Analysis of the impact of rain on perception in automated vehicle applications

IEEE Open Journal of Vehicular Technology

Received 14 February 2025; accepted 9 March 2025. Date of publication 21 March 2025;
date of current version 22 April 2025. The review of this article was coordinated by Editor Subhadip Kumar.

Digital Object Identifier 10.1109/OJVT.2025.3553718

Analysis of the Impact of Rain on Perception in Automated Vehicle Applications

TIM BROPHY [1,2,3], DARRAGH MULLINS [1,2,3], ASHKAN PARSI [1,2,3], JONATHAN HORGAN [4], ENDA WARD [4], PATRICK DENNY [1,3,5] (Member, IEEE), CIARÁN EISING [1,3,6] (Senior Member, IEEE), BRIAN DEEGAN [1,2,3] (Member, IEEE), MARTIN GLAVIN [1,2,3] (Member, IEEE), AND EDWARD JONES [1,2,3]

[1]School of Engineering, University of Galway, H91 TK33 Galway, Ireland
[2]Ryan Institute, University of Galway, H91 TK33 Galway, Ireland
[3]Lero the Research Ireland, Centre for Software, V94 NYD3 Limerick, Ireland
[4]Valeo Vision Systems, H54 Y276 Tuam, Ireland
[5]Department of Computer Science and Information Systems, University of Limerick, V94 T9PX Limerick, Ireland
[6]Department of Electronic and Computer Engineering, University of Limerick, V94 T9PX Limerick, Ireland

CORRESPONDING AUTHOR: TIM BROPHY (e-mail: tim.brophy@universityofgalway.ie).

This work was supported in part by Taighde Éireann – Research Ireland under Grant 13/RC/2094 P2 and Grant 18/SP/5942, in part by the European Regional Development Fund through the Southern and Eastern Regional Operational Programme to Lero - the Research Ireland Centre for Software, and in part by Valeo Vision Systems.

ABSTRACT The reliable performance of object detection perception algorithms in automated vehicles under adverse conditions such as rain is critical for maintaining vulnerable road user safety. Visible-spectrum cameras provide a rich source of information and are cost-effective compared with other sensors; however, their performance can degrade under adverse environmental conditions. Despite the general consensus that the object detection performance in computer vision is adversely affected by rain, there is a relative lack of research investigating this relationship in detail. This study investigates the performance of object detection under rain conditions, focusing on algorithm performance and low-level object characteristics. Using the publicly available BDD100 k dataset, this study examines object detection performance across multiple deep-learning object detection architectures, analyzing error types and image characteristics under rain and no rain conditions. In addition, statistical methods were used to compare image-level metrics to determine statistical significance. The results reveal that rain is not detrimental to object detection performance, and in some cases, better performance is observed. For some models, medium-sized objects experience improved detection and classification under rain conditions, while large objects experience a slight decline in performance. The error analysis shows an increase in localization errors and a decrease in classification errors. The object-level analysis revealed statistically significant changes in the contrast-to-noise ratio, entropy, mean pixel value, pixel variance, hue, saturation, and value, with hue and saturation experiencing the most significant changes. This study highlights the need for more detailed weather labeling in datasets to fully understand the nuances of the relationship between rain and object detection.

INDEX TERMS Adverse weather, object detection, machine learning, automated vehicles.

I. INTRODUCTION

Autonomous vehicles and advanced driver assistance systems (ADAS) have the potential to improve road safety for both Vulnerable Road Users (VRUs) and other vehicles. Despite the large amount of research conducted in this area, the influence of rain on sensor availability and its subsequent effect on algorithm performance remains an under-researched topic. If autonomous vehicles are to become practical for everyday use, they must be able to operate in a reliable and safe manner under rain conditions. There is a prevailing consensus in the literature that rain negatively impacts sensor availability and subsequently affects any further processing of collected data. However, the evidence to support this consensus is limited, particularly in the context of cameras and object detection.

Object detection has emerged as a key technology enabling automated vehicles. Object detection refers to the ability to

detect and localize the objects of interest within an image. These perception algorithms are one of the significant components of the technology stack of an automated vehicle that keeps VRUs safe. To ensure that these algorithms perform reliably, both the sensor and associated algorithm must be fully characterized across all weather conditions. Much of the initial research on autonomous vehicles was undertaken under optimal weather conditions, with little consideration given to the influence of adverse weather conditions. Therefore, understanding the limitations of the current state-of-the-art algorithms in the context of rain is an important step in ensuring the robustness and safety of autonomous vehicles. This study explores the degree to which rain influences image characteristics and, hence, object detection in an automotive environment. The focus of this study is on the inherent effects of rain, where no mitigation measures are employed.

The contributions of this study are as follows: (1) a detailed analysis at the object/pixel level of the impact of rain on camera images; (2) evaluation of object detection performance in rain using a range of deep learning architectures; (3) analysis of the different types of object detection errors.

II. RELATED WORKS

The main purpose of any object detection algorithm is to extract useful information about a scene regarding the detection, classification, and localization of targets. In the context of autonomous vehicles, special consideration is given to VRUs, such as pedestrians and cyclists; other road users, such as cars and trucks; and road infrastructure, including traffic lights, lane markings, and road signs. Several aspects of the imaging chain are vulnerable to the effects of rain [1]. These changes include changes to the road users themselves in the form of physical appearance changes [2] and sociological changes in which pedestrians behave differently in the rain [3]. Falling raindrops interact with light in a scene. The interaction between light and raindrops reflects and refracts light, causing spatiotemporal changes in the optical path [4]. Further changes occur at the lens because of the accumulation of adherent raindrops on the surface of the lens [5]. Water accumulation on the lens surface significantly distorts the image and causes occlusion [6]. All changes to the imaging chain are encoded onto the image plane, resulting in a degradation in the quality of the data collected. The accumulation of these changes renders object detection in rain a problem worthy of further consideration.

A. THE EFFECT OF RAIN ON OBJECT DETECTION

There is consensus in the literature that rain negatively affects object detection performance [7], [8]. Although the degradation of data quality in rain is unavoidable, few studies have been completed that quantify the change in the performance of an object detection algorithm under rain conditions. Hnewa and Radha [9] examined the impact of rain on object detection performance. They found that in rain, the mean Average Precision (mAP) [10] of Yolov3 [11] decreased by 2%, and that of FasterRCNN [12] decreased by 5% when compared to dry conditions. Chia et al. [13] examined the impact of rain on feature extraction using a Harris Corner Detector [14] and SURF [15]. They found that the number of detected feature points decreased by 48% and 67%, respectively. They also noted that while there was a substantial reduction in the number of detected keypoints, this reduction had a negligible effect on image stitching. Muller et al. [16] investigated the effect of simulating optical properties on the modular transfer function (MTF) and its subsequent effect on the machine learning performance. Although the optical blur model reduced the MTF in an image, this had little effect on the machine learning performance. Changes in the optical path affect traditional image quality metrics, which are typically highly sensitive to changes in an image. However, traditional image quality metrics are poor predictors of machine learning performance and changes in image quality do not always correlate with changes in machine learning performance [17].

Several studies have been completed that investigated mitigation techniques to reduce the impact of rain on an image and its subsequent effect on object detection. Amongst the most common mitigation techniques are image deraining and synthetic rain. Image deraining refers to the removal of rain from an image during post-processing, and has been the subject of several recent review papers [18], [19]. Current state-of-the-art deraining methods exhibit excellent performance from a visual perspective. However, the lack of a quantitative benchmark makes evaluating deraining algorithms from a computer-vision perspective more challenging [20]. The benefit of using deraining algorithms within the context of object detection is unclear, with Li et al. [21] noting that almost all deraining algorithms deteriorate the performance of an object detection algorithm, including FasterRCNN and Yolov3.

Another popular mitigation technique is to augment the training data with synthetic rain to attempt to make the algorithm more robust to real rain. Adding synthetic rain to an image is the reverse process of image deraining; therefore, it faces many of the same issues. The accuracy of synthetic rain relies on the use of collected image pairs, one image collected under rain, and one image collected under no rain, to train the algorithm. Even in a static scene where the objects and background are constant, illumination changes prevent the capture of an ideal image pair. Hnewa et al. [9] noted that any differences between synthetic and real rain are likely to result in a domain mismatch between training and testing.

Despite much research, the benefits of using mitigation techniques are unclear within the context of object detection. Furthermore, there is some evidence that mitigation techniques, particularly deraining, appear to have a negative impact on object detection performance with the algorithms performing better when no mitigation techniques are attempted [9].

B. ADVERSE WEATHER DATASETS

One of the main issues with current publicly available automotive datasets is the lack of detailed weather labeling included with the metadata. Although many current

TABLE 1. BDD100 k Class Breakdown

					No Rain		
Class	Total	Rain	Fog	Snow	Overcast	Cloudy	Clear
Pedestrian	129262	6720	47	9060	16649	7422	39048
Car	1021857	53511	1283	59403	118553	64374	414926
Bicycle	10229	423	7	559	1237	558	3353
Sign	343777	18517	354	21012	39148	21170	141372
Traffic Light	265906	18325	300	19310	26155	13034	112863
Bus	16505	1073	12	1155	2263	1263	5735
Truck	42963	2349	30	2708	5578	3460	13615
Train	179	11	0	18	22	24	62
Motorbike	4269	157	4	98	507	355	1176
Rider	4138	244	5	241	879	420	2349

FIGURE 1. Sample images from BDD showing various weather conditions: 'clear,' 'rain,' 'snow,' and 'fog'.

datasets have been acquired under diverse conditions, such as NuScenes [22], Waymo Open dataset [23], and Apolloscape [24], the metadata included with the collected data often contain no information regarding the weather conditions or lighting conditions at the time of data collection. BDD100k [25] is one of the few datasets that include any form of adverse weather labels. The labels are coarse and lack detailed information, although they are still useful for subdividing the dataset for training and evaluating machine-learning algorithms.

The lack of detailed labels makes fully analyzing and understanding the impact of adverse weather more challenging. Adverse weather can affect image data in a wide variety of ways, including occlusions due to adherent raindrops on the lens, raindrops distorting the optical path between the lens and objects of interest in the scene, and lens flare due to direct sunlight. Many pedestrians in the scene also appear differently from a visual perspective as they may elect to use an umbrella or wear a hood. Many of the issues with adverse weather in terms of machine learning performance will be found in edge cases, where a set of accumulating conditions can lead to a system failure. Current datasets lack the fine details regarding weather conditions to analyze when these edge cases are likely to occur.

III. METHODOLOGY

A. DATASET

To evaluate the effect of rain on object detection in the context of autonomous vehicles, a suitable dataset needs to be selected. The dataset must include rain, have detailed weather labels, and be suitable for object detection tasks. The only dataset that meets these criteria is BDD100 k (Berkeley Deep Drive) [25]. The BDD dataset contains 100,000 images collected across a diverse range of geographical, environmental, and lighting conditions. Images were collected using dashboard cameras with a resolution of 720p. The dataset comprises 210,000 objects spread across 10 classes, with approximately 50% of the objects in the dataset occluded to some degree. BDD also contains 10,000 images that have been annotated for semantic segmentation tasks.

The dataset is annotated for six weather conditions namely, 'clear,' 'overcast,' 'snowy,' 'rainy,' 'cloudy,' and 'foggy.' Table 1 shows a detailed breakdown of the number of objects assigned to each weather category for each class. Each weather category is representative of all the environmental changes that occur under these conditions. For example, the rain label may represent changes in ambient light, cloud cover, etc. as well as the presence of rain itself. The dataset was collected across three times of day — 'day,' 'night,' and 'dawn/dusk' — and across multiple locations. Fig. 1 shows several sample images taken from the BDD dataset under different weather conditions.

In this study, a subsection of the BDD dataset was used to evaluate the performance of each object detection model. To simplify the analysis, only daytime images were considered, with the assumption that the camera's Image Signal Processor (ISP) would maintain a constant target luminance value through auto-exposure and gain settings, despite minor changes in ambient lighting. Low-light or nighttime images were not considered because of the inability of the ISP to maintain a constant target luminance. The removal of low-light images isolates the effects of rain from those of changing illumination. The images were further segmented into rain and no rain categories. The no rain category was made up of the 'clear,' 'overcast,' and 'cloudy' labels in the dataset. For evaluation, the following classes were chosen: 'car,' 'pedestrian,' 'bicycle,' 'traffic light,' 'bus,' 'truck,' 'motorbike,' and 'rider' (mapped as pedestrian). These classes were chosen because of their commonality with the COCO dataset on which the object detection models were originally trained. The models were evaluated on all the available daytime rain and no rain images in the BDD dataset, producing a dataset of 30,538 images, including 2,903 rain images and 27,365 no rain images. A class breakdown of the BDD subsection is presented in Table 2. This subsection includes semantic segmentation annotations for 1,905 no rain and 210 rain images.

The data-processing pipeline used to analyze the effect of rain on object detection is shown in Fig. 2. Initially, the BDD dataset is segmented into two categories: rain and no rain.

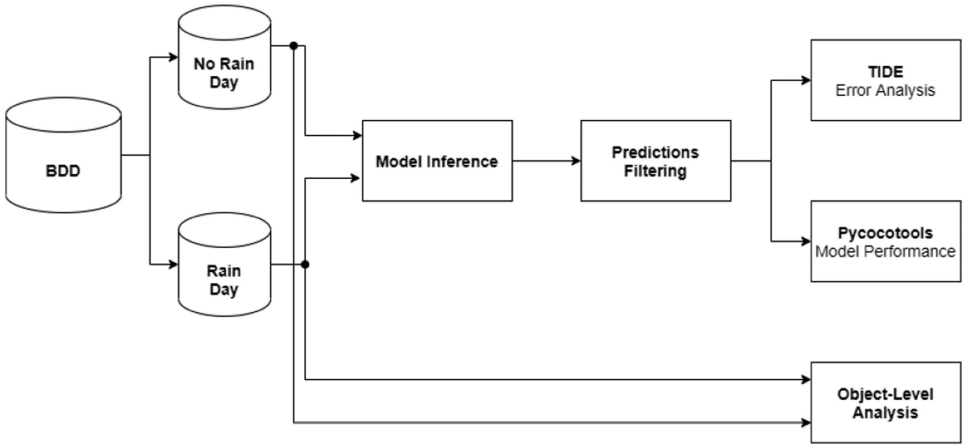

FIGURE 2. Data Processing Pipeline.

TABLE 2. BDD100 k Subsection Class Breakdown

			No Rain		
Class	**Total**	**Rainy**	**Overcast**	**Cloudy**	**Clear**
Pedestrian	45103	1716	15586	6942	20859
Car	347076	28986	102469	56366	159255
Bicycle	3504	97	1132	499	1776
Traffic Light	72314	8803	22211	11229	30071
Bus	11098	4834	1968	1124	3172
Truck	16356	297	4907	3098	8054
Train	85	8	18	22	37
Motorbike	2489	712	453	290	1034

TABLE 3. Object Detection Model Information

Models	**Year**	**Num. Parameters (millions)**	**Backbone**
DETR	2020	41.3	ResNet50
Faster RCNN	2015	49	ResNet101
RTMDet	2022	49	ResNet101
Yolov8x	2023	68.2	CSPNet
Yolvov8l	2023	43.7	CSPNet
Yolov8m	2023	25.9	CSPNet
Yolov8n	2023	3.2	CSPNet

Inference is performed for each model on the rain and no rain datasets. After the predictions are generated, several filtering steps are performed. First, predictions associated with COCO classes not labeled in the BDD dataset are removed. Then, all predictions with a confidence score below 0.25 are discarded. Finally, non-maximum suppression (NMS) is performed with a threshold of 0.75 to eliminate overlapping predictions. The filtered results from each model are analyzed using a combination of Pycocotools [26] and TIDE [27].

B. OBJECT DETECTION MODELS

In this study, several models were used to assess the impact of rain on object detection. The chosen models are representative of the full spectrum of model architecture capabilities from the literature. The selected models are representative of both single-stage and multi-stage classifiers and include multiple architectures and backbones. Table 3 summarizes these models and their primary characteristics. The models used in this study were: Yolov8 [28], FasterRCNN [12] with a ResNet101 backbone, DETR (DEtection TRansformer) [29], and RTMDet (Real-Time Models for object Detection) [30]. For the analysis, pretrained Yolov8 models from ultralytics [31] and pretrained models from MMDetction [32] were used. MMDetection provides several variations for each model. For reproducibility, the exact configurations for each MMDetection model used in this study are as follows:

- *FasterRCNN:* faster_rcnn_r101_fpn_mstrain_3x_coco
- *DETR:* detr_r50_8xb2-150e_coco
- *RTMDet:* rtmdet_l_8xb32-300e_coco

All the models were trained on the COCO dataset [26] and evaluated using the BDD100 k dataset. Note that data from the BDD100 k dataset is not used for transfer learning in this study.

The main benefit of the YOLO (You Only Look Once) architecture [33] is increased inference speed, because classification and localization are performed in a single pass at the expense of a slight decrease in overall performance. This architecture uses an anchor-free detection mechanism and introduces CSPNet [34] for efficient feature extraction. The architecture attempts to balance the speed and efficiency for object detection.

Yolov8 is available in several variations, each of which shares similar architectures but differs in their number of parameters. In this study, four variants of Yolov8 were considered:

- *Yolov8n (Nano):* The smallest Yolov8 model with 3.2 m parameters. This model has been optimized for embedded edge devices with limited resources.

- *Yolov8m (Medium):* A mid-sized Yolov8 model with 11 M parameters. The model offers a balance between speed and accuracy.
- *Yolov8l (Large):* A large model that has 46 million parameters. The model offers higher accuracy but longer inference times, and is suitable for tasks where high precision is necessary.
- *Yolov8x (Extra Large):* The largest Yolov8 model with 87 M parameters. The model is designed for tasks where accuracy is crucial; however, the increased number of parameters leads to longer inference times.

FasterRCNN is a two-stage object detection model based on RCNN [35]. FasterRCNN uses a region proposal network (RPN) to identify the coordinates of possible bounding box locations. Feature extraction is completed on the returned regions, with these features used as input to class-specific linear SVMs. FasterRCNN is known for its accuracy, particularly for detecting small objects. In this study, a version of FasterRCNN using a ResNet101 backbone was examined.

DETR uses a transformer architecture for object detection, employing self-attention mechanisms to model the relationships between objects. This model eliminates anchor boxes and uses end-to-end bipartite matching to assign object classes. The main benefit of this method is improved detection in complex scenes. Although it is a highly accurate model, the transformer architecture is computationally heavy and, therefore, has longer inference times. DETR is suitable for tasks that require a high degree of accuracy when operating in complex environments.

RTMDet is a CNN-based object detection model optimized for real-time performance. RTMDet uses lightweight convolutional networks and task-specific decoders to handle multiple objectives including object detection and semantic segmentation. To improve inference speed, the model uses an anchor-free approach and feature pyramid networks. The architecture is designed for real-time applications in embedded systems.

C. METRICS

1) OBJECT DETECTION PERFORMANCE

The most commonly used metric for evaluating machine learning performance is average precision (AP) [10]. AP is calculated from the precision-recall curve. The precision-recall curve is computed by varying the confidence threshold of a model and calculating the precision and recall values for each threshold. The area under the curve represents the AP metric. In the context of object detection models, the area of overlap between the bounding box of the prediction and that of the ground truth must also be considered when calculating true and false positives. The exact method for calculating the overlap is referred to as Intersection over Union (IoU).

$$IoU(A, B) = \frac{A \cap B}{A \cup B} \tag{1}$$

AP50, therefore, refers to the average precision with an IoU of at least 50%.

Another commonly used metric is AP50:95, which refers to the average AP value over a range of IoU thresholds. The IoU threshold is set at 5% intervals, starting at 50% and ending at 95%, with the AP value calculated at each threshold. Finally, the average of all the AP values is calculated. The main difference between AP50 and AP50:95 is that AP50:95 places greater emphasis on the accurate localization of the bounding box predicted by the model.

Average Recall (AR) is another metric commonly used to evaluate machine learning performance [36]. AR is evaluated across several IoU thresholds, typically between 50% and 100%, with the average value calculated. AR is generally evaluated for a fixed number of bounding box proposals. In this study, AR is evaluated on the top 100 proposals for each image, with the metric referred to as AR-100 in the results section below.

Another important aspect to consider when evaluating object detection performance is the size of the objects the model attempts to detect. In this study, several metrics are evaluated for objects of various sizes. The breakdown is based on the object area as follows:

$$small : area < 32 * 32 \ pixels$$

$$medium : 32 * 32 < area < 96 * 96 \ pixels$$

$$large : area > 96 * 96 \ pixels$$

All the precision and recall metrics are generated using Pycocotools [26].

To fully characterize model performance, it is essential to understand where and how errors occur. Depending on the application, some errors are more critical than others. For ADAS, false negatives are generally considered to be more dangerous than false positive errors because of safety implications. In this study, the errors are evaluated using TIDE (Toolbox for Identifying Object Detection Errors) [27]. The toolbox classifies each error from the model into one of six categories: class error, localization error, both class and localization error, duplicate, background false positive, and missed. The toolbox calculates the improvement in AP50:95 if each error type is corrected. The relative difference between the model output and corrected output, called difference in AP (dAP), is returned for each error type.

2) IMAGE-LEVEL METRICS

To gain further insight into the effects of rain on camera images, and in particular on the objects of interest in these images, several lower-level metrics are generated based on the pixel values and other characteristics of objects of interest in the scenes. The process of extracting the objects and their associated backgrounds is illustrated in Fig. 3. Background extraction is necessary, because some metrics require an assumed background for the foreground object. The delineation of objects is based on the Semantic Segmentation (SS) information provided with the dataset. The non-zero pixel values from each cropped region were used to analyze the impact

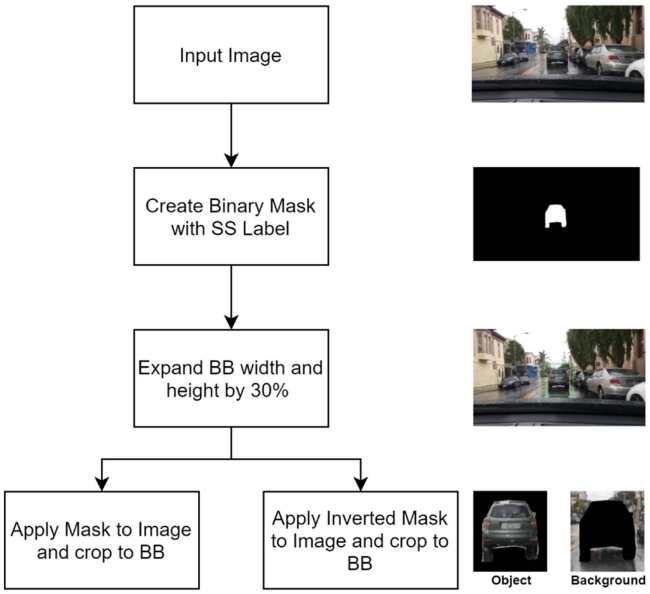

FIGURE 3. Object and Background Image Extraction Process.

of rain on individual objects and their immediate surrounding background.

Both the mean pixel values and pixel variance values in the RGB space are analyzed for each region. Each region is also converted to the HSV (Hue, Saturation, Value) color space [37]. Hue refers to the color shade or pure color and is directly linked to the visible portion of the electromagnetic spectrum. Saturation is a measure of the intensity of the color on the image plane, with a low saturation value indicating a color close to pure gray. Value refers to the relative lightness or darkness of a color. Colors with a high value reflect more light and are therefore closer to white than those with a low value. The HSV color space offers the possibility of extracting additional information that is not directly apparent in the RGB space.

As well as examining the basic pixel values the object-background contrast-to-noise ratio (CNR) is calculated. CNR is an indication of the degree to which an object stands out from the background and is calculated using the following formula:

$$CNR(O,B) = 20 * \log_{10}\left(\frac{S_o - S_b}{\sqrt{\sigma_o^2 + \sigma_b^2}}\right) \tag{2}$$

where S_o refers to the mean value of the object signal, S_b refers to the mean value of the background, σ_a^2 refers to the variance of the object signal, and σ_b^2 refers to the variance of the background signal.

In addition to CNR, Weber Contrast [38], which measures the contrast between objects and their immediate background, is also calculated. Weber contrast (also referred to as luminance contrast) differs from CNR in that pixel variance, which is an indication of noise in an image, is not included in the metric. Weber contrast is calculated using the following formula:

$$Weber\ Contrast = \frac{L_{object} - L_{background}}{L_{background}} \tag{3}$$

where L_{object} and $L_{background}$ are the average values of the luminance channel from the YUV color space of the object and background, respectively.

The Shannon Entropy [39] is also calculated for the object and background regions according to the following formula:

$$H(X) = -\sum_{i=0}^{X=255} P(Xi)\log_2(P(Xi)) \tag{4}$$

where $H(X)$ is the Shannon Entropy and $P(Xi)$ is the probability of a pixel having a given value over the range of 0 to 255 for a grayscale image. Shannon Entropy is a measure of the amount of information in an image, with the metric indicating the minimum number of bits per pixel required to represent the information in the image.

A statistics-based approach is used to compare and understand the results in terms of relative significance. Analysis of Variance (ANOVA) and Kruskal-Wallis tests are performed to compare the data collected from the rain and no rain conditions. ANOVA determines whether a statistically significant difference exists between the means of different groups. Kruskal-Wallis testing is very similar to ANOVA, with the major difference being that the Kruskal-Wallis test can be used for non-parametric data, and therefore, is useful for data that do not follow a normal distribution.

Each test returns two key values: the F-score (ANOVA) or H-score (Kruskal-Wallis) and the p-score. The F-score or H-score is a measure of the ratio of the variance between groups divided by the variance within groups. A higher value indicates a larger statistical difference between the groups. The p-score is a confidence value that indicates that a statistical difference does not exist between the groups. If the p-score is below a certain threshold, typically 5%, we can reject the null hypothesis and conclude with 95% confidence that a statistical difference exists between the groups. In the Results section of this study, the p-value of the Kruskal-Wallis test is referred to as the KW-p-score, with the AN-p-score referring to ANOVA.

IV. RESULTS

A. OBJECT DETECTION RESULTS

Table 4 summarizes object detection results, which indicate that rain does not impact detection performance in a systematic way; indeed, while certain object detection models exhibit slight performance degradation in rain, others experience small improvements, emphasizing the nuanced impact of rain on perception tasks.

At a high level, AP50 shows a very slight variation between the no rain and rain conditions across all the models. Faster R-CNN achieved the highest overall performance with a minor reduction in AP50 from 0.39 to 0.38 between no rain and rain conditions. Conversely, YOLOv8m exhibited a slight improvement in AP50 from 0.311 (no rain) to 0.313

TABLE 4. Object Detection Results for Rain and No Rain Conditions

		AP50:95	AP50	AP50:95 small	AP50:95 medium	AP50:95 large	AR-100	AR-100 small	AR-100 medium	AR-100 large
DETR	No Rain	0.154	0.322	0.028	0.168	0.382	0.245	0.073	0.262	0.513
	Rain	0.154	0.313	0.025	0.176	0.338	0.231	0.069	0.277	0.421
Faster RCNN	No Rain	0.216	0.39	0.065	0.248	0.429	0.288	0.104	0.323	0.544
	Rain	0.212	0.38	0.065	0.26	0.375	0.274	0.1	0.346	0.448
RTMDET	No Rain	0.231	0.372	0.044	0.27	0.512	0.3	0.065	0.355	0.619
	Rain	0.231	0.369	0.041	0.289	0.448	0.299	0.057	0.382	0.618
Yolov8x	No Rain	0.227	0.355	0.043	0.264	0.5	0.29	0.059	0.332	0.615
	Rain	0.224	0.348	0.039	0.272	0.439	0.278	0.05	0.346	0.59
Yolov8l	No Rain	0.213	0.332	0.04	0.252	0.495	0.273	0.053	0.309	0.605
	Rain	0.216	0.333	0.038	0.26	0.44	0.263	0.051	0.325	0.52
Yolov8m	No Rain	0.198	0.311	0.033	0.23	0.471	0.256	0.044	0.289	0.583
	Rain	0.202	0.313	0.03	0.237	0.421	0.247	0.039	0.31	0.497
Yolov8n	No Rain	0.113	0.176	0.01	0.123	0.351	0.147	0.012	0.16	0.44
	Rain	0.118	0.185	0.008	0.127	0.313	0.151	0.011	0.17	0.377

(rain). However, these differences are quantitatively quite small, suggesting that rain does not significantly affect the object detection performance.

Examining the results in more detail indicates a more nuanced effect, particularly with respect to object size. Small objects show minimal change, with a very small decline in performance with rain; for all models, for example, AP50:95 is reduced by 0.003 in the presence of rain. At the same time, there is generally a larger decline in performance for large objects, with RTMDet showing the largest drop in AP50:95, from 0.512 (no rain) to 0.448 (rain). Medium-sized objects generally show improved performance under rain across all the models. For example, YOLOv8x exhibits an increase in AP50:95 from 0.264 to 0.272 between the no rain and rain conditions.

Average recall (AR-100), a measure of the proportion of objects correctly detected, decreases by approximately 1% on average in rain, but with some variation with object size. AR-100 degrades by 6% for large objects, while for medium objects, AR-100 increases by an average of 2%. The increase in AR-100 for medium-sized objects in rain is consistent with the average precision results outlined above, suggesting that rain may enhance the detectability of certain object features, possibly due to contrast effects.

Overall, high-level object detection results indicate that while rain does affect performance in subtle ways, overall, it is not significantly detrimental. Indeed, in some cases, it reduces object detection errors, as will be discussed further below when examining the errors in more detail, as well as examining image-level metrics. This counterintuitive behavior suggests that changes in the scene caused by rain may enhance the separability of objects from their backgrounds in some cases, as highlighted by e.g. entropy measurements (see Section IV.C.).

B. ERROR ANALYSIS

TIDE [27] categorizes specific error types from either a false positive or false negative result generated by the model. The dAP for each error type is shown in Fig. 4 for each model type under rain and no rain conditions. The results show that each model behaves slightly differently under rain conditions. Given the level of variance across the models, including different architectures, number of trainable parameters, and different hyperparameter configurations, it is unlikely that any two models would perform identically in rain. Consequently, there is no universal pattern that can be applied to all examined models.

As shown in Fig. 4, there is a decrease in classification errors under rain conditions for all models. In rain, more objects are predicted as being part of their correct class. There is a slight increase in the localization error under rain conditions for all models, except for DETR. There is little change and no discernible pattern in the Miss, Background, Both, or Duplicate error categories under rain conditions.

Fig. 5 shows the effect of rain on the false positive rate (FPR) and false negative rate (FNR) for each model. The implications of false positives and false negatives are application-specific, and it is usually desirable to minimize either FPR or FNR at the expense of the other. Under rain conditions, the FPR decreases for each model, with the exception of RTMDet, by varying degrees. DETR has the largest difference, with a decrease of 2.29%. Every model also experiences a decrease in FNR. RTMDet exhibits the largest decrease in FNR, with a decrease of 4.95%.

C. IMAGE LEVEL METRIC ANALYSIS

This section presents the results of the dataset analysis using the image-level metrics described in Section III-C. Each metric was compared under rain and no rain conditions. Each object was measured independently of the surrounding background pixels, as shown in Fig. 3. The results for each metric are shown in histogram form, with the associated statistics presented in tabular form. Furthermore, the results are examined using ANOVA and Kruskal–Wallis tests to determine statistical significance.

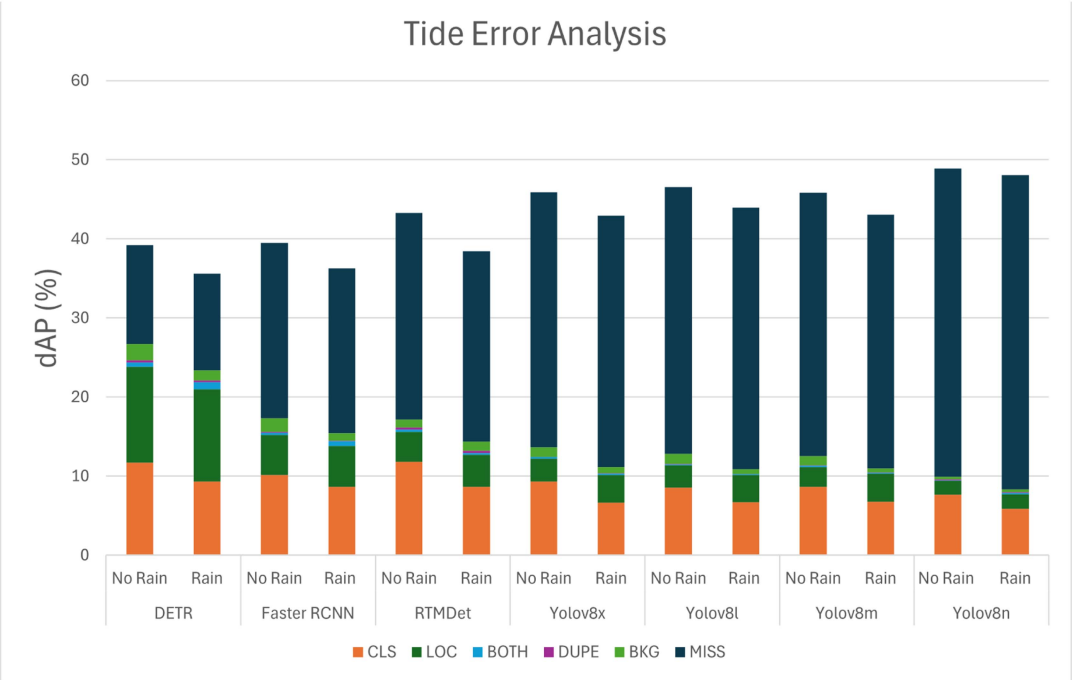

FIGURE 4. Tide Results: Error Classification.

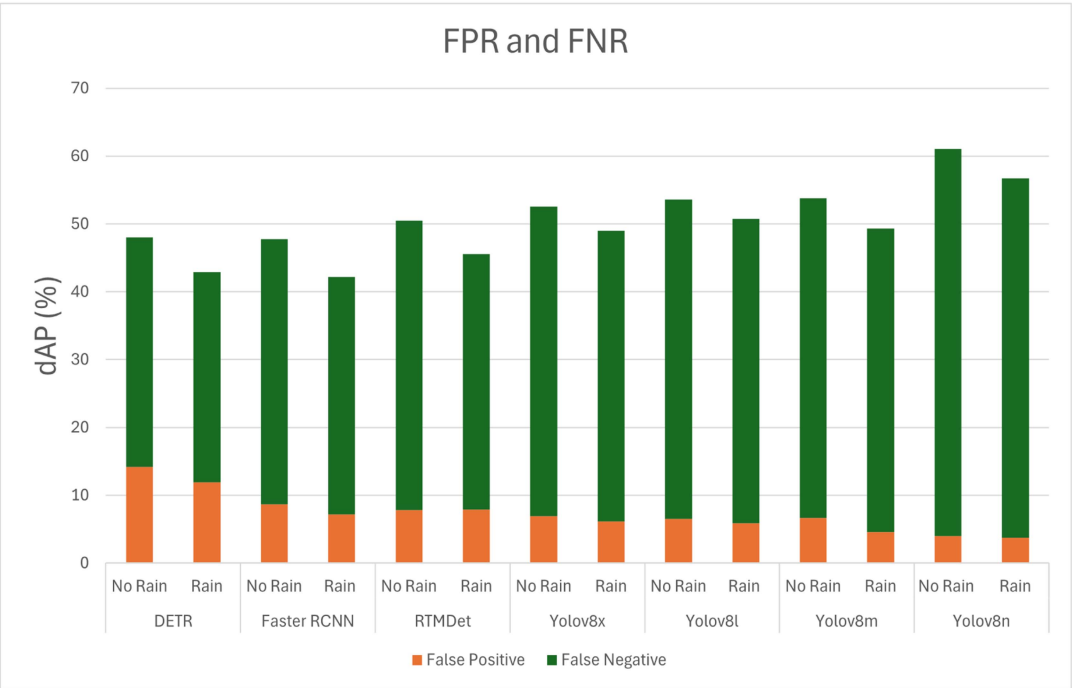

FIGURE 5. TIDE Results: False Positive Rate and False Negative Rate.

The CNR results are shown in Fig. 6 and Table 5. The results show a small decrease in the CNR under rain conditions. The results from the ANOVA test show that the difference between rain and no rain falls short of what is normally classified as statistically significant; however, the results are statistically significant according to the Kruskal-Wallis test. This slight decrease in contrast makes sense in the context of rain, with objects becoming more difficult to distinguish from the background.

The Weber Contrast results are presented in Fig. 7 and Table 6. The results show a difference between the rain and no rain conditions, with the contrast decreasing under rain

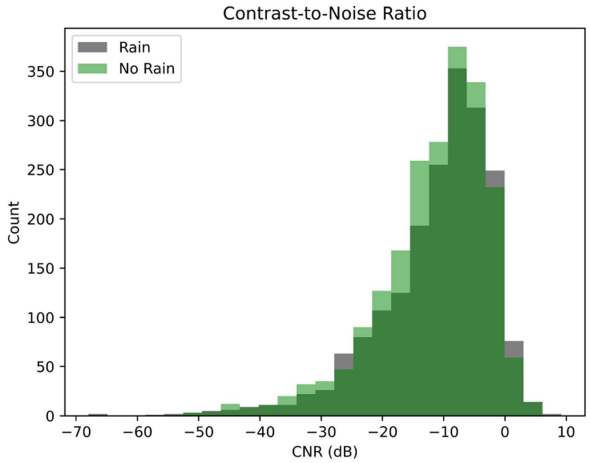

FIGURE 6. CNR Results.

TABLE 5. CNR Statistical Results

	Rain	No Rain
Mean	-11.046	-11.697
Median	-8.932	-9.794
Std. Deviation	9.355	9.269
F-score	4.924	
AN-p-score	0.027	
H-Score	8.174	
KW-p-score	0.004	

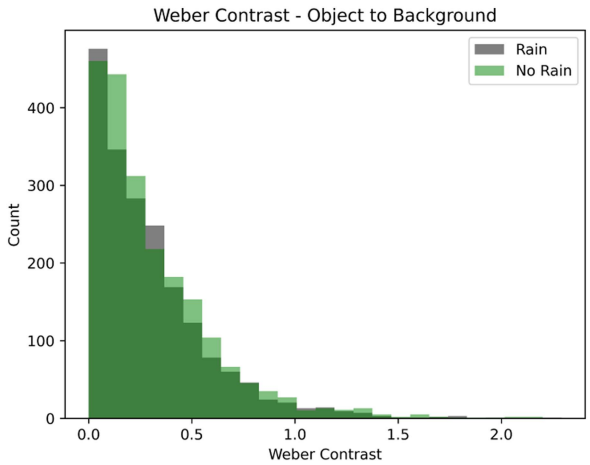

FIGURE 7. Weber Contrast Results.

TABLE 6. Weber Contrast Statistical Results

	Rain	No Rain
Mean	0.2991	0.3151
Median	0.2256	0.2233
Std. Deviation	0.3144	0.3355
F-score	3.0561	
AN-p-score	0.0805	
H-Score	1.4074	
KW-p-score	0.235	

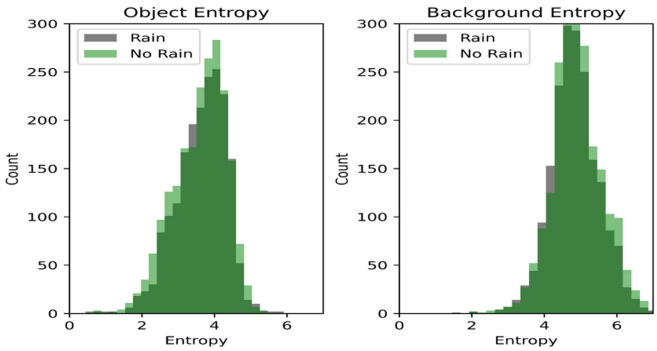

FIGURE 8. Entropy Results.

TABLE 7. Entropy Statistical Results

	Object		Background	
	Rain	**No Rain**	**Rain**	**No Rain**
Mean	3.638	3.600	4.850	4.910
Median	3.729	3.716	4.818	4.860
Std. Deviation	0.689	0.720	0.674	0.696
F-score	2.857		7.847	
AN-p-score	0.091		0.005	
H-Score	1.628		6.959	
KW-p-score	0.202		0.008	

conditions. The decrease in contrast in rain makes it more difficult for objects to be distinguished from the immediate background. While a difference exists, the difference is not statistically significant.

The entropy results are shown in Fig. 8 and Table 7. There was no significant difference in the object entropy results under rain and no rain conditions. This is reflected in the ANOVA and Kruskal-Wallis scores of less than 5.0. As there is no change in entropy, this indicates no loss of detail in the image under rain conditions. The same number of bits would be appropriate for representing pixels under both conditions, indicating no information loss has occurred.

Unlike object entropy, there is a significant difference in the background entropy under rain conditions. The results indicate a decrease in the background entropy during rain, which corresponds to less detail in the background image. Both the F-score and H-score are greater than 5.0, indicating a significant difference between the two conditions. The loss of detail in rain is an expected result and may partly explain the increase in model performance in rain. As the information loss is not uniform across the image, that is, the loss of information is only apparent in the background under rain conditions, this loss of detail may make it easier for an object detection algorithm to distinguish an object from the background.

The statistics for the pixel values are presented in Fig. 9 and Table 8. The mean pixel value refers to the average pixel value of a grayscale image. In rain, the mean pixel value increased for the objects in the dataset, with an increase of 5% compared with objects in dry conditions. F-score and H-score values of

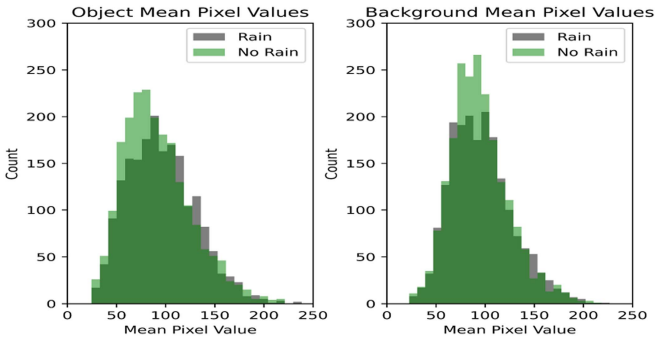

FIGURE 9. Mean Results.

TABLE 8. Mean Pixel Statistical Results

	Object		Background	
	Rain	No Rain	Rain	No Rain
Mean	95.995	91.469	96.254	94.344
Median	92.803	86.464	93.395	91.129
Std. Deviation	34.438	34.495	31.432	29.490
F-score	17.399		3.974	
AN-p-score	<0.001		0.046	
H-Score	22.247		2.245	
KW-p-score	<0.001		0.134	

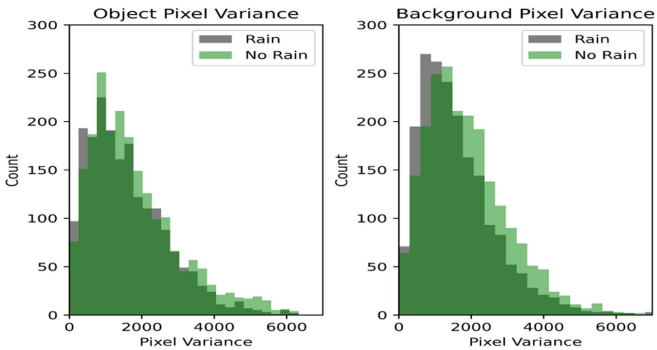

FIGURE 10. Variance Results.

TABLE 9. Pixel Variance Statistical Results

	Object		Background	
	Rain	No Rain	Rain	No Rain
Mean	1597.397	1821.534	1609.611	1879.16
Median	1392.833	1509.875	1383.267	1689.823
Std. Deviation	1092.533	1299.953	1060.986	1131.137
F-score	34.879		60.782	
AN-p-score	<0.001		<0.001	
H-Score	21.997		69.961	
KW-p-score	<0.001		<0.001	

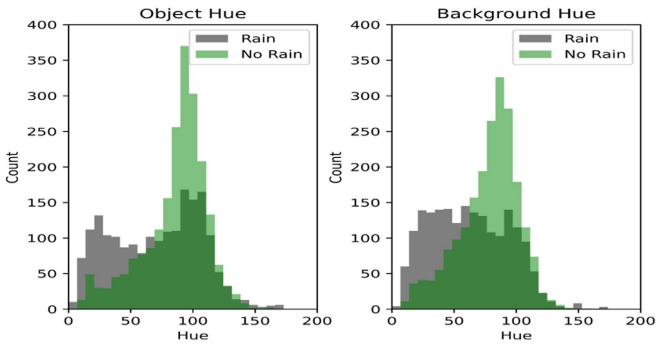

FIGURE 11. Hue Results.

TABLE 10. Hue Statistical Results

	Object		Background	
	Rain	No Rain	Rain	No Rain
Mean	70.694	85.954	62.676	78.185
Median	74.620	91.695	60.808	82.415
Std. Deviation	35.880	25.796	31.248	23.788
F-score	244.288		318.701	
AN-p-score	<0.001		<0.001	
H-Score	176.280		281.946	
KW-p-score	<0.001		<0.001	

17.4 and 22.25, respectively, reflect the significance of this increase. The increase in the mean pixel value is due to visible rain being lighter than the object in the scene, with visible streaks appearing white in an image.

The increase in the pixel value is not present in the background of each object. Both the F-score and H-score are under 5, indicating a low statistical difference. Under rain conditions, the relative increase in the mean pixel value between the object and background increases the local contrast between the object and background. The increase in local contrast is reflected in the Weber Contrast results.

Pixel variance is a good approximation of noise in an image. The pixel variance results are shown in Fig. 10 and Table 9. For both the object and background, the pixel variance is reduced in rain, with a greater decrease occurring in the background. The decrease in variance in rain conditions is indicative of less noise and texture in the image. The F-Score and H-score, of 34.8 and 22 respectively, reflect the significance of the change between rain and no rain for objects in the dataset. On average, there is a 12% reduction in pixel variance for objects in the scene. A similar change also occurs in the background of each object, with an average decrease in pixel variance of 14.4%.

The hue results presented in Fig. 11 and Table 10 refer to the hue value, which has been normalized between 0 and 180. The F-score and H-score, of 244.3 and 176.28, respectively, reflect a significant change under rain conditions with an 18% hue decrease for objects and a 20% decrease in the background. The results show a clear change in the computer vision perception of color under rain conditions. These shifts indicate changes in color perception, not intensity, with rain conditions leading to a wider, flatter hue distribution in both objects and backgrounds. This shift toward lower hue values,

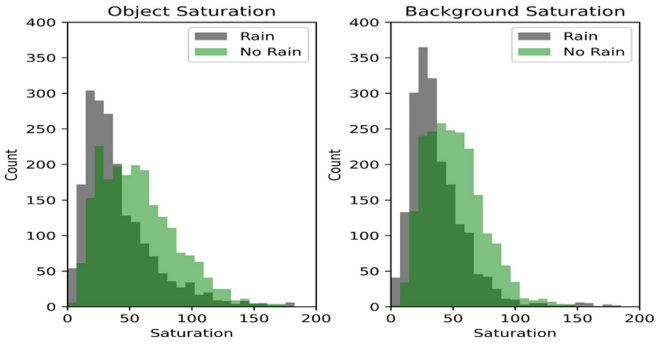

FIGURE 12. Saturation Results.

TABLE 11. Saturation Statistical Results

	Object		Background	
	Rain	No Rain	Rain	No Rain
Mean	40.644	57.000	37.771	51.223
Median	32.806	52.875	32.233	48.924
Std. Deviation	28.938	30.365	23.706	23.111
F-score	306.154		333.546	
AN-p-score	<0.001		<0.001	
H-Score	379.437		433.340	
KW-p-score	<0.001		<0.001	

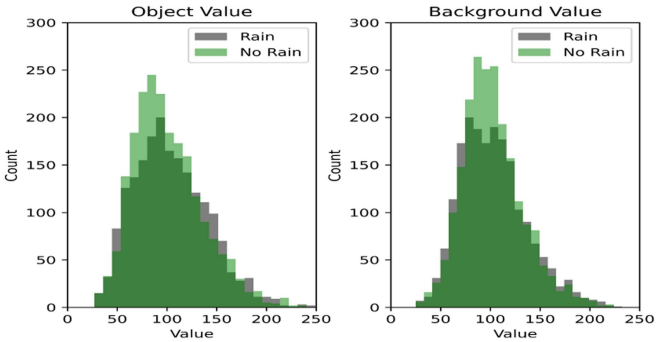

FIGURE 13. Value Results.

suggesting a move toward red, may result from additional brake light reflections in the rain.

The saturation results are presented in Fig. 12 and Table 11, where the higher the saturation value, the more vivid a color appears. For both the object and background, there is a significant decrease in saturation under rain conditions, with decreases of 20.2% and 26.4%, respectively. The large decrease is reflected in the F-score and H-scores, of 306.16 and 379.44 respectively. The distribution of the results can be characterized as being wider and flatter in no rain. In rain, the distribution is narrower, with more low saturation values present. The large decrease in saturation in rain conditions leads to less vivid colors and to images that appear to be more washed out.

The results for value are presented in Fig. 13 and Table 12. An increase in value results in a brighter color. For the objects in the dataset, there is a 4% increase in the value under rain compared with no rain. The F-score and H-score are both

TABLE 12. Value Statistical Results

	Object		Background	
	Rain	No Rain	Rain	No Rain
Mean	104.036	100.486	103.277	102.138
Median	100.137	95.117	99.804	98.884
Std. Deviation	36.85	34.95	33.415	30.475
F-score	9.886		1.285	
AN P Score	0.002		0.257	
H-Score	10.357		0.144	
KW P Score	0.001		0.705	

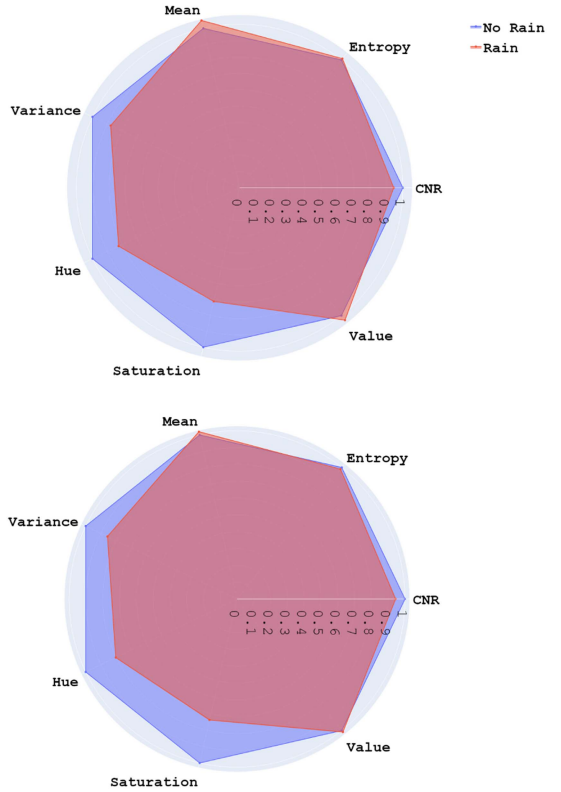

FIGURE 14. RADAR Plots showing the normalized differences between rain and no rain conditions for both the object results (top) and background results (bottom).

greater than 5, indicating a significant difference between rain and no rain, but one of much lower significance than the differences for hue and saturation. The change in value for the objects is not present in the background, which is reflected in the low F-score and H-score.

All low-level dataset analysis results are summarized in the RADAR plots in Fig. 14. The no rain values have been normalized to 1.0, with the corresponding rain results scaled accordingly. The object results are shown at the top and the background results are shown at the bottom of the plot. Both plots highlight the relative differences between rain and no rain conditions for each metric. The CNR result, which is the contrast-to-noise ratio between the object and immediate background, is included in both plots for consistency.

V. DISCUSSION

Despite being an under-explored topic of research, the prevailing consensus across the literature, which was assumed rather than proven in many cases, was that rain would negatively impact detection performance [7], [8], [9]. However, the results of this study show that rain is not systematically detrimental to object detection performance. The performance of some models decreased under rain conditions, whereas the performance of others increased. Furthermore, categorizing the results by object size shows some contradictory behaviors, where the performance of the models increases under rain conditions for medium objects and decreases for large and small objects.

Owing to the counter-intuitive nature of high-level object detection results, the scope of the study was expanded to include a detailed error analysis of the output from each algorithm. The error analysis shows a slight decrease in both FPR and FNR under rain conditions. Object detection performance in rain also generally shows a higher degree of localization errors but with fewer classification errors. Duplication errors, miss errors, and background false positive errors remained almost constant under both rain and no rain conditions.

The detailed dataset analysis revealed several changes that occur at an image level under rain conditions. When examining global performance across the dataset, the metrics with the largest statistical difference between rain and no rain were hue and saturation. Statistically significant changes were also observed in CNR, entropy, mean pixel value, pixel variance, and value. The analysis also revealed that the object and its associated background are influenced differently by rain. Direct, generalizable causal links between object-level metrics and the resultant machine learning performance cannot be easily drawn based on the experiments presented here with a single dataset. Rather, the object-level results should be viewed as average global changes that occur in the dataset under rain conditions; further analysis is needed over a broader range of datasets to determine if such relationships exist. This points to the need for further work on developing automotive datasets in rain conditions, in a wide range of scenarios; moreover, such datasets would require more detailed weather metadata to facilitate more in-depth analysis.

At the same time, it is likely that the observed discrepancies in object detection performance across different models in rain arise from a combination of model architecture factors and input data characteristics, though due to the complexity of large-scale deep learning algorithms and the relative lack of explicit interpretability mechanisms, these influences can only be hypothesized rather than definitively explained.

For example, the two-stage architecture of Faster R-CNN emphasizes localization and classification precision, which may make it more robust to the subtle effect of rain on medium-sized objects. On the other hand, YOLO-based models, which are single-stage models designed for faster inference, may benefit from the reduction in entropy and pixel variance of image backgrounds, which in turn enhances the relative contrast of medium-sized objects more than objects that are less visible. Occlusion effects or changed optical properties of a scene under rain may account for the drop in the performance with large objects. However, as noted previously, such hypotheses, while consistent with our experimental findings, remain speculative, and establishing robust causal relationships between image-level metrics and model performance requires further detailed investigation.

Furthermore, the more detailed error analysis and image-level metric analysis highlight some useful trends. The consistently improved detection performance for medium-sized objects suggests reduced background noise or enhanced relative object saliency in rain, whereas the drop in performance for large and small objects may be due to reduced contrast or occlusion effects. These observations were corroborated by detailed error analysis using TIDE, which revealed fewer classification errors but more localization errors in rain, suggesting that while rain impacts localization, it may assist classification by enhancing the contrast between object classes. The image-level analysis of pixel-level metrics further supports this. For example, the contrast-to-noise ratio (CNR), hue, and saturation exhibited statistically significant differences in rain, with saturation and hue exhibiting the largest changes. Importantly, the entropy analysis indicated a selective loss of detail in the image background in rain conditions, while the object (foreground) entropy was largely unaffected. This may explain the better performance for medium-sized objects, where models are better able to distinguish such objects from their surroundings. However, large objects may not gain the same benefit, as their large size occupies more of the image, leaving them more susceptible to the effects of rain-induced occlusions.

Future work should focus on three key areas: (1) Developing more extensive datasets with associated highly detailed weather metadata to facilitate controlled experiments and a more rigorous evaluation of particular scenarios, including edge cases. (2) More rigorous use of explainability techniques for deep-learning models, such as saliency maps or feature attribution methods, to gain further insight into model behavior. (3) Exploring synthetic datasets as a complement to real-world datasets to simulate precise rain effects and facilitate testing of specific hypotheses.

VI. DATASET LIMITATIONS

The main limitation of this study lies in the reliance of using the pre-existing weather labels of the BDD dataset in particular the 'rain' labels. While BDD is the most suitable publicly available dataset for a study such as this, the dataset's weather annotations are inherently coarse and subjective. This study analyzed what the dataset labels as rain and no rain, rather than the actual presence of rain itself. Fig. 15 shows several sample 'rain' images from the BDD dataset. Examples of light rain, heavy rain, occlusion due to adherent raindrops on the windshield, and an image in which it appears to not be raining but has some aftereffects of rain are shown. All of these images are labeled as 'rain.'

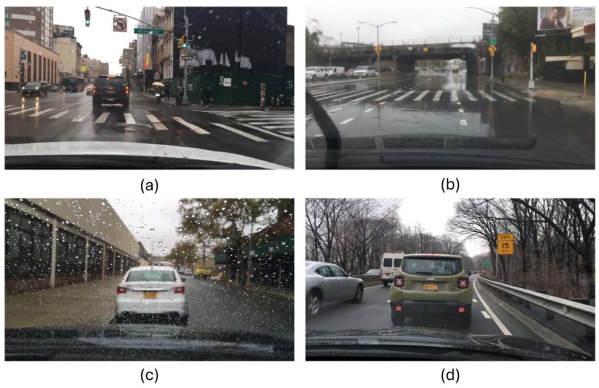

FIGURE 15. Sample rain images from the BDD dataset where (a) appears to be light rain; (b) appears to be heavy rain; (c) has heavy occlusion on the windshield due to adherent raindrops; and (d) has some after effects of rain but does not appear to be raining.

Furthermore, the BDD dataset would have been better suited for this study if it had included more granular metadata. For example, additional information such as rainfall intensity, drop size distribution, and ambient light level could significantly improve weather-centric studies. These finer details would allow for a more nuanced understanding of the impact of rain on perception tasks and advance research in this field. For instance, a light drizzle may affect object detection differently compared to heavy rainfall with a larger drop size.

VII. CONCLUSION

The primary objective of this study was to examine the impact of rain on object detection performance in automotive computer vision applications. A range of object detection models were used, covering different architectural approaches, with no fine-tuning of the models, and no rain mitigation was applied. Each algorithm was pretrained on the COCO dataset, where rain is not prevalent (some images in the COCO dataset do contain rain, but they are unlabeled and not common). Testing was performed using the BDD100 k dataset. While the BDD dataset uses real-world data, one limitation is a lack of "matched" rain and no rain image pairs of the same scene, rendering comparison between the two conditions more difficult; notwithstanding this, it is possible to draw some broad conclusions.

The results indicate that the effect of rain on object detection performance is quite subtle. Overall, there is no systematic degradation observed (which may seem counterintuitive), although closer investigation reveals subtle effects in particular scenarios, for example, different behaviors depending on object size. To gain further insight, a comprehensive error analysis of each algorithm output was completed, and the impact of rain on image-level metrics was explored.

Although the results of this study suggest that modern object detection algorithms are, on average, quite robust in the presence of rain, more detailed research is required to fully understand the relationship between changes in images at the pixel level due to rain and the subsequent knock-on effect on computer vision performance. This further emphasizes the role of comprehensive datasets with detailed weather annotations in enabling stronger conclusions to be drawn that generalize across different scenarios. This would enable a deeper understanding of the subtle relationship between rain and object detection, with precise data collection and targeted algorithmic refinements, allowing researchers to address the most problematic conditions.

REFERENCES

[1] T. Brophy et al., "A review of the impact of rain on camera-based perception in automated driving systems," *IEEE Access*, vol. 11, pp. 67040–67057, 2023.
[2] C. G. Pooley et al., "Household decision-making for everyday travel: A case study of walking and cycling in Lancaster (UK)," *J. Transport Geogr.*, vol. 19, no. 6, pp. 1601–1607, 2011.
[3] M. B. Movahhed, J. Ayoubinejad, F. N. Asl, and M. Feizbahr, "The effect of rain on pedestrians crossing speed," *Comput. Res. Prog. Appl. Sci. Eng. (CRPASE)*, vol. 6, no. 3, pp. 186–190, 2020.
[4] K. Garg and S. K. Nayar, "Vision and rain," *Int. J. Comput. Vis.*, vol. 75, pp. 3–27, 2007.
[5] H. Kurihata et al., "Rainy weather recognition from in-vehicle camera images for driver assistance," in *Proc. IEEE Intell. Veh. Symp.*, 2005, pp. 205–210.
[6] M. Uricar et al., "Let's get dirty: GAN based data augmentation for camera lens soiling detection in autonomous driving," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2021, pp. 766–775.
[7] S. Zang, M. Ding, D. Smith, P. Tyler, T. Rakotoarivelo, and M. A. Kaafar, "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.*, vol. 14, no. 2, pp. 103–111, Jun. 2019.
[8] H. Sun, M. H. Ang, and D. Rus, "A convolutional network for joint deraining and dehazing from a single image for autonomous driving in rain," in *Proc. 2019 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 962–969.
[9] M. Hnewa and H. Radha, "Object detection under rainy conditions for autonomous vehicles: A review of state-of-the-art and emerging techniques," *IEEE Signal Process. Mag.*, vol. 38, no. 1, pp. 53–67, Jan. 2021.
[10] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The PASCAL visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, pp. 303–308, 2009.
[11] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *Comput. Vision and Pattern Recognition*, vol. 1084, Berlin/Heidelberg, Germany: Springer, 2018, *arXiv:1804.02767*.
[12] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc Adv. Neural Inf. Process. Syst.*, 2015, vol. 28.
[13] W. C. Chia, L. S. Yeong, S. I. Ch'Ng, K. P. Seng, and L.-M. Ang, "The effect of rainfall on feature points extraction and image stitching," in *Proc. 2014 Int. Conf. Inf. Sci., Electron. Elect. Eng.*, 2014, vol. 3, pp. 1382–1386.
[14] C. Harris et al., "A combined corner and edge detector," in *Proc. Alvey Vis. Conf.*, Citeseer, 1988, vol. 15, no. 50, pp. 10–5244.
[15] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Comput. Vis. Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008.
[16] P. Müller and A. Braun, "Simulating optical properties to access novel metrological parameter ranges and the impact of different model approximations," in *Proc. 2022 IEEE Int. Workshop Metrol. Automot.*, 2022, pp. 133–138.
[17] P. Müller and A. Braun, "MTF as a performance indicator for ai algorithms?," *Soc. Imag. Sci. Technol.*, vol. 35, pp. 1–7, 2023.
[18] S. Du, Y. Liu, M. Zhao, Z. Shi, Z. You, and J. Li, "A comprehensive survey: Image deraining and stereo-matching task-driven performance analysis," *IET Image Process.*, vol. 16, no. 1, pp. 11–28, 2022.

[19] Z. Zhang, Y. Wei, H. Zhang, Y. Yang, S. Yan, and M. Wang, "Data-driven single image deraining: A comprehensive review and new perspectives," *Pattern Recognit.*, vol. 143, 2023, Art. no. 109740.
[20] T. Wang, X. Yang, K. Xu, S. Chen, Q. Zhang, and R. W. Lau, "Spatial attentive single-image deraining with a high quality real rain dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12270–12279.
[21] S. Li et al., "Single image deraining: A comprehensive benchmark analysis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3838–3847.
[22] H. Caesar et al., "nuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11621–11631.
[23] P. Sun et al., "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2446–2454.
[24] X. Huang et al., "The apolloscape dataset for autonomous driving," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2018, pp. 954–960.
[25] F. Yu et al., "BDD100K: A diverse driving dataset for heterogeneous multitask learning," in *Proc. IEEE/CVF Conf. Comput. Vis. pattern Recognit.*, 2020, pp. 2636–2645.
[26] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
[27] D. Bolya, S. Foley, J. Hays, and J. Hoffman, "Tide: A general toolbox for identifying object detection errors," in *Proc. Comput. Vis.–ECCV 2020: 16th Eur. Conf.*, Glasgow, U.K., 2020, pp. 558–573.
[28] R. Varghese and M. Sambath, "YOLOv8: A novel object detection algorithm with enhanced performance and robustness," in *Proc. 2024 Int. Conf. Adv. Data Eng. Intell. Comput. Syst.*, 2024, pp. 1–6.
[29] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 213–229.
[30] C. Lyu et al., "Rtmdet: An empirical study of designing real-time object detectors," 2022, *arXiv:2212.07784*.
[31] G. Jocher, J. Qiu, and A. Chaurasia, "Ultralytics yolo (version 8.0)," 2023. [Online]. Available: https://github.com/ultralytics/ultralytics
[32] M. Contributors, "Openmmlab detection toolbox and benchmark [computer software]," 2018. [Online]. Available: https://github.com/open-mmlab/mmdetection
[33] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 779–788.
[34] C.-Y. Wang, H.-Y. M. Liao, Y.-H. Wu, P.-Y. Chen, J.-W. Hsieh, and I.-H. Yeh, "CSPNet: A new backbone that can enhance learning capability of CNN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 390–391.
[35] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
[36] J. Hosang, R. Benenson, P. Dollár, and B. Schiele, "What makes for effective detection proposals?," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 4, pp. 814–830, Apr. 2016.
[37] A. R. Smith, "Color gamut transform pairs," *ACM Siggraph Comput. Graph.*, vol. 12, no. 3, pp. 12–19, 1978.
[38] E. Peli, "Contrast in complex images," *J. Opt. Soc. America A*, vol. 7, no. 10, pp. 2032–2040, 1990.
[39] C. E. Shannon, "A mathematical theory of communication," *The Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.

TIM BROPHY received the B.Eng. (Hons.) degree in 2018 from the University of Galway, Galway, Ireland, where he is currently working toward the Ph.D. degree. He is a Member of the Connaught Automotive Research group under the supervision of Prof. Edward Jones and Prof. Martin Glavin. His research interests include computer vision and sensor availability within an autonomous vehicle context.

DARRAGH MULLINS received the B.E. degree in energy systems engineering and the Ph.D. degree in electronic engineering from the University of Galway, Galway, Ireland, in 2013 and 2018, respectively. His Ph.D. research topic involved the application of imaging sensors and signal processing to wastewater treatment plant performance sensing. He was a Postdoctoral Research Fellow with the University of Galway, from 2018 to 2022, where he managed a research program and co-supervised six Ph.D. student projects involving sensors and V2X communication systems for pedestrian and vehicle monitoring from both vehicle and fixed infrastructure point-of-view. He is currently a Senior Technical Officer and Adjunct Lecturer with the School of Engineering, University of Galway. He was awarded Lero Director's Prize for Education and Public Engagement in 2020.

ASHKAN PARSI received the B.E. degree in computer software and the M.Sc. degree in artificial intelligence from the Shahrood University of Technology, Shahrood, Iran, in 2010 and 2013, respectively, and the Ph.D. degree in electrical and electronic engineering from the University of Galway, Galway, Ireland, in 2021, for which he has been awarded the Government of Ireland Postgraduate Research Scholarship (IRC). From 2012 to 2016, he was a Software Developer, Senior Researcher, and Technical Lead on several national projects with the Iran Telecommunication Research Center. He is currently leading and conducting research focusing on the development of signal and image processing technologies for advanced sensors and machine vision for in-cabin monitoring systems with Tobii AB. His research interests include signal processing, machine learning, and algorithm design.

JONATHAN HORGAN is currently a Computer Vision and Deep Learning Architecture Manager and Senior Expert with Valeo Vision Systems. He is also working on next-generation advanced computer vision and deep learning with the ultimate goal of achieving fully autonomous driving and parking. He worked in the field of computer vision for more than 16 years with a focus over the last 10 years on automotive computer vision for Advanced Driver Assistance Systems, automated parking and automated driving. He has 25 publications in peer-reviewed conferences and journals and more than 100 patents published in the field of automotive computer vision.

ENDA WARD received the B.E. degree in electronic engineering from the University of Galway, Galway, Ireland, in 1999, and the MEng.SC master's degree in research in electronic engineering in 2002, with a focus on biomedical electronics. He is responsible for defining the camera product roadmap for surround and automated driving applications with Valeo, Paris, France, and worked with key technology experts across the supply chain and with OEMs to define optimal system architectures. Lecturing for a number of years in the areas of electronics and computing systems with the Atlantic Technological University Ireland, he later moved into industry, working in the biomedical space and has spent the last 16 years in automotive ADAS design. He holds several patents in the area of automotive vision.

PATRICK DENNY (Member, IEEE) received the B.Sc. degree in experimental physics and mathematics from NUI Maynooth, Maynooth, Ireland, in 1993, the M.Sc. degree in mathematics from the University of Galway, Galway, Ireland, in 1994, and the Ph.D. degree in physics from the University of Galway (while working with GFZ Potsdam, Germany), in 2000. From 1999 to 2001, he was a radiofrequency (RF) Engineer with AVM GmbH, Germany, developing the RF hardware for the first integrated GSM/ISDN/USB modem. After working in supercomputing with Compaq-HP, from 2001 to 2002, he joined Connaught Electronics Ltd., (later Valeo), Galway, Ireland, as the Team Leader of RF Design. Over the next 20 years, he was a Lead Engineer developing novel RF and imaging systems and led the development of the first massproduction HDR automotive cameras for leading car companies, including Jaguar Land Rover, VW, Audi, BMW, and Daimler. In 2010, he became an Adjunct Professor of Engineering and Informatics with the University of Galway, and a Lecturer of Artificial Intelligence with the Department of Electronic and Computer Engineering, University of Limerick, Ireland, in 2022. In 2025, he became Senior Lecturer / Associate Professor of AI and Imaging with the Department of Computer Science and Information Systems (CSIS), University of Limerick, Limerick, Ireland. He is currently the Co-Founder and a Committee Member of the IEEE P2020 Automotive Imaging Standards Group, the AutoSens Conference on Automotive Imaging, and the IS&T Electronic Imaging Autonomous Vehicles and Machines (AVM) Conference. His research interests include automotive imaging technology and its extension into other domains, algorithm development, artificial intelligence, machine learning, applied mathematics, and the industrialization of advanced technologies.

CIARÁN (HUGHES) EISING (Senior Member, IEEE) received the degree in electronic and computer engineering and the Ph.D. degree from the University of Galway, Galway, Ireland, in 2003 and 2010, respectively. From 2009 to 2020, he was a computer vision Team Lead and Architect with Valeo Vision Systems, Tuam, Ireland, where he also held the title of Senior Expert. In 2016, he was an Adjunct Lecturer with the University of Galway. In 2020, he joined the University of Limerick, Limerick, Ireland.

BRIAN DEEGAN (Member, IEEE) received the bachelor's degree in computer engineering and the M.Sc. degree in biomedical engineering from the University of Limerick, Limerick, Ireland, in 2004 and 2005, respectively, and the Ph.D. degree in biomedical engineering from the University of Galway, Galway, Ireland, in 2011. From 2011 to 2022, he was with Valeo Vision Systems, Tuam, Ireland, as a Vision Research Engineer focusing on Image Quality. In 2022, he joined the Department of Electrical and Electronic Engineering with the University of Galway as a Lecturer and Researcher. His research focuses on the relationship between blood pressure and cerebral blood flow in humans. His research interests include high dynamic range imaging, LED flicker, topview harmonization algorithms, and the relationship between image quality and machine vision.

MARTIN GLAVIN (Member, IEEE) received the B.E. degree in electronic engineering and the Ph.D. degree in the area of algorithms and architectures for high-speed data communications systems from the University of Galway, Galway, Ireland, in 1997 and 2004, respectively, and the Higher Diploma in Third Level Education in 2007. He was a Lecturer (fixed term contract) from 1999 to 2003 and was a Permanent Member of the academic staff in 2004. He is currently the Joint Director of the Connaught Automotive Research Group, University of Galway. He is also a Funded Investigator in Lero, the Irish Software Research Centre. He has a number of Ph.D. students and Postdoctoral Researchers working in collaboration with industry in the areas of signal processing and embedded systems for automotive and agricultural applications.

EDWARD JONES received the B.E. and Ph.D. degrees in electronic engineering from the University of Galway, Galway, Ireland. His Ph.D. research topic was on the development of computational auditory models for speech processing. He was a Visiting Researcher with the Department of Electrical Engineering, Columbia University, New York, NY, USA, and was a Visiting Fellow with the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW, Australia. He is currently a Professor of electrical and electronic engineering with the School of Engineering, University of Galway. From 2009 to 2010, He also has a number of years of industrial experience in senior positions, in both start-ups and multinational companies, including Toucan Technology Ltd., PMC-Sierra Inc., Innovada Ltd., and Duolog Technologies Ltd. He also represented Toucan Technology and PMC-Sierra on international standardization groups ANSI T1E1.4 and ETSI TM6. His current research interests include signal processing for connected and autonomous vehicles, biomedical engineering, speech and audio processing. He is a Chartered Engineer and a Fellow of the Institution of Engineers of Ireland.