

IDD-AW: A Benchmark for Safe and Robust Segmentation of Drive Scenes in Unstructured Traffic and Adverse Weather

Furqan Ahmed Shaik^{*1} Abhishek Reddy Malreddy¹ Nikhil Reddy Billa¹
 Kunal Chaudhary² Sunny Manchanda² Girish Varma¹

1. Machine Learning Lab, IIT Hyderabad 2. DYSL AI, DRDO

*furqan.shaik@research.iiit.ac.in

Abstract

Large-scale deployment of fully autonomous vehicles requires a very high degree of robustness to unstructured traffic, weather conditions, and should prevent unsafe mispredictions. While there are several datasets and benchmarks focusing on segmentation for drive scenes, they are not specifically focused on safety and robustness issues. We introduce the IDD-AW dataset, which provides 5000 pairs of high-quality images with pixel-level annotations, captured under rain, fog, low light, and snow in unstructured driving conditions. As compared to other adverse weather datasets, we provide i.) more annotated images, ii.) paired Near-Infrared (NIR) image for each frame, iii.) larger label set with a 4-level label hierarchy to capture unstructured traffic conditions. We benchmark state-of-the-art models for semantic segmentation in IDD-AW. We also propose a new metric called “Safe mean Intersection over Union (Safe mIoU)” for hierarchical datasets which penalizes dangerous mispredictions that are not captured in the traditional definition of mean Intersection over Union (mIoU). The results show that IDD-AW is one of the most challenging datasets to date for these tasks. The dataset and code will be available here: <http://iddaw.github.io>.

1. Introduction

Semantic Segmentation of driving scenes is one of the major tasks in the development of autonomous driving vehicles. Even though there has been much work in terms of datasets [4, 5, 10, 17], there are still many challenges when it comes to the representation of rare events and conditions. One of the major challenges is the robustness of the model to high amounts of traffic participants and adverse weather conditions. While state-of-the-art models give high prediction accuracies, it’s unclear how they will perform in such cases. Vehicles with full autonomy cannot be deployed




Figure 1. Examples of RGB/NIR/Annotations triplets in four adverse weather conditions from IDD-AW. The color legend for annotation is given in Figure 2.

safely unless these corner cases are analyzed.

Secondly, given the focus on safety in autonomous vehicles, it is essential to check if the metrics used accurately represent the safety concerns in road scenes. Traditionally segmentation models have been evaluated based on the mIoU which is a generic metric for the quality of the segmentation. However, mIoU doesn’t capture safety concerns specific to road scenes. For example, accurate prediction of traffic participants (vehicles, living things) and roadside objects are more essential compared to far objects. Also for autonomous driving, certain types of mispredictions are more tolerable than others. For example, a car being mispredicted as a bus is more tolerable than a car being mispre-

dicted as a road. mIoU treats both these mispredictions in the same way while they represent different severity of safe mispredictions.

In this paper, we introduce the IDD-AW dataset which captures highly unstructured traffic in adverse weather conditions. Here, highly unstructured traffic refers to situations where road conditions are intricate and challenging to predict, often characterized by a diverse mix of vehicles, pedestrians, and other dynamic elements. The dataset captures driving scenarios under various adverse weather conditions, encompassing rain, fog, low light, and snow. The distribution of pixels and instances of traffic participants is higher when compared to existing datasets. Additionally, we provide paired Near Infrared (NIR) images for every frame, thereby increasing the information available to models and enhancing their ability to make precise predictions.

We also propose a new metric called "Safe mIoU (SmIoU)" which penalizes dangerous mispredictions that are not captured in the traditional mIoU metric. We propose this metric for all hierarchical datasets where the severity of the penalty is based on the tree distance. We also calculate this metric over various label sets varying from the complete set and narrowing it down to several important classes such as traffic participants and roadside objects which directly impact autonomous driving.

We benchmark models on the IDD-AW dataset and compare them with results on other similar datasets like Cityscapes [4], IDD [25], ACDC [20]. We show that appending the NIR image to the input increases the accuracy of our dataset. We also show high contrast in results when comparing mIoU and SmIoU using the same trained models. We provide qualitative and logical explanations, that the lowering of Safe mIoU as compared to mIoU exposes safety concerns in existing models.

Main Contributions.

1. We propose an adverse weather conditions dataset for robust segmentation in unstructured traffic scenes (see Section 3).
2. We also propose a new metric called "Safe mIoU" to ensure the penalizing of misclassified pixels in hierarchical datasets (see Section 4).
3. We show both qualitative and quantitative results to prove our dataset is more diverse and challenging than other existing datasets (see Section 5).

2. Related Work

Road scene segmentation datasets. The Cityscapes Dataset [4, 5] is one of the most popular datasets for driving scene segmentation. Along with it, there are other driving datasets such as KITTI [10], Pascal VOC [8], Mapil-

Dataset	Labeled images	Rain	Fog	Snow	Lowlight	Labels	NIR
Foggy Driving	101	0	101	0	0	19	
Foggy Zurich	40	0	40	0	0	19	
Nighttime Driving	50	0	0	0	50	19	
Dark Zurich	201	0	0	0	201	19	
Raincouver	326	326	0	0	95	19	
WildDash	226	13	10	26	13	19	
BDD100K	1346	213	23	345	765	19	
ACDC	4006	1000	1000	1000	1006	19	
IDD-AW	5000	1500	1500	1000	1000	30	✓

Table 1. Comparison of IDD-AW with other adverse weather condition datasets. IDD-AW provides i.) more labeled images ii.) with more (30) labels that capture unstructured driving conditions and iii.) NIR images.

lary vistas [17] etc which contain a large collection of images. However, most of these datasets are captured in structured environments and have label sets that cannot completely satisfy the unstructured scenarios. IDD [25] introduced a large-scale annotated dataset in unstructured driving conditions with more than 10K annotated images. It contains drive scenes with a larger label set when compared to Cityscapes and includes 2-wheelers, pedestrians, and roadside objects. However, none of the above datasets have images in adverse weather conditions. All the datasets have been captured in clear daylight conditions with minimum image distortion.

Adverse weather road scene datasets. Oxford Robotcar [16] was the first real-world large-scale dataset in which adverse visual conditions such as nighttime, rain, and snow were significantly represented, but it did not feature semantic annotations. Foggy Zurich [18], Foggy Cityscapes [19] have synthesized fog images generated by applying a mask over the original images. Rain1400 [9], RainyCityscapes [22] Rain100H [27], and Rain12 [15] are prime examples for synthetic rain datasets. Comprehensive Snow Dataset [3] is another example of synthetic snow images. Datasets such as Rainy Wcity [31] BDD100K [28], ACDC [20], WildDash [29], WildDash2 [30] and Ithaca365 [6] have real images which have been captured in various adverse conditions like rain, fog, snow and low light/night time. However, none of these datasets have both NIR and RGB components




Figure 2. IDD & IDD-AW hierarchical labels (x-axis) and pixel count (y-axis) comparison. IDD provides data in unstructured driving conditions and our pixel composition is similar to it. Both datasets have 30 labels and a 4-level label hierarchy. The colors used will serve as legends for all plots of annotations and predictions in this paper. Normalization is done relative to image resolution and number of images in the datasets.

which could help us give a better insight into the images in these adverse conditions. We compare these datasets with IDD-AW in Table 1.

Datasets with NIR/Multispectral Data. Nusences [2], the first dataset to carry the fully autonomous vehicle sensor suite: 6 cameras, 5 radars, and 1 lidar, all with a full 360-degree field of view. The Multi-Spectral Road Scenarios (MSRS) dataset [23] focuses on road-related scenarios by combining NIR and RGB modalities. The TNO Hyperspectral Dataset [24] provides NIR and RGB images obtained from an airborne platform. Similarly, the KAIST Multispectral Pedestrian Dataset [12] offers synchronized NIR and RGB images captured in urban settings. IDD-3D [7], which consists of multimodal data from multiple cameras and LiDAR sensors across various traffic scenarios. However, most of these datasets do not cover adverse weather images. IDD-AW dataset provides synchronized RGB and NIR images in adverse weather and unstructured driving scenarios.

Safe & Robust Segmentation. The robustness of a dataset can also be viewed in terms of the ambiguity of labels in the dataset and the misclassification of corner cases. Disambiguity of labels can be achieved by introducing a hierarchy in the dataset labels. Hierarchical semantic segmentation explains visual scenes with multi-level abstraction by considering a structured class label hierarchy.

Li et al. [14] addresses hierarchical segmentation at both scene level and object level. Shaik et al. [21] aims to accelerate semantic segmentation and other computer vision tasks on GPUs using sparsity frameworks. In the work of Li et al. [13], the authors aim to enhance hierarchical semantic

segmentation by modifying the cross-entropy loss and penalizing it over various class hierarchies. In the research by Bar et al. [1], the ‘teacher-student framework’ is employed, where a pre-trained model guides the training of another model, a common practice to improve segmentation. In the study by Kamann et al. [11], the authors examine the robustness of semantic segmentation networks concerning a wide range of image corruptions based on their architecture.

IDD-AW contains images in unstructured driving scenarios which are also captured in various adverse weather conditions. This signifies the robustness of the dataset along various conditions. Additionally, we introduce a metric called Safe-mIoU which penalizes unsafe mispredictions more than traditional mIoU. This helps in improving the safety of the models especially for hierarchical driving datasets.

3. The IDD-AW Dataset

Data Collection or Acquisition. The dataset primarily focuses on capturing traffic scenarios in adverse weather conditions such as rain, fog, snow, and low light, with an emphasis on unstructured driving scenarios. Our dataset comprises a total of 202 driving sequences, providing a comprehensive collection of challenging driving scenarios. The dataset is captured at different locations in India comprising urban, rural, and even hilly areas to ensure diversity across various conditions.

For data capture, we used JAI’s FS-3200D-10GE camera. This is a 2-CMOS multispectral prism camera that captures images in different spectral bands simultaneously. Specifically, it provides a visible color channel from 400-670 nm and a near-infrared (NIR) channel from 740-1000 nm. To ensure a clear separation between different adverse




Figure 3. Comparison of ACDC & IDD-AW: Pixel wise (left), traffic participants instances count (right). Pixel counts are normalized by resolution and the number of images in the datasets. IDD-AW has more pixels (left) as well as more instances per image (right) of traffic participants (TP) which includes all vehicles and living things that represent unstructured traffic. Note that there are over 300 images in IDD-AW with more than 20 instances of TP while ACDC has only around 10.

conditions, each capture was made under only one adverse scenario to make sure the images were diverse between various conditions.

Frame Selection. The frames captured by the camera are filtered using a 3-second threshold to eliminate any repetitive frames. The goal of this threshold is to ensure that unique frames are used for subsequent processing. After the filtering step, the RGB-NIR pairs are matched based on their time stamps. However, to ensure accuracy, the matches are also manually verified by reviewing all frames. From the matched pairs, visually appealing frames that feature a diverse range of traffic participants and showcase adverse driving conditions are shortlisted. This process allows for the selection of high-quality frames for the dataset.

Label hierarchy and Annotation. The label set used is the same as in IDD [25]. This label set has a hierarchical structure with four levels, consisting of 7 labels at level 1 and 30 labels at level 4 (see x-axis of Figure 2). This adds a higher level of complexity to our dataset when compared to existing datasets like Cityscapes, and even when compared to adverse weather datasets like ACDC or Foggy Cityscapes. For labelling the dataset, we had a team of highly skilled annotators. The annotation process for each image typically requires 1.5 to 2.5 hours to complete, encompassing the initial annotation and quality check. To ensure the correctness of the annotations, multiple steps are taken, such as verifying annotations against predefined criteria and conducting a final quality check by experienced annotators.

Dataset Splits. IDD-AW is split into four sets corresponding to the examined conditions. We manually selected 1500 rainy, 1500 foggy, and 1000 lowlight and 1000 snowy

images from the recordings for dense pixel-level semantic annotation, for a total of 5000 adverse-condition images. The selection process aimed at maximizing the complexity and diversity of captured scenes.

We used a thorough train-test split technique with set limits to permit robust evaluations. All images from a drive sequence are either completely in train or in test to ensure that the test is truly unseen and distinct. For each weather condition, the test set includes drive sequences in a ratio of 0.18 to 0.22 of total sequences, ensuring adequate representation of varied driving scenarios. The average number of frames per drive sequence in the test set was kept within the ratio of 0.9 to 1.2 of the average frames per drive sequence of the entire dataset. Class-specific fairness is maintained by ensuring that the average instances per image in the test set vary between 0.8 and 1.2 of the overall dataset. For precise pixel-level accuracy evaluation, the test dataset must meet strict criteria: an average pixel ratio of 0.8 to 1.2 for at least 18 classes, and 0.7 to 1.3 for at least 22 classes.

Comparison with Other Datasets. In Table 1, we compare IDD-AW to existing datasets that address semantic segmentation under adverse conditions. Most of these datasets are specialized in a single condition and are relatively small in scale. IDD and ACDC are the two primary large-scale datasets, focusing on unstructured driving scenes and adverse weather, respectively. IDD-AW is unique in its focus on drive scenes that are both unstructured and affected by adverse weather. Additionally, we provide a Near-Infrared (NIR) component, recognizing its potential to enhance segmentation accuracy and safety. NIR images tend to exhibit higher contrast and sharpness, while RGB images provide color and luminescence information that enhances the visual appeal.

Figure 2 illustrates the pixel-wise comparison between

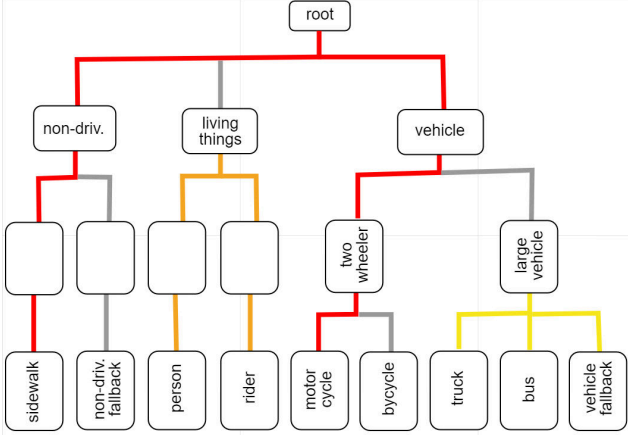


Figure 4. Example of tree distances between labels. Tree distance captures the severity of the misprediction using the label hierarchy. The $td(\text{sidewalk}, \text{motorcycle}) = 3$ since the length of the path is 6 and td is $\text{length}/2$. Similarly $td(\text{person}, \text{rider}) = 2$ and $td(\text{truck}, \text{bus}) = 1$.

IDD and IDD-AW datasets. Although our dataset features images captured in adverse weather conditions, most of the classes exhibit pixel counts similar to those of the IDD dataset which represents unstructured traffic.

Additionally, Figure 3 provides a pixel-wise comparison, as well as traffic participant, counts when compared with ACDC. Our dataset notably contains more pixels for crucial traffic participants such as roads, pedestrians, riders, and cars. Similarly, when comparing instance counts for traffic participants per image, nearly every image in our dataset has more instances than those in ACDC. Moreover, over 300 images in our dataset contain more than 20 instances, while ACDC only includes approximately 10-20 images with the same number of instances. This observation highlights that our dataset features numerous dense images in terms of both pixels per class and traffic participants per image.

4. Safe-mIoU: A metric for safe segmentation

Traditionally, mIoU (mean of Intersection over Union) is a widely used metric for semantic segmentation tasks. It reflects the quality of segmentation by considering both false positives (pixels mistakenly classified as a class) and false negatives (pixels of the class not detected). It captures the balance between precision and recall. However, while mIoU is an excellent general-purpose metric, its limitations become apparent when attempting to quantify the safety of driving scenes. However, mIoU’s strength lies in its generality, and this generality becomes its limitation in contexts where class importance as well as the severity of the mispredictions varies. In driving scenes, the safety of pedestrians, vehicles, traffic signs, and other critical elements holds paramount importance. Misclassifying these safety-critical

classes can lead to dire consequences. Also, certain misclassifications are dangerous, while others are tolerable. For example, it is tolerable to misclassify one vehicle as another type, while it is dangerous to misclassify a pedestrian or the vehicle as the road (see Figure 5). Unfortunately, traditional mIoU does not inherently capture these safety concerns. It treats all classes as equal entities, disregarding their impact on real-world driving scenarios.

To address this gap, a more refined metric is needed – one that can not only assess the quality of segmentation results but also consider the severity of unsafe mispredictions. We propose a new **Safe mIoU (SmIoU)** metric. The calculation of Safe mIoU involves selecting a subset of critical/important classes C_{imp} from the entire class set C . These chosen classes typically encompass those that have a direct impact on safety within driving scenes, such as pedestrians, vehicles, and traffic signs. Furthermore, we use the distance between the predicted and the actual label in the hierarchical label tree of IDD-AW, to measure the severity of misclassifications in C_{imp} . It aims to provide a more holistic evaluation that aligns with the priorities of driving scene applications and the overall safety of the environment.

The essence of SmIoU lies in the introduction of hierarchical penalty, a strategy that takes into account the semantic relationships between classes. Misclassifications within critical classes, and non-critical classes classified as critical, are penalized based on their distance in the class hierarchy. The *tree distance* (td) between a pair of labels is the length of the shortest path in the class hierarchy tree divided by 2 (see Figure 4). So the td between person and rider is 1, while sidewalk and motorcycle is 3 (see the class hierarchy in 2), denoting that the latter is a more severe misprediction than the former. The calculation of SmIoU involves computing the I_c^{safe} for each class as shown in equation 1. The safe IoUs are the IoUs with a penalty for misclassification of that class weighted by the tree distance. The final SmIoU score is obtained by taking the mean of these I_c^{safe} as shown in equation 2, resulting in a metric that accurately measures both the segmentation quality and the safety-related considerations of driving scene applications.

Let C be the set of all classes at the bottom level of the hierarchy, $d(c, s)$ be the tree distance between class c, s , and n be the number of levels in the hierarchy, gt_c the set of pixels in the ground truth with label c and $pred_s$ the set of pixels in the prediction with label s . We define the following quantities:

$$I_{c,s}^{\text{safe}} = \frac{|gt_c \cap pred_s|}{|gt_c \cup pred_s|} \quad I_{c,c} = \frac{|gt_c \cap pred_c|}{|gt_c \cup pred_c|} \quad (\text{a})$$

Now, we define SmIoU as follows:

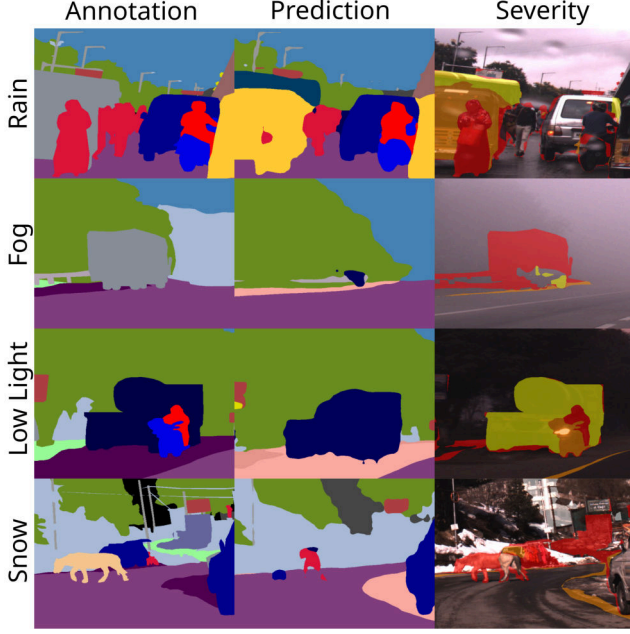


Figure 5. Examples from IDD-AW, showing a pedestrian in the rain, a truck in the fog, a rider/bike in low light, and an animal in the snow are mispredicted respectively. In the Severity Column, Red, orange, and yellow colors denote the level of danger posed by mispredictions. In the third row, a bike being mispredicted as a truck is less dangerous (in orange) than a rider being predicted as a vehicle (in red). mIoU considers these different colored mispredictions in the same way, while the proposed safe mIoU gives a higher penalty for orange and red mispredictions. In these images, the proposed safe mIoU is less than mIoU by more than 20%.

$$I_c^{\text{safe}} = \begin{cases} I_{c,c} - \sum_{s \in C, s \neq c} \frac{d(c,s)}{n} I_{c,s}^{\text{safe}} & \text{if } c \in C_{\text{imp}} \\ I_{c,c} - \sum_{s \in C_{\text{imp}}} \frac{d(c,s)}{n} I_{c,s}^{\text{safe}} & \text{else.} \end{cases} \quad (1)$$

$$\text{SmIoU} = \frac{\sum_{c \in C} I_c^{\text{safe}}}{|C|} \quad (2)$$

Note that the definition of SmIoU requires the definition of important classes C_{imp} . When SmIoU is mentioned without specifying C_{imp} , we take it to be the union of traffic participants and roadside object classes (all classes except Far Objects and Sky in Figure 2).

Properties of SmIoU. The following are the important properties of the proposed SmIoU when compared with traditional mIoU, which show their relationship and motivate the definition of SmIoU for capturing unsafe mispredictions.

Dataset	CS	ACDC	IDD	Rain	Fog	LL	Snow	IDD-AW
Test →								
Train ↓								
CS RGB	83	-	-	46	45	42	43	46
ACDC RGB	-	75	-	47	51	42	38	48
IDD RGB	-	-	73	52	55	50	33	54
IDD-AW RGB	49	51	51	62	64	62	53	64
IDD-AW NIR	-	-	-	61	58	57	51	61
IDD-AW NIR+RGB	-	-	-	66	65	63	53	67

Table 2. Comparison of mIoU scores (%) of InternImage-b model trained on CityScapes (CS), ACDC, IDD, and IDD-AW datasets (RGB, NIR, and Combined) for IDD-AW test set in individual conditions and jointly for all conditions. The model trained on IDD-AW NIR+RGB gives the highest accuracy which is 20% more than CS, ACDC, and 14% more than IDD. The NIR+RGB model gives 3% more accuracy compared to the RGB model, which indicates that the NIR image adds useful information for prediction. Also, note that the model trained on IDD-AW gives 3% better accuracy when tested on CS and ACDC as compared to the models trained on the respective datasets tested on IDD-AW (in grey). This trend is reversed in IDD since its training set is double the size of IDD-AW (in darker grey).

1. $\text{SmIoU} = \text{mIoU}$, when $C_{\text{imp}} = \emptyset$. When there are no important classes, SmIoU matches traditional mIoU. This is because the penalty becomes zero from the above formula, resulting in SmIoU becoming equal to the original mIoU.
2. $\text{SmIoU} \in [-1, 1]$. As we are penalizing the mIoU with the tree distance divided by the number of levels in the hierarchy, it always ranges between -1 to 1 inclusive of both extremes.
3. $\text{SmIoU} = 1$ if and only if all pixels in an image are accurately classified into their respective classes.
4. $\text{SmIoU} = -1$ if and only if all pixels in an image are misclassified with the farthest classes in the tree hierarchy.

5. Results

Training on CS, IDD, and ACDC. We have performed several experiments on Cityscapes, ACDC, IDD, and IDD-AW datasets using the InternImage [26] model. InternImage was chosen due to its state-of-the-art performance currently in semantic segmentation benchmarks. We have taken InternImage-b backbones of the InternImage [26]. Regarding the results from Table 2, we can see the mIoU when trained on Cityscapes does not even cross 45% mIoU when tested on IDD-AW. As the underlying dataset does not have any images with different weather conditions in it, the performance of a CS-trained model on IDD-AW is very low.

eval →	cross				same		
	Rain	Fog	LL	Snow	mIoU	SmIoU (tp)	SmIoU
Test →							
Train ↓							
IDD	52	55	50	33	-	-	-
IDD-AW	-	-	-	-	64	60	51
Rain	-	55	40	29	64	58	48
Fog	51	-	53	29	64	58	47
LL	52	57	-	30	62	58	48
Snow	35	38	33	-	53	43	28

Table 3. Cross-evaluation of InternImage-b condition experts on the various conditions of IDD-AW are shown in the first 4 columns. Comparison of mIoU (%) with SmIoU (%) metric at different levels and label sets for various adverse weather conditions shown in the last 3 columns. Here, SmIoU refers to Safe mIoU and tp refers to setting the important classes to be just the traffic participants (ie Living Things \cup Vehicles in Figure 2).

Similarly, a model trained on ACDC does not cross 50%, which shows the dataset is that much harder and has a larger label set. Even though ACDC is also a driving dataset that is captured in multiple weather conditions, there’s still not much improvement from Cityscapes. This already shows that our dataset is much harder than both Cityscapes and ACDC datasets.

When using pre-trained models of the IDD dataset, it performs much better than Cityscapes and ACDC despite testing with a larger label set. This is due to the unstructured images being present in the IDD dataset. The IDD pre-trained models performed efficiently in well-lit conditions like rain and fog, nonetheless the performance did not cross the 50% mIoU mark in low-light conditions. Also, the major concern is the performance in snow conditions, where the best IDD pre-trained model, gives only 32% mIoU which is a significant drop and is not usable for this condition at all. So even though the final mIoU reaches around 53%, it cannot be generalized over each condition and overly fails to perform in certain conditions.

Training and Evaluation on IDD-AW. We perform the training in three different ways on IDD-AW. First, we train the network only on the RGB images of our dataset. We can observe that training on IDD-AW gives more than 10% mIoU score when compared to IDD and almost 20% when taking the best performing Cityscapes or the ACDC pre-trained model. This is because of the variety of images in our dataset in various weather conditions which are not present in regular driving datasets like IDD or Cityscapes, as well as the unstructured nature of images and varying traffic when compared to other adverse datasets like ACDC. From Figure 6, we can see qualitative results from each pre-

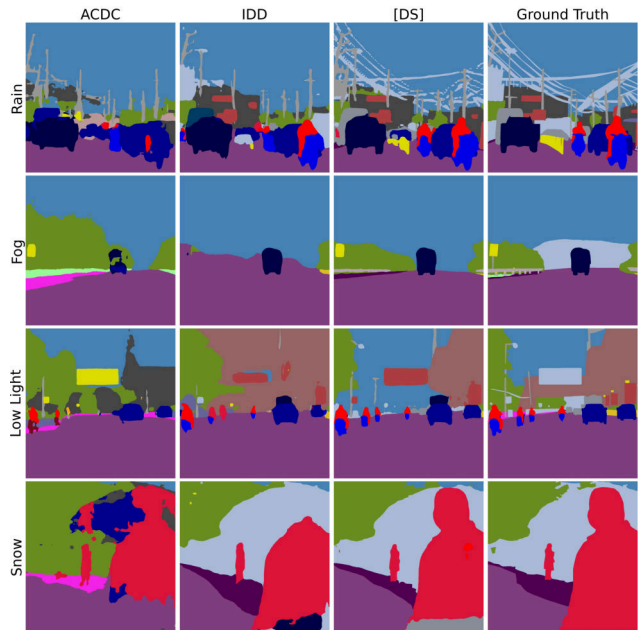


Figure 6. Qualitative examples from each condition with predictions using models trained on ACDC, IDD, and IDD-AW Datasets and the ground truth at the end. The mIoUs (%) for each pre-trained model on the IDD-AW test set are 46.57, 53.4, and 64.5 using ACDC, IDD, and IDD-AW pre-trained models respectively.

trained model on our dataset.

We conducted a similar experiment using only the NIR components of the images, yielding a mIoU of over 60% across various models and highlighting the quality of both the RGB and NIR components. Furthermore, we include training with stacked RGB+NIR images as input. Notably, the combination of RGB and NIR images through stacking led to a significant increase in mIoU by more than 3% when compared to using just RGB or NIR components alone (refer to Table 2), underscoring the critical role of NIR components within our dataset.

Condition Expert Training. In condition expert training, we train specifically on each condition separately and then test on each condition to assess the quality of images across each set. In Table 3, we can see the differences between training on one condition and testing on each of them. We can see that training on any other condition except Snow, does not yield any significant mIoU increase in the snow images.

Safe mIoU on IDD-AW test set. In Table 3, we can see the drop in SmIoU when we consider the important classes. For the InternImage-b framework, the SmIoU again drops more than 15% at least for each individual condition. This clearly shows the discrepancy and the dangers of traditional

Metric	Condition	Condition																		
		road	drivable fallback	sidewalk	person	rider	bike	bicycle	rickshaw	car	truck	bus	vehicle fallback	curb	wall	guard rail	billboard	traffic sign	traffic light	
mIoU %	All	95	51	48	76	72	68	5	83	85	74	76	45	78	52	74	62	58	52	
	Rain	96	49	49	48	73	68	2	86	87	79	52	47	81	47	46	64	54	52	
	Fog	97	64	24	62	75	73	8	67	91	82	80	49	79	57	78	61	69	47	
	LL	95	60	54	73	73	68	38	75	80	69	86	29	65	50	72	58	35	53	
	Snow	85	42	-	80	62	40	0	-	82	58	70	48	23	56	64	60	37	-	
SmIoU %	All	92	32	16	64	58	52	-22	77	81	68	70	21	69	32	61	42	40	27	
	Rain	94	28	25	15	59	54	-7	81	84	74	46	23	74	22	20	46	32	26	
	Fog	95	51	-48	41	62	58	-63	43	87	78	77	28	71	52	69	45	53	12	
	LL	92	46	22	60	59	52	14	70	76	61	82	2	48	26	52	35	15	32	
	Snow	79	20	-	70	44	-2	-99	-	76	48	56	23	-46	23	42	38	5	-	

Table 4. Comparison of class-wise labels for important classes between mIoU vs SmIoU for InternImage-b model on segmentation.




Figure 7. Distribution of Safe mIoU (%) vs mIoU (%) for IDD-AW test set. y-axis represents the number of images with a particular value of SmIoU/mIoU. The SmIoU distribution is shifted to the lower side, indicating that it is finding a significant amount of dangerous mispredictions that are not accounted for in mIoU.

mIoU in hierarchical autonomous driving datasets. In Figure 7, we can see the distribution of SmIoU vs mIoU over the test set has shifted toward the lower side which some images exhibiting even negative SmIoU.

Notice that $SmIoU = mIoU$ if $C_{imp} = \emptyset$. When we add more classes to C_{imp} the metric value decreases. This can be observed in the last 3 columns of Table 3. They correspond to $C_{imp} = \emptyset$, traffic participants and traffic participants \cup roadside objects respectively. The SmIoU (tp) has values between mIoU and SmIoU. This shows our new metric is flexible and very effective at identifying the images in which the important classes are misclassified.

When comparing class-wise mIoUs and SmIoUs for the InternImage-b framework as in Table 4, we can see several classes showing great disparities between mIoU and SmIoU. There is a significant difference between things

like bicycles, traffic signals, and sidewalks. It can be dangerous to misclassify these classes, especially when driving. When compared to conventional mIoU, our suggested SmIoU metric more accurately measures this risk. SmIoU also shows several classes with negative values showing dangerous misclassification, especially bicycle, sidewalk, and curb. The danger level associated with misclassification that could be brought on by these objects in the drive scenes is represented by classes like person in rain, sidewalk, and rickshaw in fog, vehicle fallback in lowlight, bike, and curb in snow that is well captured by SmIoU but missed by the traditional mIoU.

6. Conclusion

We have presented IDD-AW, a large-scale dataset and a benchmark suite for semantic driving scene understanding in adverse weather and unstructured driving conditions. We also present a new metric called Safe mIoU which incorporates safety concerns in the definition of mIoU. We benchmark state-of-the-art models for semantic segmentation in IDD-AW and also show the differences between traditional mIoU and safe mIoU while considering important classes. Finding appropriate loss functions, which can better optimize safe mIoU more efficiently is an interesting direction for future work.

7. Acknowledgement

We acknowledge the funding from DYSL-AI, DRDO for this project. The dataset was collected in collaboration with [iHub Data](#), IIIT Hyderabad.

References

- [1] Andreas Bar, Marvin Klingner, Serin Varghese, Fabian Huger, Peter Schlicht, and Tim Fingscheidt. Robust se-

- mantic segmentation by redundant networks with a layer-specific loss contribution and majority vote. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 332–333, 2020. 3
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscnets: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 3
- [3] Wei-Ting Chen, Hao-Yu Fang, Cheng-Lin Hsieh, Cheng-Che Tsai, I Chen, Jian-Jiun Ding, Sy-Yen Kuo, et al. All snow removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4196–4205, 2021. 2
- [4] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 1, 2
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset. In *CVPR Workshop on the Future of Datasets in Vision*, volume 2. sn, 2015. 1, 2
- [6] Carlos A. Diaz-Ruiz, Youya Xia, Yurong You, Jose Nino, Junan Chen, Josephine Monica, Xiangyu Chen, Katie Luo, Yan Wang, Marc Emond, Wei-Lun Chao, Bharath Hariharan, Kilian Q. Weinberger, and Mark Campbell. Ithaca365: Dataset and driving perception under repeated and challenging weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21383–21392, June 2022. 2
- [7] Shubham Dokania, A. H. Abdul Hafez, Anbumani Subramanian, Manmohan Chandraker, and C. V. Jawahar. IDD-3D: Indian driving dataset for 3d unstructured road scenes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 4482–4491, January 2023. 3
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International journal of computer vision*, 88:303–338, 2010. 2
- [9] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3855–3863, 2017. 2
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1, 2
- [11] Christoph Kamann and Carsten Rother. Benchmarking the robustness of semantic segmentation models with respect to common corruptions. *International journal of computer vision*, 129:462–483, 2021. 3
- [12] Chengyang Li, Dan Song, Ruofeng Tong, and Min Tang. Multispectral pedestrian detection via simultaneous detection and segmentation. *arXiv preprint arXiv:1808.04818*, 2018. 3
- [13] Liulei Li, Tianfei Zhou, Wenguan Wang, Jianwu Li, and Yi Yang. Deep hierarchical semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1246–1257, 2022. 3
- [14] Qing Li, Aihua Liang, and Hongzhe Liu. Hierarchical semantic segmentation of image scene with object labeling. *EURASIP Journal on Image and Video Processing*, 2018:1–10, 2018. 3
- [15] Yu Li, Robby T Tan, Xiaojie Guo, Jiangbo Lu, and Michael S Brown. Rain streak removal using layer priors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2736–2744, 2016. 2
- [16] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017. 2
- [17] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 4990–4999, 2017. 1, 2
- [18] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2
- [19] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, Sep 2018. 2
- [20] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10765–10775, 2021. 2
- [21] Furqan Ahmed Shaik, Thejasvi Konduru, Girish Varma, and Kishore Kothapalli. Accelerating computer vision tasks on gpus using ramanujan graph product framework. In *Proceedings of the 6th Joint International Conference on Data Science & Management of Data (10th ACM IKDD CODS and 28th COMAD)*, pages 113–117, 2023. 3
- [22] Vishwanath A Sindagi, Poojan Oza, Rajeev Yasarla, and Vishal M Patel. Prior-based domain adaptive object detection for hazy and rainy conditions. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 763–780. Springer, 2020. 2
- [23] Linfeng Tang, Jiteng Yuan, Hao Zhang, Xingyu Jiang, and Jiayi Ma. Piafusion: A progressive infrared and visible image fusion network based on illumination aware. *Information Fusion*, 83:79–92, 2022. 3
- [24] Alexander Toet. The tno multiband image data collection. *Data in Brief*, 15:249–251, 2017. 3

- [25] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1743–1751. IEEE, 2019. [2](#), [4](#)
- [26] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14408–14419, 2023. [6](#)
- [27] Wenhan Yang, Robby T Tan, Jiashi Feng, Zongming Guo, Shuicheng Yan, and Jiaying Liu. Joint rain detection and removal from a single image with contextualized deep networks. *IEEE transactions on pattern analysis and machine intelligence*, 42(6):1377–1393, 2019. [2](#)
- [28] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. [2](#)
- [29] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. Wilddash - creating hazard-aware benchmarks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [30] Oliver Zendel, Matthias Schörghuber, Bernhard Rainer, Markus Murschitz, and Csaba Beleznai. Unifying panoptic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21351–21360, 2022. [2](#)
- [31] Xian Zhong, Shidong Tu, Xianzheng Ma, Kui Jiang, Wenxin Huang, and Zheng Wang. Rainy wcity: A real rainfall dataset with diverse conditions for semantic driving scene understanding. In *International Joint Conference on Artificial Intelligence*, pages 1743–1749, 2022. [2](#)