

A Comprehensive Analysis of Object Detectors in Adverse Weather Conditions

Vatsa S. Patel, Kunal Agrawal, Tam V. Nguyen

Department of Computer Science

University of Dayton

Dayton, Ohio, US

Email: {patelv20, agrawalk2, tamnguyen}@udayton.edu

Abstract— In this paper, we meticulously examine the robustness of computer vision object detection frameworks within the intricate realm of real-world traffic scenarios, with a particular emphasis on challenging adverse weather conditions. Conventional evaluation methods often prove inadequate in addressing the complexities inherent in dynamic traffic environments—an increasingly vital consideration as global advancements in autonomous vehicle technologies persist. Our investigation delves specifically into the nuanced performance of these algorithms amidst adverse weather conditions like fog, rain, snow, sun flare, and more, acknowledging the substantial impact of weather dynamics on their precision. Significantly, we seek to underscore that an object detection framework excelling in clear weather may encounter significant challenges in adverse conditions. Our study incorporates in-depth ablation studies on dual modality architectures, exploring a range of applications including traffic monitoring, vehicle tracking, and object tracking. The ultimate goal is to elevate the safety and efficiency of transportation systems, recognizing the pivotal role of robust computer vision systems in shaping the trajectory of future autonomous and intelligent transportation technologies.

Keywords—Robustness Evaluation, Object Detection, Adverse Weather Conditions, Computer Vision, Deep Learning

I. INTRODUCTION

In recent years, the evolution of computer vision algorithms has significantly transformed the landscape of traffic surveillance and control systems. These algorithms, designed to automate the detection, tracking, and identification of objects in traffic scenarios, hold immense potential for revolutionizing transportation systems. The incorporation of deep learning methodologies has been particularly instrumental in propelling this progress, resulting in substantial improvements in accuracy and overall performance.

Despite these strides, deploying robust and reliable computer vision systems in real-world traffic environments remains a formidable challenge. The inherent intricacies and uncertainties in such settings, marked by dynamic variations in lighting, weather conditions, traffic flow, and the presence of occlusions and obstacles, pose significant hurdles to the accuracy of object detection algorithms. Existing evaluation methods, while contributing to algorithmic development, often fall short of encapsulating the complexities and uncertainties of real-world traffic scenarios.



Figure 1. YOLOv5n [1] object detector fails to detect objects accurately due to foggy weather.

This challenge assumes critical significance as the demand for autonomous driving technology burgeons, poised to reshape contemporary travel and commuting norms. Addressing this pivotal issue necessitates the formulation of novel evaluation methods capable of accommodating the multifaceted challenges associated with real-world traffic scenarios. These challenges encompass variations in weather and lighting, occlusions, and other obstacles. This research project undertakes the imperative task of evaluating the robustness of computer vision algorithms in the context of traffic videos, aiming to delineate the strengths and weaknesses of existing models in confronting intricate real-world scenarios.

In this paper, our primary focus lies in the meticulous evaluation of the robustness exhibited by eight distinctive YOLO (You Only Look Once) [2] object detection models when confronted with adverse weather conditions. These challenging scenarios are meticulously crafted through the utilization of a carefully curated dataset comprising 9700 images sourced from three diverse repositories: KITTI [3], Udacity [4], and IDD [5]. This amalgamated dataset presents an extensive array of real-world scenarios, from which a comprehensive test set is randomly assembled. The adverse weather effects, ranging from rain, fog, shadow, sun flare, and snow to intricate combinations like fog and rain, fog and snow, sun flare, and shadow, are systematically applied to this dataset, aptly christened as the Urban Weather Diversity Dataset (UWDD). Figure 1 shows an exemplary image of the YOLOv5n [1] object detection method failing to detect the object accurately due to adverse fog conditions, mistakenly identifying a building as a bus and some objects as people. The nomenclature not only reflects the varied meteorological conditions introduced but also emphasizes the dataset's

intent—to encapsulate the diverse weather challenges prevalent in urban environments.

Our objective is to scrutinize the performance of these object detection models under adverse weather conditions, comparing their accuracy against normal conditions. This evaluation aims to unravel the efficacy of each model in adverse scenarios and provide insights into its comparative strengths and weaknesses. The findings of this research endeavor contribute to the ongoing development of more resilient and dependable computer vision systems tailored for application in traffic surveillance and control.

II. RELATED WORKS

Recently, significant research efforts have been dedicated to developing robust computer vision algorithms for traffic monitoring and management in images. However, existing evaluation methods often fall short of capturing the intricacies and uncertainties prevalent in real-world traffic environments. This limitation can lead to inaccurate or inconsistent algorithm performance, especially in challenging scenarios like low-light conditions, adverse weather, and occlusions.

A. Datasets and Object Detection Techniques

The landscape of computer vision datasets has played a crucial role in supporting the development and evaluation of algorithms, particularly in autonomous driving and related applications. Datasets such as BDD100K [6], Waymo Open Dataset [7], ApolloScape [8], Udacity Self-Driving Car Dataset [4], Cityscapes [9], KITTI [10], COCO [11], and PASCAL VOC [12] have become pivotal resources in this field. While each dataset brings unique strengths, including variations in sensors, data collection locations, and annotation approaches, they also have limitations and potential biases. Notably, many of these datasets exhibit constraints related to limited weather and lighting conditions or focus on specific scenarios. Nguyen *et al.* [13] addressed traffic issues using a simulator. To overcome these limitations, our dataset aims to provide a comprehensive video dataset, encompassing diverse scenarios and weather conditions. This approach addresses the need for a more versatile dataset for object detection tasks, a domain that gained prominence since the advent of neural networks for image classification by Krizhevsky *et al.* [14].

Various object detection methods have been proposed for computer vision tasks, including Faster R-CNN [15], YOLOv5 [1], and SSD [16]. While two-stage methods like Faster R-CNN [15], YOLOv5 [1], and SSD [16]. Faster R-CNN [17] and Libra R-CNN [18], CARAFE [19], CenterNet [19], MTDL [21], and ATSS [22] achieve state-of-the-art performance, they are relatively slow due to their two-stage pipeline. On the other hand, one-stage methods like YOLOv5 [6] and SSD [16] are faster but may compromise accuracy for speed, especially in small object detection scenarios. However, a common challenge among all these methods is their susceptibility to adverse weather conditions, which can




Figure 2: Example of snow effect rendering on an input frame using our weather synthesizer.

lead to degraded performance and even failure in object detection. In our research work, we only focus on Yolo object detection models.

To address this limitation, our research will fine-tune and compare the results of image-adaptive object detection methods capable of adapting to changes in weather conditions. Evaluation metrics such as precision, recall, mean average precision (mAP), and IoU (Intersection over Union) thresholds will be employed to comprehensively assess the accuracy of object detection across various categories. This approach aims to provide a more nuanced understanding of algorithm performance under diverse and challenging conditions, contributing valuable insights to the ongoing enhancement of computer vision systems for real-world traffic scenarios.

B. Evaluation Studies on Deteriorating Weather

Numerous studies have demonstrated the substantial impact of adverse weather conditions, such as rain, fog, and snow, on various computer vision algorithms, specifically object detection and tracking. For instance, Michaelis *et al.* [22] conducted a robustness assessment of six object detection models using a dataset of traffic images captured in diverse weather conditions. The findings revealed a significant decrease in accuracy for all models under adverse weather conditions. Similarly, Mirza *et al.* [23] investigated the robustness of multiple object detectors in varying weather conditions, highlighting a notable reduction in performance during heavy rain and snow.

In addressing the impact of adverse weather on object detection models, various preprocessing techniques have been proposed. Classical dehazing or image enhancement methods, as suggested in [24–28], offer a straightforward approach to improve the models' robustness to rain. Halder *et al.* [29] introduced a physics-based rendering approach, generating synthetic rain images for improved object detection in adverse weather conditions. Additionally, benchmark datasets like CADP by Shah *et al.* [30] provide a standardized platform for evaluating object detection algorithms under degrading weather conditions.

To enhance the training and evaluation of object detection models in adverse weather conditions, Von Bernuth *et al.* [31] and Liu *et al.* [26] proposed a simulation approach, generating synthetic snow and fog images. The use of multi-modal sensors and fusion techniques, combining visual and non-visual data sources such as thermal and radar imaging, has also been explored (Dong *et al.* [32]) to improve object detection and tracking in inclement weather.




Figure 3: The figure illustrates the UWDD dataset, depicting various weather effects. From (a) to (j): input image with no effect, ground-truth, rain, shadow, snow, fog, sun flare, fog and rain, fog and snow, and sun flare and shadow, respectively.

Despite these approaches, existing studies often overlook the limited variations and dataset sizes [6 – 10] present in most datasets. Our benchmark, however, seeks to address this limitation by enabling the evaluation of a larger dataset with diverse situations and weather conditions. This approach aims to thoroughly assess the true robustness and potential of object detection techniques in adverse weather scenarios.

III. PROPOSED WORK

In this section, we introduce the Urban Weather Diversity Dataset (UWDD) and delve into the intricacies of its creation. We aim to provide a comprehensive overview of the methodology employed in generating this dataset, shedding light on the techniques utilized to craft a collection comprising 9,700 images.

A. Urban Weather Diversity Dataset (UWDD) Collection

The cornerstone of our research lies in the creation of the Urban Weather Diversity Dataset (UWDD). Building this dataset was a significant endeavor, given the challenge of assembling a comprehensive collection of dashcam videos depicting real traffic scenes in diverse weather conditions. To address this, we strategically amalgamated three distinct datasets—Kitti, Udacity, and IDD.

- Kitti Dataset [3]: Originating from the Karlsruhe Institute of Technology and Toyota Technological Institute in Chicago, USA, the Kitti dataset is a reservoir of autonomous driving images. Out of its total of 12,919 images, we randomly selected 3,465 for our dataset.
- Udacity Dataset [4]: Comprising a whopping 404,916 video frames for training, the Udacity dataset offered a rich resource. From this extensive pool, we handpicked 3,353 images to contribute to our dataset.
- IDD Dataset [5]: Hailing from Hyderabad, India, the Indian Driving Dataset (IDD) boasts a substantial collection of 10,003,182 images. Our dataset drew from this source, selecting 3,182 images to add a global diversity dimension.

The cornerstone of our research lies in the creation of the Urban Weather Diversity Dataset (UWDD). Building this dataset was a significant endeavor, given the challenge of assembling a comprehensive collection of dashcam videos depicting real traffic scenes in diverse weather conditions. To address this, we strategically consider three available datasets, namely, Kitti, Udacity, and IDD.

1) Single Effects

In the exploration of Single Effects, we meticulously delved into five distinct weather phenomena: rain, fog, snow, sun flare, and shadow. Each of these conditions was crafted with precision to simulate the challenges posed by adverse weather scenarios. To achieve this, we employed sophisticated augmentation techniques, leveraging the *imgaug* library [33] for the introduction of fog, snow, and rain effects. Additionally, we utilized *albumentations* [34] to authentically replicate sun flare and shadow effects, ensuring a realistic portrayal of each weather type. Figure 2 explains the way we rendered image augmentation techniques on the dataset images. This comprehensive approach to single effects was instrumental in diversifying our dataset, capturing the essence of individual weather challenges with accuracy.

2) Additive Effects

The Additive Effects category amalgamates various weather conditions, offering a holistic perspective on challenges encountered in real-world scenarios. These combinations, including fog and rain, fog and snow, and sun flare and shadow, are strategically derived from the Single Effects category. The rationale behind these combinations is to present a nuanced spectrum of possible weather conditions encountered in diverse traffic scenarios. By seamlessly merging these distinct weather types, we aim to create a dataset authentically mirroring the complexity of adverse conditions faced by object detection models in real-world settings.

Figure 3 illustrates an exemplary image with eight different effects. Through simulated adverse weather scenarios, we seek to unravel the intricacies of model robustness under varying and challenging conditions, contributing valuable insights to the advancement of computer vision systems in real-world traffic environments.

Table 1. Results of object detection models on weather effects.

	Yolov5n	Yolov5x	Yolov6l6	Yolov6n6	Yolov7	Yolov7-e6e	Yolov8n	Yolov8x
Fog	0.380	0.611	0.612	0.384	0.596	0.630	0.418	0.601
Rain	0.6118	0.600	0.598	0.416	0.587	0.611	0.425	0.649
Shadow	0.897	0.897	0.640	0.520	0.610	0.643	0.594	0.598
Sun flare	0.682	0.682	0.557	0.402	0.530	0.561	0.453	0.534
Snow	0.650	0.591	0.592	0.418	0.581	0.601	0.472	0.591
Fog and Rain	0.258	0.560	0.548	0.293	0.548	0.577	0.291	0.577
Fog and Snow	0.230	0.512	0.510	0.251	0.509	0.531	0.266	0.520
Shadow and Sun flare	0.640	0.640	0.547	0.383	0.525	0.558	0.434	0.516

B. Robustness Evaluation of Deep Learning Models

In assessing the robustness of computer vision algorithms using the Urban Weather Diversity Dataset (UWDD), our approach involves a systematic series of steps. Initially, we preprocess the images by resizing and normalizing them to the input size required by pre-trained models. In particular, we consider YOLO (You Only Look Once) [2] family object detection methods for benchmarking. The reasons are three-fold. First, YOLO is known for its real-time object detection capabilities, making it suitable for applications such as traffic monitoring where quick and efficient analysis is essential. Second, YOLO models, especially the later versions like YOLOv4 [35] and YOLOv5 [1], have demonstrated competitive accuracy and precision in object detection tasks. These models are trained on large datasets and have advanced architectures to handle various object classes. Last but not least, YOLO models often come with pre-trained weights on large datasets. This can be advantageous as pre-training helps the model to learn general features and patterns, potentially improving its performance on specific tasks like traffic object detection.

Utilizing the YOLO pre-trained models, we focus on seven specific classes—bicycle, bus, car, motorcycle, person, train, and truck—common throughout the dataset. Testing the dataset with the pre-trained model of YOLOv5s [1] without augmented images establishes a ground truth for subsequent comparisons. Having established the ground truth, we proceed to evaluate the dataset's performance with augmentation using various deep learning object detection models, including YOLOv5n [1], YOLOv5x [1], YOLOv6l6 [36], YOLOv6n6 [36], YOLOv7 [37], YOLOv7-e6e [37], YOLOv8n [38], and YOLOv8x [38]. Our comparison methodology centers on assessing the accuracy of model predictions against ground truth labels, employing the Intersection over Union (IoU) metric to measure bounding box overlap. This involves calculating IoU while reading the object-detected file from the model and ground truth data from generated text files, as shown in Figure 4.

The comparison extends to evaluating the correspondence of classes between ground truth and model output data. Additionally, we scrutinize the alignment of output data and ground truth bounding boxes, leading to the computation of

true positive (TP), false positive (FP), and false negative (FN) counts for each of the seven specified classes. Iterating through testing files, we accumulate class-wise statistics and overall accuracy. Concluding this evaluation, we present detailed results encompassing class accuracy, precision, recall, and F1 score. The average accuracy is calculated to provide an overarching view of model performance. The results are then meticulously saved to a file, facilitating a comprehensive comparison across all eight deep-learning object detection models.

IV. DISCUSSION

This section conducts an in-depth analysis of model performance based on the Intersection over Union (IoU) scores detailed in Table 1. The IoU metric serves as a pivotal indicator, gauging the accuracy of object detection, with higher values indicating superior performance. The YOLO-based models under examination encompass Yolov5n [1], Yolov5x [1], Yolov7 [37], Yolov7-e6e [37], Yolov8n [38], Yolov8x [38], Yolov6l6 [36], and Yolov6n6 [36].

Notably, Yolov5x emerges as the top performer, consistently achieving competitive IoU scores across diverse weather conditions. Conversely, Yolov6n6 demonstrates less robust performance, establishing itself as the least effective method with comparatively lower scores. Although all models excel in Shadow conditions, challenges surface in additive scenarios involving Fog and Snow.

As shown in Figure 4, when scrutinizing single weather effects, the worst-performing method varies by condition. Yolov5n exhibits the lowest effectiveness in fog, while Yolov6n6 struggles the most in rainy conditions. Yolov6n6 faces challenges in sun flare scenarios, emerging as the least accurate, and records the least effective performance in snowy conditions.

Contrastingly, under Shadow and Sun Flare conditions, all models perform similarly to normal conditions, with Yolov5n and Yolov5x excelling and Yolov6n6 lagging. In conclusion, this assessment underscores the diverse capabilities of YOLO-based models across various weather conditions. Yolov5x consistently emerges as a top performer, showcasing notable excellence, while Yolov6n6 encounters hurdles, particularly in challenging conditions like fog, snow, and their combinations.




Figure 4. Visualization of the object detection results in different weather conditions. For optimal viewing, please zoom to 400%.

Despite commendable proficiency in Shadow conditions, challenges arise in adverse scenarios like Fog and Snow or Fog and Rain, where heightened occlusion impedes accurate object detection. These insights are pivotal for refining object detection algorithms, emphasizing the need for weather-aware models to enhance the resilience of autonomous systems in diverse and challenging environmental conditions.

V. CONCLUSION AND FUTURE WORK

In this paper, we extensively investigate the performance dynamics of YOLO-based object detection models in a diverse array of challenging weather conditions, offering crucial insights for advancing computer vision algorithms in real-world traffic scenarios. The core of our investigation revolves around the thorough evaluation of YOLOv5n, YOLOv5x, YOLOv7, YOLOv7-e6e, YOLOv8n, YOLOv8x, YOLOv6l6, and YOLOv6n6, with the robust Urban Weather Diversity Dataset (UWDD) as the cornerstone of our analysis.

YOLOv5x consistently stands out as the top performer, showcasing resilience and accuracy across various weather scenarios. On the contrary, YOLOv6n6 encounters challenges, particularly in adverse conditions like fog, snow, and their combinations. A notable observation is the commendable proficiency displayed by all models in Shadow conditions, where minimal occlusion preserves the majority of the original

image. However, these models face difficulties in scenarios of Fog and Snow or Fog and Rain, where heightened occlusion poses significant obstacles to precise object detection.

These findings underscore the crucial role of weather-aware models in fortifying the resilience of autonomous systems, especially given their operation in diverse and challenging environmental conditions. The necessity for adaptive algorithms capable of navigating the intricacies of real-world traffic scenarios becomes evident. By illuminating the performance variations under different weather conditions, this research significantly contributes to the ongoing refinement of object detection algorithms, laying the foundation for the development of more dependable and robust computer vision systems tailored for traffic surveillance and control.

In a future dominated by autonomous technologies, the insights derived from this research serve as a foundational cornerstone for advancements in the safety and efficiency of transportation systems. By bridging the gap between algorithmic capabilities and real-world challenges, this study represents a pivotal step in the continuous evolution of computer vision technologies. It propels us closer to the realization of intelligent, adaptive, and seamlessly integrated autonomous transportation systems capable of navigating the complexities of our ever-changing urban landscapes.

ACKNOWLEDGMENT

This research was supported by the National Science Foundation (NSF) under Grant 2025234. The first author was also supported by the Graduate Student Summer Fellowship at University of Dayton.

REFERENCES

- [1] YOLO Version 5, Github: <https://github.com/ultralytics/yolov5>. [Accessed: December 1, 2023]
- [2] Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi. "You only look once: Unified, real-time object detection." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 779-788. 2016.
- [3] Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2013). Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research, 32(11), 1231-1237.
- [4] Gonzalez, E, Virgo, M, Lovekesh, J, Jensen, J, Kirksey, C. Annotated Driving Dataset. <https://github.com/udacity/self-driving-car> [Accessed: December 1, 2023].
- [5] Varma, G., Subramanian, A., Namboodiri, A., Chandraker, M., & Jawahar, C. V. (2019, January). IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1743-1751). IEEE.
- [6] Yu, B., Yin, K., Christy, M., Du, Y., Liu, Z., & Darrell, T. (2020). BDD100K: A diverse driving video database with scalable annotation tooling. arXiv preprint arXiv:2005.05618.
- [7] Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, T., & Chrabaszcz, P. (2020). Scalability in perception for autonomous driving: Waymo open dataset. arXiv preprint arXiv:1912.04838.
- [8] Huang, X., Cheng, X., Yang, Y., Wan, J., & Prokhorov, D. (2018). The ApolloScape open dataset for autonomous driving and its application. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (pp. 1626-1635).
- [9] Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., & Schiele, B. (2016). The cityscapes dataset for semantic urban scene understanding. In Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp. 3213-3223.
- [10] Geiger, A., Lenz, P., & Urtasun, R. (2012). Are we ready for autonomous driving? KITTI vision benchmark suite. In Conference Computer Vision and Pattern Recognition Workshops
- [11] Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., ... & Zitnick, C. L. (2014). Microsoft COCO: Common objects in context. In European conference on computer vision (pp. 740-755). Springer, Cham.
- [12] Everingham, M., Van Gool, L., Williams, C. K., Winn, J., & Zisserman, A. (2010). The PASCAL Visual Object Classes challenge. International journal of computer vision, 88(2), 303-338.
- [13] Nguyen, Tam V., Thanh Ngoc-Dat Tran, Viet-Tham Huynh, Bao Truong, Minh-Quan Le, Mohit Kumavat, Vatsa S. Patel, Mai-Khiem Tran, and Minh-Triet Tran. "Data-Driven City Traffic Planning Simulation." In 2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct), pp. 859-864, 2022.
- [14] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Proceedings of the 25th International Conference on Neural Information Processing Systems (NIPS), Lake Tahoe, Nevada, USA, pp. 1097-1105.
- [15] Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. In Advances in neural information processing systems (pp. 91-99).
- [16] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. Y., & Berg, A. C. (2016). Ssd: Single shot multibox detector. In ECCV pp. 21-37, 2016.
- [17] Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. In Advances in neural information processing systems (pp. 91-99).
- [18] Z. Li, C. Peng, G. Yu, X. Zhang, Y. Deng, and J. Sun. "Libra R-CNN: Towards Balanced Learning for Object Detection." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019.
- [19] J. Ding, X. Xue, Z. Long, Y. Xia, and W. Lu. "Context-Aware Reassembling of Features for Saliency Detection." Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020.
- [20] K. Law and J. Deng. "Cornersnet: Detecting objects as paired keypoints." Proceedings of the European Conference on Computer Vision. 2018.
- [21] Patel, Vatsa S., Zhongliang Nie, Trung-Nghia Le, and Tam V. Nguyen. "Masked face analysis via multi-task deep learning." Journal of Imaging 7, no. 10 (2021): 204.
- [22] Michaelis, Claudio, et al. "KITTI-360: A dataset for immersive driving scenarios." IEEE Robotics and Automation Letters 4.2 (2019): 1336-1343.
- [23] Mirza, Mohsin, et al. "RoadDamageDetector: A large scale dataset for road damage detection." IEEE Transactions on Intelligent Transportation Systems (2021).
- [24] Guo, Qiang, et al. "Building extraction from VHR remote sensing images using multiscale spatial spectral features." ISPRS Journal of Photogrammetry and Remote Sensing 164 (2020): 177-194.
- [25] He, Kaiming, Jian Sun, and Xiaoou Tang. "Single image haze removal using dark channel prior." IEEE transactions on pattern analysis and machine intelligence (2010): 2341-2353.
- [26] Liu, Yichao, et al. "Deep learning for generic object detection: A survey." International Journal of Computer Vision (2019): 1-42.
- [27] Hang, Zhao, et al. "A survey on deep learning-based methods for semantic segmentation of 3D LiDAR point clouds." Sensors 20.23 (2020): 6621.
- [28] Qin, Junjie, et al. "A comprehensive survey of deep learning for image captioning." ACM Transactions on Multimedia Computing, Communications, and Applications (2020): 1-25.
- [29] Halder, S. S., Seshadrinathan, K., & Chakrabarti, A. (2019). Physics-based rendering for improving robustness to rain. In Proceedings of the IEEE International Conference on Computer Vision (pp. 9404-9413).
- [30] A. P. Shah, J. B. Lamare, T. Nguyen-Anh, and K. W. Ng, "CADP: A novel dataset for CCTV traffic camera based accident analysis," in 2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), 2018, pp. 1-6.
- [31] Von Bernuth, A., Volk, G., & Krauß, T. (2019). Simulating photo-realistic snow and fog on existing images for enhanced CNN training and evaluation. 2019 IEEE Intelligent Vehicles Symposium (IV), 2113-2118. doi: 10.1109/IVS.2019.8814168
- [32] Dong, Zhen, et al. "Deep learning for object detection: A comprehensive review." Neurocomputing 396 (2020): 242-256.
- [33] Imgaug [augments.weather](https://imgaug.readthedocs.io/en/latest/source/overview/weather.html), [online] <https://imgaug.readthedocs.io/en/latest/source/overview/weather.html>. [Accessed: December 1, 2023]
- [34] Augmentations in Albumentations [Online] https://albumentations.ai/docs/examples/example_weather_transforms/ [Accessed: December 1, 2023]
- [35] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934.
- [36] Li, C., Li, L., Jiang, H., Weng, K., Geng, Y., Li, L., ... & Wei, X. (2022). YOLOv6: A single-stage object detection framework for industrial applications. arXiv preprint arXiv:2209.02976. Github: <https://github.com/meituan/YOLOv6> [Accessed: December 1, 2023]
- [37] Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 7464-7475). Github: <https://github.com/WongKinYiu/yolov7> [Accessed: December 1, 2023]
- [38] YOLO Version 8, Github: <https://github.com/ultralytics/ultralytics>. [Accessed: December 1, 2023].