ENHANCEMENT OF AUTONOMOUS VEHICLES DURING EXTREME WEATHER

Pandian Anand
Inspirit AI
Keller, Texas, United States

ABSTRACT

Autonomous vehicles are self-driving vehicles that operate without human assistance. Self-driving vehicles, which include Teslas, BMWs, Waymos, Fords, and Mercedes, struggle to detect objects under extreme weather conditions, such as rain and fog, due to a lack of visibility, thereby significantly increasing the likelihood of collisions. Although autonomous vehicles have proven to be exceptional in nominal conditions, their poor performance in stormy, foggy, and snowy conditions currently limits their deployment on real roads today. This is important to know because AI can be used to improve the safety of vehicles using algorithms like deep learning via neural networks. In order to answer the question, it involved training and fine-tuning a neural network called the YOLOv8n model. We loaded a pretrained model, called yolov8n. This model was inaccurate at predicting cars when the vision isn't clear. However, from the DAWN dataset from Kaggle + YOLOV8, we fine-tuned the model and got it to predict vehicles when the vision wasn't clear. Initially, the model failed to predict other vehicles in extreme weather conditions, or often predicted other vehicles and objects with low confidence levels. However, after the model was fine-tuned, it was able to increase its prediction confidence level by 40%, also being able to predict 91% of the presented objects, way higher compared to its predicting 22% of its presented objects per image. However, during normal conditions, the fine-tuned model predicted other vehicles at a lower confidence level compared to the original Yolov8n model due to it being trained during extreme weather and not nominal conditions. It was concluded that the fine-tuning of YOLOv8n contributed to the feasibility of object detection during extreme weather conditions, with a 64% increase in objects detected with it being able to detect 94.87% of vehicles and pedestrians, but led to a 10% decrease in confidence level and <1% decrease in object detection rate under nominal conditions and failed to detect traffic lights, which the non-tuned model could do. This makes it necessary for self-driving providers to build a car with at least two different models: one non-tuned YOLOv8n model, operating during nominal conditions and during traffic lights, and the other being a tuned YOLOv8n model that operates during extreme weather.

INTRODUCTION

Autonomous vehicles continue to be built to reduce the occurrences of accidents caused by human error, while also promoting hands-free driving and comfort. However, as of right now, self-driving vehicles operate poorly under extreme weather conditions, far worse even than human driving. According to the article Nature, AVs are more commonly involved in accidents during rainy conditions, accounting for 11% of such accidents, compared to human-driven vehicles, which experience these conditions in only 5% of accidents ("Top Five Dangers of Self-Driving Cars"). According to Bizzieri Law, those accidents are caused by improper object detections - including a truck being classified as a boat, but more significantly, when there are 3 cars on the road and only 1 being detected, potentially leading to fatal accidents due to the inability to detect the vehicles around it (Bizzieri Law Offices). That is why research was conducted to resolve the self-driving cars' inability to accurately detect objects in rainy or foggy conditions. The models of self-driving vehicles are currently reliable in that they can detect objects correctly under nominal weather; however, they still fail in edge cases like dangerous weather that causes the vehicle to see poorly on the road. The model performs less flawlessly under extreme weather conditions like rain, sandstorms, snow, and other times when the vision isn't clear. This makes the action of improving vehicles under extreme weather conditions more necessary. The overall goal of this project is to make sure that it classifies the vehicles and objects confidently across conditions, and also raises the confidence level of the classifications. At the completion of the project, it is expected that there will be an increase in objects detected during extreme conditions like rain and fog, as well as an expected decrease in objects detected during nominal conditions, including when it is sunny and during daytime. This side effect can be mitigated by allowing an autonomous vehicle to contain the original YOLOv8n model as well, allowing it to be in use during nominal conditions.

Figure 1: The YOLOv8n model's failure to detect the rest of the cars because of poor visibility

BACKGROUND

Various sources have provided insights into the potential risks of self-driving vehicles. In Why weather is a problem for autonomous vehicle safety | Geotab, Bizzieri Law discusses how autonomous vehicles currently struggle to see during bad weather, leading to an increased likelihood of accidents and collisions (Geotab). The article revealed that frequent computer vision failure was the cause of autonomous vehicles struggling to detect objects and vehicles in extreme weather, leading to accidents and collisions, including hydroplaning. Moreover, from How automated vehicles run in bad weather with Finnish software, David Fusiek discusses how self-driving cars can't see and react quickly enough to objects in bad weather, showcasing the issue of computer vision failure (Fusiek). In order to address the issue of computer vision, the idea for a fine-tuning of a YOLOv8n model - a model that is the fastest version of the YOLOv8 object detection model and is capable of detecting multiple objects during extreme weather. - was sparked. The fine-tuning of the YOLOv8n model would also make it more accurate for the vehicle to detect objects during extreme weather, making the YOLOv8n model suitable for autonomous vehicles during dangerous weather.

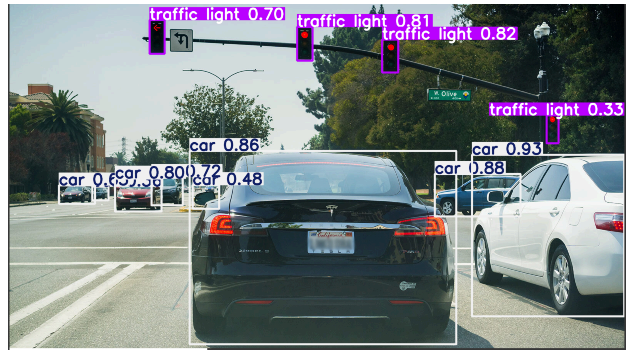

Figure 2: The capability of YOLOv8n in object detection

Figure 3: The YOLOv8n model's object detection failure in extreme weather

DATASET

The dataset that was used for the project was the DAWN dataset (Detection in Adverse Weather Nature), a dataset found on Kaggle. This is a vision dataset that shows all the images - stored as .jpg - along with the labels - stored as .txt - that allow for the training of the YOLOv8n model. This dataset reveals 1025 images along with the road vision in conditions like sand storm (287 images), fog (104 images), snow (203 images), and rain (200 images), along with extras like mist (82 images) and haze (114 images). The images found in the dataset consist of real-world images collected under various adverse weather conditions, along with diverse traffic environments (like rural and urban areas). While this dataset can be beneficial for the enhancement of autonomous vehicles during unusual weather, its bias towards extreme weather instead of nominal weather will create a tradeoff, improving safety in extreme conditions but lowering confidence during nominal weather. Given that the overall goal of the project is to improve self-driving cars during extreme weather, the lowered confidence level during nominal weather doesn't automatically favor other datasets that consist of images that include both nominal and extreme weather images. Furthermore, the preprocessing of the dataset was done through first importing the Ultralytics module and also loading YOLOv8n to detect sample images before the dataset was preprocessed. This paved the way for the option to view the uselessness of object detections by the non-fine-tuned model: only 30.77% of objects were accurately detected during extreme weather by autonomous vehicles. Later, it involved going into the cleaning and filtering by having all label files read and checked for specific classes, which include truck, person, bicycle, car, motorcycle, bus, and original images and labels were copied into the new-archive folder. Thirdly, the splitting of the data into training (80%) and testing (20%) was done by initially creating a dataframe with these two columns: an image path - the path to the image- and label - the dominant class in that image.

Figure 4: Image found from the dataset that enhanced YOLOv8n in object detection.

METHODOLOGY/MODELS

A pretrained ViT model (Vision Transformer) - a deep learning model that recognizes images and is used for image classification - was initially used in order to solve the research problem of autonomous vehicles struggling to perform efficiently and safely under extreme weather like snow, rain, and sandstorms. While image classification is used for detecting the whole image - a visual representation composed of object(s)- as a whole, the ViT model wasn't properly trained in detecting objects like cars, buses, bikes, pedestrians, and other objects found on the road. If a ViT model were to be implemented in the self-driving industry, that would lead to failures like collisions and accidents. Even if the ViT model was trained on the detection of cars, buses, pedestrians, and all the objects on the road, it wouldn't be sufficient due to its ability to only classify images and not objects. In other words, the ViT model is more suitable for image classification, which includes the following questions: Is this an image of a car or a bus? Are those vehicles classified as a truck or a bus? Knowing that the necessary model for autonomous vehicles is a model that can detect multiple objects in the same image, it led to the conclusion that the ViT model is insufficient and has no relevance to the self-driving industry, as it also displayed images used in everyday contexts and not specifically on the road. Those images include an office, chair, windows, vacuum, all features that an autonomous vehicle wouldn't need to detect. After this realization, a new method, which is safer and more tailored towards autonomous vehicles, has been found. The method that was used to solve the research problem included fine-tuning a YOLOv8n model - a deep learning model that is developed by Ultralytics and optimized for real-time object detection. The tuning of the model involved presenting the model with 1025 images, which were tailored to extreme weather, and 1025 labels, ensuring that the YOLOv8n model is trained to recognize objects - including cars, buses, trucks, pedestrians, motorcycles, bikes, and all objects found on the road during extreme weather - found on the road during extreme weather. Since the job of the necessary model was to be able to only detect vehicles and objects found on the road, YOLOv8n was more suitable for the research project compared to the pretrained ViT model. However, since YOLOv8n is more accurate and efficient than ViT at object detection, particularly with the ability to detect multiple objects and not just give one word for the classification of the image and looking at the entire image at once, YOLOv8n was chosen, particularly to be able to answer the following results: There's a car here, and a person there, and another car over there. Another reason YOLOv8n was used instead of ViT is due to its ability to handle blurry, low-quality images, which includes times during extreme weather like foggy and rainy conditions. Machine learning algorithms, like a training and testing dataset, were used in order to allow the model to learn on vehicles - truck, bicycle, car, motorcycle, bus, and pedestrians during different conditions - including sandstorm, rain, snow, and fog. Using YOLOv8n, to train the model, using StratifiedShuffleSplit, the training and testing data were able to be split with training being 80% and testing being 20%. Even though more epochs, like

50-100, could have led the fine-tuned model to greater accuracy, potentially from 95% to 99%+, 20 epochs were chosen for the training of the YOLOv8n model. 50-100 epochs wasn't chosen because, while an increase in accuracy can occur, more epochs increase the chance of overfitting - making the model memorize the training dataset instead of generalizing it to new images under extreme weather. Furthermore, on the T4 GPU, 20 epochs took 15 minutes for the model to tune. If 50-100 epochs were run, knowing that it takes 3 quarters of a minute per epoch, it would take 37.5 minutes to 75 minutes.. The training of 20 epochs involved having the training and test images originally split into separate folders, enabling ease of initial training. All object annotations were converted into the YOLOv8n format. It involved making all directories start with no labels or images. The model was tested by taking a sample of the images from the detected objects - 11 from the non-tuned model and 11 from the tuned model. This involved finding the accuracy of the non-tuned model (objects detected/objects presented) and the accuracy of the tuned model (objects detected/objects presented). An increase in confidence was also observed by taking the average of the old confidence levels and comparing it to the average of the new confidence levels.

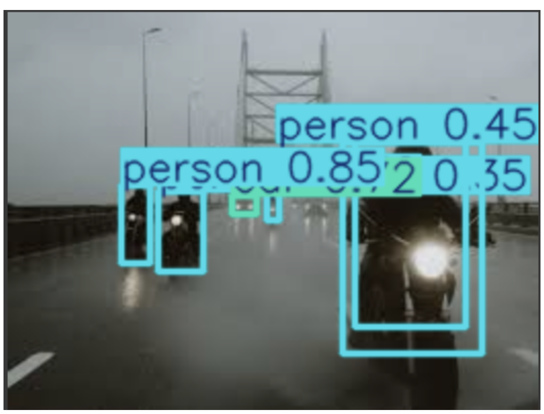

Figure 5: The capability of YOLOv8n (fine-tuned model) to detect multiple objects in one image.

RESULTS AND DISCUSSION

The research findings indicate that the non-tuned YOLOv8n model was accurate at object and vehicle detection in sunny conditions, with the ability to detect traffic lights. During sunny or nominal conditions, with it being almost 100% accurate in detecting or noticing all vehicles and pedestrians, its classification of the objects was with an accuracy of 92%. The other 8% was off due to misclassifications, which involves, for example, classifying a truck as a bus with an average confidence level of .7. Other examples of misclassification of images involve a car being classified as a truck with a confidence level of .32, a van being classified as a truck with a confidence level of .69, and a group of ground lights being detected as a car with a confidence level of .68, two distant cars being classified as a boat with a confidence level of .3. During accurate and inaccurate predictions, the fluctuation of the confidence levels was due to the brightness of day: brighter times would indicate a high confidence level of .9, while darker days would result in low confidence levels of .5. However, while the non-tuned model was able to detect almost all vehicles and pedestrians in nominal conditions, it could only detect 83% of traffic lights at an average confidence level of .6, indicating that the model was better off in vehicle and pedestrian classification. Furthermore, in extreme conditions, it failed to predict most of the vehicles, indicating that it struggles in object detection in abnormal conditions. During extreme weather conditions, which make the computer vision unclear, making it harder for the model to detect vehicles and pedestrians, the model could only detect 21% of the objects found on the road, with an average confidence level of .45. Another issue with the non-tuned YOLOv8n model was that it included misclassification of a few objects. For example, a car's tail light was predicted to be a traffic light at the .32 confidence level. More examples include when the distant trucks and cars were classified as boats, ranging from .3 to .37 confidence level. This indicates that the non-tuned YOLOv8n model isn't safe during extreme weather due to its failure to detect all objects and the low confidence levels. On the other hand, the tuned YOLOv8n model was able to predict most of the vehicles, people, and bikes under extreme conditions. Out of 23 test vehicles, the tuned model was able to predict 21 of the images with an average confidence of .85. This confidence level ranged from .28 to .96, with very low confidence levels being due to the big distance. Furthermore, instances of over-detections - more objects detected- were observed during the testing phase of the model. For example, one truck - detected with a confidence level of .76 - involved its detection of the tailight as a car with a confidence level of .37. The tuned model was also able to predict vehicles in normal conditions; however, the confidence levels of the predictions were lower than those of the original YOLOv8n models, and the tuned model couldn't predict traffic lights. This indicates that the original YOLOv8n model must be used in non-extreme weather conditions, along with times when there exists a traffic light, but the tuned YOLOv8n model can be used in extreme weather conditions and shouldn't be used to predict traffic lights. In other words, the fine-tuning of the YOLOv8n model significantly improves the performance of autonomous vehicles' accuracy of object detection by 64% and confidence by 40%, contributing to a safe autonomous vehicle mobility environment during adverse conditions. However, this results in the cost of lower confidence in nominal conditions and the failure to detect traffic lights, indicating that a dual-model approach is necessary for fully autonomous vehicles.

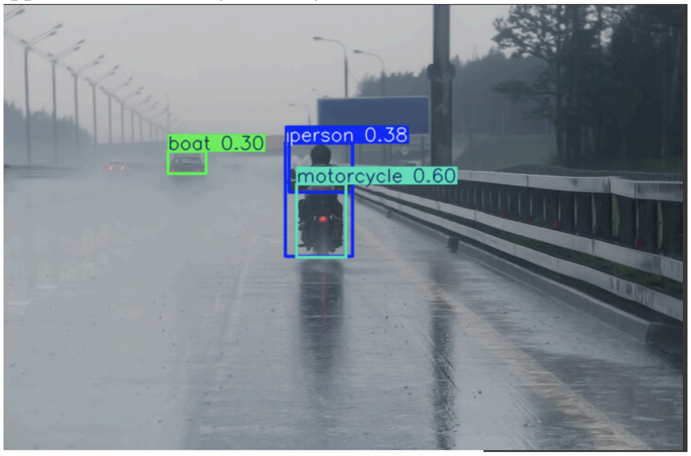

Figure 6: YOLOv8n's (non-tuned) misclassification of detecting a distant car as a boat

Figure 7: Example of the fine-tuned YOLOv8n's tendency to overdetect objects

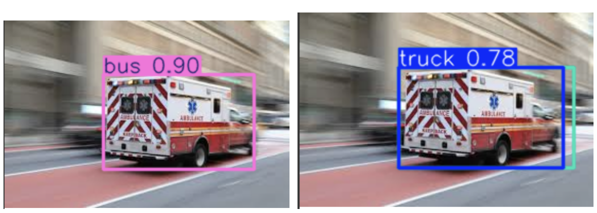

Figure 8: Example of the fine-tuned YOLOv8n (right image) detecting an object at a lower confidence compared to the non-tuned model (left image).

CONCLUSION

The goal of the research paper was to improve self-driving vehicles regarding issues like high collision rates during extreme weather, making the project focus on enhancing self-driving vehicles so that they can operate well under unfavorable, extreme weather. This goal is significant because with this enhancement, fatality rates will decrease, crash rates will decrease as poor operations with autonomous vehicles cause vehicular collisions, and insurance rates won't be so high, leading customers to be satisfied with the self-driving industry. The goal of improving autonomous vehicles during extreme weather was all done through training a YOLOv8n model, resulting in an increased accuracy of 64% - from ~31% to ~95% - objects detected during extreme weather, with also an increase in the correction of detected objects by 13.04% and the confidence level increase from .5 to .9. The fine-tuned model encountered an increase in an accuracy of 64% compared to the non-tuned model due to the amount of epochs being 20 and the vast number of images - 1025 - that were able to tune the model. The next steps worth taking with the research are for it to be able to detect traffic lights and their colors, along with deploying the distance it can see during foggy conditions, and calculating how much the vehicle should slow down during extreme weather, potentially using more epochs. This would allow for self-driving vehicles to be able to use the non-tuned YOLOv8n model for nominal conditions and the tuned YOLOv8n model for extreme weather, instead of the non-tuned model being present in order to detect traffic lights when the tuned YOLOv8n model is detecting vehicles and pedestrians. The pros of two or more perception models, each specialized for different conditions, include increased accuracy in object detection and more accurate localization of objects, and therefore enhanced overall safety. While the one con includes decreased efficiency and speed - caused by the switch from one model to another -the pros outweigh the cons because the overall goal is to ensure that autonomous vehicles are safe. Another step worth taking is improving another model during extreme weather.

Image	Objects Detected	Objects presented
#1	1	3
#2	3	7
#3	1	3
#4	0	2
#5	1	1
#6	3	10
#7	2	5
#8	1	7
#9	0	1
Total	12	39
Accuracy	30.77%	

Chart 1: The performance of the non-tuned YOLOv8n model in dangerous weather, achieving a 30.77% accuracy

Image	Objects Detected	Objects presented
#1	3	3
#2	7	7
#3	2	3
#4	1	2
#5	1	1
#6	10	10
#7	5	5
#8	7	7
#9	1	1
Total	37	39
Accuracy	94.87%	

Chart 2: The performance of the tuned YOLOv8n model in dangerous weather (achieved a 94.87% accuracy)

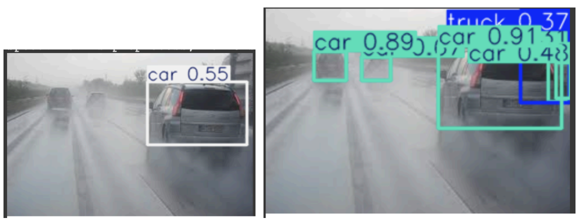

Figure 9: Image #1 showing an increased number of detected objects. The left image shows the performance of the non-fine-tuned model and the right image shows the performance of the fine-tuned model.

Figure 10: Image #2 showing an increased number of detected objects. The top image shows the performance of the non-fine-tuned model and the bottom image shows the performance of the fine-tuned model.

Figure 11: Image #3 showing an increased number of detected objects. The top image shows the performance of the non-fine-tuned model detecting ⅓ of the presented objects and the bottom image shows the performance of the fine-tuned model detecting ⅔ of the presented objects.

ACKNOWLEDGEMENTS

I, Pandian Anand, thank the Inspirit AI team for allowing me to work with Udgam Goyal, who I also thank for helping me through the project and suggesting the correct topic of autonomous vehicles.

REFERENCES

https://www.nature.com/articles/s41467-024-48526-4

Autonomous Vehicle Accidents in Chicago: Weather-Related Liability Concerns | Bizzieri Law

https://arxiv.org/abs/2104.13395

https://www.eib.org/en/stories/finnish-automated-vehicles-weather

Works Cited

Bizzieri Law Offices. “Understanding How Weather Conditions Affect Autonomous Vehicle Safety.” Bizzieri Law, https://bizzierilaw.com/2024/11/understanding-how-weather-conditions-affect-autonomous-vehicle-safety/.

Fusiek, David. “European Investment Bank.” How automated vehicles run in bad weather with Finnish software, 2022, https://www.eib.org/en/stories/finnish-automated-vehicles-weather. Accessed 25 July 2025.

Geotab. “Why weather is a problem for autonomous vehicle safety | Geotab.” Geotab, https://www.geotab.com/blog/autonomous-vehicles-safety-2/. Accessed 25 July 2025.

“Top Five Dangers of Self-Driving Cars.” Adam S. Kutner, Injury Attorneys, https://www.askadamskutner.com/auto-accident/self-driving-cars/. Accessed 25 July 2025