Assessing the Impact of Fog on Autonomous Vehicle Perception Systems

Alexandre Ray da Silva
Department of Computer Engineering and Digital Systems
University of Sao Paulo, Polytechnic School
Sao Paulo, Brazil
alexandre.ray.silva@usp.br

João Batista Camargo Junior
Department of Computer Engineering and Digital Systems
University of Sao Paulo, Polytechnic School
Sao Paulo, Brazil
joaocamargo@usp.br

***Abstract*—This study addresses the impact of adverse weather, specifically fog, on the perception systems of autonomous vehicles, which are critical for detecting and responding to traffic scenarios. Using over 10,000 images, an object recognition model was developed with Roboflow and YOLOv8, while fog disturbances were generated with GANs. The research simulates various traffic scenarios, comparing system performance under clear and foggy conditions. Results show that training models with a wider range of conditions enhances accuracy, highlighting the importance of diverse training for safe autonomous vehicle operation. This work offers insights for improving perception systems in autonomous vehicles.**

***Index Terms*—Autonomous Vehicles; Perception System; Adverse Weather Conditions; Object Recognition; Generative Adversarial Networks (GAN)**

I. Introduction

In recent years, autonomous vehicles have emerged as a transformative technology with the potential to revolutionize the way we move and transport. With advancements in artificial intelligence, sensor technology, and connectivity, autonomous vehicles promise to increase road safety, reduce congestion, and enhance mobility for everyone [1, 13]. However, to fully realize the potential benefits of autonomous vehicles, ensuring their safety under various environmental conditions becomes paramount [7, 8, 14].

A critical aspect that requires meticulous attention is the detection and adaptation to adverse weather conditions, such as fog. Foggy weather presents significant challenges for autonomous vehicles, as it dramatically affects visibility and introduces uncertainties that can compromise the safe operation of these vehicles [9, 11, 12]. Therefore, establishing robust and reliable methods to detect and respond to fog conditions is crucial to ensure the safety of autonomous vehicles and the passengers, pedestrians, and infrastructure around them [2].

The adoption of autonomous vehicles has immense potential to mitigate traffic accidents caused by human errors, improve transportation efficiency, and reduce environmental impact. However, widespread acceptance and integration of autonomous vehicles into our daily lives require ensuring their ability to navigate safely through various weather conditions. Fog, being a common occurrence in many regions, represents a significant challenge for the reliable operation of autonomous vehicles. [1, 6]

II. Related Work

The impact of adverse weather conditions on autonomous vehicle (AV) perception systems has been widely studied, with a particular focus on object detection challenges. Prior research highlights the detrimental effects of fog on visibility and AV performance, driving the need for adaptive solutions. For example, Appiah and Mensah (2024) examine object detection challenges in poor weather, presenting approaches for enhanced accuracy in low-visibility conditions [2]. Similarly, Liu et al. (2023) introduced GCANet, a model leveraging feature fusion to improve detection accuracy under fog, suggesting that specialized network architectures can better handle adverse weather disturbances [11].

III. Methodology

The study aims to demonstrate how such adverse conditions can negatively affect the accuracy of object detection, which is essential for the safe operation of autonomous vehicles.

A. Data Preparation

BDD100K, developed by the Berkeley DeepDrive Center at UC Berkeley, is one of the most comprehensive datasets available for advancing research in autonomous driving. It includes 100,000 videos recorded under diverse geographical, climatic, and temporal conditions, providing a rich variety of real-world driving scenarios [15]. The dataset features detailed annotations for tasks such as object detection, semantic segmentation, and traffic condition assessment, covering essential categories like vehicles, pedestrians, traffic signs, and traffic lights, along with driving conditions like road type, weather, and time of day [15].

To manage and process this data, Roboflow was chosen for its user-friendly interface and powerful pre-processing tools, despite its limitation of processing up to 10,000 images for free [4]. Given this constraint, a reduced version of BDD100K was used, consisting of 9,900 images, with 7,892 for training and 2,008 for validation [3]. These images represent traffic scenarios with varying lighting and weather conditions and contain a total of 185,995 annotations, averaging 18.6 annotations per image.

The images in the dataset have a resolution of 1280x720 pixels, chosen to provide sufficient visual detail for accurate analysis while maintaining efficient storage and processing requirements. This resolution is widely recognized for balancing clarity and performance in computer vision applications.

B. Foggy Image Generation

This research involves generating foggy images to assess the performance of computer vision algorithms in adverse weather conditions. The Foggy-CycleGAN project was utilized, employing a machine learning approach to transform clear images into foggy ones. In Figure 1, a comparison is shown, with the original clear urban scene on the left and the foggy version on the right, illustrating the effectiveness of the CycleGAN algorithm in simulating foggy conditions.

Fig. 1. On the left, we have the original image and on the right, we have the image with the application of the Foggy-Cycle-GAN algorithm transforming it into a foggy image. Source: BDD100K Dataset.

Six different experiments were conducted by re-training the model with distinct combinations of training and validation sets, each representing different weather scenarios. The Clear/Clear experiment serves as the baseline, offering a reference point to evaluate the model's performance under altered weather conditions.

IV. Results and Discussion

The number of epochs was chosen as 10 and 100 in the experiments to observe the model's performance across both shorter and longer training durations, allowing for a comparison of how quickly the model converges and whether additional training improves or plateaus the results.

A. For 10 epochs

The results of different tests using the Foggy-CycleGAN method are shown respectively for IoU (Intersection over Union) of 0.5 and for 10 epochs with a batch size of 16. The average execution time of each experiment was 1 hour.

Exp. #	Car	Pedestrian	Traffic Light	Traffic Signs
Clear/Clear	77.5%	55.4%	53.5%	62.7%
Clear/Fog	23.1%	15.3%	0.29%	1.13%
Fog/Fog	61.5%	32.5%	34.7%	36%
Fog/Clear	38.9%	27.1%	0.66%	2.84%
Fog + Clear/Clear	75.9%	51.4%	49.9%	57.9%
Fog + Clear/Fog	56.7%	30.6%	29.5%	32.6%

TABLE I
Experiment results (by class) for 10 epochs

The Table I details the mAP-0.5 metrics for different object classes: Cars, Pedestrians, Traffic Lights, and Traffic Signs.

- **Cars:** Detected with high precision in most experiments, with Experiment Clear/Clear showing the best result (77.5%) and Experiment Fog/Clear the worst (38.9%).
- **Pedestrians:** Detection was moderately successful, with the best result in Experiment Clear/Clear (55.4%) and the worst in Experiment Clear/Fog (15.3%).
- **Traffic Lights:** Results vary significantly, from 0.29% in Experiment Clear/Fog to 53.5% in Experiment Clear/Clear, possibly indicating some inconsistency in the model or data quality.
- **Traffic Signs:** Show considerable variation but less drastic than traffic lights, with the best result in Experiment Clear/Clear (62.7%) and the worst in Experiment Clear/Fog (1.13%).

B. For 100 epochs

The results of different tests using the FoggyCycle GAN method are shown respectively for IoU (Intersection over Union) of 0.5 and for 100 epochs with a batch size of 16. The average execution time of each experiment was 13 hours.

Exp. #	Car	Pedestrian	Traffic Light	Traffic Signs
Clear/Clear	81%	60.6%	62.2%	67.4%
Clear/Fog	24.7%	18.5%	0.5%	1.73%
Fog/Fog	68.7%	39.1%	46%	45.1%
Fog/Clear	38.2%	33.8%	1.61%	4.78%
Fog + Clear/Clear	80%	58.3%	60.1%	65.1%
Fog + Clear/Fog	66.9%	38.9%	42.9%	43%

TABLE II
Experiment Results

The Table II details the mAP-0.5 metrics for different object classes: Cars, Pedestrians, Traffic Lights, and Traffic Signs.

- **Clear/Clear:** High performance across all categories; Cars had the highest accuracy (81%), indicating effective car detection under ideal conditions. Pedestrians (60.6%), Traffic Lights (62.2%), and Traffic Signs (67.4%) also performed well.
- **Clear/Fog:** Significant drop in accuracy for all categories in foggy conditions; Traffic Lights (0.5%) and Traffic Signs (1.73%) were especially low, highlighting the impact of fog on detecting smaller or detailed objects.
- **Fog/Fog:** Improvement in all categories under foggy conditions compared to Experiment Clear/Fog, but still moderate to low results; Cars (68.7%) had the highest

accuracy, followed by Traffic Lights (46%), Traffic Signs (45.1%), and Pedestrians (39.1%).
- **Fog/Clear:** Low accuracy across all categories when trained in fog and validated in clear weather; notably poor performance for Traffic Lights (1.61%) and Traffic Signs (4.78%), suggesting that adverse condition training does not well prepare the model for ideal conditions.
- **Fog + Clear/Clear:** High performance for Cars (80%) and good results for other categories when trained with mixed data (Clear + Fog) and validated in clear conditions; Pedestrians (58.3%), Traffic Lights (60.1%), and Traffic Signs (65.1%).
- **Fog + Clear/Fog:** Decrease in accuracy for all categories when validated in foggy conditions, using the same dataset as Experiment Fog + Clear/Clear; Cars (66.9%) had the highest accuracy, followed by Traffic Lights (42.9%) and Traffic Signs (43%). Pedestrians were at 38.9%, showing the challenges of detecting people in fog.

The data suggests that training with a combination of conditions (Clear + Fog) tends to provide better generalization, but foggy conditions consistently lower detection accuracy across all classes, with Traffic Lights and Traffic Signs being the most affected.

C. For 10 and 100 epochs

The results of different tests using the FoggyCycle GAN method are shown respectively for IoU (Intersection over Union) of 0.5 and for 10 epochs and 100 epochs with a batch size of 16 for both. The average execution time for each experiment was 1 hour and 13 hours, respectively.

Exp. #	Training Data	Validation Data	mAP-0.5	
			10 epochs	100 epochs
Clear/Clear	Clear	Clear	62.3%	67.8%
Clear/Fog	Clear	Fog	9.97%	11.4%
Fog/Fog	Fog	Fog	41.2%	49.7%
Fog/Clear	Fog	Clear	17.4%	19.6%
Fog + Clear/Clear	Fog + Clear	Clear	58.8%	65.9%
Fog + Clear/Fog	Fog + Clear	Fog	37.4%	47.9%

TABLE III
COMPARISON BETWEEN 10 EPOCHS AND 100 EPOCHS

Table III compares the performance of object detection models under different training and validation conditions, evaluated by the global mAP-0.5 metric after 10 and 100 epochs. The mAP-0.5 metric represents the average precision for an Intersection over Union (IoU) of 0.5. Analyzing the improvement of mAP-0.5 from 10 to 100 epochs for each experiment:

- **Clear/Clear:** There was an improvement of 5.5 percentage points, from 62.3% to 67.8%, indicating that the model benefits from a higher number of training epochs under consistent clear image conditions.
- **Clear/Fog:** The improvement was only 1.47 percentage points, from 9.97% to 11.4%, suggesting that training under clear conditions has limited transfer to validation under foggy conditions, even with more training epochs.
- **Fog/Fog:** There was an increase of 8.5 percentage points, from 41.2% to 49.7%, showing that training and validation under consistent foggy conditions benefit from more training epochs.
- **Fog/Clear:** The increment was 2.2 percentage points, from 17.4% to 19.6%. This modest increase suggests that training only in fog does not generalize well to clear conditions, even with more training.
- **Fog + Clear/Clear:** The mAP-0.5 increased by 7.1 percentage points, from 58.8% to 65.9%, indicating that a diverse training dataset improves generalization to clear conditions when the model is trained for longer.
- **Fog + Clear/Fog:** An improvement of 10.5 percentage points, from 37.4% to 47.9%, was seen here, suggesting that a diverse training dataset also improves detection under foggy conditions with more training epochs.

Increasing training epochs usually enhances model performance, but the benefit depends on training conditions. Exposure to diverse conditions during training helps models generalize better, improving their robustness with additional epochs.

Currently, the model performs adequately under certain weather scenarios but falls short for reliable deployment in autonomous vehicles operating across different weather conditions. Nonetheless, the data and methods applied in this study lay a strong foundation for future enhancements. Incorporating more diverse data and advanced modeling techniques could help improve the model's reliability across a broader spectrum of conditions.

Though the study has not fully achieved its intended goal, the results point to a viable path forward. The research data and methods provide critical insights necessary for model refinement, ultimately working toward ensuring safe and effective operation under diverse weather conditions.

To be considered safe for use in autonomous vehicles, a computer vision model generally needs to achieve a minimum mAP of 85-90% under representative testing conditions [5, 10]. Additionally, practical implementation must include redundancies, rigorous validations, and comply with specific regulatory standards to ensure the safety and reliability of the system.

V. CONCLUSION

Evaluating the Foggy-CycleGAN method over 10 and 100 epochs provides insights into performance under various conditions. While 10 epochs showed variability, with Cars achieving the highest accuracy and Traffic Lights showing inconsistency, extending to 100 epochs improved accuracy across all categories. This highlights the benefits of longer training, although foggy conditions remained challenging.

Training with diverse data, covering both clear and foggy conditions, is essential for robust performance. Models trained with mixed conditions generalized better and achieved higher accuracy in both scenarios, emphasizing the value of varied training data to enhance model robustness.

Despite improvements, the model's performance in foggy conditions is still suboptimal, highlighting the need for more data augmentation and methodological advances. For autonomous vehicle applications, the model should ideally achieve an mAP of 85-90% under representative conditions. Achieving this will require continued investment in diverse data, advanced modeling, and strict safety standards to ensure reliable real-world operation.

Acknowledgment

I would like to express my sincere gratitude to the Universidade de São Paulo for their invaluable support throughout the duration of this research. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

References

[1] Guarino Alfonso, Rocco Zaccagnino, Emiliano Del Gobbo, Divya Garikapati, and Sneha Sudhir Shetiya. Autonomous Vehicles: Evolution of Artificial Intelligence and the Current Industry Landscape. *Big Data and Cognitive Computing 2024, Vol. 8, Page 42*, 8(4):42, 4 2024. ISSN 2504-2289. doi: 10.3390/BDCC8040042. URL https://www.mdpi.com/2504-2289/8/4/42/htmhttps://www.mdpi.com/2504-2289/8/4/42.

[2] Emmanuel Owusu Appiah and Solomon Mensah. Object detection in adverse weather condition for autonomous vehicles. *Multimedia Tools and Applications*, 83(9):28235–28261, 3 2024. ISSN 15737721. doi: 10.1007/S11042-023-16453-Z/METRICS. URL https://link.springer.com/article/10.1007/s11042-023-16453-z.

[3] Pedro Azevedo. BDD100k Dataset. *Roboflow Universe*, 2023.

[4] B. Dwyer, J. Nelson, T. Hansen, and et al. Roboflow (Version 1.0), 2024.

[5] Mark Everingham, Luc Van Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 6 2010. ISSN 09205691. doi: 10.1007/S11263-009-0275-4/METRICS. URL https://link.springer.com/article/10.1007/s11263-009-0275-4.

[6] Caner Filiz and Caner Filiz. Can Autonomous Vehicles Prevent Traffic Accidents? *Accident Analysis and Prevention*, 6 2020. doi: 10.5772/INTECHOPEN.93020. URL https://www.intechopen.com/chapters/72643undefined/chapters/72643.

[7] André Hartwecker, Steffen Müller, and Christian Schyr. Safety of Use Analysis for Autonomous Driving Functions Under Laboratory Conditions. *Lecture Notes in Mechanical Engineering*, pages 1183–1192, 2022. ISSN 2195-4364. doi: 10.1007/978-3-031-07305-2{_}110. URL https://link.springer.com/chapter/10.1007/978-3-031-07305-2_110.

[8] Philip Koopman, Uma Ferrell, Frank Fratrik, and Michael Wagner. A Safety Standard Approach for Fully Autonomous Vehicles. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11699 LNCS:326–332, 2019. ISSN 1611-3349. doi: 10.1007/978-3-030-26250-1{_}26. URL https://link.springer.com/chapter/10.1007/978-3-030-26250-1_26.

[9] Jinlong Li, Runsheng Xu, Jin Ma, Qin Zou, Jiaqi Ma, and Hongkai Yu. Domain Adaptive Object Detection for Autonomous Driving under Foggy Weather. URL https://github.com/jinlong17/DA-Detect.

[10] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng Yang Fu, and Alexander C. Berg. SSD: Single shot multibox detector. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9905 LNCS:21–37, 2016. ISSN 16113349. doi: 10.1007/978-3-319-46448-0{_}2/FIGURES/5. URL https://link.springer.com/chapter/10.1007/978-3-319-46448-0_2.

[11] Zhaohui Liu, Shiji Zhao, and Xiao Wang. Research on Driving Obstacle Detection Technology in Foggy Weather Based on GCANet and Feature Fusion Training. *Sensors 2023, Vol. 23, Page 2822*, 23(5):2822, 3 2023. ISSN 1424-8220. doi: 10.3390/S23052822. URL https://www.mdpi.com/1424-8220/23/5/2822/htmhttps://www.mdpi.com/1424-8220/23/5/2822.

[12] Xianglin Meng, Yi Liu, Lili Fan, and Jingjing Fan. YOLOv5s-Fog: An Improved Model Based on YOLOv5s for Object Detection in Foggy Weather Scenarios. *Sensors 2023, Vol. 23, Page 5321*, 23(11):5321, 6 2023. ISSN 1424-8220. doi: 10.3390/S23115321. URL https://www.mdpi.com/1424-8220/23/11/5321/htmhttps://www.mdpi.com/1424-8220/23/11/5321.

[13] Yi Lu Murphey, Ilya Kolmanovsky, and Paul Watta. Advances, Opportunities and Challenges in AI-enabled Technologies for Autonomous and Connected Vehicles. *Lecture Notes in Intelligent Transportation and Infrastructure*, Part F1376:1–16, 2023. ISSN 2523-3459. doi: 10.1007/978-3-031-06780-8{_}1. URL https://link.springer.com/chapter/10.1007/978-3-031-06780-8_1.

[14] Noor Ul Ain Tahir, Zuping Zhang, Muhammad Asim, Junhong Chen, and Mohammed ELAffendi. Object Detection in Autonomous Vehicles under Adverse Weather: A Review of Traditional and Deep Learning Approaches. *Algorithms 2024, Vol. 17, Page 103*, 17(3):103, 2 2024. ISSN 1999-4893. doi: 10.3390/A17030103. URL https://www.mdpi.com/1999-4893/17/3/103/htmhttps://www.mdpi.com/1999-4893/17/3/103.

[15] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: A Diverse Driving Video Database with Scalable Annotation Tooling. 5 2018. URL http://arxiv.org/abs/1805.04687.