. For feature set 5 the best performing models are random forest and decision trees both achieving 98.8 % Accuracy and 98.8% F1 score. As the number of features increase, both rforest and dtrees perform significantly better on the same dataset.

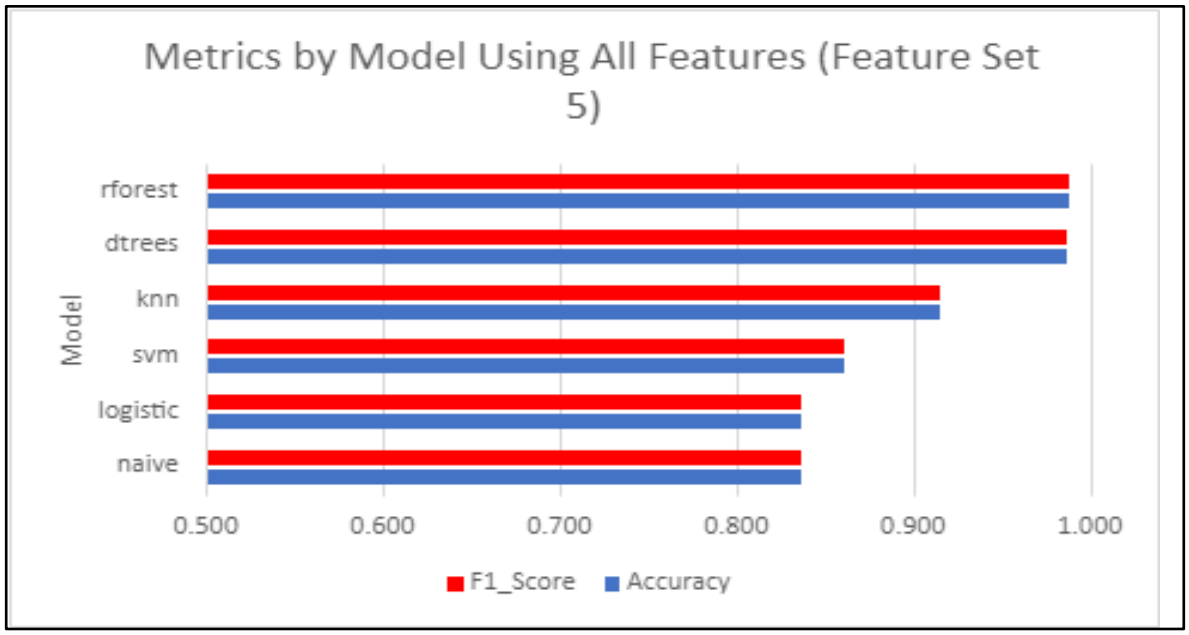

Figure 15. Comparison of Accuracy, Precision, Recall, and F1 score by model for feature set 5.

So, to summarize the results, both the image-level features and weather sensor input are equally important as shown in Figure 2.6, 2.7, 2.8, and 2.9. A critical advantage of the image data is that is that it is precisely local to the car, although the weather sensor provides excellent area-wide information that may impact road visibility, the image data from the vehicle can be used to accurately determine, with input from the general weather data, what the road conditions are in the current location of the vehicle. Adding easily available weather data from existing infrastructure is a highly effective means of improving our ability to estimate local road conditions.

3.4 Conclusions

In this study we derived a method of estimating the snow coverage on the road using a single camera sensor and infrastructure weather data inputs using ML. Firstly, data was collected

using the instrumented research vehicle along arterial roads in Kalamazoo, MI. This data was then processed and cleaned for model development. Additionally, infrastructure-based weather sensor data such as snow visibility was acquired from ASOS. Features were extracted from the processed camera data and ASOS dataset to further create different features sets. These feature sets were used as inputs to the different supervised ML models. In total we had 35 different model-feature sets combinations. We compared and analyzed the performance of all models based on metrics such as Accuracy, Precision, Recall, and F1 score. The best-performing model using all image-level features (feature set 2) yielded an accuracy of 52.8% whereas the best-performing model with both image-level features and weather data feature (feature set 5) had an accuracy of 98.8%. This demonstrates that both image-level features and weather sensor inputs equally improve the performance of the models.

Overall, this study demonstrates that we can estimate the snow coverage on the roads using a custom dataset with just one camera sensor and infrastructure weather data. Categorizing snow coverage will enable ADAS products to operate in inclement weather conditions. This study lays the foundation for broadening the ODD of AVs which will also positively impact the operation of AVs, minimizing crash injuries and fatalities. Additionally, higher resolution on-vehicle weather sensor data as inputs in conjunction with image data would further enhance the model's performance. We could get accurate local weather information from an on-vehicle weather sensor such as the MARWIS which provides us with dynamic road condition information [68]. Adding additional features available in the ASOS dataset along with the on-vehicle weather sensor data such as friction, ice percent, road condition, water film height, and precipitation would help in

improving the model's performance. Future work for this study will include estimating snow coverage using data from both infrastructure and on-vehicle sensor data and using DL models.

Our second research question this “Can we utilize on-vehicle camera sensor and infrastructure weather sensor data inputs to estimate the amount of snow coverage on the road local to the vehicle?”. In Chapter 2, we developed a methodology that outputs the snow coverage in regions of snow occluded lanes using a single on-vehicle camera and infrastructure-based weather sensor data inputs. The work established in this chapter complements the work in Chapter 1 by providing the subsystems with accurate local road conditions. The methodology established in Chapter 2 can further help expand the ODD of ADAS in snow conditions.

Future work would include expansion of this study using CNN’s for predicting the road condition. We would also use an on-vehicle weather sensor known as the Mobile Advanced Road Weather Information Sensor (MARWIS) to obtain various weather parameters needed for further model development.

CHAPTER IV

CONCLUSIONS

The methodologies outlined in this document represent a novel expansion of Advanced ADAS under adverse weather conditions, offering significant contributions to the field of vehicle engineering. The areas of focus include enhancing drivable region detection for ADAS in snowy weather using an on-vehicle camera sensor, along with accurately estimating the extent of snow coverage on roads through the integration of the on-vehicle camera sensor and infrastructure weather sensor.

The first study details a comprehensive end-to-end pipeline that addresses the research gap in drivable region detection on snow-covered roads using a single camera sensor. The study also covers an extensive data collection and labelling process. This effort aims to expand the ODD of ADAS in inclement weather conditions, when their assistance is most crucial. By applying the outlined methodology of data collection, data processing, machine learning and deep learning, results demonstrated a 93% Intersection over Union (IoU) score on training set and 88% IoU score on unseen validation set. This study addressed the stated research question, proving that current vehicle technology can indeed be used to broaden the ODD of ADAS in snowy weather conditions to detect the drivable region and enable safe lane navigation using purpose-built deep learning models.

The second study demonstrates the effective utilization of image data in conjunction with existing infrastructure weather sensor information to estimate the extent of snow coverage on roads proximate to the vehicle. At first, we gathered a specific dataset tailored for this study. We then

assessed this data, assigning subjective snow levels based on image features. Afterward, we extracted essential characteristics to ensure the best performance from our machine learning models. A total of six distinct feature sets were created and evaluated on five different machine learning models resulting in a total of 35 different machine learning model combinations. The results showed that using all image features along with infrastructure weather sensor data inputs achieved an accuracy of ~ 98.8% in classifying the road condition. This second study successfully addressed the associated research question by describing and demonstrating a novel methodology that expands the ODD of ADAS in snowy weather conditions by providing on-vehicle road classification using camera and infrastructure weather sensor data inputs. This study can be further used for active sensor toggling and improved feedback for vehicle safety systems.

What tools are needed to expand the ODD of ADAS and how might these work? We learned from work established in Chapter 2 that lane detection serves as the backbone for ADAS and lane detection in snow occluded lanes is the means to expand the ODD of ADAS in such conditions. In order to put in place purpose-built perceptions systems, we need to first estimate the road conditions and this has been in established in Chapter 3. The tools developed in these studies form a basis for the future improvement of ADAS in inclement weather conditions and thus expanding the ODD of automated vehicles in snowy weather conditions. Figure 3 shows an overall system's level diagram to show how they would tie together in the full vehicle system. This has been filed as a patent through the WMU Tech Transfer Office.

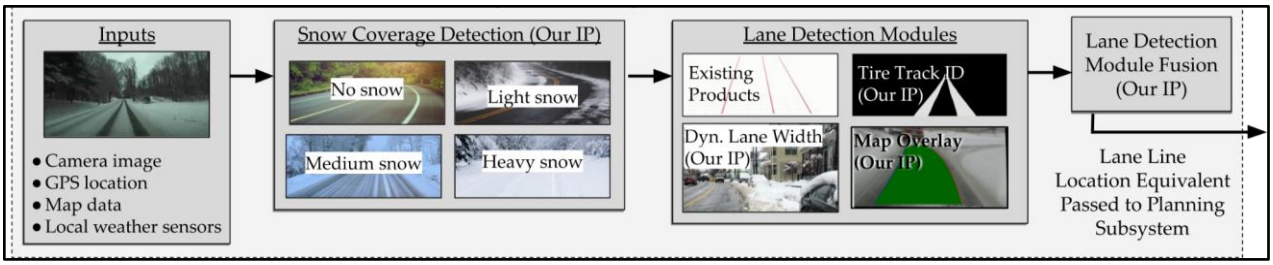

Figure 16. Systems level diagram for work presented in Chapter 1 and 2. Filed for patent.

CHAPTER V

FUTURE WORK

Much of this work is currently being actively pursued under the guidance of my advisor Dr. Zachary Asher and my committe members towards my Doctoral dissertation. This work is being funded by the National Sceince Foundation - Partnership for Innovation. Chapter 1 and 2 would be expanded with the help of expansive data collection that would include a varied dataset. I am also assisting on the Infrastructure Enabled Energy Efficient Autonomous Vehicles project funded by the Department of Energy (DOE). The future chapters of this work would consist of work under the two projects and would appear in my Doctoral dissertation.

REFERENCES

1. Goberville, N.A., Kadav, P., and Asher, Z.D., “Tire Track Identification: A Method for Drivable Region Detection in Conditions of Snow-Occluded Lane Lines,” SAE Technical Paper, 2022.

2. Kadav, P., Goberville, N., Motallebiaraghi, F., Fong, A., and Asher, Z.D., “Tire track identification: Application of U-net deep learning model for drivable region detection in snow occluded conditions,” *Intelligent Transportation Systems World Congress*.

3. Svancara, A.M., Horrey, W.J., Tefft, B., and Benson, A., “Potential Reduction in Crashes, Injuries and Deaths from Large-Scale Deployment of Advanced Driver Assistance Systems,” *AAA Foundation for Traffic Safety.*, 2018.

4. Varghese, J.Z. and Boone, R.G., “Overview of autonomous vehicle sensors and systems,” *International Conference on,* 2015.

5. Sternlund, S., Strandroth, J., Rizzi, M., Lie, A., and Tingvall, C., “The effectiveness of lane departure warning systems—A reduction in real-world passenger car injury crashes,” *Traffic Inj. Prev.* 18(2):225–229, 2017.

6. Kusano, K.D., Gabler, H., and Gorman, T.I., Fleetwide Safety Benefits of Production Forward Collision and Lane Departure Warning Systems, *SAE International Journal of Passenger Cars - Mechanical Systems* 7(2):514–527, 2014, doi:10.4271/2014-01-0166.

7. Kusano, K.D. and Gabler, H.C., “Comparison of Expected Crash and Injury Reduction from Production Forward Collision and Lane Departure Warning Systems,” *Traffic Inj. Prev.* 16 Suppl 2:S109–14, 2015.

8. Real-world benefits of Crash Avoidance technologies, Insurance Institute for Highway Safety & Highway Loss Data Institute, 2020.

9. Advanced driver assistance systems: global revenue growth 2020-2023, https://www.statista.com/statistics/442726/global-revenue-growth-trend-of-advanced-driver-assistance-systems/, Feb. 2022.

10. Jiménez, F., Naranjo, J.E., Anaya, J.J., García, F., Ponz, A., and Armingol, J.M., “Advanced Driver Assistance System for Road Environments to Improve Safety and Efficiency,” *Transportation Research Procedia* 14:2245–2254, 2016.

11. Asher, Z.D., Tunnell, J.A., Baker, D.A., Fitzgerald, R.J., Banaei-Kashani, F., Pasricha, S., and Bradley, T.H., “Enabling Prediction for Optimal Fuel Economy Vehicle Control,” SAE Technical Paper, 2018.

12. Motallebiaraghi, F., Yao, K., Rabinowitz, A., Hoehne, C., Garikapati, V., Holden, J., Wood, E., Chen, S., Asher, Z., and Bradley, T., “Mobility energy productivity evaluation of prediction-based vehicle powertrain control combined with optimal traffic management,” 2022-01-0141, SAE Technical Paper, 2022.

13. Kadav, P. and Asher, Z.D., “Improving the Range of Electric Vehicles,” *2019 Electric*

Vehicles International Conference (EV), 1–5, 2019.

14. Rabinowitz, A., Araghi, F.M., Gaikwad, T., Asher, Z.D., and Bradley, T.H., "Development and Evaluation of Velocity Predictive Optimal Energy Management Strategies in Intelligent and Connected Hybrid Electric Vehicles," *Energies* 14(18):5713, 2021.

15. Mahmoud, Y.H., Brown, N.E., Motallebiaraghi, F., Koelling, M., Meyer, R., Asher, Z.D., Dontchev, A., and Kolmanovsky, I., "Autonomous Eco-Driving with Traffic Light and Lead Vehicle Constraints: An Application of Best Constrained Interpolation," *IFAC-PapersOnLine* 54(10):45–50, 2021.

16. How do weather events impact roads?, https://ops.fhwa.dot.gov/weather/q1_roadimpact.htm, Feb. 2022.

17. Gern, A., Moebus, R., and Franke, U., "Vision-based lane recognition under adverse weather conditions using optical flow," *Intelligent Vehicle Symposium, 2002. IEEE*, 652–657 vol.2, 2002.

18. Brandon, S., "SENSOR FUSION: A COMPARISON OF CAPABILITIES OF HUMAN HIGHLY AUTOMATED."

19. Lei, Y., Emaru, T., Ravankar, A.A., Kobayashi, Y., and Wang, S., Semantic Image Segmentation on Snow Driving Scenarios, *2020 IEEE International Conference on Mechatronics and Automation (ICMA)*, 2020, doi:10.1109/icma49215.2020.9233538.

20. Rawashdeh, N.A., Bos, J.P., and Abu-Alrub, N.J., “Drivable path detection using CNN sensor fusion for autonomous driving in the snow,” *Autonomous Systems: Sensors, Processing, and Security for Vehicles and Infrastructure 2021*, SPIE: 36–45, 2021.

21. Rjoub, G., Wahab, O.A., Bentahar, J., and Bataineh, A.S., “Improving Autonomous Vehicles Safety in Snow Weather Using Federated YOLO CNN Learning,” *Mobile Web and Intelligent Information Systems*, Springer International Publishing: 121–134, 2021.

22. Shetty, S.H., Shetty, S., Singh, C., and Rao, A., Supervised machine learning: Algorithms and applications, *Fundamentals and Methods of Machine and Deep Learning* 1–16, 2022, doi:10.1002/9781119821908.ch1.

23. Chourasiya, S. and Jain, S., A Study Review On Supervised Machine Learning Algorithms, *International Journal of Computer Science and Engineering* 6(8):16–20, 2019, doi:10.14445/23488387/ijcse-v6i8p104.

24. Osisanwo, F.Y., Akinsola, J.E.T., Awodele, O., Hinmikaiye, J.O., Olakanmi, O., and Akinjobi, J., “Supervised machine learning algorithms: classification and comparison,” *International Journal of Computer Trends and Technology (IJCTT)* 48(3):128–138, 2017.

25. Seif, G., “Deep Learning vs Classical Machine Learning,” https://towardsdatascience.com/deep-learning-vs-classical-machine-learning-9a42c6d48aa, 2018.

26. Long, J., Shelhamer, E., and Darrell, T., “Fully convolutional networks for semantic

segmentation," *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440, 2015.

27. Albawi, S., Mohammed, T.A., and Al-Zawi, S., "Understanding of a convolutional neural network," *2017 International Conference on Engineering and Technology (ICET)*, 1–6, 2017.

28. Why Convolutional Neural Networks Are The Go-To Models In Deep Learning, https://analyticsindiamag.com/why-convolutional-neural-networks-are-the-go-to-models-in-deep-learning/, 2018.

29. Sankesara, H., "UNet," https://towardsdatascience.com/u-net-b229b32b4a71, 2019.

30. Ronneberger, O., Fischer, P., and Brox, T., U-Net: Convolutional Networks for Biomedical Image Segmentation, *arXiv [cs.CV]*, 2015.

31. Baldi, P. and Sadowski, P., "Understanding Dropout," https://proceedings.neurips.cc/paper/2013/file/71f6278d140af599e06ad9bf1ba03cb0-Paper.pdf, Feb. 2022.

32. Brownlee, J., "A Gentle Introduction to Dropout for Regularizing Deep Neural Networks," https://machinelearningmastery.com/dropout-for-regularizing-deep-neural-networks/, 2018.

33. Classification: Accuracy, https://developers.google.com/machine-learning/crash-course/classification/accuracy, Feb. 2022.

34. Duque-Arias, D., Velasco-Forero, S., Deschaud, J.-E., Goulette, F., Serna, A., Decencière, E., and Marcotegui, B., "On power jaccard losses for semantic segmentation," *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, SCITEPRESS - Science and Technology Publications, ISBN 9789897584886, 2021, doi:10.5220/0010304005610568.

35. Godoy, D., "Understanding binary cross-entropy / log loss: a visual explanation," https://towardsdatascience.com/understanding-binary-cross-entropy-log-loss-a-visual-explanation-a3ac6025181a, 2018.

36. Le, Q.V., Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., and Ng, A.Y., On optimization methods for deep learning, 2011.

37. Kingma, D.P. and Ba, J., Adam: A Method for Stochastic Optimization, *arXiv [cs.LG]*, 2014.

38. Yaqub, M., Jinchao, F., Zia, M.S., Arshid, K., Jia, K., Rehman, Z.U., and Mehmood, A., "State-of-the-Art CNN Optimizer for Brain Tumor Segmentation in Magnetic Resonance Images," *Brain Sci* 10(7), 2020, doi:10.3390/brainsci10070427.

39. People - All Victims, https://www-fars.nhtsa.dot.gov/People/PeopleAllVictims.aspx, Sep. 2022.

40. How Do Weather Events Impact Roads?, https://ops.fhwa.dot.gov/weather/q1_roadimpact.htm, Oct. 2022.

41. Neumeister, D.M., Pape, D.B., and Battelle Memorial Institute, “Automated vehicles and adverse weather: Final report,” FHWA-JPO-19-755, United States. Department of Transportation. Intelligent Transportation Systems Joint Program Office, 2019.

42. Walker, C.L., Boyce, B., Albrecht, C.P., and Siems-Anderson, A., “Will Weather Dampen Self-Driving Vehicles?,” *Bull. Am. Meteorol. Soc.* 101(11):E1914–E1923, 2020.

43. Traffic Safety Facts: Overview: 2008 Data, *PsycEXTRA Dataset*, 2009, doi:10.1037/e626152012-001.

44. Hearst Autos Research, “ADAS: Everything you need to know,” https://www.caranddriver.com/research/a31880412/adas/, 2020.

45. Motallebiaraghi, F., Yao, K., Rabinowitz, A., Holden, J., Wood, E., Chen, S., Asher, Z., Bradley, T., and Others, “Mobility Energy Productivity Evaluation of Prediction-based Vehicle Powertrain Control Combined with Optimal Traffic Management,” SAE Technical Paper, 2022.

46. Zang, S., Ding, M., Smith, D., Tyler, P., Rakotoarivelo, T., and Kaafar, M.A., “The Impact of Adverse Weather Conditions on Autonomous Vehicles: How Rain, Snow, Fog, and Hail Affect the Performance of a Self-Driving Car,” *IEEE Veh. Technol. Mag.* 14(2):103–111, 2019.

47. Rasshofer, R.H., Spies, M., and Spies, H., “Influences of weather phenomena on automotive laser radar systems,” *Adv. Radio Sci.* 9:49–60, 2011.

48. LiDAR, https://www.synopsys.com/glossary/what-is-lidar.html, Oct. 2022.

49. Kulemin, G.P., “Influence of propagation effects on a millimeter-wave radar operation,” *Radar Sensor Technology IV*, SPIE: 170–178, 1999.

50. Wallace, H.B., “Millimeter-wave propagation measurements at the Ballistic Research Laboratory,” *IEEE Trans. Geosci. Remote Sens.* 26:253–258, 1988.

51. Society of Automotive Engineers, “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles,” J3016_202104, 2021.

52. ©Alaska Department of Transportation, Facilities, P., and Reserved, A.R., “[No title],” https://roadweather.alaska.gov/gis, Oct. 2022.

53. (4) Classification of road conditions: From camera images and weather data, https://www.researchgate.net/publication/261019430_Classification_of_road_conditions_From_camera_images_and_weather_data, Oct. 2022.

54. Qian, Y., Almazan, E.J., and Elder, J.H., “Evaluating features and classifiers for road weather condition analysis,” *2016 IEEE International Conference on Image Processing (ICIP)*, 4403–4407, 2016.

55. Kadav, P., Sharma, S., Araghi, F.M., and Asher, Z.D., “Development of Computer Vision Models for Drivable Region Detection in Snow Occluded Lane Lines,” *Machine Learning and Optimization Techniques for Automotive Cyber-Physical Systems*, Springer Nature.

56. ZED 2 - AI Stereo Camera, https://www.stereolabs.com/zed-2/, May 2022.

57. Wasserman, S.E. and Monte, D.J., "A Relationship between Snow Accumulation and Snow Intensity as Determined from Visibility," *J. Appl. Meteorol. Climatol.* 11(2):385–388, 1972.

58. aum-toc.pdf.

59. Landolt, S.D., Lave, J.S., Jacobson, D., Gaydos, A., DiVito, S., and Porter, D., "The Impacts of Automation on Present Weather–Type Observing Capabilities across the Conterminous United States," *J. Appl. Meteorol. Climatol.* 58(12):2699–2715, 2019.

60. Martinaitis, S.M., Cocks, S.B., Qi, Y., Kaney, B.T., Zhang, J., and Howard, K., "Understanding Winter Precipitation Impacts on Automated Gauge Observations within a Real-Time System," *J. Hydrometeorol.* 16(6):2345–2363, 2015.

61. Snow, https://glossary.ametsoc.org/wiki/Snow, Oct. 2022.

62. Brooks, D., "Lightning appears to be the culprit as Concord airport weather station goes offline," *Concord Monitor*, 2019.

63. Duboue, P., "The Art of Feature Engineering: Essentials for Machine Learning," Cambridge University Press, ISBN 9781108571647, 2020.

64. Motallebiaraghi, F., Rabinowitz, A., and Holden, J., "High-Fidelity Modeling of Light-Duty Vehicle Emission and Fuel Economy Using Deep Neural Networks," *SAE Technical*, 2021.

65. Sen, P.C., Hajra, M., and Ghosh, M., "Supervised Classification Algorithms in Machine

Learning: A Survey and Review," *Emerging Technology in Modelling and Graphics*, Springer Singapore: 99–111, 2020.

66. Singh, A., Thakur, N., and Sharma, A., "A review of supervised machine learning algorithms," *2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom)*, 1310–1315, 2016.

67. Pedregosa, Varoquaux, and Gramfort, "Scikit-learn: Machine learning in Python," *Of Machine Learning …*.

68. Road & Runway Sensors - MARWIS - Mobile Advanced Road Weather Information Sensor, https://www.lufft.com/products/road-runway-sensors-292/marwis-umb-mobile-advanced-road-weather-information-sensor-2308/, Oct. 2022.