



**HAL**  
open science

# MUAD: Multiple Uncertainties for Autonomous Driving, a benchmark for multiple uncertainty types and tasks

Gianni Franchi, Xuanlong Yu, Andrei Bursuc, Angel Tena, Rémi Kazmierczak, Séverine Dubuisson, Emanuel Aldea, David Filliat

## ► To cite this version:

Gianni Franchi, Xuanlong Yu, Andrei Bursuc, Angel Tena, Rémi Kazmierczak, et al.. MUAD: Multiple Uncertainties for Autonomous Driving, a benchmark for multiple uncertainty types and tasks. 33rd British Machine Vision Conference, Nov 2022, London, United Kingdom. hal-04454911

**HAL Id: hal-04454911**

**<https://hal.science/hal-04454911>**

Submitted on 13 Feb 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# MUAD: Multiple Uncertainties for Autonomous Driving, a benchmark for multiple uncertainty types and tasks

Gianni Franchi<sup>1†</sup>

gianni.franchi@ensta-paris.fr

Xuanlong Yu<sup>1,2†</sup>

xuanlong.yu@universite-paris-saclay.fr

Andrei Bursuc<sup>3</sup>

andrei.bursuc@valeo.com

Angel Tena<sup>4</sup>

angel.tena@anyverse.ai

Rémi Kazmierczak<sup>1</sup>

remi.kazmierczak@ensta-paris.fr

Séverine Dubuisson<sup>5</sup>

severine.dubuisson@lis-lab.fr

Emanuel Aldea<sup>2</sup>

emanuel.aldea@universite-paris-saclay.fr

David Filliat<sup>1</sup>

david.filliat@ensta-paris.fr


<sup>1</sup> U2IS, ENSTA Paris, IP Paris

<sup>2</sup> SATIE, Paris-Saclay University

<sup>3</sup> valeo.ai

<sup>4</sup> Anyverse

<sup>5</sup> Aix Marseille University



(a) No perturbation

(b) With OOD instances

(c) Level 1 perturbation

(d) Level 2 perturbation

Figure 1: **Snapshots from the MUAD dataset** showing different types of adverse conditions and events to evaluate perception models (OOD : Out Of Distribution, i.e., not seen during training).

## Abstract

Predictive uncertainty estimation is essential for safe deployment of Deep Neural Networks in real-world autonomous systems. However, disentangling the different types and sources of uncertainty is non trivial for most datasets, especially since there is no ground truth for uncertainty. In addition, while adverse weather conditions of varying intensities can disrupt neural network predictions, they are usually under-represented in both training and test sets in public datasets. We attempt to mitigate these setbacks and

introduce the MUAD dataset (Multiple Uncertainties for Autonomous Driving), consisting of 10,413 realistic synthetic images with diverse adverse weather conditions (night, fog, rain, snow), out-of-distribution objects and annotations for semantic segmentation, depth estimation, object and instance detection. MUAD allows to better assess the impact of different sources of uncertainty on model performance. We conduct a thorough experimental study of this impact on several baseline Deep Neural Networks across multiple tasks, and release our dataset to allow researchers to benchmark their algorithm methodically in adverse conditions. More visualizations and the download link for MUAD are available at <https://muad-dataset.github.io/>.

## 1 Introduction

In recent years, Deep Neural Networks (DNNs) have achieved remarkable results in various computer vision tasks [9, 57, 59]. This has turned DNNs into an essential tool for effective automatic perception. Although DNNs achieve outstanding performance across benchmarks and tasks, there are still a few major bottlenecks to solve before a widespread deployment. One of the most frequent and known criticisms of DNNs is related to their lack of reliability under varying levels of shifts in the data distribution, and it became crucial to address this limitation. To achieve this, we focus on studying the uncertainties of the DNN predictions for computer vision tasks. The predictive uncertainty of a DNN stems from two main types of uncertainty [52]: *aleatoric* and *epistemic*. The former is related to randomness of the world and of the sensing system, typically instantiated as noise in the data. The latter concerns finite size training datasets. The epistemic uncertainty captures the uncertainty in the DNN parameters and their lack of knowledge on the model that generated the training data. In spite of their simple and intuitive definitions, the sources of uncertainty are notoriously hard to separate in most datasets, where data are typically curated and various outlier or noisy samples are removed before annotation.

For autonomous driving, uncertainty estimation and reliability are essential for safely deploying DNNs in real-world conditions. Here, DNNs are expected not only to reach high predictive performance and real-time inference speed, but also to deal effectively with the two types of uncertainty under various forms (noise, distribution shift, out-of-distribution samples, sensor degradation, etc.). In the last years, numerous works have moved the needle towards more reliable predictive uncertainty for DNNs [0, 9, 18, 22, 36, 38, 46, 47, 59]. However, evaluating such methods is not obvious as there is no ground truth for uncertainty and the different sources of uncertainty are conflated due to prior data curation.

We introduce a new dataset to study uncertainty estimation methods for perception in autonomous vehicles. While most datasets aim to improve the predictive performance of DNNs [13, 24, 49, 72], only recently datasets addressed the robustness of DNNs under unseen weather conditions [15, 59, 60] or objects [8, 7, 30]. However, these datasets are either limited to only one task, typically semantic segmentation, or only focus on a single type of uncertainty, or are not being precise enough in the different levels of uncertainties. We address these limitations in our dataset that allows to quantify all levels of uncertainty in the same conditions. Our dataset, MUAD (Multiple Uncertainties for Autonomous Driving) is composed of 3,420 images for training, 492 for validation, and 6,501 for testing.

To summarize, our contributions are as follows: (1) We introduce MUAD: a new automotive dataset with annotations for multiple tasks and multiple uncertainty sources. (2) We perform a wide range of benchmarks on MUAD dataset for multiple computer vision tasks and settings (semantic segmentation, depth estimation, object detection) to further support

research in this area. (3) We conduct an extensive study on uncertainty quantification for 2D output tasks for recent Transformer-based architectures.

## 2 Related work

### 2.1 Datasets

A variety of real-world datasets for autonomous driving have been recently released [8, 9, 13, 24, 33, 53, 54, 56, 72]. They have enabled tremendous progress in the area but they typically focus on a single task, e.g., semantic segmentation [13, 53, 72], object detection [8, 24, 54], motion prediction [8, 33] and do not have evaluation tracks for uncertainty and out-of-distribution detection. Synthetic datasets, e.g., GTA-V [56], SYNTHIA [57], virtual KITTI [20] can provide abundant training data alleviating the need for costly annotation of real images as well as privacy preservation concerns in the case of real data. Currently, they are mostly designed and used for domain adaptation, typically imitating the content and classes from a given real dataset. Several datasets have emerged towards meeting the reliability requirement for self-driving vehicles [8, 9, 30, 54] and evaluate the performance of semantic segmentation DNNs when facing out-of-distribution objects (OOD). Other datasets investigate the robustness against different weather conditions, e.g., night [14, 15, 50], rain [50, 54], fog [58, 50], however they are often acquired in different locations and conditions leading to a performance drop that overlaps with the one from the difficult weather conditions.

In order to provide images of the same locations, to address the lack of diversity in real environments and to evaluate better the impact on the epistemic uncertainty, some works promoted inpainting of virtual objects [30] or synthesised weather conditions [53]. In this setting however, questions may be raised about the veracity of the result. Therefore, the recent ACDC dataset [50] is composed entirely of real images taken from the same locations, and includes multiple sources of aleatoric uncertainty. However, not having any control on the noise level makes it harder to quantify the link between noise and uncertainty. Acquiring images with uncertainty corner cases is problematic as these cases are rare (long tail) and also costly to annotate, e.g., 3.3 hours/image [50]. Given this scarcity, such images are better used for validation as a small test set to assess the reliability of DNNs before deployment. These system validation stages can be seen as stress tests with corner cases to mirror challenging real-world conditions. It is thus interesting even from a more applied standpoint to have a synthetic dataset that mimics these rare conditions with some good fidelity constraint to quantify the robustness of DNNs. Synthetic data is abundant and can allow us to measure finer drifts in the input distribution. In addition, most such datasets mainly focus on semantic segmentation, while we propose to address multiple tasks (semantic segmentation, monocular depth, object detection, and instance segmentation).

In Table 1 we provide a summary of the main existing uncertainty datasets. In this work, we propose a fully synthetic dataset, called MUAD, integrating different weather conditions with various intensities, and suitable for a multitude of vision tasks and for the comprehensive characterisation of their uncertainty.

### 2.2 Uncertainty

Several works address the two types of uncertainty, in particular for the classification task. Most approaches build upon Bayesian learning, frequently using Bayesian Neural Networks (BNNs) [9, 16, 18, 19, 51, 43, 70], which estimate the posterior distribution of the DNN

Dataset	Adversarial annotations	Weather				Classes	Out of distribution	Depth	Object detection 2D/3D	Instance segmentation
		Fog	Night	Rain	Snow					
Foggy Driving [13]	101	✓	-	-	-	19	-	-	✓	-
Foggy Zurich [14]	40	✓	-	-	-	19	-	-	-	-
Nighttime Driving [15]	50	-	✓	-	-	19	-	-	-	-
Dark Zurich [16]	201	-	✓	-	-	19	-	-	-	-
Raincouver [17]	326	-	✓	✓	-	3	-	-	-	-
WildDash [18]	226	✓	✓	✓	✓	19	-	-	-	-
BDD100K [19]	1346	✓	✓	✓	✓	19	-	-	-	-
ACDC [20]	4006	✓	✓	✓	✓	19	-	✓	✓	-
Virtual KITTI 2 [21]	21260	✓	-	✓	-	14	-	✓	✓	✓
Fishyscapes [22]	373	-	-	-	-	19+2	✓	-	-	-
LostAndFound [23]	1203	-	-	-	-	19+9	✓	-	-	-
RoadObstacle21 [24]	327	-	✓	-	✓	19+1	✓	-	-	-
RoadAnomaly21 [25]	100	-	-	-	✓	19+1	✓	-	-	-
Streetazard [26]	6625	-	-	-	-	13+250	✓	-	-	-
BDD anomaly [27]	810	✓	✓	✓	✓	17+2	✓	-	-	-
MUAD	10413	✓	✓	✓	✓	16+9	✓	✓	✓	✓

Table 1: **Comparative overview** of the different datasets for uncertainty on autonomous driving.

weights to marginalize the likelihood distribution at inference time. Yet most BNNs are difficult to train and scale to complex computer vision tasks that have been addressed, so far, by fewer uncertainty estimation methods. Ensembles [58] and pseudo-ensembles [19, 22, 23] achieve state-of-the-art performance on various tasks, at the high cost of multiple training and/or multiple forward passes at inference. Some approaches [56] formalize DNNs to output a parametric distribution, and their goal is to estimate the distribution parameters. These approaches can be applied to optical flow [54] and object detection [12], yet they mainly focus on aleatoric uncertainty. Besides the additional challenges posed by complex computer vision tasks, progress on uncertainty estimation in this area is hindered by the lower number of datasets for properly assessing both the quality of the predictive uncertainty and the predictive performance. With MUAD we hope to encourage research in this essential area for practical applications with annotations and benchmarks for multiple tasks.

### 3 Multiple Uncertainties for Autonomous Driving benchmark (MUAD)

According to the categorization of the uncertainty in line with the current works of the community [23] (summarized in the Supplementary Material), we propose to use the dataset to better evaluate the results and uncertainty estimations given by the DNNs in the context of autonomous driving. Let us link the two main types of uncertainty - aleatoric and epistemic - to the specific context of our application. In the scenario of autonomous driving, we believe that the aleatoric uncertainty of the DNNs will occur due to different weather conditions than the ones present in the training set. The epistemic uncertainty of the DNNs should arise when the class or the appearance of objects in the picture differ from those of the data provided in the training set. The design of MUAD dataset is based on this hypothesized relationship between uncertainty and autonomous driving scenarios. In the remainder of this section, we will detail the composition of MUAD dataset.

The goal of MUAD is to confront DNNs to uncertain environments and to characterize numerically their robustness in adverse conditions, more specifically in the presence of

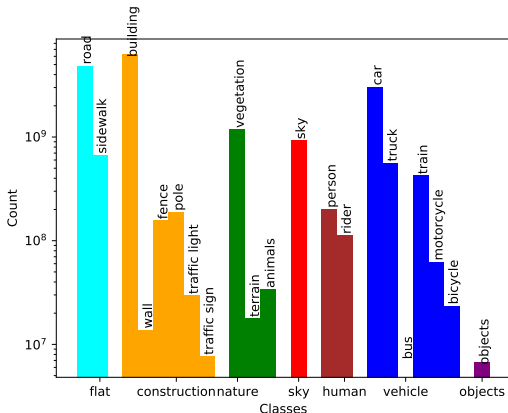


Figure 2: Number of annotated pixels per class in MUAD.

rain, fog and snow. Photorealism is essential for guaranteeing that synthetic datasets are challenging with respect to real-world conditions, and also for keeping them relevant for use in industrial applications. This is particularly important for accommodating weather artifacts [40, 63, 68]. Our dataset is generated using a physics based synthetic image rendering engine to produce high-quality realistic images and sequences. The engine uses an accurate light transport model [44, 57] and provides a physics description of lights, cameras and materials. This allows for a detailed simulation of the amount of light that is reaching the camera sensor. The camera sensor itself is simulated converting the energy coming from the scene in the form of photons into electrons. Electrons are finally converted into a voltage that is digitized to produce the digital values that represent the color image. We provide the photorealistic rendering descriptions for different weather conditions in Section 3.3. For each sample in MUAD, the corresponding ground truth information contains the semantic segmentation, the depth map, and for some specific classes (pedestrian, car, van, traffic light, traffic sign) the instance segmentation with the corresponding bounding boxes. We follow the standard data split strategy, however the training and validation set contain only images with normal weather conditions and without some specific classes which are denoted as OOD. The test set is organized into seven subsets following the intensity of the adverse weather conditions:

- **normal set:** images without OOD objects nor adverse conditions, as in Figure 1a.
- **normal set overhead sun:** images without OOD objects nor adverse conditions, in which we simulate the sun with a zenith angle of  $0^\circ$ , that we denote for the sake of simplicity as overhead sun.
- **OOD set:** images with OOD objects and without adverse conditions, as in Figure 1b.
- **low adv. set:** images with medium intensity adverse conditions (fog, rain or snow).
- **high adv. set:** images containing high intensity adverse conditions (fog, rain or snow).
- **low adv. with OOD set:** images containing both OOD objects and medium intensity adverse conditions (fog, rain or snow), as in Figure 1c.
- **high adv. with OOD set:** images containing both OOD objects and high intensity adverse conditions (fog, rain or snow), as in Figure 1d.

In Figure 3 and 4 we illustrate the instance segmentation and the semantic segmentation of 3 images. The adverse weather conditions are realistic and challenging as they bring a mix of difficult (unknown during training) environment conditions and perturbation of the




Figure 3: **Illustration of semantic segmentation images of MUAD dataset.** The first row is composed of the original images of the **high adv. set**. The second row is their corresponding ground truth.

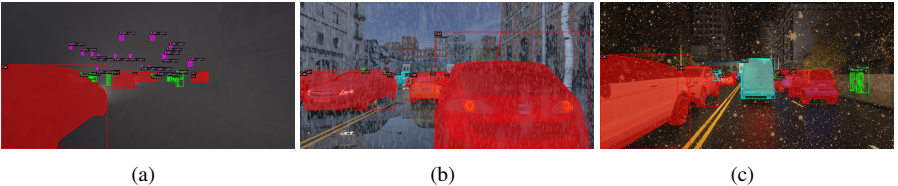


Figure 4: **Illustration of instance segmentation images of MUAD dataset.** The three images are selected from the **high adv. set**. We illustrated fog (4(a)), rain (4(b)), and snow (4(c)) conditions.

visibility in the scene. We argue that such settings are helpful for autonomous driving since the autonomous system must face and be robust against a variety of weather conditions and situations.

### 3.1 MUAD statistics

Our dataset contains 3,420 images in the train set, and 492 in the validation set. The test set is composed of 6,501 images divided as follows: 551 in the **normal set**, 102 in the **normal set no shadow**, 1,668 in the **OOD set**, 605 in the **low adv. set**, 602 in the **high adv. set**, 1,552 in the **low adv. with OOD set** and 1,421 in the **high adv. with OOD set**. All of these sets cover day and night conditions with 2/3 of day images and 1/3 of night images. Test datasets address diverse weather conditions (rain, snow, and fog with different levels), and various OOD objects. The resolution of all images is  $1024 \times 2048$ .

The dataset aims to provide a general and consistent coverage for a typical urban and suburban environment under different times of day and weather conditions. Ego-vehicle poses are drawn randomly within a complex environment, and in a second stage the field of view is populated stochastically with dynamic objects of interest following distributions in compliance with their expected behaviour. The pose and context changes as well as the variation of the models for the objects of interest ensure that content diversity is high, in addition to images being photorealistic. The simulator makes use of approximately 300 different person models and 150 different vehicle models, which are sampled while varying their visual characteristics.

Cityscapes classes	MUAD classes	nb. of images with the annotations
Road	Bots, Tram Tracks, Crosswalk, Parking Area, Garbage - Road, Road Lines, Sewer Longitudinal Crack, Transversal Crack, Road, Asphalt hole, Polished Aggregate, Vegetation - Road, Sewer - Road, Construction Concrete	9,055
Sidewalk	Lane Bike, Kerb Stone, Sidewalk, Kerb Rising Edge	8,948
Building	House, Construction Scaffold, Building, Air Conditioning, Construction Container, TV Antenna, Terrace, Water Tank, Pergola Garden, Stairs, Dog House, Sunshades, Railings, Construction Stock, Marquees, Hangar Airport	9,089
Wall	Wall	1,101
Fence	Construction Fence, Fences	8,622
Pole	Traffic Signs Poles or Structure, Traffic Lights Poles, Street lights, Lamp	8,984
Traffic light	Traffic Lights Head, Traffic Cameras, Traffic Lights Bulb (red, yellow, green)	8,222
Traffic sign	Traffic Signs	2,672
Vegetation	Vegetation	9,072
Terrain	Terrain, Tree Pit	8,377
Sky	Sky	8,591
Person	Walker, All colors of Construction Helmet, All colors of Safety Vest, Umbrella, People	8,843
Rider	Cyclist, Biker	3,470
Car	Car, Beacon Light, Van, Ego Car	9,026
Truck	Truck	5,533
Bus	Bus	0
Train	Train, Subway	2,240
Motorcycle	Motorcycle, Segway, Scooter Child	2,615
Bicycle	Bicycle, Kickbike, Tricycle	2,816
Animals	Cow, Bear, Deer, Moose	603
Objects anomalies	Food Stand, Trash Can, Garbage Bag	352
Background	Others	-

Table 2: Overview of annotated classes

## 3.2 Class labels

The class ontology of MUAD is presented in Table 2. MUAD comprises 155 different classes that we have regrouped into 21 classes. The first 19 classes are similar to the CityScapes classes [13], then we added object anomalies and animals to have more diversity in the anomalies. In addition to ensuring high content diversity, this ontology facilitates the mapping of MUAD to specific environments which require or impose a lower number of more generic classes. Consequently, trained models are easily transferable for existing datasets, and we provide the mapping towards the 21 classes widely used by the community, e.g., [10, 13, 56, 57]. The dataset statistics for the 21 classes are presented in Figure 2. For the evaluation of OOD detection, we have excluded nine classes (train, motorcycle, bicycle, bears, cow, deers, moose, food stand, garbage bags) from the training and validation sets. These classes are present in the test set as OOD objects. DNNs that process samples belonging to one of these nine classes are expected to have a low confidence score.

## 3.3 Photorealistic rendering

Our physically based approach simulates the weather conditions taking into consideration the amount of ozone, the humidity, among other factors. **Regarding the sky**, the renderer uses a physical model of the light coming from the sky. The amount of ozone and humidity in the atmosphere changes the emissive spectral profile of the sky, impacting the color of the objects in the scene. Apart from ozone and humidity, there are other factors that the render takes into account, for instance, turbidity and scattering asymmetry. **Regarding the rain and the snow**, the simulation of every raindrop allows us to model physical dispersion. For improved realism we choose the falling speed and size of raindrops according to observed real rain [10, 48]. For snow, the same principle applies, but changing in this case the material and the



Methods	Architectures	normal set				low adv. without OOD set				high adv. without OOD set			
		mIoU ↑	ECE ↓	AUROC ↑	FPR ↓	mIoU ↑	ECE ↓	AUROC ↑	FPR ↓	mIoU ↑	ECE ↓	AUROC ↑	FPR ↓
Baseline (MCP) [49]	DeepLab v3+ [40]	68.90%	0.0138	38.77%	0.3238	22.51%	0.4567						
Baseline (MCP) [49]	SegFormer-B0 [77]	69.04%	0.0135	48.67%	0.1004	30.0%	0.2396						
MC-Dropout [42]	DeepLab v3+ [40]	65.33%	0.0172	42.08%	0.2587	27.68%	0.3846						
MC-Dropout [42]	SegFormer-B0 [77]	68.55%	0.0119	45.01%	<b>0.0758</b>	26.59%	<b>0.1594</b>						
Deep Ensembles [45]	DeepLab v3+ [40]	69.80%	0.0129	42.81%	0.2444	23.91%	0.4500						
Deep Ensembles [45]	SegFormer-B0 [77]	<b>70.00%</b>	<b>0.0115</b>	<b>49.10%</b>	0.0837	<b>31.67%</b>	0.3167						

Methods	Architectures	OOD set				low adv. with OOD set				high adv. with OOD set						
		mIoU ↑	ECE ↓	AUROC ↑	FPR ↓	mIoU ↑	ECE ↓	AUROC ↑	FPR ↓	mIoU ↑	ECE ↓	AUROC ↑	FPR ↓			
Baseline (MCP) [49]	DeepLab v3+ [40]	57.32%	0.0607	0.8624	0.2604	31.84%	0.3078	0.6349	0.1185	0.6746	18.94%	0.4356	0.6023	0.1073	0.7547	
Baseline (MCP) [49]	SegFormer-B0 [77]	58.91%	0.06465	0.8578	0.21576	40.22%	0.1544	0.7448	0.1361	0.5876	<b>27.51%</b>	0.6564	0.6564	0.1071	0.7011	
MC-Dropout [42]	DeepLab v3+ [40]	55.62%	0.0645	0.8439	0.2225	0.4575	33.38%	0.1329	0.7506	0.1545	0.5807	20.77%	0.3809	0.6864	0.1185	0.6751
MC-Dropout [42]	SegFormer-B0 [77]	58.81%	<b>0.0574</b>	0.8811	0.2535	0.3435	39.64%	<b>0.1172</b>	<b>0.7698</b>	<b>0.1557</b>	<b>0.5498</b>	26.52%	<b>0.1771</b>	<b>0.6965</b>	<b>0.1237</b>	<b>0.6633</b>
Deep Ensembles [45]	DeepLab v3+ [40]	58.29%	0.0588	0.871	<b>0.2802</b>	0.3760	34.91%	0.2447	0.6543	0.1212	0.6425	20.19%	0.4227	0.6101	0.1162	0.7212
Deep Ensembles [45]	SegFormer-B0 [77]	<b>59.50%</b>	0.05928	<b>0.8843</b>	0.2611	<b>0.3342</b>	40.00%	0.1400	0.6933	0.1198	0.6290	25.89%	0.3305	0.5939	0.0959	0.7287

Table 3: **Comparative results for semantic segmentation on MUAD.** The mIoU is related to the main task performance, while the rest of the metrics evaluate the uncertainty quality when the model is confronted with different types of perturbations.

dynamics. **Regarding the fog**, we use a full volumetric approach for the simulation where scattering effects are considered. **Regarding the level of noise**, to the best of our knowledge, there is no standard procedure to measure the intensity of adverse weather conditions for driving scenarios. We empirically selected the number of raindrops, snowflakes, and fog intensity from a human point of view. All the efforts mentioned above improved our dataset realism. A study [45] was performed that confirmed that our render enhances the realism of MUAD compared to SYNTHIA.

## 4 Experiments

### 4.1 Semantic segmentation experiments

Our semantic segmentation study consists of two experiments. Firstly, we evaluate on MUAD the uncertainty quantification of three benchmarks (MCP [49], Deep Ensembles [45], MC Dropout [42]), by taking advantage of the OOD/adverse weather splits. The second experiment evaluates the quality of transfer learning from MUAD to Cityscapes [43] and the quality of the uncertainty quantification on Cityscapes. We aim here to verify whether MUAD can be used for unsupervised domain adaptation.

For the first experiment, we train a DeepLabV3+ [40] network with ResNet50 encoder [28] and a SegFormer-B0 [77] on MUAD. Table 3 shows the results of our three baselines. The first criterion we use is the mIoU [45], and the second criterion is the expected calibration error (ECE) [26] that measures how the confidence score predicted by a DNN is related to its accuracy. Finally, we use the AUPR, AUROC, and the FPR-95-TPR defined in [49] that evaluate the quality of a DNN to detect OOD data. We can see that Deep Ensembles outperform other strategies, especially when mixed with Transformers. Yet, MC Dropout seems to have better performance on more complicated sets. Hence MUAD is well suited for quantifying the uncertainty evaluation of different DNNs.

For the second experiment, we train a DeepLabV3+ [40] segmentation network on MUAD and evaluate it on Cityscapes. Results reported in Table 4 show that models trained on MUAD images modified with simple histogram matching [62]<sup>1</sup> with Cityscapes images achieve the same performance as models trained on the much larger GTA dataset [56].

<sup>1</sup>We use the scikit-image [63] implementation: [https://scikit-image.org/docs/dev/api/skimage.exposure.html#skimage.exposure.match\\_histograms](https://scikit-image.org/docs/dev/api/skimage.exposure.html#skimage.exposure.match_histograms)

Training set	mIoU $\uparrow$
Baseline trained on Cityscapes	76.84%
Baseline trained on MUAD	16.71%
Baseline trained on MUAD with histogram eq.	32.12%
Baseline trained on GTA [24]	32.85%
Baseline trained on SYNTHIA [24]	29.45%

Table 4: **Comparative results for semantic segmentation simple domain adaptation from MUAD to Cityscapes.** First row is the original baseline, the second row is the performance of the model trained directly on MUAD and the third row is the performance of the model trained on MUAD with histogram matching technique.

Methods	normal set					low adv. without OOD set					high adv. without OOD set					normal set overhead sun					
	Depth results			Uncertainty results		Depth results			Uncertainty results		Depth results			Uncertainty results		Depth results			Uncertainty results		
	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	
Baseline	0.922	0.114	3.357	-	-	0.786	<b>0.147</b>	5.005	-	-	0.632	<b>0.207</b>	6.989	-	-	0.951	0.090	3.646	-	-	-
Deep Ensembles [35]	<b>0.929</b>	<b>0.111</b>	<b>3.199</b>	<b>0.291</b>	0.060	0.767	0.156	4.892	<b>0.740</b>	0.105	0.566	0.243	7.498	<b>1.182</b>	0.153	<b>0.955</b>	<b>0.083</b>	3.479	<b>0.336</b>	0.055	
MC Dropout [22]	0.919	0.119	3.209	0.634	0.061	<b>0.798</b>	0.151	<b>4.580</b>	1.063	0.098	<b>0.657</b>	<b>0.207</b>	<b>6.278</b>	1.382	<b>0.128</b>	0.948	0.092	<b>3.407</b>	0.786	0.058	
Single-PU [24]	0.905	0.132	3.230	0.313	0.081	0.773	0.159	4.865	0.789	0.112	0.571	0.248	7.680	1.740	0.171	0.946	0.105	3.546	0.358	0.079	
SLURP [23]	0.922	0.114	3.357	0.467	<b>0.048</b>	0.786	<b>0.147</b>	5.005	1.167	<b>0.090</b>	0.632	<b>0.207</b>	6.989	1.707	<b>0.128</b>	0.951	0.090	3.646	0.525	<b>0.033</b>	

Methods	OOD set					low adv. with OOD set					high adv. with OOD set				
	Depth results			Uncertainty results		Depth results			Uncertainty results		Depth results			Uncertainty results	
	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$	d1 $\uparrow$	AbsRel $\downarrow$	RMSE $\downarrow$	AUSE RMSE $\downarrow$	AUSE AbsRel $\downarrow$
Baseline	0.896	0.125	3.616	-	-	0.713	2.637	4.764	-	-	0.555	0.459	6.916	-	-
Deep Ensembles [35]	<b>0.903</b>	<b>0.114</b>	3.447	<b>0.427</b>	0.074	0.709	<b>1.810</b>	4.707	<b>0.692</b>	<b>0.129</b>	0.521	<b>0.331</b>	7.411	<b>1.072</b>	<b>0.151</b>
MC Dropout [22]	0.893	0.145	<b>3.432</b>	0.724	0.080	<b>0.744</b>	3.925	<b>4.364</b>	0.927	0.206	<b>0.610</b>	0.545	<b>6.176</b>	1.245	0.314
Single-PU [24]	0.888	0.132	3.463	0.447	0.095	0.714	4.349	4.716	0.744	0.482	0.529	0.351	7.627	1.347	0.156
SLURP [23]	0.896	0.125	3.616	0.721	<b>0.068</b>	0.713	2.637	4.764	1.072	0.212	0.555	0.459	6.916	1.564	<b>0.151</b>

Table 5: **Comparative results for monocular depth on MUAD.** We use NeWCRFs [24] as the based DNN for monocular depth task.

## 4.2 Monocular depth experiments

We provide results for monocular depth using NeWCRFs [24], which is one of the SOTA on KITTI dataset [24]. NeWCRFs does not output uncertainty by default. Similarly to [34, 36, 50], we modify the DNN to output the parameters of a Gaussian distribution (i.e., the mean and variance). We denote the result as single predictive uncertainty (Single-PU). Based on this modification we train a Deep Ensembles [35] with 3 DNNs. We also provide the results from SLURP [23], which needs 2 DNNs to predict the depth and the uncertainty respectively, and MC-Dropout [22]. For depth evaluation, we use the same metrics as Eigen et al. [17] which are used in many following works [39, 24]. For uncertainty quality evaluation, we follow the implementation of Poggi et al. [52]. More details on implementation and evaluation criteria are provided in the Supplementary Material.

Table 5 lists some of the depth and uncertainty results of the above techniques on our dataset due to the space limit. We observe that in the presence of OOD, the uncertainty results of Deep Ensembles are comparatively better, while MC-Dropout provides more robust depth estimations under different perturbation. Additionally, we provide self-supervised monocular depth results using left-right image consistency [25], along with the full supervised results in Supplementary Material. We also propose a baseline method for depth domain adaptation from MUAD to KITTI and report its performance in Table 6. Compared to the direct adaptation from Virtual KITTI 2 [20], which is specifically designed based on the target dataset KITTI [24], the model trained on MUAD can achieve competitive performance.

## 4.3 Object detection experiments

For the object detection task, we trained a Gaussian YOLOV3 [12] and a Faster-RCNN [54] on the training data. The Faster R-CNN are trained with ResNet101 and ResNet50

Training set	KITTI							
	d1↑	d2↑	d3↑	Abs Rel↓	Sq Rel↓	RMSE↓	RMSE log↓	SILog↓
KITTI [10]	0.975	0.997	0.999	0.052	0.148	2.072	0.078	6.9859
Virtual KITTI 2 [11]	0.835	0.957	0.989	0.129	0.706	4.039	0.177	15.534
MUAD	0.731	0.927	0.983	0.187	1.059	4.754	0.227	18.581

Table 6: **Comparative results for monocular depth estimation simple domain adaptation from MUAD to KITTI eigen-split [10].** First row is the original baseline, the second and the third rows are the performance of the model trained directly on Virtual KITTI 2 [11] and MUAD respectively.

Evaluation data	normal set			low adv. without OOD set			high adv. without OOD set			OOD set			low adv. with OOD set			high adv. with OOD set		
	mAP↑	AP50↑	PDQ↑	mAP↑	AP50↑	PDQ↑	mAP↑	AP50↑	PDQ↑	mAP↑	AP50↑	PDQ↑	mAP↑	AP50↑	PDQ↑	mAP↑	AP50↑	PDQ↑
Faster R-CNN (ResNet101)	<b>39.91%</b>	<b>54.91%</b>	<b>16.88%</b>	25.00%	36.89%	<b>8.61</b>	<b>13.97%</b>	<b>22.01%</b>	<b>0.041</b>	<b>35.85%</b>	<b>48.9%</b>	<b>14.33%</b>	<b>24.73%</b>	<b>35.70%</b>	<b>8.49%</b>	<b>12.41%</b>	<b>19.66%</b>	<b>3.86%</b>
Faster R-CNN (ResNet50)	38.43%	53.13%	15.02%	<b>25.19%</b>	<b>37.38%</b>	8.18%	13.29%	21.53%	0.0389%	34.52%	47.63%	12.96%	23.93%	34.51%	7.95%	12.11%	19.46%	3.64%
Gaussian YOLOV3 [12]	20.81%	32.84%	2.22%	8.79%	16.40%	0.57%	3.28%	6.30%	0.22%	17.44%	28.16%	1.52%	10.80%	18.71%	0.64%	3.21%	6.15%	0.26%

Table 7: **Comparative results for object detection on MUAD.** The first criteria are the mAP AP50 [13] related to the accuracy, and the second criterion is the PDQ [14] that measures how well detectors probabilistically localise objects in an image.

backbones with FPN [15]. All the results are presented on Table 7. We can see that Faster R-CNN’s performance drops with the adversarial conditions, which confirms that considering the adversarial behavior is important when designing algorithms.

## 4.4 Discussion

The experiments show that the best main task OOD contender might not always be the most suited against different sources of uncertainty, thus it is important to test thoroughly and adapt the processing pipeline to the expected type of perturbations. The similar ranking of methods on our synthetic dataset and on real data (see [20, 30]) is encouraging as it allows us to generalize the analysis performed on MUAD to actual scenarios. An additional benefit of synthetic datasets is related to the reduced data privacy concerns and regulations that typically affect real world datasets, in particular in urban settings that include pedestrians. All these traits allow for faster validation of new algorithms before their deployment in in real-world settings. Finally, a potential different usage of MUAD concerns unsupervised domain adaptation from synthetic to real domains. Our preliminary results are encouraging.

## 5 Conclusion

Previous research in deep learning and autonomous cars has established that it is essential to robustify DNNs. In this paper, we present MUAD, a synthetic but highly realistic dataset incorporating multiples sources of uncertainties for autonomous driving, that provides insight into the robustness of DNNs for various applications. Based on MUAD, we provide a set of baselines for three fundamental computer vision tasks. Uncertainty is related to events that occur rarely; synthetic data is very valuable for dealing with infrequent events. We hope that our dataset can improve the reliability of DNNs, especially in autonomous driving scenarios. We are the first, to our knowledge, to provide a dataset with such noise dichotomies present in automotive applications. Our extensive benchmarks show the greater than ever importance of considering uncertainty quantification in addition to accuracy, for decision making systems in sensitive applications.

## Acknowledgement

We gratefully acknowledge the support of DATAIA Paris-Saclay institute which supported the creation of the dataset (ANR-17-CONV-0003/RD42). This work was performed using HPC resources from GENCI-IDRIS (Grant 2020 - AD011011970) and (Grant 2021 - AD011011970R1) and Saclay-IA computing platform.

## References

- [1] baranidesign. Rain drop size and speed of a falling rain drop. <https://www.baranidesign.com/faq-articles/2020/1/19/rain-drop-size-and-speed-of-a-falling-rain-drop>, 2020. [Online; accessed 06-October-2022].
- [2] Victor Besnier, Andrei Bursuc, David Picard, and Alexandre Briot. Triggering failures: Out-of-distribution detection by learning from local adversarial attacks in semantic segmentation. In *ICCV*, 2021.
- [3] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. Fishyscapes: A benchmark for safe semantic segmentation in autonomous driving. In *ICCV Workshops*, 2019.
- [4] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *ICML*, 2015.
- [5] Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual KITTI 2. *arXiv:2001.10773*, 2020.
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [7] Robin Chan, Krzysztof Lis, Svenja Uhlemeyer, Hermann Blum, Sina Honari, Roland Siegwart, Mathieu Salzmann, Pascal Fua, and Matthias Rottmann. Segmentmeifyoucan: A benchmark for anomaly segmentation. In *NeurIPS Datasets and Benchmarks*, 2021.
- [8] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019.
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *TPAMI*, 2017.
- [10] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018.

- [11] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *ICCV*, 2017.
- [12] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *ICCV*, 2019.
- [13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [14] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *ITSC*, 2018.
- [15] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *IJCV*, 2020.
- [16] Michael Dusenberry, Ghassen Jerfel, Yeming Wen, Yian Ma, Jasper Snoek, Katherine Heller, Balaji Lakshminarayanan, and Dustin Tran. Efficient and scalable bayesian neural nets with rank-1 factors. In *ICML*, 2020.
- [17] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014.
- [18] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Séverine Dubuisson, and Isabelle Bloch. Encoding the latent posterior of bayesian neural networks for uncertainty quantification. *arXiv:2012.02818*, 2020.
- [19] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Séverine Dubuisson, and Isabelle Bloch. TRADI: Tracking deep neural network weight distributions. In *ECCV*, 2020.
- [20] Gianni Franchi, Nacim Belkhir, Mai Lan Ha, Yufei Hu, Andrei Bursuc, Volker Blanz, and Angela Yao. Robust semantic segmentation with superpixel-mix. In *BMVC*, 2021.
- [21] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016.
- [22] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- [23] Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv:2107.03342*, 2021.
- [24] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [25] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.

- [26] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *ICML*, 2017.
- [27] David Hall, Feras Dayoub, John Skinner, Haoyang Zhang, Dimity Miller, Peter Corke, Gustavo Carneiro, Anelia Angelova, and Niko Sünderhauf. Probabilistic object detection: Definition and evaluation. In *WACV*, 2020.
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [29] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017.
- [30] Dan Hendrycks, Steven Basart, Mantas Mazeika, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. A benchmark for anomaly segmentation. *arXiv:1911.11132*, 2019.
- [31] José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *ICML*, 2015.
- [32] Stephen C Hora. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *RESS*, 1996.
- [33] John Houston, Guido Zuidhof, Luca Bergamini, Yawei Ye, Long Chen, Ashesh Jain, Sammy Omari, Vladimir Iglovikov, and Peter Ondruska. One thousand and one hours: Self-driving motion prediction dataset. *arXiv:2006.14480*, 2020.
- [34] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *ECCV*, 2018.
- [35] Paul Jaccard. The distribution of the flora in the alpine zone. *New phytologist*, 1912.
- [36] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017.
- [37] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [38] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*, 2017.
- [39] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv:1907.10326*, 2019.
- [40] Kunming Li, Yu Li, Shaodi You, and Nick Barnes. Photo-realistic simulation of road scene for data-driven methods in bad weather. In *ICCVW*, 2017.
- [41] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.

- [42] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [43] Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. In *NeurIPS*, 2019.
- [44] Marcus A Magnor, Oliver Grau, Olga Sorkine-Hornung, and C Thebalt. *Digital Representations of the Real World*. CRC Press, 2015.
- [45] Candela Mateos. All synthetic data is not made equal. <https://anyverse.ai/synthetic-data/synthetic-data-object-detection/>, 2020. [Online; accessed 06-October-2022].
- [46] Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip HS Torr, and Puneet K Dokania. Calibrating deep neural networks using focal loss. In *NeurIPS*, 2020.
- [47] Jishnu Mukhoti, Andreas Kirsch, Joost van Amersfoort, Philip HS Torr, and Yarin Gal. Deterministic neural networks with appropriate inductive biases capture epistemic and aleatoric uncertainty. *arXiv preprint arXiv:2102.11582*, 2021.
- [48] Ken’ichiro Muramoto, Toru Shiina, Tatsuo Endoh, Hiroyuki Konishi, Koh’ichi Kitano, et al. Measurement of snowflake size and falling velocity by image processing. 1989.
- [49] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kontschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *CVPR*, 2017.
- [50] David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *ICNN*, 1994.
- [51] Peter Pinggera, Sebastian Ramos, Stefan Gehrig, Uwe Franke, Carsten Rother, and Rudolf Mester. Lost and found: detecting small road hazards for self-driving vehicles. In *IROS*, 2016.
- [52] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *CVPR*, 2020.
- [53] Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *CVPR*, 2018.
- [54] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [55] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [56] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016.
- [57] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, 2016.

- [58] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *IJCV*, 2018.
- [59] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Map-guided curriculum domain adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. *arXiv:2005.14553*, 2020.
- [60] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, 2021.
- [61] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020.
- [62] Panos E Trahanias and Anastasios N Venetsanopoulos. Color image enhancement through 3-d histogram equalization. In *ICPR*, 1992.
- [63] Maxime Tremblay, Shirsendu Sukanta Halder, Raoul de Charette, and Jean-François Lalonde. Rain rendering for evaluating and improving robustness to bad weather. *IJCV*, 2021.
- [64] Frederick Tung, Jianhui Chen, Lili Meng, and James J Little. The Raincouver scene parsing benchmark for self-driving in adverse weather and at night. *RAL*, 2017.
- [65] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image processing in python. *PeerJ*, 2014.
- [66] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *WACV*, 2019.
- [67] Eric Veach. *Robust Monte Carlo methods for light transport simulation*. Stanford University, 1998.
- [68] Alexander Von Bernuth, Georg Volk, and Oliver Bringmann. Simulating photo-realistic snow and fog on existing images for enhanced cnn training and evaluation. In *ITSC*, 2019.
- [69] Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In *ICLR*, 2020.
- [70] Andrew Gordon Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *arXiv:2002.08791*, 2020.
- [71] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021.



- [72] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020.
- [73] Xuanlong Yu, Gianni Franchi, and Emanuel Aldea. SLURP: Side learning uncertainty for regression problems. In *BMVC*, 2021.
- [74] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. NeW CRFs: Neural window fully-connected CRFs for monocular depth estimation. In *CVPR*, 2022.
- [75] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. Wilddash-creating hazard-aware benchmarks. In *ECCV*, 2018.