IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

IEEE VEHICULAR TECHNOLOGY SOCIETY SECTION

Received 4 April 2023, accepted 18 June 2023, date of publication 28 June 2023, date of current version 7 July 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3290143

SURVEY

A Review of the Impact of Rain on Camera-Based Perception in Automated Driving Systems

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

[1]School of Engineering, University of Galway, Galway, H91 TK33 Ireland
[2]Ryan Institute, University of Galway, Galway, H91 TK33 Ireland
[3]Valeo Vision Systems, Tuam, Galway, H54 Y276 Ireland
[4]Department of Electronic and Computer Engineering, University of Limerick, Limerick, V94 T9PX Ireland

Corresponding author: Tim Brophy (t.brophy2@universityofgalway.ie)

This work was supported, in part, by Science Foundation Ireland under Grants 18/SP/5942 and 13/RC/2094 P2, and co-funded under the European Regional Development Fund through the Southern & Eastern Regional Operational Program to Lero–the Science Foundation Ireland Research Centre for Software (www.lero.ie), and in part by Valeo Vision Systems.

ABSTRACT Automated vehicles rely heavily on image data from visible spectrum cameras to perform a wide range of tasks from object detection, classification, and avoidance to path planning. The availability and reliability of these sensors in adverse weather is therefore of critical importance to the safe and continuous operation of an automated vehicle. This review paper presents a data communication-inspired Image Formation Framework that characterizes the data flow from object through channel to sensor, and subsequent processing of the data. This framework is used to explore the degree to which adverse weather conditions affect the cameras used in automated vehicles for sensing and perception. The effects of rain on each element of the model are reviewed. Furthermore, the prevalence of these rain-induced changes in publicly available open-source datasets is reviewed. The degree to which synthetic rain generation techniques can accurately capture these changes is also examined. Finally, this paper offers some suggestions on how future adverse weather automotive datasets should be collected.

INDEX TERMS Adverse weather, automated vehicles, computer vision, rain, sensor availability.

I. INTRODUCTION

Automated vehicles is an area of research that has garnered huge amounts of interest in recent years, in both academia and industry. Despite the vast amount of publicly available research that has been conducted, the impact of adverse weather conditions on automated vehicle performance remains an under-explored and under-researched area. Much existing research into automated vehicles focused on the operation in optimal weather conditions, as a result, the influence of adverse weather conditions on sensor performance has yet to be comprehensibly considered. For autonomous vehicles to be suitable for widespread use, they need to be able to operate reliably across all weather conditions. For level 5 automation [1], where a vehicle can safely operate in all conditions with no driver involvement or monitoring, the negative effects of adverse weather on the sensors used on a vehicle need to be fully understood to allow optimal and robust sensor configuration [2].

This review aims to explore how a camera is affected by adverse weather conditions in the context of automated vehicles. The review focuses on visible spectrum cameras as these are of critical importance to autonomous vehicle operation and are heavily involved in the majority of current autonomous vehicle functionality. To fully explore the problem space, a data communications-style Image Formation Framework for data transmission from the object being imaged to the sensor and subsequent image processing is

The associate editor coordinating the review of this manuscript and approving it for publication was Hassan Omar.

proposed. This framework is used to segment the problem space, with a focus placed on the changes that occur in each component of the framework under adverse conditions. Given the passive nature of cameras, any changes to the optical path are likely to result in changes to the data on the image plane. The accurate characterization of camera performance in adverse conditions will help to provide guidance for system design to facilitate safe, reliable autonomous vehicle operations.

The contributions of the paper are as follows: (1) A data-communications-inspired Image Formation Framework that characterizes data flow from objects in the scene through to the image sensor, and the subsequent processing of the collected data is presented; (2) a detailed review of the literature on the impact of rain on camera images is presented, using the Image Formation Framework as a basis for the review; (3) the degree to which current automotive datasets capture rain-induced changes is reviewed; (4) some guidelines on how future adverse weather automotive datasets should be collected are presented.

II. CHARACTERISTICS OF RAIN

Rain is one of the most common adverse weather conditions. Across Europe, rainfall varies significantly with location and season, with North-Western Europe experiencing the most rain [3]. The mean rainfall intensity, including periods of no rain, is between 0.05-0.3 mm/h [3], while the total annual precipitation ranges from 100 to 1300 mm. In Ireland, it rains up to 300 days per year depending on location [4], with an average rainfall of approximately 1000 mm per year [5] and an average rainfall rate of approximately 2 mm/h. Rainfall can be broadly classified into three categories based on the rainfall intensity [6], as follows:

- Slight: < 0.5 mm/h
- Moderate: 0.5 to 4 mm/h
- Heavy: > 4 mm/h

In tropical climates, the rainfall rate can exceed 50 mm/h; however, this typically occurs in short sharp bursts rather than on a continuous basis over an extended period of time.

Rain is often described as a stochastic process [7]. It can be viewed as a series of falling water particles that act as a series of near-spherical lenses, creating fluctuations over time in the optical path at varying spatial locations. The interactions between raindrops, with each other and the environment, produce 'random' fluctuations in lighting across a scene. These changes in the optical path are encoded on the image plane when capturing a scene with a camera. Several features of rainfall, including drop size and speed, can be measured using a laser disdrometer. The rainfall features discussed in the remainder of this section are typically measured using a disdrometer.

Several studies have investigated different features of rainfall, the results of which indicate that raindrops are uniformly distributed in space [9], [10]; however, this is not the case for drop size distribution. Laws and Parsons [8] noted that the drop size distribution is correlated with rainfall intensity.

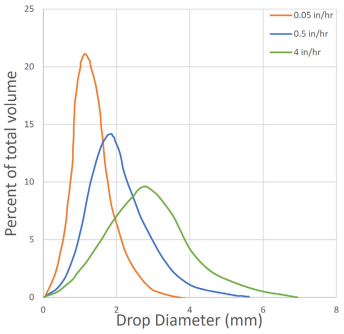

FIGURE 1. Drop size distributions for rainfall of varying intensities based on [8].

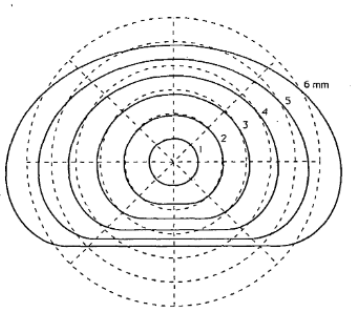

FIGURE 2. Drop shape variation for drops of varying size (given as drop radius) [11].

As rainfall intensity increases, so too does the average drop size. This increase in drop size is also characterized by a wider distribution, meaning that as the rainfall intensity increases, so does the range of the droplet size distribution. Figure 1 shows the distribution of raindrop size for three rainfall intensities. Typical drop sizes range from just above 0 mm to 6 mm for rain falling at 12.5 mm/h (0.5 in/h). The shape of a raindrop is also correlated with its size. As raindrops increase in size, their shape changes from spherical toward that of an elongated ellipsoid. The base of a raindrop experiences the most significant distortion owing to the effect of air resistance [11]. This effect is illustrated in Figure 2.

As the drop size increases, the terminal velocity of a given drop also increases [12]. Several studies have produced different models for the terminal drop speed [12], [13], [14], [15], [16], [17], [18]. The results of these studies are summarized in Figure 3, with a high level of correlation demonstrated between the studies, the exception being the model proposed by Garg [16], which estimates higher terminal velocities than the other models for the same drop size. Typical terminal velocities are in the range of 2 to 10 m/s.

Drop size is one of the key parameters that affects the visibility of a drop on an image plane [16]. Mason and Andrews [19] noted that drop size is also dependent on the type of rain. They also noted that broadly applicable universal relationships between rain parameters are unlikely to hold true given the possible sources of variation, including wind shear and inhomogeneities in the system. These sources

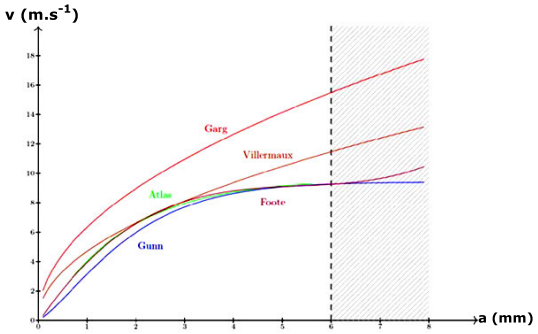

FIGURE 3. Raindrop terminal velocity models versus drop radius [20].

of variation, which introduce errors into any mathematical model of rain, need to be considered when attempting to simulate the effects of rain.

III. IMAGE FORMATION FRAMEWORK

To fully assess the impact of rain on a camera, a data communications-style framework is used to segment the problem. The Image Formation Framework is shown in Figure 4. This framework describes the ''data channel'' from objects in 3D space to the image plane and the subsequent processing of the acquired data. This section explores how each part of the proposed framework is influenced by rain. The components of the framework are defined as follows:

1) Target: Any object of interest in a scene including, but not limited to, pedestrians, vehicles, road signs, traffic lights, and road markings. In this paper, special consideration is given to vulnerable road users (VRUs).
2) Path: The space between the 'Target' and the sensor enclosure i.e. the lens of a camera. This space is characterized by atmospheric and environmental changes over time. The path is where the majority of adverse weather conditions occur.
3) Lens: An optical lens fitted to the camera to focus light onto the image sensor. The accumulation of water droplets on a lens, due to adhesion, leads to occlusions on the image plane.
4) Camera:
   a) Sensor: An image sensor is an electronic device that detects light and converts the detected light to an image. The type of image collected depends on the color filter array, for example, the Bayer Filter [21].
   b) ISP: An image signal processor (ISP) is a digital signal processing unit optimized for images. ISPs have a wide range of functions including auto-exposure, auto-gain, white balance, noise reduction, sharpening, and demosaicing.
5) Algorithm: Further processing is completed on the collected image. Processing is completed with a view to extracting useful information from the image plane. This includes image enhancement (e.g., rain removal), object detection, and classification.

While the analogy is not strict, these components of the framework can be considered to correspond to elements of a typical data communication system. The Target is similar to the data transmitter (source of the signal), while the Path is similar to the transmission channel. Lens distortion can be considered similar in nature to a ''noise source'' or additional distortion in the channel or receive filter, while the camera is a form of ''demodulation'' system that produces output data. Additional processing e.g. image enhancement could be viewed in the same manner as channel coding (used for error detection and correction in a data communications system).

A. TARGET

Under rain conditions, several factors affect how targets are viewed by a given sensor. These factors include changes to the characteristics of the target, such as color, contrast, and reflectivity. Each of these properties is highly dependent on the others and they overlap significantly. VRUs also experience sociological factors such as changes in behavior and changes in physical appearance due to rain. As water accumulates on a target, the interaction between light and the target is affected by scattering, absorption, and attenuation. Depending on the material properties of the target, water can either accumulate on the surface of the target through adhesion or be absorbed by the target. Generally, only porous materials, such as concrete and fabric absorb water.

1) TARGET CHARACTERISTICS

It is well known that as porous materials become wet, they appear darker in color. Twomey et al. [22] attributed this change, in multiple scattering media, to a decrease in relative refractive index as the material becomes surrounded by water instead of air. Ångström [23] proposed that colors appear darker when wet due to total internal reflection in a film of water coating the material. Rough surfaces, including roads, are more prone to this effect owing to the diffuse reflection of light. The orientation of the target is also likely to affect the likelihood of this occurring due to surface run-off. This work was later verified and refined by Lekner and Dorf [24], who also investigated how absorption decreases the relative refractive index. The effect can also be seen in fabrics; as a fabric moves from a dry state to a wet state, the color darkens. The change in appearance is due to less light being reflected by the fabric; therefore, the color of the fabric appears darker. This is true across the full spectrum of visible light [25]. Lee et al. [26] investigated the lightness value, the perceived luminance of an object, and found that the lightness of a textile sample decreases linearly when plotted against the absorption weight of water up to a value of 60%. The change in moisture accounted for a decrease of approximately 14% in the lightness value of the HSL colorspace.

Nayar and Narasimhan derived a model to describe the chromatic effects of atmospheric scattering in fog or haze [27]. The model referred to as 'The Dichromatic Atmospheric Scattering Model,' proposes that the color of a scene

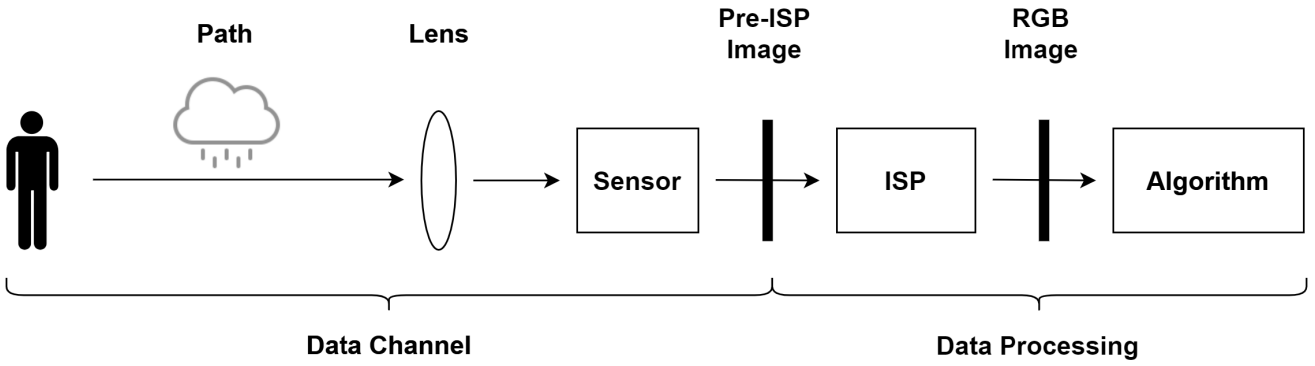

FIGURE 4. The data communications-inspired Image Formation Framework used to segment the problem space.

point (or target) under fog or haze conditions is a linear combination of the direct transmission color (as measured at a time of low atmospheric scattering) and airlight color. Airlight is a phenomenon that causes the atmosphere to act as a source of light. This concept was first introduced by Koschmieder [28] and recently reevaluated by Lee [29], who noted that Koschmieder's model of visibility is akin to identifiability but not necessarily to detectivity.

The visibility of a target is directly related to its reflectivity. The more reflective a target is the more light that returns from the target to the sensor. Higher reflectivity improves the contrast of a given target relative to the background. Retroreflective objects, which redirect light back toward the illumination source, are used to boost visibility in low-light conditions. Common retroreflective objects that are typically found on the roadside include traffic cones, certain road signs, and road studs (commonly referred to as cat's eyes). Coulomb et al. [30] measured the spectral response, in the wavelength range of 350 to 2450 nm, of several common targets in a road environment. Of particular interest in these results is the spectral response of a pedestrian wearing a blue fleece. The spectral response to visible light (400 - 700 nm) is very poor, with the maximum response being 10% at 400 nm. With some exceptions, such as [30], how the spectral response of automotive targets changes in adverse weather is an underexplored area; however, some studies have investigated the impact of water on reflectance for specific surfaces and material types.

Hautiere et al. [31] noted that there is a change in the reflectance of wet materials, in particular, the surface of a road. They noted that the road can appear darker or more specular, depending on the viewing angle of the observer. The increased specularity generally occurs with low light sources such as the sun, at dawn or dusk, or the headlights of oncoming vehicles. The mirror-like surface of the road leads to increased glare and high dynamic range scenes, which are challenging to image. They also attributed the reduction in reflectivity to water on the surface of the road, which leads to more internal reflections at the water-air interface, and to water under the surface of the road, which leads to more scattering, and consequently, more absorption. Senthilkumar et al. [25] measured the reflectance of wet and dry fabric and found a minimum reflectance decrease of approximately 10% across the full visible light spectrum. While exploring the spectral response of wet and dry sand, Nolet et al. [32] found a distinct difference in the reflectance between wet and dry conditions; however, after the initial reduction, the reflectance remains relatively constant as the percentage of water content in the sand varies. This implies that the surface moisture of the sand is a key factor governing the change in reflectance. It is currently unknown how this effect translates onto a porous road surface.

Under adverse weather conditions, especially rain, targets are harder to distinguish from the background scene. This is due to a reduction in contrast across the scene. In very wet conditions, the light reflected from targets is severely scattered by the atmosphere. The resulting change in contrast decreases with the depth of scene, according to an exponential function [33]. Therefore, targets positioned farther from the camera are more difficult to find in a given scene. Burghardt et al. [34] investigated the contrast change of road markings in rain. They found that the contrast ratio in rain (Weber Contrast, also referred to as luminance contrast [35]) decreased by approximately 70% compared with with weather. They noted that high retroreflected luminance (RL) values are strongly correlated with machine vision performance, with higher RL values leading to better contrast ratios, which machine vision algorithms can take advantage of.

2) BEHAVIORAL FACTORS

Under adverse weather conditions, VRUs, especially pedestrians, undergo significant changes to their physical appearance. When it rains, many pedestrians elect to carry an umbrella or wear a jacket with a hood. This wide variation in pedestrian appearance, coupled with lower visibility, makes pedestrian detection in rain particularly challenging. Rasouli et al. [132] investigated which pedestrian attributes contribute to pedestrian detection performance, and noted that the presence of an umbrella, which is rarely seen in datasets, is associated with high algorithm failure rates. In a behavioral study examining the decision-making of households on whether to walk or cycle, the importance of having suitable clothing to protect from the weather was noted. This

clothing included waterproof coats and trousers, hats of various types, and shoes [37]. The protective clothing combined with an umbrella significantly changes the appearance of a pedestrian, which can lead to a domain mismatch for an algorithm attempting to detect pedestrians.

Coupled with a change in appearance, VRUs also experience a change in behavior under inclement weather conditions. VRUs tend to try to limit their exposure to adverse conditions. This manifests itself in more erratic behavior, with VRUs being more willing to take risks than in better weather conditions. Movahhed et al. [38] investigated pedestrian crossing speed in rain and found that pedestrian crossing speed increased by up to 20% in rain conditions. A study examining pedestrians' risk-taking behavior in rain concluded that pedestrians take more risks [39]. Kulmala and Salusjarvi measured the gaps between an oncoming vehicle and a pedestrian who chose to cross the road. In rain, they found that the gap allowed by a pedestrian was on average 0.4 seconds shorter (an 8% reduction). They also noted that pedestrians without adequate protection from the rain accepted shorter gaps while crossing.

Although VRUs undergo significant changes in appearance and behavior in rain, it should also be noted that several studies indicate that the overall number of VRUs is likely to decrease in rain. de Montigny et al. [40] found that on days when there was precipitation the number of pedestrians dropped by 32% to 42% depending on which city the pedestrians were located in. One limitation of this study was that all meteorological data was measured with a resolution of one day. Other studies reported similar reductions: Clifton and Livi [41] noted that 40% of people reported a reduction in walking activities in inclement weather, and Aultman-Hall et al. [42] found a reduction of approximately 13% in hourly pedestrian volumes under precipitation conditions. Miranda-Moreno and Nosa [43] investigated the impact of weather on cycling in Montreal and found that precipitation significantly reduces the number of cyclists on the roads. They also observed a lagging effect, where the number of cyclists was reduced if there had been precipitation in the previous three hours.

The variation in pedestrian behavior makes the accurate sensing and perception of pedestrians more challenging. More erratic pedestrian behavior leads to shorter crossing gaps with oncoming vehicles, requiring quicker reaction times for automated vehicles. The change in pedestrian appearance can lead to a domain mismatch for detection algorithms, resulting in increased algorithm error rates. Despite the more challenging conditions for perception, there is typically a reduction in the overall number of VRUs that need to be monitored.

B. PATH

Rain is complex and dynamic in nature. The path of the Image Formation Framework is arguably the component that is most affected by rain. The 'path' is characterized by the optical path of visible light between the target and the lens. In clear weather conditions, light undergoes atmospheric attenuation while traveling between the target and the sensor, which can be modeled according to a power law [44]. However, given the relatively short range of an automotive camera, the effects of this attenuation are negligible, therefore, we can assume a clean optical path under clear conditions. Unlike other adverse weather conditions such as fog and haze, which are relatively static, rain is more dynamic in nature. As rain falls it causes random spatial and temporal changes to the optical path. As light travels between the target and the sensor along the path, changes to the light pattern are encoded onto the image plane.

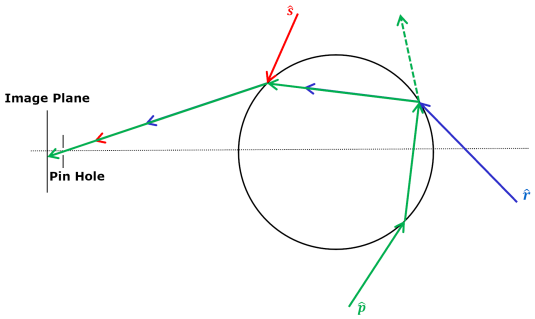

FIGURE 5. Simplified Raindrop Model showing light rays interacting with a raindrop, inspired by [16].

1) INTERACTION OF RAIN AND LIGHT

Every falling raindrop has the potential to interact with light rays in a scene. This interaction usually occurs in the form of reflection or refraction, where a given light ray is redirected. This redirection means light that ordinarily would have reached a camera may not get there and rays that would not ordinarily have traveled to the camera will be received. This leads to a misrepresentation of the data on the image plane, which takes the form of fluctuations in pixel values. Figure 5 shows three light rays interacting with a raindrop, with each ray experiencing a different phenomenon. Path $\hat{r}$ experiences refraction, $\hat{s}$ experiences specular reflection, and $\hat{p}$ experiences internal reflection from the drop. Each phenomenon redirects a light ray toward the camera; as a result, the appearance of the drop on the image plane is a complex mapping of the scene radiance.

Adding more drops to the model shown in Figure 5 leads to a scenario in which each light ray is affected by multiple drops. The extra raindrops add a further layer of complexity to the model, and as a given ray undergoes more reflection and refraction, the effect of an individual drop contributes less to the overall effect. In such scenarios, scattering theory and geometric optics can be used to model the effect of raindrops on light rays. Two widely used scattering models are Rayleigh scattering [45] and Mie scattering [46]. The choice of model depends on the size of the scattering particles interacting with light. Rayleigh scattering is suitable for use when the

particle size is significantly smaller than the wavelength of light. For visible light, this would place an upper limit of approximately 70 nm on the particle size (assuming that the maximum particle size is limited to 1/10 of the wavelength of light [47]). Particles that fit this description include small dust particles and cigarette smoke; however, it is unsuitable for significantly larger raindrops. Mie scattering is used to model scattering in particles of sizes up to the wavelength of light, up to approximately 700 nm for visible light. Particles of this scale include water droplets in fog and clouds; however, raindrops that are typically larger than 700 nm, are still too large for Mie scattering to be accurate.

For larger particles such as raindrops, geometric optics can be used to model their interactions with light. Geometric optics, also called ray optics, refers to modeling the propagation of light in terms of rays [48]. Inakage [49] noted that geometric optics is suitable for modeling the refraction of light by raindrops, albeit with some simplifications. While defining a 'Rainbow Model,' Inakage assumed that all raindrops individually act as refractive prisms and take the form of spherical lenses. This assumption leads to a simple optics problem that can be modeled using Snell's law of refraction. The issue of the refraction of light with multiple raindrops remains an unmodeled effect. In addition to the obvious simplification of the raindrop shape, this model also fails to account for raindrop oscillation as drops fall. Tokay and Beard [50] noted that all raindrops over 1 mm in diameter oscillate as they fall.

2) AFTER-EFFECTS OF RAIN

The path is also susceptible to several static effects during and after rain in the form of changes to the scene. The accumulation of water produces puddles that in turn produce reflections. These reflections are essentially a source of noise from the perspective of the camera, as they distort the appearance of the scene and can cause several issues depending on their size and location relative to the camera. At best, reflections have no effect on the camera; however, they can lead to false positives in terms of object detection [51] and the occlusion of lane markings [52]. Resnick et al. [53] identified 'Cloudy Dark Puddles' as a challenging area for object detection algorithms.

Puddles can also produce dynamic effects, such as spray, in a scene. Yu and Sun [54] investigated the effects of spray caused by high-speed vehicles. They found that the effects of spray are similar to that of heavy rain when the depth of water is at least 5 mm. On very wet roads, vehicles tend to spray water backward, often toward the front-view camera of an autonomous vehicle. Given the direction of the spray, this often leads to heavy lens occlusion.

3) SCENE LIGHTING

Another important factor to consider is the effect of lighting in the scene. Garg and Nayar [16] proposed a photometric raindrop model to describe the effects of refraction and reflection of light by a single drop of rain. The model illustrates how light from different sources (and different directions) can reach the lens of a camera. This effect increases the radiance of the drop, thereby increasing its visibility on the image plane. Typically, lighting phenomena in urban scenes are complex with multiple sources of light, further compounded by reflections from objects in the scene; therefore, the lighting is very unlikely to be uniform or symmetrical. The dominant lighting source will have the largest influence on the visibility of raindrops in the path. During daylight hours, this is likely to be the sun, and at night, this is likely to be from either street lighting or the headlamps of other vehicles. At night, the problem of uneven scene lighting is more severe given the lower relative levels of ambient lighting. Garg and Nayar [10] a the strong dependence of the appearance of a water streak on the illumination direction.

The presence of rain in the 'Path' scatters light resulting in fluctuating pixel values on the image plane. The after-effects of rain, including puddles and spray, can increase the number of reflections in the scene and create complex occlusions on the lens. The brightness and dynamic range of the scene will influence the degree to which rain in the 'Path' is visible on the image plane.

C. LENS

For automated driving, the choice of lens fitted to a camera is of critical importance. Several lens parameters significantly influence the quality of data collected by a camera. One of the key lens parameters is the field of view (FOV). As the FOV increases, each pixel on the image sensor is mapped to a larger real-world area, thereby reducing the pixel density on a real-world target. The reduction in pixel density relates the maximum object detection distance to the FOV. For autonomous vehicles, different FOVs are typically used depending on the location of the camera on a vehicle. To avoid blind spots, camera systems are designed such that the rear-view and sometimes the side-view lenses are generally wide-angle fish-eye lenses with a FOV of approximately 180°. Multiple side-view cameras with narrower FOVs can also be used. Front-view cameras tend to have a considerably narrower FOV. The current industry trend is to have a FOV in the region between 30° and 100°, and occasionally, more than one camera is located at the front of the vehicle. A narrower FOV allows objects to be detected over a much longer range, which is important when traveling at higher speeds. Figure 6 shows the typical camera locations and FOVs of an autonomous vehicle.

Automated vehicles typically use fixed-focus lenses, also called focus-free lenses, instead of autofocusing lenses. Fixed-focus lenses are designed such that every object in a scene, beyond a certain point, is in focus. For a fixed-focus lens, the f-number is particularly important. The f-number refers to the ratio of the focal length to the diameter of the aperture. For automated vehicles, an f-number between F1.6 and F2.4 is typically used [55]. Given that the focal

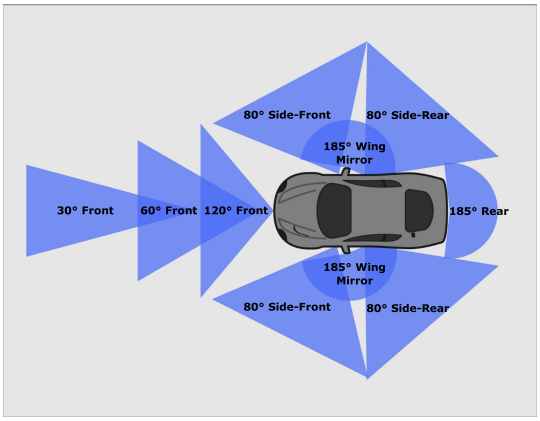

FIGURE 6. A typical camera setup for an autonomous vehicle.

length is directly related to the FOV, which is already constrained, changing the f-number requires a change to the aperture of the lens. This directly controls the amount of light that passes through the lens to the image sensor. The low f-number typically observed in automated vehicle cameras implies a wide aperture. This has the added benefit of reducing noise as more light reaches the image sensor, which requires lower gain values to achieve target image brightness.

1) LENS OCCLUSION

In autonomous vehicles, the lens, lens covering, or windshield (if the camera is located inside the car) is vulnerable to the impact of adherent raindrops. These raindrops stick to the lens and form occlusions in the image plane. This area of research is being actively investigated, and there is a strong overlap with lens soiling [56], where dirt accumulates on the surface of the lens, causing occlusions on the image plane. An example image with adherent raindrops is shown in Figure 7. In the image, several adherent drops of varying sizes are shown, although each drop does not cause a total occlusion but rather acts as a secondary lens heavily distorting the light passing through. In the case of a front-view camera located behind a windshield, the impact of lens occlusion is likely to differ from that of other cameras positioned around the car. The presence of an active wiper system will periodically remove adherent drops, and the system will also benefit from aerodynamic self-cleaning. Both of these factors limit the impact of adherent drops on cameras located behind the windshield.

You et al. studied the effect of a single adherent raindrop. They found that the effect of a raindrop is similar to that of a fish-eye lens and that the scene is contracted through the drop [58]. This effect is illustrated in Figure 8. It was also noted in the same study that motion viewed through an adherent raindrop is significantly slower than that in other parts of an image, on the order of 20 to 30 times.

Several attempts have been made to model adherent raindrops. The determining factor of the geometry of an adherent water drop is wetting. Wetting refers to the ability of a liquid to remain in contact with a solid surface [59]. One

FIGURE 7. The impact of adherent raindrops [57].

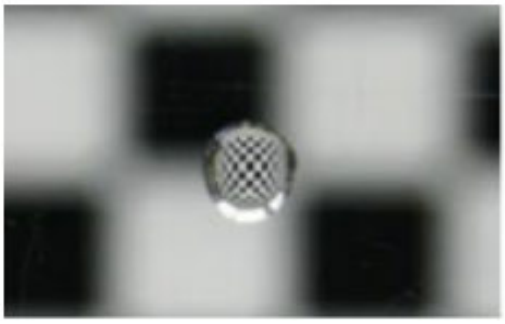

FIGURE 8. The fish-eye effect of adherent raindrops [58].

of the major challenges of modeling adherent raindrops is that their appearance is heavily dependent on the background of a given image. The ability to model adherent raindrops accurately is a key step when attempting to remove them from an image. Roser and Geiger [60], attempted to improve image registration in rainy scenes by modeling raindrops as sections of spheres. They also noted that the occlusion caused by adherent raindrops can be filled in using pixel intensity values from subsequent frames, provided that the image registration is accurate. Roser et al. [61] later proposed using Bézier curves to more accurately capture the shape than using a simple spherical model. They found that a cubic Bézier curve is sufficient to model the curve of an adherent drop. One advantage of this approach is that the model is able to describe the shape of drops that are not on a flat surface. Kurihata et al. [62] proposed an image feature extraction method based on principal component analysis, to detect adherent drops in an image. By creating a raindrop template, referred to as an 'eigendrop,' they were able to perform template matching. Three key properties of adherent raindrops were identified:

- Edges that feature a raindrop outline.
- Blurry edges behind raindrops.
- Refraction of light by raindrops.

The main disadvantage of this approach is a high false positive rate due to variations in the background of the image. Eigen et al. [63] attempted to use a convolutional neural

network (CNN) to try to predict the 'clean' image from an input rainy image. While their algorithm performed better than the median filtering used as a benchmark, the simple structure of the network struggles to capture the complex nature of adherent drops.

Yamashita et al. have published several studies in this area [64], [65], [66], [67], all of which focus on the detection and removal of adherent raindrops. Reference [64] detects areas of interference in an image and replaces these regions with pixel values from another camera. This work was extended to a stereo camera pair in [65]; after transforming the images to a parallel stereo pair, disparities between each image within the common FOV can be easily detected. The main drawback of both techniques is that multiple similarly aligned cameras are required to detect and remove adherent drops. Reference [67] makes use of pan-tilt cameras to achieve the same outcome; two images taken from different reference angles are compared, with high noise regions being detected. As the camera rotates, adherent raindrops occlude different regions in the two images. Following image registration, the regions behind the adherent drop in the initial image can be replaced with the relevant pixel values from the second image. This technique assumes that adherent drops are stationary and is therefore unsuitable for use in heavy rain conditions. Finally, [66] uses spatio-temporal images collected from a single rotating camera. As a result of the projective transform, which describes what happens to the perceived positions of objects as the perspective of the observer changes, the trajectories of a static background become vertical in the cross-section of a spatio-temporal image with the trajectories of adherent noise regions forming curves. These curved tracks in the spatio-temporal image are used to identify regions of adherent noise. Once identified, the regions can be filled in using the image restoration technique described in [68].

2) IMPACT OF LENS PARAMETERS ON RAIN VISIBILITY

Garg and Nayar [69] conducted extensive research on how different camera parameters, including some lens parameters, affect the visibility of rain. Among other parameters, they noted that reducing the f-number of a camera has a significant effect on reducing the visibility of rain. A lower f-number narrows the depth of field, leading to fewer raindrops in focus on the image plane. Garg and Nayar also noted that rain visibility decreases linearly with scene brightness. This effect is likely echoed by increasing the aperture diameter. As previously discussed, autonomous vehicle cameras typically use a relatively low f-number, which helps to limit the impact of rain. Finally, Garg and Nayar noted that rain visibility is also related to the depth of field and the focal length of a lens. However, in the case of autonomous vehicles, the focal length tends to be fixed in order to set the FOV. Notwithstanding the work of [69], the influence of lens parameters on the appearance of rain in an image remains underresearched. The work completed by Garg and Nayer used traditional low dynamic range cameras; when using high dynamic range (HDR) cameras, the visibility of rain is likely to be still dependent on scene brightness; however, due to contrast compression [70], the effect is unlikely to be linear. The visibility of drops is also likely to be dependent on the local tone mapping, which in turn is highly dependent on the local content in an image.

D. CAMERA

1) SENSOR

The key component of any imaging system is the image sensor. For an autonomous vehicle, several key design features need to be considered when selecting an image sensor. For automotive sensors, a dynamic range of 100 to 120 dB, and a resolution of 2 to 4 MP provide a good indication of current industry standards. Another important factor that must be considered is the sensitivity (signal-to-noise ratio) of the sensor. In an ideal world, each of these parameters (dynamic range, resolution, and SNR) would be as high as possible, but inevitable trade-offs between performance and cost place upper limits on one or more of the parameters. The current trend for dynamic range is still 120 dB while 8 MP cameras are currently being explored in practical systems [71].

Pixel size is usually in the range of 2-3um [72], [73]. A smaller pixel size means that the image sensor will not have the necessary dynamic range or signal-to-noise ratio, whereas a larger pixel size means that the sensor will lose spatial resolution. Pixel bit-depth is typically 10-12 bits; this value sets the number of unique shades of color that can be represented. Automotive image sensors are almost exclusively CMOS sensors, which are preferred over CCD sensors for a number of reasons:

- Lower cost
- Higher potential frame rates
- Lower power consumption.

The majority of automotive image sensors still use the Bayer color filter array [21]. Very little, if any, research exists that demonstrates how the primary design parameters of an image sensor influence the impact of rain on the sensor. The noise profile of the sensor will likely affect the visibility of rain in an image, with rain being less visible in regions of high noise.

2) ISP

The role of the ISP is to adjust the data collected from the sensor to optimize image quality. Traditionally, ISPs have been configured and optimized for the human visual system, with efforts focused on making images as perceptually realistic as possible. However for computer vision tasks, especially in the automotive sector, the concept of ''image quality'' is poorly defined, with a single definition yet to be agreed upon [74]. ISPs are used for a wide variety of tasks including denoising, demosaicing, sharpening, and automatic white balancing. A sample pipeline for an ISP is shown in Figure 9.

The impact of rain on an ISP is an underexplored topic, with a relatively small number of studies examining

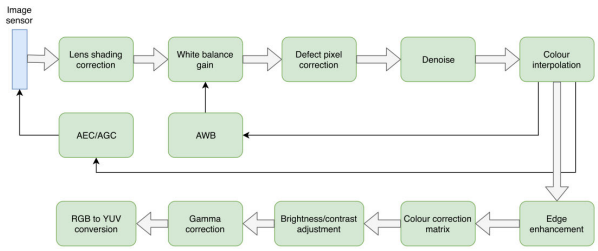

FIGURE 9. A representative ISP pipeline [75].

individual elements of the ISP. Garg and Nayar [69] explored how exposure time is related to the visibility of rain in an image. They found that significantly increasing the exposure time reduces the visibility of rain in an image. However, this comes at the expense of motion blur of a given drop, resulting in streaks and blurring in an image. A long exposure time is also infeasible in the automotive space owing to increased motion blur caused by the motion of the vehicle. The frame rate of the camera also imposes an upper limit on the maximum permissible exposure time. To the best of our knowledge, no other study exists that attempts to tune an ISP to minimize the effects of rain. Under normal circumstances, an ISP adapts to mitigate the effects of adverse changes in the signal path. However, the degree to which a standard ISP configuration can achieve this is unknown because of the large number of parameters that a typical ISP can control, and the confounding effects of these parameters operating simultaneously. There is potential to optimize the configuration and parameter tuning for a given task, for example, object detection, for use in wet conditions. Certain ISP parameters are likely to be more sensitive to the presence of rain which could be leveraged to detect rain in an image.

E. ALGORITHM

1) ALGORITHM PERFORMANCE IN RAIN

The main purpose of any image processing algorithm utilized in autonomous vehicles is to extract useful information from an image. This is often in the form of detecting and identifying an object of interest, classification, and locating where the object is relative to a vehicle. These are often subdivided into smaller tasks, such as pedestrian detection, where a single algorithm focuses on a small task while achieving better results. The majority of state-of-the-art algorithms use machine learning techniques; however, more traditional approaches can also be used, particularly for feature extraction. At its simplest, object detection can be divided into two stages: feature extraction and classification. A wide variety of feature extraction techniques exist, including gradient-based features such as histogram of oriented gradients (HOG) [76], motion-based features [77], shape-based features [78], and texture-based features [79]. Following feature extraction, the extracted features are used to classify detected objects. Chia et al. [80] examined the performance of feature point extraction in rain; they found that the performance of SURF (Speeded Up Robust Feature) and the Harris corner detector decreased by 67% and 48%, respectively, in rainy conditions compared to dry conditions.

Deep learning is the most widely-used approach for modern object detection algorithms, with convolutional neural networks (CNN) very commonly used for object detection in images [81], [82], [83], [84], [85]. It is well known that CNNs fail to generalize outside of the training domain [86], [87], [88], which presents a problem for systems trained in dry conditions but used in rain conditions. One of the most common metrics used to evaluate object detection performance is Average Precision (AP). AP refers to the area under the precision-recall curve. When calculating the true positive and false positive rates, the area of overlap between the prediction and the ground truth must be considered. The overlap between a prediction and the ground truth is referred to as Intersection over Union (IoU). AP50 refers to the area under the curve with an IOU of at least 50%. Mean Average Precision (mAP) refers to the average AP50 score across all classes of objects in a dataset.

Hnewa and Radha [89] investigated the performance of Yolov3 [90] and FasterRCNN [82] under rain conditions. When trained using only 'clear' weather images, the mAP decreased for both algorithms by 3.3% and 5.4%, respectively. Figure 10 and Figure 11 provide simple illustrations of the issues that algorithms have in rain conditions. In Figure 10, when simulated rain is added, there is a large discrepancy between the performance of the two algorithms in terms of object location. Figure 11 shows the performance of Faster RCNN when synthetic rain is added to an image. As the synthetic rainfall rate increases, the mAP decreases, with a maximum decrease of approximately 40% at 80 mm/h.

Dodge and Karam [91] investigated the effect of Gaussian noise on the performance of deep neural networks. They concluded that deep neural networks are susceptible to the effects of blur and noise but are less susceptible to changes in contrast. This is of particular interest for images subjected to rain, as rain streaks blur the background and introduce noise. The loss of contrast caused by rain may be less of an issue. Da Costa et al. [92] investigated how different types of noise impact image classification, and arrived at a conclusion similar to that of Dodge and Karam. Additionally, [92] noted that while the denoising of distorted images improves performance, the performance with the original image, without noise, is still significantly better. They attributed this to the blurring effect of denoising techniques.

Several techniques have been used to attempt to improve the performance of CNNs in rain. Perhaps the most obvious technique is to simply include an adequate number of rain images in the training data. Several adverse weather datasets exist, such as [95], which contain a large number of images captured while it was raining. However, the ratio of wet to clear images required by a machine learning algorithm is an under-explored topic, and the potential improvement in performance, due to the use of a broad training set, is also unknown. On the other hand, if rain is viewed as a form of

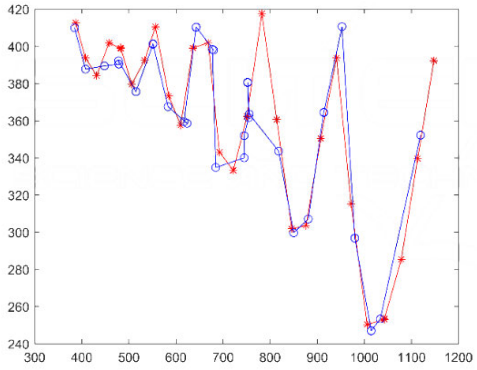

FIGURE 10. The detected position of a bouncing ball in clear (red) and synthetic rain (blue) conditions [93].

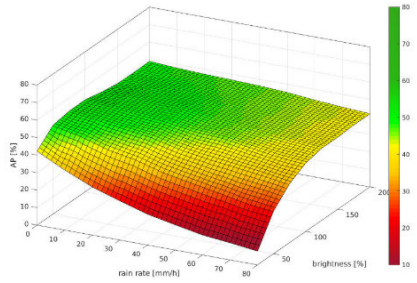

FIGURE 11. The performance of Faster RCNN against synthetic rain [94].

image "corruption," simply removing the corruption from the image should improve object detection performance. Several attempts have been made to 'derain' an image [96], [97]. Image deraining is discussed in more detail below. Another valid approach is to artificially add corruption to the training data. This approach is usually referred to as synthetic rain or simulated rain. Essentially, the training data, or a portion of it, is corrupted with synthetic rain to improve the robustness of the network to real rain. This technique is discussed in more detail in Section III-E4.

2) DETECTING/MEASURING RAIN

When viewing an image without any prior knowledge of the weather conditions, it can often be difficult to determine if it was raining when the image was captured. This presents a significant challenge for machine learning algorithms. As discussed in [69], the visibility of rain is heavily reliant on camera parameters. While having little impact from a human visual perspective, the impact of rain can heavily influence the performance of CNN-based algorithms. Knowing if it is raining is, therefore, an important initial step, as this knowledge can be used to optimize subsequent image processing and computer vision algorithms.

Several studies have been published using a variety of techniques concerning the automatic detection of rain. Barnum et al. [98] demonstrated that despite their apparent chaotic nature in the image domain, rain and snow have a predictable global effect in the frequency domain. Xue et al. [99] made use of the wavelet transform and bilateral filtering to create a detailed edge map containing rain. Several studies have focused on detecting adherent drops on windshields [60], [62], [100], [101]. All these studies used an in-vehicle camera located behind a windshield, with a variety of image processing and machine learning techniques used to process the subsequent images. Applying some of these techniques to common datasets, such as [102], could potentially return coarse labels describing the current rain conditions in an image. This could potentially alleviate the problem of a lack of labeled weather conditions in open-source datasets. Another valid approach to detect rain is to leverage the changes that take place to 'Targets,' especially pedestrians, in a scene. The indirect observables brought about by rain, such as a pedestrian using an umbrella, may prove to be stronger heuristics of rain than any direct measurement obtained from an image.

It is also worth noting that rain may also be detected through the use of on-vehicle hardware. Rain sensors can be used to detect the presence of rain to control smart wiper systems, for example, [103] reviews several different rain sensor types. Optics-based rain sensors can be positioned inside the windshield. Such sensors use the changing reflection and refraction properties of two surfaces in contact to detect rain on the windshield. On-vehicle hardware can accurately detect the presence of rain; however, not all vehicles are equipped with such hardware. The focus of this paper is on the most general situation where no assumptions are made regarding the presence of specific rain-detecting hardware on the vehicle. In such situations, using an algorithmic approach to detect the presence of rain is the only viable option.

3) DERAINING

The goal of image deraining is to remove rain from an image to be left with a clear image, free from the interference caused by rain. This is an active topic of research and has been the subject of several recent reviews [104], [105], [106]. Image deraining is not a trivial problem. In their review, Yang et al. [105] identified several key problems:

1) Rendering physically correct rain, to add to previously acquired images, is very complicated. Owing to the complex nature of rain, a large number of factors, such as shape, size, and orientation need to be considered. The appearance of rain also depends on the background of a scene.
2) Decoupling rain information from background and lighting information is very difficult. Given that an image comprises a matrix of pixel values, all rain, object, background, and lighting information is fused into a single value for each pixel.
3) Rain and background information often overlap in the feature space. Background textures found in images can closely resemble those of rain. Without prior knowledge of the background textures, these textures can result in false detections.

FIGURE 12. The RESCAN algorithm with the rainy input on the left and the derained output on the right [107].

4) Nearly all machine learning training methods for image deraining rely on image pairs, one with rain and one without rain. Gathering comprehensive, suitable, real-world data for this task is nearly impossible. Even for a static scene in which all objects in the background remain in place, illumination changes in the scene prevent the collection of an ''ideal'' image pair for training purposes. Inaccurate rain models or synthetic rain used to generate image pairs can result in the development of inaccurate deraining algorithms.

Given the detailed nature of the reviews already published in this area, the rest of this section focuses on the merits of the current state-of-the-art deraining algorithms and their effect on object detection algorithms. Although there is no doubt about the effectiveness of current deraining algorithms for human viewing, there is little evidence to suggest that deraining algorithms can improve the performance of object detection in rain conditions. One of the state-of-the-art deraining algorithms, RESCAN (Recurrent SE Context Aggregation Net), is based on a recurrent neural network (RNN) [107]. An example illustrating the performance of RESCAN is shown in Figure 12.

Despite their excellent visual performance, deraining algorithms have not yet been optimized for object detection. Wang et al. [108] noted the lack of a quality benchmark to evaluate the performance of image-deraining algorithms. Many deraining algorithms are evaluated using subjective metrics or image quality metrics, including PSNR and SSIM. These algorithms work as intended from the perspective of these metrics; however, this does not necessarily correlate with improved object detection performance. In their review, Yang et al. performed a series of tests to evaluate the effectiveness of several deraining algorithms on pretrained object detection algorithms. They used YOLOv4 [81] and faster R-CNN [82] to evaluate several deraining algorithms. At best, the performance of each algorithm was no worse after deraining than before deraining; however, in most cases, a small decrease in mAP was observed. Li et al. [109] carried out a similar experiment and formed a similar conclusion. They found that ''almost all existing deraining algorithms will deteriorate the detection performance compared to directly using the rainy images.''

Pei et al. [110] performed an empirical analysis of various de-hazing techniques on both real and simulated haze images. De-hazing is a similar problem to de-raining and faces many of the same challenges. They concluded that de-hazing algorithms had little positive effect on improving object detection performance. Furthermore, they noted that there is little correlation between the current image-level de-hazing metrics, including pixel-wise errors and local structural similarities, and object detection results. A major concern for deraining techniques is that the deraining process may remove true discriminatory information from the image. Perhaps a fairer experiment would be to fine-tune an object detection algorithm with images that have already been derained. This approach would allow an algorithm to generalize to the imperfections of a deraining algorithm and boost the overall object detection performance.

4) SYNTHETIC RAIN

Augmenting existing datasets with synthetic rain has become a popular technique to attempt to limit the degradation in the performance of neural networks when presented with rain images. Adding synthetic rain is essentially the reverse process of de-raining, and therefore faces many of the same issues outlined above. Given the difficulty in collecting ideal image pairs, it is challenging to design accurate synthetic rain models. Hnewa and Radha [89] noted that the inaccuracy of synthetic rain models leads to a domain mismatch.

A variety of techniques are often used to generate synthetic rain. Broadly speaking, these techniques can be divided into three categories: (i) physics-based models; (ii) data-driven models; and (iii) a combination of the two. Garg and Nayar were among the first to explore the factors affecting the appearance of rain [16] and to model physics-based photorealistic streaks [10]. Halder et al. [111] used a physics-based simulator [112], to generate synthetic rain images in an attempt to improve object detection in rain. Other techniques for rendering rain streaks include a frequency-based model [113] and a technique that uses ray tracing, as described in [114].

Ni et al. [115] proposed the Rain Intensity Controlling Network (RICNet) which is capable of adding or removing rain from a given image. The bi-directionality of the network is unique, and allows the user to set the rain intensity to the desired level. The network was trained in a supervised manner making use of synthetic rain datasets. Wei et al. [116] designed a CycleGAN network capable of synthesizing rain in an unsupervised manner. Two generators create rain-free and rainy images, respectively. The major advantage of this approach is that the need to use rain and rain-free image pairs for training is negated, thereby allowing the use of unmatched rain and rain-free images. Several synthetic rain datasets exist that are commonly used to train rain rendering and rain removal networks. Synthetic rain datasets often include images from more common automotive datasets, such as [102] and [117], which have been augmented using a variety of rain rendering techniques. Among the most common synthetic rain datasets found in the literature are Rain12000 [118], Rain800 [119], Rain200H [96], and SPAData [108].

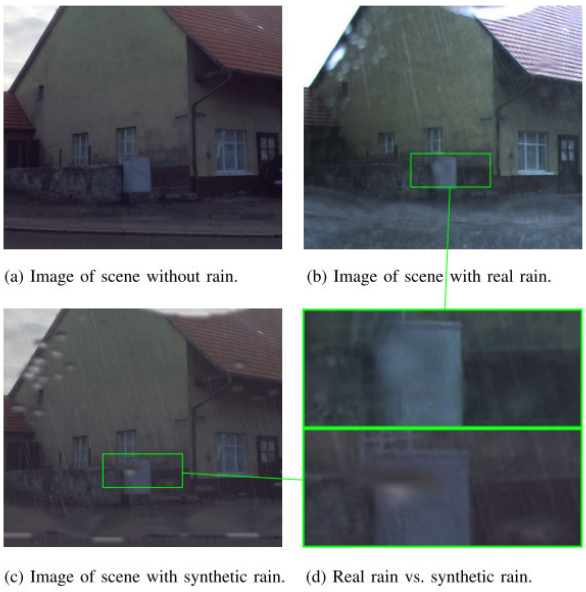

FIGURE 13. Comparison between synthetic and real rain [94].

Figure 13 shows an example of synthetic rain which has been added to an image. The results are visually impressive; however, as in the case of image deraining, there is no standardized objective quality benchmark to evaluate performance against. Many rain rendering algorithms are evaluated subjectively using surveys that ask participants to rate the realism of rendered rain. One of the major issues with this evaluation technique is demonstrated in [120], where only slightly over 75% of the participants agreed that a control image with real rain was representative of real rain, whereas the remaining 25% failed to agree that the rain appeared to be real. As with rain removal algorithms, the subjective measure from a visual perspective has little bearing on computer vision interpretation.

Despite the lack of qualitative evaluation metrics for synthetic rain, some other limitations also need to be considered. Given that synthetic rain is overlaid on top of an image taken in dry conditions, the current techniques fail to account for several of the changes previously outlined in the data communication-based model presented earlier in this review. Perhaps the most important change unaccounted for is the change in object (pedestrian) behavior. Changes in pedestrian behavior and appearance can have a large effect on algorithm performance, and these changes are not accounted for using the current state-of-the-art synthetic rain generation techniques. These behavioral changes are among the most difficult to model.

Rain has the potential to negatively influence the performance of computer vision and machine learning algorithms. Techniques to limit the impact of rain, such as augmenting training data with synthetic rain or removing rain from the image plane, do not appear to have the required mitigating effect. In order to fully characterize the impact of rain on computer vision and machine learning performance, datasets with detailed adverse weather labeling are required.

IV. DESIGNING A DATASET

The selection of data for training a neural network for any application is of critical importance. If the data are not an accurate representation of the real scenario, the performance of any network will degrade significantly when tested under realistic conditions [86]. Similarly, the amount of available training data is also of particular importance. The ideal amount of data required strongly correlates with the network depth and the number of trainable parameters in the network. Using too much training data will cause the network to overfit the training data and fail to generalize on similar scenarios. Conversely, when using too little data, the network fails to capture the relationships between the input data and the output of the network.

Within the context of this review, for an automotive dataset to accurately represent the complex effects of rain, all the factors discussed in previous sections should be included. Several of the changes discussed are difficult to model. Therefore, the most appropriate way to accomplish this is to capture real data in a realistic environment. The importance of labeling data should not be underestimated. Providing detailed weather and lighting labels may not improve the overall performance of a network; however, they will provide more details about the edge cases where the network fails. Ideally, as much environmental data as possible should be captured alongside the sensor data.

A. EXISTING DATASETS

Table 1 presents a collection of major publicly released automotive datasets used for object detection and semantic segmentation. Each of the datasets has been analyzed in terms of the diversity of conditions, namely Day, Night, Dawn/Dusk, and Rain, and the presence or absence of detailed weather labels and ambient light information. Other basic details of the datasets have also been included in Table 1.

When presented in tabular form, the trends across many major datasets become clear. The general diversity of many datasets is high, with several datasets containing images across all times of day and in a variety of rain conditions e.g. [95] and [124]. All of these datasets captured images in real-world conditions and are therefore assumed to have accurately captured the changes between wet and dry scenes described in Section III of this paper.

The table reveals a lack of weather labels included with the sensor data. The Berkley Deep Drive (BDD) dataset [95] is the only large-scale dataset that contains any weather information. These labels are coarse and lack detail of specific conditions; however, they are still useful for segmenting the dataset into various weather conditions. To the best of our knowledge, no major publicly available automotive dataset contains any labels describing the ambient light conditions. The NightOwls dataset [131] attempts to infer ambient light

TABLE 1. Adverse weather datasets for autonomous vehicles review. OD: Object detection, SS: Semantic segmentation, D: Day, N: Night, DD: Dawn/Dusk, R: Rain, BB: Bounding box, WL: Weather labels, AL: Ambient light, SR: Synthetic rain, ✓*: Coarse label estimated from pixel intensities.

Name	Date	Citations	Frames (k)	Classes		Conditions				BB		WL	AL	SR
				OD	SS	D	N	DD	R	2D	3D			
PandaSet [121]	2021	1	60	25	37	✓	✓	✓	✗	✓	✓	✗	✗	✗
RaidaR [122]	2021	3	5	00	17	✓	✓	✗	✓	✗	✗	✗	✗	✗
NuScenes [123]	2020	1103	40	23	00	✓	✓	✓	✓	✓	✓	✗	✗	✗
Waymo Open [124]	2020	469	200	04	00	✓	✓	✓	✓	✓	✓	✗	✗	✗
A2D2 [125]	2020	92	12.5	03	38	✓	✗	✗	✓	✗	✓	✗	✗	✗
Woodscape [126]	2019	124	10	07	40	✓	✗	✗	✗	✓	✗	✗	✗	✗
H3D [127]	2019	95	27	08	00	✓	✗	✗	✗	✗	✓	✗	✗	✗
Boxy [128]	2019	13	200	01	00	✓	✗	✓	✓	✓	✓	✗	✗	✗
BLVD [129]	2019	23	120	03	00	✓	✓	✗	✗	✗	✓	✗	✗	✗
EuroCity [130]	2019	111	47	07	00	✓	✓	✓	✓	✓	✗	✗	✗	✗
NightOwls [131]	2019	35	279	04	00	✗	✓	✓	✓	✓	✗	✗	✓*	✗
Pie [132]	2019	89	293	01	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
SPAData [108]	2019	3	29.5	00	00	✓	✗	✗	✓	✗	✗	✗	✗	✓
KAIST Multi-Spectral [133]	2018	107	8.9	03	00	✓	✓	✓	✗	✓	✗	✗	✗	✗
BDD [95]	2018	361	100	10	40	✓	✓	✓	✓	✓	✗	✓	✗	✗
ApolloScape [134]	2018	306	147	00	35	✓	✓	✗	✓	✗	✗	✗	✗	✗
Rain12000 [118]	2018	3	5	00	00	✓	✗	✗	✗	✗	✗	✗	✗	✓
Bosch Small Traffic Lights [135]	2017	152	13	15	00	✓	✗	✗	✓	✓	✗	✗	✗	✗
CityPersons [136]	2017	459	5	06	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
Raincouver [137]	2017	22	0.326	00	03	✓	✓	✓	✓	✗	✗	✗	✗	✗
Mapillary Vistas Dataset [138]	2017	641	25	00	66	✓	✓	✓	✓	✗	✗	✗	✗	✗
JAAD [139]	2016	67	88	01	00	✓	✓	✓	✓	✓	✗	✗	✗	✗
Tsinghua-Daimler Cyclist [140]	2016	100	15	06	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
CityScapes [117]	2016	6280	5	00	30	✓	✗	✗	✗	✗	✗	✗	✗	✗
KAIST Multispectral Pedestrian [141]	2015	481	95	03	00	✓	✓	✓	✗	✓	✗	✗	✗	✗
PascalRaw [142]	2014	5	4.259	03	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
KITTI [102]	2013	5020	15	08	00	✓	✗	✗	✗	✓	✓	✗	✗	✗
German Traffic Sign [143]	2013	592	5	43	00	✓	✗	✓	✓	✓	✗	✗	✗	✗
TME Motorway [144]	2012	168	30	02	00	✓	✗	✓	✗	✓	✗	✗	✗	✗
TUD-Brussels Pedestrian [145]	2010	495	1.6	01	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
Caltech Pedestrian Dataset [146]	2009	1446	250	01	00	✓	✗	✗	✗	✓	✗	✗	✗	✗
Daimler Pedestrian [147]	2009	1544	21	01	00	✓	✗	✗	✗	✓	✗	✗	✗	✗

conditions by estimating a value from pixel intensity values. At best, this approximation only provides an estimate of the actual ambient light level.

B. CONSIDERATIONS FOR BUILDING A DATASET

Section III-A to Section III-D of this paper, which describes the 'data channel' of the Image Formation Framework, summarized the changes that take place in the environment. The changes that occur to the 'targets,' in the 'path,' and on the lens of the camera should all be captured when collecting data. Ideally, data should be collected under varying conditions and at varying times of year. The more adverse weather training data, and the more variation within the data that is collected, the more robust an algorithm trained on the data is likely to become. Current state-of-the-art synthetic rain generation techniques are not a suitable replacement for gathering real rain data because they fail to account for the full scope of changes that occur under rain conditions.

The major problem with open-source autonomous vehicle datasets, from the perspective of adverse weather, is the lack of detailed weather labels. The lack of labels makes characterizing the effects of adverse conditions more challenging. Adding weather labels to a dataset as a post-processing step is challenging and unlikely to be accurate. When designing a dataset suitable for characterizing sensor performance under adverse conditions, as many labels as possible should be recorded at the time of data acquisition. The more labels that are recorded and the greater the granularity within these labels, the easier the task of characterizing sensor performance becomes. In urban environments, owing to occlusion by buildings, local conditions at the vehicle change rapidly; therefore, weather and environmental data labels should be regularly updated and reviewed during data collection. The addition of sensors to monitor ambient conditions, such as lux meters, can make the process of updating labels more reliable.

Categorizing rainfall into discrete categories is desirable, and categorization should be completed from the perspective of a machine learning algorithm. If a large volume of data with labeled continuous rainfall information becomes available, this process could be completed using contribution value plots, that is, plotting the rainfall rate (or some other feature of rainfall) versus an algorithm performance metric. Identifying key regions in the plot, such as knee points, would reveal the boundaries that should be used for the categories. The plots should be repeated for various features of rainfall, as the rainfall rate may not be the feature that has the largest impact on machine learning performance.

Another approach to categorize rainfall, which would require less detailed rainfall labeling, is through the use of generative adversarial networks (GANs) [148]. A similar

concept is explored in [149], to examine how lens soiling is viewed from the perspective of a network. The discriminator in the GAN is asked to identify whether a lens is soiled or not. The performance of the discriminator, coupled with human intuition, indicates the degree of separability between the classes. Poor discriminator performance indicates low discernibility between classes, with the GAN struggling to distinguish between the categories. Starting with arbitrarily assigned rainfall categories, a GAN could be used to measure the separability between each category from a network's perspective. If the classes are not discernible from the perspective of a network, they do not add useful information to the dataset.

V. CONCLUSION

This review has presented a data communications-inspired Image Formation Framework to explore how adverse weather conditions affect camera images collected for use in an autonomous vehicle application. In particular, the effect of rain on cameras is examined. The framework segments the problem space, and the changes that occur under rain conditions in each section of the framework are reviewed. The subsequent effect of using rain-degraded data as input for subsequent data processing is also reviewed. A focus was placed on object detection with the related areas of synthetic rain generation and rain removal included. Finally, this study reviews publicly available datasets suitable for automotive applications.

Generally, across all aspects explored using the Image Formation Framework, there is a lack of research that characterizes the effects of rain. This lack of research is especially evident in the areas concerning the ISP and the image sensor. The subsequent effects on processing algorithms, such as object detection, are also underexplored. The current state of open-source datasets with respect to comprehensive coverage of adverse weather conditions is relatively limited. While many datasets contain a high level of diversity, in terms of weather conditions, a lack of detailed labeling and other metadata limits the potential of this data. Without gathering more detailed weather labels at source when collecting data, fully understanding the impact of rain is not possible. Another cause for concern lies in the fact that uncommon edge cases are likely to be among the most difficult challenges of adverse weather to overcome. Identifying the conditions that lead to edge cases, where the failure of an autonomous vehicle is most common, requires significantly more detailed metadata.

In summary, rain impacts a wide variety of aspects of an autonomous vehicle environment. Many of these aspects remain under-explored, and further research is required. In order to conduct this research, datasets with high levels of diversity and detailed labeling will be required.

REFERENCES

[1] *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, Standard J3016, 2021. [Online]. Available: https://www.sae.org/standards/content/j3016_202104/

[2] C. Michaelis, B. Mitzkus, R. Geirhos, E. Rusak, O. Bringmann, A. S. Ecker, M. Bethge, and W. Brendel, "Benchmarking robustness in object detection: Autonomous driving when winter is coming," 2019, *arXiv:1907.07484*.

[3] A. Overeem, E. V. D. Besselaar, G. V. D. Schrier, J. F. Meirink, and E. V. Der, "EURADCLIM: The European climatological high-resolution gauge-adjusted radar precipitation dataset," *Open Access Earth Syst. Sci. Data Discuss.*, vol. 2022, pp. 1–34, Nov. 2022.

[4] S. Walsh, "New long-term rainfall averages for Ireland," in *Proc. Irish Nat. Hydrol. Conf.*, 2012, pp. 3–12.

[5] D. Fitgerald and F. Forrestal, "Monthly and annual averages of rainfall for Ireland 1961–1990," 1996. [Online]. Available: http://hdl.handle.net/2262/70486

[6] S. Jebson, "Fact sheet number 3: Water in the atmosphere," *Nat. Meteorological Library Arch.*, pp. 1–27, 2007. [Online]. Available: https://www.metoffice.gov.uk/binaries/content/assets/metofficegovuk/pdf/research/library-and-archive/library/publications/factsheets/factsheet_3-water-in-the-atmosphere.pdf

[7] L. Le Cam, "A stochastic description of precipitation," in *Proc. 4th Berkeley Symp. Math. Statist. Probab.* Berkeley, CA, USA: Univ. of California, Berkeley, 1961, pp. 165–186.

[8] J. O. Laws and D. A. Parsons, "Relationship of raindrop size and intensity," *Eos, Trans. Amer. Geophys. Union*, vol. 24, no. 2, pp. 452–460, 1943.

[9] P. W. M. Marshall, "The distribution of raindrop size," *J. Atmos. Sci.*, vol. 5, no. 4, pp. 148–162, 1948.

[10] K. Garg and S. K. Nayar, "Photorealistic rendering of rain streaks," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 996–1002, Jul. 2006.

[11] K. V. Beard and C. Chuang, "A new model for the equilibrium shape of raindrops," *J. Atmos. Sci.*, vol. 44, no. 11, pp. 1509–1524, Jun. 1987.

[12] R. Gunn and G. D. Kinzer, "The terminal velocity of fall for water droplets in stagnant air," *J. Meteorol.*, vol. 6, no. 4, pp. 243–248, Aug. 1949.

[13] G. G. Rodgers, G. Poots, J. K. Page, and W. M. Pickering, "Theoretical predictions of raindrop impaction on a slab type building," *Building Sci.*, vol. 9, no. 3, pp. 181–190, Sep. 1974.

[14] M. N. Chowdhury, F. Y. Testik, M. C. Hornack, and A. A. Khan, "Free fall of water drops in laboratory rainfall simulations," *Atmos. Res.*, vol. 168, pp. 158–168, Feb. 2016.

[15] G. B. Foote and P. S. D. Toit, "Terminal velocity of raindrops aloft," *J. Appl. Meteorol.*, vol. 8, no. 2, pp. 249–253, Apr. 1969.

[16] K. Garg and S. K. Nayar, "Vision and rain," *Int. J. Comput. Vis.*, vol. 75, no. 1, pp. 3–27, Jul. 2007.

[17] K. Garg and S. K. Nayar, "Detection and removal of rain from videos," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Washington, DC, USA, 2004, p. 1.

[18] E. Villermaux and F. Eloi, "The distribution of raindrops speeds," *Geophys. Res. Lett.*, vol. 38, no. 19, pp. 10–13, 2011.

[19] B. J. Mason and J. B. Andrews, "Drop-size distributions from various types of rain," *Quart. J. Roy. Meteorological Soc.*, vol. 86, no. 369, pp. 346–353, Jul. 1960.

[20] M. Colomb, P. Duthon, and S. Laukkanen, "Characteristics of adverse weather conditions," DENSE, CER, Brussels, Belgium, 2017. [Online]. Available: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5b0632e8d&appId=PPGMS

[21] B. E. Bayer, "Color imaging array," U.S. Patent 3 971 065, Jul. 20, 1976. [Online]. Available: https://patents.google.com/patent/US3971065

[22] S. A. Twomey, C. F. Bohren, and J. L. Mergenthaler, "Reflectance and albedo differences between wet and dry surfaces," *Appl. Opt.*, vol. 25, no. 3, pp. 431–437, 1986.

[23] A. Ångström, "The albedo of various surfaces of ground," *Geografiska Annaler*, vol. 7, no. 4, pp. 323–342, Dec. 1925.

[24] J. Lekner and M. C. Dorf, "Why some things are darker when wet," *Appl. Opt.*, vol. 27, no. 7, pp. 1278–1280, Apr. 1988.

[25] M. Senthilkumar and N. Selvakumar, "Parameters involved in assessment of dry state colour of fabrics from their wet state colour," *Indian J. Fibre Textile Res.*, vol. 32, no. 1, pp. 126–134, 2007.

[26] J. H. Lee, S. H. Kim, K. J. Lee, D. Y. Lim, and H. Y. Jeon, "Determining the absorption properties of split-type microfiber fabrics by measuring the change in color depth," *Textile Res. J.*, vol. 74, no. 3, pp. 271–278, Mar. 2004.

[27] S. K. Nayar and S. G. Narasimhan, "Vision in bad weather," in *Proc. 7th IEEE Int. Conf. Comput. Vis.*, Sep. 1999, pp. 820–827.

[28] H. Koschmieder, "Theory of horizontal visibility," *Contribution Phys. Free Atmos.*, vol. 12, pp. 33–53, 1924.

[29] Z. Lee and S. Shang, "Visibility: How applicable is the century-old Koschmieder model?" *J. Atmos. Sci.*, vol. 73, no. 11, pp. 4573–4581, Nov. 2016.
[30] M. Colomb, P. Duthon, and F. Bernardin, "Spectral reflectance characterization of the road environment to optimize the choice of autonomous vehicle sensors," in *Proc. IEEE Intell. Transp. Syst. Conf. (ITSC)*, Oct. 2019, pp. 1085–1090.
[31] N. Hautiere, E. Dumont, R. Bremond, and V. Ledoux, "Review of the mechanisms of visibility reduction by rain and wet road," in *Proc. 8th Int. Symp. Automot. Lighting*, 2009, pp. 445–455.
[32] C. Nolet, A. Poortinga, P. Roosjen, H. Bartholomeus, and G. Ruessink, "Measuring and modeling the effect of surface moisture on the spectral reflectance of coastal beach sand," *PLoS ONE*, vol. 9, Nov. 2014, Art. no. e112151.
[33] S. G. Narasimhan and S. K. Nayar, "Contrast restoration of weather degraded images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 6, pp. 713–724, Jun. 2003.
[34] T. E. Burghardt, R. Popp, B. Helmreich, T. Reiter, G. Böhm, G. Pitterle, and M. Artmann, "Visibility of various road markings for machine vision," *Case Stud. Construct. Mater.*, vol. 15, Dec. 2021, Art. no. e00579.
[35] E. Peli, "Contrast in complex images," *J. Opt. Soc. Amer. A, Opt. Image Sci.*, vol. 7, no. 10, pp. 2032–2040, 1990.
[36] A. Rasouli, I. Kotseruba, and J. K. Tsotsos, "It's not all about size: On the role of data properties in pedestrian detection," in *Proc. ECCV* (Lecture Notes in Computer Science), vol. 11129, 2018, pp. 210–225.
[37] C. G. Pooley, D. Horton, G. Scheldeman, M. Tight, T. Jones, A. Chisholm, H. Harwatt, and A. Jopson, "Household decision-making for everyday travel: A case study of walking and cycling in Lancaster (UK)," *J. Transp. Geography*, vol. 19, no. 6, pp. 1601–1607, Nov. 2011.
[38] 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.
[39] R. Kulmala and M. Salusjarvi, "The effects of weather conditions on pedestrians' risk taking," *Central Org. Traffic Saf.*, vol. 87, 1988. [Online]. Available: https://trid.trb.org/view/283179
[40] L. de Montigny, R. Ling, and J. Zacharias, "The effects of weather on walking rates in nine cities," *Environ. Behav.*, vol. 44, no. 6, pp. 821–840, Nov. 2012.
[41] K. J. Clifton and A. D. Livi, "Gender differences in walking behavior, attitudes about walking, and perceptions of the environment in three Maryland communities," *Res. Women'S Issues Transp.*, vol. 2, pp. 79–88, Feb. 2005.
[42] L. Aultman-Hall, D. Lane, and R. R. Lambert, "Assessing impact of weather and season on pedestrian traffic volumes," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2140, no. 1, pp. 35–43, Jan. 2009.
[43] L. F. Miranda-Moreno and T. Nosal, "Weather or not to cycle: Temporal trends and impact of weather on cycling in an urban environment," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2247, no. 1, pp. 42–52, Jan. 2011.
[44] A. B. R. Anthonisamy and A. V. S. James, "Formulation of atmospheric optical attenuation model in terms of weather data," *J. Opt.*, vol. 45, no. 2, pp. 120–135, Jun. 2016.
[45] A. T. Young, "Rayleigh scattering," *Appl. Opt.*, vol. 20, no. 4, pp. 533–535, 1981.
[46] R. M. Drake and J. E. Gordon, "Mie scattering," *Amer. J. Phys.*, vol. 53, no. 10, pp. 955–962, 1985.
[47] D. Jackel and B. Walter, "Modeling and rendering of the atmosphere using Mie-scattering," *Comput. Graph. Forum*, vol. 16, no. 4, pp. 201–210, Oct. 1997.
[48] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 7th ed. Cambridge, U.K.: Cambridge Univ. Press, 1999.
[49] M. Inakage, "An illumination model for atmospheric environments," in *New Advances in Computer Graphics*. Tokyo, Japan: Springer, 1989, pp. 533–548.
[50] A. Tokay and K. Beard, "A field study of raindrop oscillations," *J. Appl. Meteorol. Climatol.*, vol. 148, pp. 148–162, Jan. 1996.
[51] D. Abe, E. Segawa, O. Nakayama, M. Shiohara, S. Sasaki, N. Sugano, and H. Kanno, "Robust small-object detection for outdoor wide-area surveillance," *IEICE Trans. Inf. Syst.*, vol. 91, no. 7, pp. 1922–1928, Jul. 2008.
[52] N. Hashimoto, S. Thompson, and Y. Takinami, "Evaluation experiments for lane keeping assistance system under several conditions," in *Proc. Int. Conf. Comput. Sci. Comput. Intell. (CSCI)*, Dec. 2018, pp. 442–445.
[53] C. Resnick, O. Litany, A. Kar, K. Kreis, J. Lucas, K. Cho, and S. Fidler, "Causal BERT: Improving object detection by searching for challenging groups," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2021, pp. 2972–2981.
[54] B. Yu and Y. Sun, "Simulation of impact of water-film spray on visibility," *J. Transp. Eng., B, Pavements*, vol. 144, no. 4, Dec. 2018, Art. no. 04018054.
[55] F. Sahin, "Long-range, high-resolution camera optical design for assisted and autonomous driving," *Photonics*, vol. 6, no. 2, p. 73, Jun. 2019.
[56] M. Uricár, G. Sistu, H. Rashed, A. Vobecký, V. R. Kumar, P. Krížek, F. Burger, and S. Yogamani, "Let's get dirty: GAN based data augmentation for camera lens soiling detection in autonomous driving," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2021, pp. 766–775.
[57] F. Al Machot, M. Ali, A. H. Mosa, C. Schwarzlmuller, M. Gutmann, and K. Kyamakya, "Real-time raindrop detection based on cellular neural networks for ADAS," *J. Real-Time Image Process.*, vol. 16, no. 4, pp. 931–943, Aug. 2019.
[58] S. You, R. T. Tan, R. Kawakami, Y. Mukaigawa, and K. Ikeuchi, "Adherent raindrop modeling, detection and removal in video," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 9, pp. 1721–1733, Sep. 2016.
[59] L. Gao and T. J. McCarthy, "Wetting 101°," *Langmuir*, vol. 25, no. 24, pp. 14105–14115, Dec. 2009.
[60] M. Roser and A. Geiger, "Video-based raindrop detection for improved image registration," in *Proc. IEEE 12th Int. Conf. Comput. Vis. Workshops*, Sep. 2009, pp. 570–577.
[61] M. Roser, J. Kurz, and A. Geiger, "Realistic modeling of water droplets for monocular adherent raindrop recognition using Bézier curves," in *Proc. Asian Conf. Comput. Vis.* Cham, Switzerland: Springer, 2010, pp. 235–244.
[62] H. Kurihata, T. Takahashi, I. Ide, Y. Mekada, H. Murase, Y. Tamatsu, and T. Miyahara, "Rainy weather recognition from in-vehicle camera images for driver assistance," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2005, pp. 205–210.
[63] D. Eigen, D. Krishnan, and R. Fergus, "Restoring an image taken through a window covered with dirt or rain," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 633–640.
[64] A. Yamashita, M. Kuramoto, T. Kaneko, and K. T. Miura, "A virtual wiper-restoration of deteriorated images by using multiple cameras," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2003, pp. 3126–3131.
[65] A. Yamashita, Y. Tanaka, and T. Kaneko, "Removal of adherent waterdrops from images acquired with stereo camera," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Aug. 2005, pp. 400–405.
[66] A. Yamashita, I. Fukuchi, T. Kaneko, and K. T. Miura, "Removal of adherent noises from image sequences by spatio-temporal image removal of adherent noises from image sequences by spatio-temporal," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2008, pp. 2386–2391.
[67] A. Yamashita, T. Harada, T. Kaneko, and K. T. Miura, "Removal of adherent noises from images of dynamic scenes by using a pan-tilt camera," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2004, pp. 437–442.
[68] M. Bertalmio, L. Vese, G. Sapiro, and S. Osher, "Simultaneous structure and texture image inpainting," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 882–889, Aug. 2003.
[69] K. Garg and S. K. Nayar, "When does a camera see rain?" in *Proc. 10th IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2005, pp. 1067–1074.
[70] Y. Monobe, H. Yamashita, T. Kurosawa, and H. Kotera, "High dynamic range compression for digital video camera using local contrast enhancement," in *Dig. Tech. Papers, Int. Conf. Consum. Electron.*, 2005, pp. 217–218.
[71] V. Singh, S. K. S. Hari, T. Tsai, and M. Pitale, "Simulation driven design and test for safety of AI based autonomous vehicles," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2021, pp. 122–128.
[72] (2018). *IMX490: CMOS Image Sensor for Automotive*. Sony. [Online]. Available: https://www.sony-semicon.com/files/62/pdf/p-15_IMX490.pdf
[73] (2022). *Isocell-3b6: Automotve Image Sensor*. Samsung. [Online]. Available: https://semiconductor.samsung.com/image-sensor/automotive-image-sensor/isocell-3b6/
[74] D. W. Hertel and E. Chang, "Image quality standards in automotive vision applications," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2007, pp. 404–409.

[75] L. Yahiaoui, J. Horgan, B. Deegan, S. Yogamani, C. Hughes, and P. Denny, "Overview and empirical analysis of ISP parameter tuning for visual perception in autonomous driving," *J. Imag.*, vol. 5, no. 10, p. 78, Sep. 2019.

[76] T. Surasak, I. Takahiro, C. Cheng, C. Wang, and P. Sheng, "Histogram of oriented gradients for human detection in video," in *Proc. 5th Int. Conf. Bus. Ind. Res. (ICBIR)*, May 2018, pp. 172–176.

[77] P. Viola, M. J. Jones, and D. Snow, "Detecting pedestrians using patterns of motion and appearance," *Int. J. Comput. Vis.*, vol. 63, no. 2, pp. 153–161, Jul. 2005.

[78] Y. Liu, S. Shan, W. Zhang, X. Chen, and W. Gao, "Granularity-tunable gradients partition (GGP) descriptors for human detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 1255–1262.

[79] L. Yeffet and L. Wolf, "Local trinary patterns for human action recognition," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 492–497.

[80] W. C. Chia, L. S. Yeong, S. I. Ch'ng, K. P. Seng, and L. Ang, "The effect of rainfall on feature points extraction and image stitching," in *Proc. Int. Conf. Inf. Sci., Electron. Electr. Eng.*, vol. 3, Apr. 2014, pp. 1382–1386.

[81] A. Bochkovskiy, C.-Y. Wang, and H.-Y. Mark Liao, "YOLOv4: Optimal speed and accuracy of object detection," 2020, *arXiv:2004.10934*.

[82] 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.*, vol. 28, 2015, pp. 1–9.

[83] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 145–151, May 2012.

[84] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.

[85] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. 3rd Int. Conf. Learn. Represent.*, 2015, pp. 1–14.

[86] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, and E. K. Oermann, "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study," *PLoS Med.*, vol. 15, no. 11, pp. 1–17, 2018.

[87] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *Proc. 2nd Int. Conf. Learn. Represent.*, 2014, pp. 1–10.

[88] M. A. Alcorn, Q. Li, Z. Gong, C. Wang, L. Mai, W. Ku, and A. Nguyen, "Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4840–4849.

[89] 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.

[90] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," 2018, *arXiv:1804.02767*.

[91] S. Dodge and L. Karam, "Understanding how image quality affects deep neural networks," in *Proc. 8th Int. Conf. Quality Multimedia Exper. (QoMEX)*, Jun. 2016, pp. 1–6.

[92] G. B. P. Da Costa, W. A. Contato, T. S. Nazare, J. E. S. B. Neto, and M. Ponti, "An empirical study on the effects of different types of noise in image classification tasks," 2016, *arXiv:1609.02781*.

[93] R. Song, J. Wetherall, S. Maskell, and J. Ralph, "Weather effects on obstacle detection for autonomous car," in *Proc. 6th Int. Conf. Vehicle Technol. Intell. Transp. Syst.*, 2020, pp. 331–341.

[94] G. Volk, S. Muller, A. V. Bernuth, D. Hospach, and O. Bringmann, "Towards robust CNN-based object detection through augmentation with synthetic rain variations," in *Proc. IEEE Intell. Transp. Syst. Conf. (ITSC)*, Oct. 2019, pp. 285–292.

[95] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, "BDD100K: A diverse driving dataset for heterogeneous multitask learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 2633–2642.

[96] W. Yang, R. T. Tan, J. Feng, J. Liu, Z. Guo, and S. Yan, "Deep joint rain detection and removal from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1685–1694.

[97] X. Zhang, H. Li, Y. Qi, W. Leow, and T. Ng, "Rain removal in video by combining temporal and chromatic properties," in *Proc. IEEE Int. Conf. Multimedia Expo*, Jul. 2006, pp. 461–464.

[98] P. C. Barnum, S. Narasimhan, and T. Kanade, "Analysis of rain and snow in frequency space," *Int. J. Comput. Vis.*, vol. 86, nos. 2–3, pp. 256–274, Jan. 2010.

[99] X. Xue, X. Jin, C. Zhang, and S. Goto, "Motion robust rain detection and removal from videos," in *Proc. IEEE 14th Int. Workshop Multimedia Signal Process. (MMSP)*, Sep. 2012, pp. 170–174.

[100] C. Lai and C. G. Li, "Video-based windshield rain detection and wiper control using holistic-view deep learning," in *Proc. IEEE 15th Int. Conf. Autom. Sci. Eng. (CASE)*, Aug. 2019, pp. 1060–1065.

[101] S. Gormer, A. Kummert, S.-B. Park, and P. Egbert, "Vision-based rain sensing with an in-vehicle camera," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2009, pp. 279–284.

[102] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.

[103] A. Vasile, I. Vasile, A. Nistor, L. Vladareanu, M. Pantazica, F. Caldararu, A. Bonea, A. Drumea, and I. Plotog, "Rain sensor for automatic systems on vehicles," *Proc. SPIE*, vol. 7821, pp. 509–514, Dec. 2010.

[104] 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, Nov. 2023, Art. no. 109740.

[105] W. Yang, R. T. Tan, S. Wang, Y. Fang, and J. Liu, "Single image deraining: From model-based to data-driven and beyond," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 4059–4077, Nov. 2021.

[106] 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, Jan. 2022.

[107] X. Li, J. Wu, Z. Lin, H. Liu, and H. Zha, "Recurrent squeeze-and-excitation context aggregation net for single image deraining," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 254–269.

[108] T. Wang, X. Yang, K. Xu, S. Chen, Q. Zhang, and R. W. H. Lau, "Spatial attentive single-image deraining with a high quality real rain dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 12262–12271.

[109] S. Li, I. B. Araujo, W. Ren, Z. Wang, E. K. Tokuda, R. H. Junior, R. Cesar-Junior, J. Zhang, X. Guo, and X. Cao, "Single image deraining: A comprehensive benchmark analysis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3833–3842.

[110] Y. Pei, Y. Huang, Q. Zou, Y. Lu, and S. Wang, "Does haze removal help CNN-based image classification?" in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 682–697.

[111] S. Halder, J. Lalonde, and R. D. Charette, "Physics-based rendering for improving robustness to rain," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 10202–10211.

[112] R. de Charette, R. Tamburo, P. C. Barnum, A. Rowe, T. Kanade, and S. G. Narasimhan, "Fast reactive control for illumination through rain and snow," in *Proc. IEEE Int. Conf. Comput. Photography (ICCP)*, Apr. 2012, pp. 1–10.

[113] Y. Weber, V. Jolivet, G. Gilet, and D. Ghazanfarpour, "A multiscale model for rain rendering in real-time," *Comput. Graph.*, vol. 50, pp. 61–70, Aug. 2015.

[114] P. Rousseau, V. Jolivet, and D. Ghazanfarpour, "Realistic real-time rain rendering," *Comput. Graph.*, vol. 30, no. 4, pp. 507–518, Aug. 2006.

[115] S. Ni, X. Cao, T. Yue, and X. Hu, "Controlling the rain: From removal to rendering," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 6324–6333.

[116] Y. Wei, Z. Zhang, Y. Wang, M. Xu, Y. Yang, S. Yan, and M. Wang, "DerainCycleGAN: Rain attentive CycleGAN for single image deraining and rainmaking," *IEEE Trans. Image Process.*, vol. 30, pp. 4788–4801, 2021.

[117] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.

[118] H. Zhang and V. M. Patel, "Density-aware single image de-raining using a multi-stream dense network," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 695–704.

[119] H. Zhang, V. Sindagi, and V. M. Patel, "Image de-raining using a conditional generative adversarial network," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 11, pp. 3943–3956, Nov. 2020.

[120] M. Tremblay, S. S. Halder, R. de Charette, and J.-F. Lalonde, "Rain rendering for evaluating and improving robustness to bad weather," *Int. J. Comput. Vis.*, vol. 129, no. 2, pp. 341–360, Feb. 2021.

[121] P. Xiao, Z. Shao, S. Hao, Z. Zhang, X. Chai, J. Jiao, Z. Li, J. Wu, K. Sun, K. Jiang, Y. Wang, and D. Yang, "PandaSet: Advanced sensor suite dataset for autonomous driving," in *Proc. IEEE Int. Intell. Transp. Syst. Conf. (ITSC)*, Sep. 2021, pp. 3095–3101.

[122] J. Jin, A. Fatemi, W. M. P. Lira, F. Yu, B. Leng, R. Ma, A. Mahdavi-Amiri, and H. Zhang, "RaidaR: A rich annotated image dataset of rainy street scenes," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Oct. 2021, pp. 2951–2961.

[123] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "NuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11618–11628.

[124] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, and V. Vasudevan, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 2443–2451.

[125] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Muhlegg, S. Dorn, T. Fernandez, M. Jänicke, S. Mirashi, C. Savani, M. Sturm, O. Vorobiov, M. Oelker, S. Garreis, and P. Schuberth, "A2D2: Audi autonomous driving dataset," 2020, *arXiv:2004.06320*.

[126] S. Yogamani, C. Hughes, J. Horgan, G. Sistu, S. Chennupati, M. Uricar, S. Milz, M. Simon, K. Amende, C. Witt, H. Rashed, S. Nayak, S. Mansoor, P. Varley, X. Perrotton, D. Odea, and P. Pérez, "WoodScape: A multi-task, multi-camera fisheye dataset for autonomous driving," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9307–9317.

[127] A. Patil, S. Malla, H. Gang, and Y. Chen, "The H3D dataset for full-surround 3D multi-object detection and tracking in crowded urban scenes," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 9552–9557.

[128] K. Behrendt, "Boxy vehicle detection in large images," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 840–846.

[129] J. Xue, J. Fang, T. Li, B. Zhang, P. Zhang, Z. Ye, and J. Dou, "BLVD: Building a large-scale 5D semantics benchmark for autonomous driving," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 6685–6691.

[130] M. Braun, S. Krebs, F. Flohr, and D. M. Gavrila, "EuroCity persons: A novel benchmark for person detection in traffic scenes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 1844–1861, Aug. 2019.

[131] L. Neumann, M. Karg, S. Zhang, C. Scharfenberger, E. Piegert, S. Mistr, O. Prokofyeva, R. Thiel, A. Vedaldi, and A. Zisserman, "Nightowls: A pedestrians at night dataset," in *Proc. Asian Conf. Comput. Vis.* Cham, Switzerland: Springer, 2018, pp. 691–705.

[132] A. Rasouli, I. Kotseruba, T. Kunic, and J. Tsotsos, "PIE: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6261–6270.

[133] Y. Choi, N. Kim, S. Hwang, K. Park, J. S. Yoon, K. An, and I. S. Kweon, "KAIST multi-spectral day/night data set for autonomous and assisted driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 3, pp. 934–948, Mar. 2018.

[134] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, "The ApolloScape dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 954–960.

[135] K. Behrendt, L. Novak, and R. Botros, "A deep learning approach to traffic lights: Detection, tracking, and classification," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2017, pp. 1370–1377.

[136] S. Zhang, R. Benenson, and B. Schiele, "CityPersons: A diverse dataset for pedestrian detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4457–4465.

[137] F. Tung, J. Chen, L. Meng, and J. J. Little, "The raincouver scene parsing benchmark for self-driving in adverse weather and at night," *IEEE Robot. Autom. Lett.*, vol. 2, no. 4, pp. 2188–2193, Oct. 2017.

[138] G. Neuhold, T. Ollmann, S. R. Bulo, and P. Kontschieder, "The mapillary vistas dataset for semantic understanding of street scenes," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5000–5009.

[139] I. Kotseruba, A. Rasouli, and J. K. Tsotsos, "Joint attention in autonomous driving (JAAD)," 2016, *arXiv:1609.04741*.

[140] X. Li, F. Flohr, Y. Yang, H. Xiong, M. Braun, S. Pan, K. Li, and D. M. Gavrila, "A new benchmark for vision-based cyclist detection," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2016, pp. 1028–1033.

[141] S. Hwang, J. Park, N. Kim, Y. Choi, and I. S. Kweon, "Multispectral pedestrian detection: Benchmark dataset and baseline," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1037–1045.

[142] A. Omid-Zohoor, D. Ta, and B. Murmann, "PASCALRAW: Raw image database for object detection," Stanford Digit. Repository, 2015. [Online]. Available: http://purl.stanford.edu/hq050zr7488

[143] S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing, and C. Igel, "Detection of traffic signs in real-world images: The German traffic sign detection benchmark," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Aug. 2013, pp. 1–8.

[144] C. Caraffi, T. Vojír, J. Trefný, J. Sochman, and J. Matas, "A system for real-time detection and tracking of vehicles from a single car-mounted camera," in *Proc. 15th Int. IEEE Conf. Intell. Transp. Syst.*, Sep. 2012, pp. 975–982.

[145] C. Wojek, S. Walk, and B. Schiele, "Multi-cue onboard pedestrian detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 794–801.

[146] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: A benchmark," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 304–311.

[147] M. Enzweiler and D. M. Gavrila, "Monocular pedestrian detection: Survey and experiments," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 12, pp. 2179–2195, Dec. 2009.

[148] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Commun. ACM*, vol. 63, no. 11, pp. 139–144, 2020.

[149] M. Uricar, P. Krizek, D. Hurych, I. Sobh, S. Yogamani, and P. Denny, "Yes, we GAN: Applying adversarial techniques for autonomous driving," 2019, *arXiv:1902.03442*.

TIM BROPHY received the B.Eng. degree from the University of Galway, in 2018, where he is currently pursuing the Ph.D. degree. He is also a member of the Connaught Automotive Research (CAR) Group, under the supervision of Prof. Edward Jones and Prof. Martin Glavin. His research interests include computer vision and sensor availability with 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, 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 has managed a research program and has co-supervised six Ph.D. student projects, which involved 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 an Adjunct Lecturer with the School of Engineering, University of Galway. He received the 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, Ireland, in 2021. From 2012 to 2016, he was a Software Developer, a Senior Researcher, and a Technical Project Manager of several national projects with the Iran Telecommunication Research Center (ITRC). He is currently leading and conducting research in the development of signal and image processing for advanced sensors and

machine vision technologies in a range of problems in connected and autonomous vehicles with Xperi Company. He received the Government of Ireland Postgraduate Research Scholarship (IRC) for his studies. His research interests include signal processing, machine learning, and algorithm design.

JONATHAN HORGAN is a Manager of computer vision and deep learning architecture and a Senior Expert with Valeo Vision Systems. He has worked in the field of computer vision for over 16 years with a focus over the last ten years on automotive computer vision for advanced driver assistance systems (ADAS), automated parking, and automated driving. He is currently working on next-generation advanced computer vision and deep learning with the ultimate goal of achieving fully autonomous driving and parking. He has 25 publications in peer-reviewed journals and conferences and over 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, in 1999, and the M.Eng.Sc. degree in research in electronic engineering, with a focus on biomedical electronics, in 2002.

He is responsible for defining the camera product roadmap for surround and automated driving applications with Valeo. He was with key technology experts across the supply chain and within OEMs to define optimal system architectures. He was a Lecturer of electronics and computing systems with Atlantic Technological University, Ireland, for a number of years. Later, he moved to 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, Ireland, in 1993, and the M.Sc. degree in mathematics and the Ph.D. degree in physics from the University of Galway, Ireland, in 1994 and 2000, respectively.

He was researching electromagnetic planetary physics with GFZ Potsdam, Germany. From 1999 to 2001, he was an RF Engineer with AVM GmbH, Germany, designed and developed the hardware for the first integrated GSM/ISDN/USB modem. After working in supercomputing development at Compaq/Hewlett Packard, from 2001 to 2002, he joined Connaught Electronics Ltd. (subsequently Valeo), Tuam, Ireland, as a Team Leader of radiofrequency design. Over the next 20 years, he is a Senior Expert with Valeo, designing and developing novel radiofrequency and imaging systems, including the first mass production high dynamic range automotive camera, for leading car companies, including BMW, Daimler, and Land Rover. In 2010, he became an Adjunct Professor of automotive electronics with the University of Galway. In 2022, he became a Lecturer of artificial intelligence with the Department of Electronic and Computer Engineering, University of Limerick, Ireland. His research interests include automotive imaging technology and its extension into other domains, algorithmic design, artificial intelligence, applied mathematics, and the industrialization of advanced technologies.

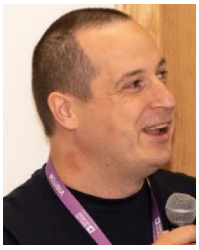

CIARÁN (HUGHES) EISING (Senior Member, IEEE) received the degree in electronic and computer engineering and the Ph.D. degree from the University of Galway, in 2003 and 2010, respectively. From 2009 to 2020, he was a Computer Vision Team Lead and an Architect with Valeo Vision Systems, where he also held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

BRIAN DEEGAN received the bachelor's degree in computer engineering from the University of Limerick, in 2004, the M.Sc. degree in biomedical engineering from the University of Limerick, in 2005, and the Ph.D. degree in biomedical engineering from the University of Galway, in 2011. The focus of his research was the relationship between blood pressure and cerebral blood flow in humans. From 2011 to 2022, he was with Valeo Vision Systems, as a Vision Research Engineer, focused on image quality. In 2022, he joined the Department of Electrical and Electronic Engineering, University of Galway, as a Lecturer and a Researcher. His research interests include high dynamic range imaging, LED flicker, top-view 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 algorithms and architectures for high-speed data communications systems from the University of Galway, Ireland, in 1997 and 2004, respectively, and the Higher Diploma degree in third level education, in 2007. He was a lecturer (fixed term contract), from September 1999 to December 2003, and became a permanent member of the academic staff, in January 2004. He is currently the Joint Director of the Connaught Automotive Research (CAR) Group, University of Galway. He is also a Funded Investigator of Lero, the Irish Software Research Centre. He has a number of Ph.D. students and postdoctoral researchers in collaboration with industry in the areas of signal processing and embedded systems for automotive and agricultural applications.

EDWARD JONES (Senior Member, IEEE) received the B.E. and Ph.D. degrees in electronic engineering from the University of Galway, Ireland. His Ph.D. research topic was on the development of computational auditory models for speech processing. He is currently a Professor of electrical and electronic engineering with the School of Engineering, University of Galway. From 2009 to 2010, he was a Visiting Researcher with the Department of Electrical Engineering, Columbia University, New York, NY, USA; and a Visiting Fellow with the School of Electrical Engineering and Telecommunications, The University of New South Wales, Sydney, Australia. He 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 DSP algorithm development and embedded implementation for applications in biomedical engineering, speech and audio processing, and image processing. He is a Chartered Engineer and a fellow of the Institution of Engineers of Ireland.