

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# FishEye8K: A Benchmark and Dataset for Fisheye Camera Object Detection

Munkhjargal Gochoo<sup>1,2</sup> Munkh-Erdene Otgonbold<sup>1,2</sup> Erkhembayar Ganbold<sup>1,2</sup> Jun-Wei Hsieh<sup>3</sup> Ming-Ching Chang<sup>4</sup> Ping-Yang Chen<sup>5</sup> Byambaa Dorj<sup>6</sup> Hamad Al Jassmi<sup>1,2</sup> Ganzorig Batnasan<sup>1</sup> Fady Alnajjar<sup>1</sup> Mohammed Abduljabbar<sup>1</sup> Fang-Pang Lin<sup>7</sup>

<sup>1</sup> Department of Computer Science and Software Engineering, United Arab Emirates University, UAE
<sup>2</sup> Emirates Center for Mobility Research, United Arab Emirates University, UAE
<sup>3</sup> College of AI and Green Energy, National Yang Ming Chiao Tung University, Taiwan
<sup>4</sup> University at Albany — State University of New York, NY, USA
<sup>5</sup> Department of Computer Science, National Yang Ming Chiao Tung University, Taiwan
<sup>6</sup> Mongolian University of Science and Technology, Mongolia
<sup>7</sup> National Center for High-Performance Computing, Taiwan

mgochoo@uaeu.ac.ae, omunkuush@uaeu.ac.ae, eganbold@uaeu.ac.ae, jwhsieh@nctu.edu.tw, mchang2@albany.edu, pingyang.cs08@nycu.edu.tw, dorj@must.edu.mn, h.aljasmi@uaeu.ac.ae, gbatnasan@uaeu.ac.ae, fady.alnajjar@uaeu.ac.ae, 201970087@uaeu.ac.ae, fplin@narlabs.org.tw

# Abstract

With the advance of AI, road object detection has been a prominent topic in computer vision, mostly using perspective cameras. Fisheye lens provides omnidirectional wide coverage for using fewer cameras to monitor road intersections, however with view distortions. To our knowledge, there is no existing open dataset prepared for traffic surveillance on fisheye cameras. This paper introduces an open FishEye8K benchmark dataset for road object detection tasks, which comprises 157K bounding boxes across five classes (Pedestrian, Bike, Car, Bus, and Truck). In addition, we present benchmark results of State-of-The-Art (SoTA) models, including variations of YOLOv5, YOLOR, YOLO7, and YOLOv8. The dataset comprises 8,000 images recorded in 22 videos using 18 fisheye cameras for traffic monitoring in Hsinchu, Taiwan, at resolutions of 1080×1080 and  $1280 \times 1280$ . The data annotation and validation process were arduous and time-consuming, due to the ultra-wide panoramic and hemispherical fisheye camera images with large distortion and numerous road participants, particularly people riding scooters. To avoid bias, frames from a particular camera were assigned to either the training or test sets, maintaining a ratio of about 70:30 for both the number of images and bounding boxes in each class. Experimental results show that YOLOv8 and YOLOR outperform on input sizes  $640 \times 640$  and  $1280 \times 1280$ , respectively. The dataset will be available on the GitHub link with PASCAL VOC, MS COCO, and YOLO annotation formats. The Fish-*Eye8K benchmark will provide significant contributions to* the fisheye video analytics and smart city applications.



Figure 1. Sample of the 5 classes in the FishEye8K dataset: Pedestrian (all visible people on the streets), Bike (people riding bicycles, motorcycles, or scooters), Car (light vehicles such as sedans, SUVs, Vans, *etc.*), Bus, and Truck (dump-truck, semi-trailers, *etc.*)

# **1. Introduction**

Fisheye lenses have gained popularity owing to their natural, wide, and omnidirectional coverage, which traditional cameras with narrow fields of view (FoV) cannot achieve. In traffic monitoring systems, fisheye cameras are advantageous as they effectively reduce the number of cameras required to cover broader views of streets and intersections.

Dataset	Frame	Boxes	Task	Vehicles	Pedestrian	Weather	Occlusion	Altitude	View	Classes	Location	Туре
MIT-Car 2000 [15]	1.1K	1.1K	D	+						-	Surveillance	2D
KITTI-D 2014 [4]	15K	80.3K	D	+	+		+			3	Car	2D
UA-DETRAC 2015 [21]	140K	1210K	D,T	+		+	+			4	Surveillance	2D
Detection in LLC 2017 [9]	7.5K	15K	D	+		+				12	Car	2D
CARPK 2017 [5]	1.5K	90K	D	+						-	Drone	2D
UAVDT 2017 [2]	80K	841.5K	D,T	+		+	+	+	+	-	Drone	2D
NEXET 2017 [6]	50K	-	D	+		+				5	Car	2D
BDD100k 2018 [22]	5.7K	-	D,T	+	+	+				10	Car	2D
AAU RainSnow 2018 [1]	2.2K	13297	D,Seg	+		+					Surveillance	RGB&Thermal
MIO-TCD CCTV 2018 [12]	113K	200K	D	+		+				5	Surveillance	2D
BDD100k Adas 2018 [24]	100K	250K	D,Seg	+		+				10	Car	2D
Woodscape 2018/2019 [23]	10K	-	D,3D,T	+		+				7	Car	Fish-Eye
CityFlow2D 2021 [14]	-	313.9K	D,T	+							Surveillance	2D
FishEye8K 2023 [our]	8K	157.0K	D	+	+				+	5	Surveillance	Fish-Eye

Table 1. Summary of existing road traffic datasets. The second and third columns  $(1K = 10^3)$  indicate the number of images containing at least one object on them and the unique object bounding boxes. Remaining columns: additional attributes for each dataset, i.e., "D": target is a detection task, "3D": target is a three-dimensional detection task, "T": target is a tracking task, and the "Seg": target is a segmentation task.

Despite these benefits, fisheye cameras present distorted views that necessitate a non-trivial design for image undistortion and unwarping or a dedicated design for handling distortions during processing. It is worth noting that, to the best of our knowledge, there is no open dataset available for fisheye road object detection for traffic surveillance applications. The WoodScape dataset [23] was collected using an in-car fisheye dash camera; however, it was intended for self-driving scenarios.

In this paper, we present a new open FishEye8K benchmark dataset for the training and evaluation of 2D road object detection tasks. The FighEye8K dataset consists of 8,000 image frames with 157K bounding box annotations of 5 object classes, namely, Pedestrian, Bike, Car, Truck, Bus, and Truck; see Figure 1. A total of 22 short (8 to 20 minutes) videos were extracted from many hour-long videos collected from 35 fisheye cameras. These traffic surveillance cameras are properties of the police department of Hsinchu City, Taiwan, and our data collection is free from user consent agreements or license issues. However, efforts are performed in blurring out visible faces and license plates in the video frames. The dataset comprises different traffic patterns and conditions, including urban highways, road intersections, various illumination, and shooting angles of the five road object classes in various scales.

The labeling of objects of interest is meticulous. Specifically, we labeled all visible and recognizable objects even if they are located far away. The FishEye8K sample images are split into the training and test sets, with a ratio of about 70:30. Efforts are made to keep a similar ratio for each class of road objects. To avoid bias, the train and test sets do not share frames from the same camera. Annotations are provided in several standard formats, including Pascal-VOC [3], MS COCO [11], and YOLO [18].

We also provide benchmarking results of the latest State-of-The-Art (SoTA) two-stage object detection models, including YOLOv5x [7], YOLOR [19], YOLOv7 [20], and YOLOv8, and report in standard metrics including *Precision*, *Recall*, *mAPs*, *AP*<sub>S</sub>, *AP*<sub>M</sub>, *AP*<sub>L</sub>, *F*1–*score*, and their inference time.

The FishEye8K benchmark dataset will be available at https://github.com/MoyoG/FishEye8K upon paper acceptance.

## 2. Related Works

Road datasets. High-resolution, diverse, and large-scale road datasets play a critical role in advancing and enhancing traffic monitoring systems. In the last decade, the number of open road datasets [1,2,4-6,9,12,14,15,21-24] for 2D and 3D road object detection, single and multiple object tracking, object segmentation tasks have significantly increased. Table 1 provides a summary of popular road datasets that are used in both model development as well as for benchmarking and public contests. In terms of camera locations, the following datasets are captured using fixed surveillance cameras: MIT-Car [15], UA-DETRAC [21], AAU Rain-Snow [1], MIO-TCD [12], and AI-City [14] datasets. The CARPK [5] and UAVDT [2] dataasets are captured using drones. The KITTI [4], Detection in LLC [9], NEXET [6], BDD100K [22], and Woodscape [23] datasets are captured using in-dash cameras mounted on a car. In terms of FoV, all the datasets were constructed using standard perspective cameras, with the drawback of narrow FoV. The only exception is the WoodScape dataset [23] that are captured using an in-dash 180° fisheye camera. To our knowledge, the proposed FishEye8K dataset is the first of the kind among the open datasets, that are designed and constructed specifically for the development and evaluation of road object detection using fisheye traffic surveillance cameras.

Fixed perspective traffic camera-based datasets. Table 1 shows that most datasets are captured using fixed,



Figure 2. Sample images of FishEye8K dataset: (Top) the original unlabelled images, (Middle) the labeled ground truths, (Bottom) the YOLOv5x6 [7] detected objects. The columns illustrate several viewing angles, time of day, various intersections and road participants in the dataset.

perspective cameras, which are limited by the narrow FoV. All the datasets have annotations for 2D road object detection task; on top of it, a few datasets [2, 14] have multiple objects tracking annotation, and one [1] has segmentation mask annotation. In 2000, MIT-Car dataset [15] was publicly offered as a flagship dataset pioneering the road automation research field. The dataset has 1.1K frames, including 1.1K bounding boxes for the vehicle detection task. In 2016, UA-DETRAC [21] dataset was offered with 140K frames, including rich annotations of illumination, vehicle type, occlusion, and 1210K bounding boxes. The dataset has four classes (car, van, bus, and others) for detection and multiple object detection tasks. In the same year, similarly, MIO-TCD CCTV [12] dataset is offered with 113K frames, including 200K bounding boxes for the detection task. In 2018, the AAU RainSnow [1] dataset was offered as a benchmark for evaluating state-of-the-art rain removal algorithms. The dataset has 22 five-minute realworld camera video sequences collected from 7 urban intersections covering various weather conditions, i.e., snow, rain, haze, and fog. They have extracted 100 frames from each five-minute video to construct 2200 frames, including 13297 bounding boxes. Recently, in 2021, AI-City Challenge [14] was held, including vehicle detection and reidentification on CityFlowV2-ReID dataset and multi-target multi-camera vehicle tracking challenge on CityFlow2D dataset. CityFlow2D dataset has 313.9K bounding boxes for 880 distinct vehicles.

**Drone based datasets.** Lately, drone road datasets have been publicly offered in the literature, namely CARPK [5] and UAVDT [2]. Both datasets were captured from a high altitude with a viewing angle of the top by narrow FOV cameras for the drone-based road monitoring systems. Thus

they are not suitable for fixed surveillance camera-based traffic monitoring.

## 3. The FishEye8K Dataset

We provide detailed information on the new FishEye8K road object detection dataset. The dataset consists of 8,000 annotated images with 157K bounding boxes of five object classes. Figure 2 shows sample images of the wide-angle fisheye views, which provide new opportunities for large coverage, but also new challenges of large distortions of the road objects.

#### 3.1. Video Acquisition

We have acquired a total of 35 fisheye videos captured using 20 traffic surveillance cameras at 60 FPS in Hsinchu City, Taiwan. Among them, the first set of 30 videos (**Set 1**) was recorded by the cameras mounted at Nanching Hwy Road on July 17, 2018, with  $1920 \times 1080$  resolution, and each video lasts about 50-60 minutes. The second set of 5 videos (**Set 2**) was recorded at  $1920 \times 1920$  resolution, and each video lasts about 20 minutes.

All cameras are the property of the local police department, so there is no issue of user consent or license issues. All images in the dataset will be made available to the public for academic and R&D use.

#### **3.2. Dataset Preparation and Characteristics**

**Sampling.** We chose 18 videos from the recorded footage, with 15 videos coming from Set 1. These were cropped into shorter videos, each lasting approximately 8 to 10 minutes, except for one that lasted 16 minutes. Using a sampling method of one frame per 50 and 200 frames for Set 1 and Set 2 videos, respectively, we extracted over



Figure 3. The class distributions of objects in terms of (a) Splits for FishEye8K dataset; (b) Illumination; and (c) Scale.

10,000 frames. The resulting images were then resized to  $1080 \times 1080$  and  $1280 \times 1280$  for Set 1 and Set 2, respectively.

To incorporate a wide range of perspectives on road conditions, we carefully selected videos for our dataset that feature diverse camera angles, including side-view and frontview shots, as well as varying video quality. The dataset also includes images from different intersection types, such as T-junctions, Y-junctions, cross-intersections, midblocks, pedestrian crossings, and non-conventional intersections. The videos were captured under various lighting conditions, including morning, afternoon, evening, and night, and diverse traffic congestion levels ranging from freeflowing to steady and busy. Figure 2 illustrates some of the wide-ranging road conditions with ground truth annotations of road objects and detection results obtained from YOLOv5x6 [7].

**Object classes:** We annotate 5 major classes for road objects, namely, **Pedestrian** (all visible people on the streets), **Bike** (riders on bicycles, motorcycles, or scooters), **Car** (light vehicles such as sedans, SUVs, vans, *etc.*), **Bus**, and **Truck** (dump-truck, semi-trailers, *etc.*).

**Distant objects:** The wide fisheye lens creates a wide FoV but also results in a panoramic hemispherical image that is notably distorted with a barrel effect. Additionally, the camera has a tendency to produce blurred images of objects located around the edges of the lens. As a consequence, distant objects can appear minuscule and indistinct. Annotating these distant objects can be an arduous or even impossible task due to their lack of clarity.

**Illumination:** Four categories of illumination conditions were identified, namely morning (sunrise), afternoon (sunny), evening (sunset), and night. The distribution of video sequences based on their respective illumination attributes is illustrated in Figure 3(b), with the majority of bounding boxes falling under the afternoon category. Nighttime sequences follow in second place, with morning and evening categories trailing behind respectively. Notably, the distribution of classes across all times of day is remarkably similar

**Object scale:** We define the scale of the bounding boxes of road participants based on their size (length and width) in pixels. The MS COCO evaluator is employed for small

and medium, and large scaled objects. However, as the size of the image grows toward  $1080 \times 1080$  or  $1280 \times 1280$ , respectively for Sets 1 and 2, we doubled the size of standard scales, i.e., *small* (pixels  $\leq 64 \times 64$ ), *medium* ( $64 \times 64$ < pixels  $\leq 192 \times 192$ ), and *large* (pixels  $> 192 \times 192$ ). The distribution of road participants in the dataset in terms of scale is presented in Figure 3 (c), where small and mediumscaled objects make the most of the dataset. Bus and Truck classes have a similar number of small and medium scaled objects. On the contrary, other classes have a comparatively high number of small-scaled objects than medium and large-scale objects.

## 3.3. Annotation

Annotation Rule. The road participants were annotated based on their clarity and recognizability to the annotators, regardless of their location. In some cases, distant objects were also annotated based on this criterion.

**Annotation.** Two researchers/annotators manually labeled over 10,000 frames using the DarkLabel annotation program over a period of one year. After cleaning the dataset, a total of 8,000 frames containing 157012 bounding boxes remained. Unsuitable frames were removed, including those featuring road participants outside the five classes of interest.

The distribution of objects per class for each video is depicted in Figure 4. Notably, the night video captured by Camera 3 has the highest number of objects. In this dataset, the dominant classes are Bike (88,373) and Car (50,597), which can be attributed to the semi-tropical location of the country where the videos were recorded. On the other hand, the classes of Truck (3,317) and Bus (2,982) have the lowest number of objects, rendering the dataset highly imbalanced. Figure 1 displays a selection of samples from all classes, showcasing various scales. Furthermore, the distributions of classes are depicted as bar graphs in Figure 3.

For the sake of convenience, we provide three different formats for the annotations of FishEye8K datasets, i.e., Pascal-VOC [3], MS COCO [11], and YOLO [18].

#### 3.4. Validation

Given the complexity and effort required for the labeling task, human errors were inevitable, and it was necessary to

Train Set																	
Camera #	3		5	6	8	9	10	11	12	13	14	15	16	17	18		•⁄
Parts of Day	Α	Ν	Α	Α	Α	Α	Α	М	Α	Α	Α	Α	Α	Α	Α	All	70
Bus	186	161	178	49	365	66	225	264	38	378	98	11	24	9	0	2052	68.8
Bike	12615	13461	1173	243	3869	9668	3991	1943	345	6457	5026	1236	377	1642	22	62068	70.2
Car	6123	6894	2093	427	2678	1254	1804	1575	97	7778	2310	808	690	1873	69	36473	72.1
Pedestrian	1216	1130	0	0	2124	452	849	18	20	1569	1108	109	483	33	0	9111	77.6
Truck	21	40	291	82	62	87	396	23	0	729	73	128	17	121	45	2115	63.8
					Т	otal Bo	unding	Boxes								111819	71.2
			Test S	Set													
Camera #	1	2		4	L .		7		0/								
Parts of Day	Α	Α	М	Α	Е	Ν	Α	All	%								
Bus	385	240	29	0	10	46	220	930	31.19			М	Mornin	g			
Bike	8803	2056	6388	7	3073	2969	3009	26305	29.77			Α	Afterno	on			
Car	3619	1322	3811	31	2375	1325	1641	14124	27.91			E	Evenin	3			
Pedestrian	1288	680	267	1	61	248	87	2632	22.41			Ν	Night				
Truck	49	63	589	3	238	22	238	1202	36.24								
	٦	otal Bou	Inding I	Boxes				45193	28.78								

Figure 4. Heat maps represent the number of extracted objects per class from all 22 short videos recorded by 18 cameras for training and test sets of the FishEye8K dataset. For the training set, the darkest blue color refers to 13461 labeled bikes from the video recorded at night with Camera 3.

correct them to avoid inaccurate results. Therefore, in order to minimize human error, we employed two semi-automatic approaches to validate all bounding boxes.

In the case of mislabeled objects, we followed a two-step approach. Firstly, we cropped and copied the objects based on their respective bounding boxes into the corresponding directories. Secondly, our annotators manually verified if the objects were correctly placed in their designated directories through simple inspection, which is highly accurate and requires less time and effort. However, this approach is blind to objects that were not labeled in the first place, which is known as a missing label error. To address this issue, we inspected the False Positives generated by the YOLOv7 model [20] trained on FishEye8K, which helped identify numerous missing label errors. This approach was especially effective in identifying errors in distant areas and regions with high traffic density of vehicles and bikes.

## 3.5. Dataset Splits

In order to minimize dataset bias, we ensured that frames from the same camera were not included in both the train and test sets. Specifically, all frames from a given camera were assigned to either the train or test set. Figure 4 illustrates the heat maps of 22 videos (captured during morning, afternoon, evening, and night) recorded by Cameras 1-18, from which all images were extracted to create the Fish-Eye8K dataset. To satisfy the criteria, we selected Cameras 1, 2, 4, and 7 for the test set and the remaining cameras for the training set. This division resulted in a training set that constitutes 66.07% of the dataset, while the test set constitutes 33.93%.

In order to maintain a roughly 70:30 ratio of objects for each class, the training set was composed of 111,835 objects and the test set contained 45,193 objects, which correspond to 71.28% and 28.78% of all objects, respectively. The classes Bike, Bus, and Car follow this ratio in both sets.

#### 3.6. Data Anonymization

The identification of road participants such as people's faces and vehicle license plates from the dataset images was found to be unfeasible due for various reasons. The cameras used for capturing the images were installed at a higher ground level, making it difficult to capture clear facial features or license plates, especially when they are far away. Additionally, the pedestrians are not looking at the cameras, and license plates appear too small when viewed from a distance. However, to maintain ethical compliance and protect the privacy of the road participants, we blurred the areas of the images containing the faces of pedestrians and the license plates of vehicles, whenever they were visible.

## 4. Benchmark

## 4.1. One-Stage 2D Object Detection Methods

In order to assess the performance of 2D object detection methods, particularly for pedestrian and vehicle detection, we conducted a benchmark of the latest state-of-the-art onestage detectors. Our selection process involved reviewing the literature and identifying the best-performing models, including YOLOv5 [7], YOLOR [19], YOLOv7 [20], and the latest YOLOv8. One-stage detectors predict bounding boxes on images without requiring a region proposal step, which results in faster processing times and makes them suitable for real-time applications. However, these detectors prioritize inference speed and may not perform as well for recognizing irregularly shaped objects or groups of small objects. Table 2 presents the results of our benchmark of the one-stage detectors.

## 4.2. Training Procedure

We utilized several frameworks and platforms, i.e., Darknet [17], Pytorch [16], and PaddlePaddle [13], for the model training.

**Hyperparameters.** All YOLO variations were pretrained on MS COCO [11] dataset. Among the models, we trained four models (YOLOv7 [20], YOLOv7-X [20], YOLOv8l, and YOLOv8x on the input size  $640 \times 640$ . Six models (YOLOv5x6 [7], YOLOv5l6 [7], YOLOR-W6 [19], YOLOR-P6 [19], YOLOv7-D6 [20], YOLOv7-E6E [20]) on the input size  $1280 \times 1280$ . All models have trained with the same training procedures for 250 epochs, Adam [8] optimizer with the momentum of 0.937 except for YOLOv5 [7] which employed SGD optimizer. The confidence and NMS (Non Max Suppression) IoU (Intersection over Union) thresholds were both 0.5, and a learning rate of 0.01.

**Data preprocessing.** For the purpose of training and testing, the input images were resized to  $640 \times 640$  and  $1280 \times 1280$  for particular models, see Table 2.

**Loss Objective.** We employed the Focal loss [10] as it is commonly used in the multi-object detection and multilabel image classification domain. The loss function is defined as:

$$FL(p_t) = -\alpha_t (1 - p_t)^{\gamma} log(p_t), \tag{1}$$

where by default  $\gamma = 0.5$  and  $\alpha = 0.5$ ,  $p_t$  is the predicted probability for the object indexed by t.

#### 4.3. Metrics

All models are analyzed and evaluated with the same metrics, i.e., *Precision*, *Recall*, *mAPs*, *AP*<sub>S</sub>, *AP*<sub>M</sub>,  $AP_L$ , F1 - score, and their inference time.

**F1-score** metric measures the balance between *Precision* and *Recall*. When both *Precision* and *Recall* are high, the F1 score is high as well, indicating good model performance. On the other hand, a low F1 score indicates that the *Precision* and *Recall* values are imbalanced, and the model is not performing well. The F1 score is calculated as below:

$$F_1 = \frac{2 \times Precision \times Recall}{Precision + Recall}$$
(2)

Average Precision (AP) represents all *Precision* and *Recall* values into a single score. The AP is calculated according to:

$$AP = \sum_{k=0}^{n-1} [Recall_{(k+1)} - Recall_{(k)}] * Precision_{(k+1)},$$

where k is an index of the frame, and n is the number of frames for a given class.

**Intersection over Union (IoU).** The model predicts the bounding boxes of the detected objects; however, it is expected that the predicted box will not match exactly the ground truth box. Intersection over Union (IoU) is employed to quantify the measure to score how the ground truth and predicted boxes match:  $IoU = \frac{Intersection Areaa}{Union Area}$ .

**Normalized Confusion Matrix** is used to determine the prediction quality of the model by each class. A confusion matrix is made up of 4 components, namely, True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN).

Mean Average Precision (mAPs) is the mean of the APs for all classes. The mAP of the object detection model is calculated according to:

$$mAP = \frac{1}{n} \sum_{k=1}^{n} AP_k, \tag{3}$$

where n is the number of classes in the dataset and AP(k) is the average precision (AP) for a given class k.

### 4.4. Performance

In this subsection, we report the experimental results of variations of YOLOv5 [7], YOLOR [19], YOLOv7 [20], and YOLOv8.

Table 2 presents two sets of models that were trained on the FishEye8K dataset, with input sizes of  $1280 \times 1280$  and  $640 \times 640$ .

#### 4.4.1 Results on Input Size 640 × 640

For input size  $640 \times 640$ , the highest two  $mAP_{0.5}$ s of 0.6146 and 0.612 are achieved by YOLOv8x and YOLOv8l, respectively. The lowest  $mAP_{0.5}$ s of 0.4235 is result of YOLOv7 [20]. In terms of *F1-score* and *Recall*, YOLOv7-X achieved the highest performance with 0.5794 and 0.4888, respectively. Further, in terms of object scale, YOLOv7-X outperformed on all three scales (small, medium, and large) as well.

The confusion matrix for the best-performing model, YOLOv8x, on the input size of  $640 \times 640$ , is presented in Figure 5, and Table 3 tabulates the results. The Car class achieved the highest  $mAP_{0.5}$  score of 0.749, followed by Bus, Bike, Truck, and finally Pedestrian with a score of 0.4596. Surprisingly, the Bike class had the highest FPrate of 0.82, with many objects mispredicted as Bike on the background. Additionally, a significant portion of objects across all classes were undetected, with normalized FNsranging from 0.45 to 0.84. However, the model performed significantly well in terms of *Precision* for all classes, with values ranging from 0.74 to 0.94. The Pedestrian class had

Model	Version	Input Size	Precision	Recall	mAP <sub>0.5</sub>	mAP.595	F1-score	APS	$AP_M$	$AP_L$	Inference [ms]
VOI 0.5 [7]	YOLOv516	1280×1280	0.7929	0.4076	0.6139	0.4098	0.535	0.1299	0.434	0.6665	22.7
	YOLOv5x6	1280×1280	0.8224	0.4313	0.6387	0.4268	0.5588	0.133	0.452	0.6925	43.9
YOLOR [19]	YOLOR-W6	1280×1280	0.7871	0.4718	0.6466	0.4442	0.5899	0.1325	0.4707	0.6901	16.4
	YOLOR-P6	$1280 \times 1280$	0.8019	0.4937	0.6632	0.4406	0.6111	0.1419	0.4805	0.7216	13.4
VOL 0v7 [20]	YOLOv7-D6	1280×1280	0.7803	0.4111	0.3977	0.2633	0.5197	0.1261	0.4462	0.6777	26.4
10101/[20]	YOLOv7-E6E	1280×1280	0.8005	0.5252	0.5081	0.3265	0.6294	0.1684	0.5019	0.6927	29.8
YOLOv7 [20]	YOLOv7	640×640	0.7917	0.4373	0.4235	0.2473	0.5453	0.1108	0.4438	0.6804	4.3
	YOLOv7-X	640×640	0.7402	0.4888	0.4674	0.2919	0.5794	0.1332	0.4605	0.7212	6.7
	YOLOv81	640×640	0.7835	0.3877	0.612	0.4012	0.5187	0.1038	0.4043	0.6577	8.5
102010	YOLOv8x	640×640	0.8418	0.3665	0.6146	0.4029	0.5106	0.0997	0.4147	0.7083	13.4

Table 2. Results of state-of-the-art models trained on FishEye8K datasets. The table consists of two groups of various versions of YOLO object detection models for input sizes  $1280 \times 1280$  and  $640 \times 640$ .



Figure 5. Normalized Confusion Matrix of YOLOv8x model on the input size  $640 \times 640$ .

$YOLOv8x-640 \times 640$										
Classes	Precision	Recall	$mAP_{0.5}$	mAP_595	F1-score					
Bus	0.9331	0.4796	0.7156	0.5419	0.6335					
Bike	0.8035	0.377	0.6062	0.3208	0.5132					
Car	0.9493	0.5331	0.749	0.5208	0.6827					
Pedestrian	0.7785	0.1402	0.4596	0.2168	0.2376					
Truck	0.7444	0.3028	0.5424	0.4141	0.4304					
All	0.8418	0.3665	0.6146	0.4029	0.5106					

Table 3. Results of YOLOv8x model on the input size  $640 \times 640$ .

the lowest normalized *TP* rate at 0.14, indicating incorrect predictions of this class as others, mainly as Background which has the maximum normalized *FN* rate at 0.76.

#### 4.4.2 Results on Input Size 1280 × 1280

Table 2 shows that for an input size of  $1280 \times 1280$ , YOLOR-P6 [19] and YOLOR-W6 [19] achieved the highest  $mAP_{0.5}$  scores of 0.6632 and 0.6466, respectively. In contrast, YOLOv7-D6 [20] yielded the lowest  $mAP_{0.5}$ score of 0.3977. YOLOv7-E6E [20] demonstrated the highest performance in terms of *F1-score* and *Recall*, with values of 0.6294 and 0.5252, respectively.

Furthermore, with regard to object scale, YOLOv7-E6E



Figure 6. Normalized Confusion Matrix of YOLOR-P6 model on input size  $1280 \times 1280$ .

YOLOR-P6-1280×1280 [19]										
Classes	Precision	Recall	mAP <sub>0.5</sub>	mAP.595	F1-score					
Bus	0.9429	0.6753	0.8161	0.6271	0.7869					
Bike	0.8537	0.4316	0.6553	0.3725	0.5733					
Car	0.9473	0.6062	0.7876	0.5575	0.7393					
Pedestrian	0.4903	0.2014	0.3621	0.2007	0.2855					
Truck	0.7753	0.5541	0.695	0.4451	0.6462					
All	0.8019	0.4937	0.6632	0.4406	0.6111					

Table 4. Results of YOLOR-P6 model on the input size  $1280 \times 1280$ .

[20] exhibited higher performance over the other models in detecting small and medium-sized objects, achieving *APs* of 0.1684 and 0.5019, respectively. In contrast, YOLOR-P6 [19] demonstrated exceptional accuracy in detecting large objects, with an  $AP_L$  of 0.7216.

Figure 6 shows the confusion matrix and Table 4 tabulates the results provided by the best-performing model YOLOR-P6 [19] on the input size of  $1280 \times 1280$ . The most accurately predicted class is Bus with an  $mAP_{0.5}$  of 0.8161 followed by Car, Truck, Bike and finally Pedestrian with  $mAP_{0.5}$  of 0.3621.

The Bike has the maximum normalized FP rate at 0.65 when the background is incorrectly detected as Bike. Ad-

ditionally, a substantial fraction of objects in each class remains undetected, as indicated by their normalized FN rates varying between 0.29 to 0.72. Despite this, the model demonstrates comparatively good performance in terms of *Precision* across all classes, with values ranging from 0.77 to 0.95, with the exception of the Pedestrian class, which displays a significantly low *Precision* of 0.49.

#### 4.4.3 Inference Time

The inference time for each model was measured on a workstation featuring an  $11^{th}$  Gen i7 CPU and an Nvidia RTX 3080 GPU, and the results are presented in Table 2. The outcomes demonstrate that all models perform efficiently on this high-end computer, with inference times varying between 4.3 ms to 43.9 ms.

## 5. Discussions

The majority of the dataset, consisting of images from Cameras 1-15, were derived from fisheye surveillance camera footage captured on a single day in July 2018 in Taiwan. Although the dataset contains images of 5 major road participants captured from varying angles and under different illumination conditions, it lacks diversity in terms of weather conditions, such as fog, rain, snow, and storms. Additionally, the dataset is imbalanced, with the class Bike having the highest number of objects at 88K, while the Bus class has the lowest number at 2.98K.

**Hard cases** of the best-performing YOLOR-W6 [19] are represented by few samples in Figure 7.

In Figure 7(a), several examples of false positives are shown where the labeled objects are not detected. These instances can be categorized into two groups: parked/stationary vehicles and road participants in motion. In the top left, only two out of nine scooters parked in a row on the sidewalk are correctly detected. On the top right, two partially visible cars parked in a garage are not detected. The presence of numerous parked vehicles in the dataset and the misdetection of such vehicles contribute to the high false negative rates observed across all classes.

The second type of false positives involves road participants in motion, such as a truck, a pedestrian, and a bus shown in the three crops at the bottom of Figure 7(a)

The examples shown in Figure 7(b) illustrate instances where the background is misclassified as one of the object classes, resulting in higher false positive rates. Specifically, in the top left, a road sign is incorrectly detected as a Pedestrian, while in the bottom left, a yellow building is misclassified as a Bus. In the center, a building pillar is erroneously labeled as a Pedestrian, and on the right, a horizontal road sign is detected as a Bike.

In Figure 7(c), we can observe cases where classes are misclassified as other classes. The four images, from the



Figure 7. Some samples of hard cases of YOLOR-P6 detections on input size  $1280 \times 1280$ .

bottom to the top, show how the predictions change as Pedestrians walk away from the camera. We can see that misclassification occurs when the size of the objects gets smaller. Specifically, the objects were initially correctly detected as Pedestrians when they were closer to the camera, but as they moved away and became smaller, they were misclassified as Bikes.

# 6. Conclusions

We presented the FishEye8K benchmark dataset along with the evaluation of the SoTA one-stage object detectors for the use of fisheye cameras for road object detection. This dataset fills the gap in the lack of a fisheye surveillance camera dataset for road 2D object detection tasks. The anonymized dataset includes 8000 frames with 157K bounding boxes of 5 different road participants and various aspects of road conditions. Our evaluation results show that YOLOv8 and YOLOR models [19], which are pretrained on MS-COCO [11], outperforms the other models. Therefore the FishEye8K dataset will be a significant contribution to the fisheye video analytics and smart city applications.

**Future work** includes the creation of a large and more balanced dataset with more diverse street object categories that can be used for object re-identification model training and evaluation.

Acknowledgements. Emirates Center for Mobility Research (EMCR) provided support for our research through Grant 12R012, while SciDM and National Center for High-performance Computing (NCHC) provided necessary storage resources.

# References

- Chris H Bahnsen and Thomas B Moeslund. Rain removal in traffic surveillance: Does it matter? *IEEE Transactions on Intelligent Transportation Systems*, 20(8):2802–2819, 2018.
  2, 3
- [2] Dawei Du, Yuankai Qi, Hongyang Yu, Yifan Yang, Kaiwen Duan, Guorong Li, Weigang Zhang, Qingming Huang, and Qi Tian. The unmanned aerial vehicle benchmark: Object detection and tracking. In *ECCV*, pages 370–386, 2018. 2, 3
- [3] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html. 2, 4
- [4] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361. IEEE, 2012. 2
- [5] Meng-Ru Hsieh, Yen-Liang Lin, and Winston H Hsu. Dronebased object counting by spatially regularized regional proposal network. In *ICCV*, pages 4145–4153, 2017. 2, 3
- [6] Itay Klein, Nexar Blog. NEXET the largest and most diverse road dataset in the world, 2017. https: //data.getnexar.com/blog/nexet-the-largestand-most-diverse-road-dataset-in-the-world, Last accessed on 2021-10-24. 2
- [7] Glenn Jocher, Alex Stoken, Jirka Borovec, NanoCode012, ChristopherSTAN, Liu Changyu, Laughing, tkianai, yxNONG, Adam Hogan, lorenzomammana, AlexWang1900, Ayush Chaurasia, Laurentiu Diaconu, Marc, wanghaoyang0106, ml5ah, Doug, Durgesh, Francisco Ingham, Frederik, Guilhen, Adrien Colmagro, Hu Ye, Jacobsolawetz, Jake Poznanski, Jiacong Fang, Junghoon Kim, Khiem Doan, and Lijun Yu. ultralytics/yolov5: v4.0 - nn.SiLU() activations, Weights & Biases logging, PyTorch Hub integration, Jan. 2021. 2, 3, 4, 5, 6, 7
- [8] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. 6
- [9] Roman Kvyetnyy, Roman Maslii, Volodymyr Harmash, Ilona Bogach, Andrzej Kotyra, aklin, Aizhan Zhanpeisova, and Nursanat Askarova. Object detection in images with low light condition. In *Photonics Applications in Astronomy, Communications, Industry,* and High Energy Physics Experiments 2017, volume 10445, page 104450W. International Society for Optics and Photonics, 2017. 2
- [10] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *CoRR*, abs/1708.02002, 2017. 6
- [11] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 2, 4, 6, 8
- [12] Zhiming Luo, Frederic Branchaud-Charron, Carl Lemaire, Janusz Konrad, Shaozi Li, Akshaya Mishra, Andrew Achkar, Justin Eichel, and Pierre-Marc Jodoin. Mio-tcd: A new benchmark dataset for vehicle classification and localization. *IEEE Transactions on Image Processing*, 27(10):5129–5141, 2018. 2, 3
- [13] Yanjun Ma, Dianhai Yu, Tian Wu, and Haifeng Wang. PaddlePaddle: An open-source deep learning platform from industrial practice. *Frontiers of Data and Domputing*, 1(1):105, 2019. 6
- [14] Milind Naphade, Shuo Wang, David C Anastasiu, Zheng Tang, Ming-Ching Chang, Xiaodong Yang, Yue Yao, Liang Zheng, Pranamesh Chakraborty, Christian E Lopez, et al. The 5th ai city challenge. In *CVPR*, pages 4263–4273, 2021. 2, 3

- [15] Constantine Papageorgiou and Tomaso Poggio. A trainable system for object detection. *International journal of computer vision*, 38(1):15–33, 2000. 2, 3
- [16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In Advances in Neural Information Processing Systems 32, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [17] Joseph Redmon. Darknet: Open source neural networks in C. http://pjreddie.com/darknet/, 2013-2016. 6
- [18] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. *CoRR*, abs/1506.02640, 2015. 2, 4
- [19] Chien-Yao Wang, I-Hau Yeh, and Hong-Yuan Mark Liao. You only learn one representation: Unified network for multiple tasks. *CoRR*, abs/2105.04206, 2021. 2, 5, 6, 7, 8
- [20] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. arXiv preprint arXiv:2207.02696, 2022. 2, 5, 6, 7
- [21] Longyin Wen, Dawei Du, Zhaowei Cai, Zhen Lei, Ming-Ching Chang, Honggang Qi, Jongwoo Lim, Ming-Hsuan Yang, and Siwei Lyu. DETRAC: A new benchmark and protocol for multi-object tracking. *CoRR*, abs/1511.04136, 2015. 2, 3
- [22] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. In *CVPR*, pages 2174–2182, 2017. 2
- [23] Senthil Yogamani, Ciarán Hughes, Jonathan Horgan, Ganesh Sistu, Padraig Varley, Derek O'Dea, Michal Uricár, Stefan Milz, Martin Simon, Karl Amende, et al. WoodScape: A multi-task, multi-camera fisheye dataset for autonomous driving. In *ICCV*, pages 9308–9318, 2019. 2
- [24] 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*, pages 2636–2645, 2020. 2