

Investigating the Performance of Real-time Object Detection Frameworks YOLOv8, YOLOv9 and YOLOv10 for Object Detection in adverse weather conditions

“How do the real-time object detection frameworks; YOLOv8, YOLOv9, and YOLOv10, comparatively perform in terms of model efficiency in detecting various objects in adverse weather conditions? “

A Computer Science Extended Essay

Word Count: 3,984
Session: November 2024

CS EE World
<https://cseeworld.wixsite.com/home>
30/34 (A)
November 2024

Submitter info:
Hey, Nethra here :)
feel free to contact me at [neythrasubr \[at\] gmail \[dot\] com](mailto:neythrasubr@gmail.com) if you want me to review your EE or something.

Table Of Contents

| | |
|-------------------------------------|-----------|
| 1 Introduction..... | 3 |
| 1.1 Context and Scope | 3 |
| 1.2 Research Question | 4 |
| 1.3 Related Work | 5 |
| 2 Background | 5 |
| 2.1 Neural Networks | 5 |
| 2.1.1 Definitions..... | 5 |
| 2.2 You Only Look Once (YOLO) | 9 |
| 2.2.1 YOLOv8 | 10 |
| 2.2.2 YOLOv9 | 12 |
| 2.2.3 YOLOv10 | 14 |
| 3 Methodology | 16 |
| 3.1 Dataset..... | 16 |
| 3.2 Variables | 18 |
| 3.3 Evaluation Metrics | 19 |
| 3.4 Model Training..... | 21 |
| 3.4.1 YOLOv8 | 21 |
| 3.4.2 YOLOv9 | 22 |
| 3.4.3 YOLOv10 | 23 |
| 3.5 Hypothesis..... | 25 |
| 4 Results and Analysis | 25 |
| 4.1 Tabular Representation | 25 |
| 4.2 Graphical Representation..... | 26 |
| 4.2.1 F1-Confidence..... | 27 |
| 4.2.2 Precision on Recall | 28 |
| 4.2.3 Precision on Confidence | 28 |

| | |
|---|-----------|
| 4.2.4 Recall on Confidence | 29 |
| 4.2.5 Confusion Matrix | 30 |
| 4.3 Analysis..... | 32 |
| 5 Conclusion | 35 |
| 6 Further Scope | 35 |
| 7 Limitations and Challenges..... | 36 |
| 8 Bibliography | 37 |
| 9 Appendices..... | 42 |
| 9.1 Appendix A: Terminology | 42 |
| 9.2 Appendix B: Background | 42 |
| 9.2.1 Artificial Neural Networks (NNs) | 42 |
| 9.2.2 Convolutional Neural Networks (CNNs)..... | 43 |
| 9.3 Appendix D: Investigation Code | 46 |

1 Introduction

1.1 Context and Scope

There is an extensive use of Computer Vision (CV) in today's day of age, and the proliferation of its use has advanced. CV is often used to substitute manual labour in organisations and enterprises. One such interesting domain is Rescue Robots. Integrating CV into Search and Rescue missions (SAR) facilitates their deployment in disaster relief and search operations (*Flynn et al*).

CV is made up of Neural Networks (NNs). NNs are capable of processing raw data to output valuable information (*"Computer Vision"*). Datasets, which consist of labelled data, are fed into these networks to facilitate the creation of scenarios. Numerous other factors influence the prediction's accuracy such as frameworks which the model is built upon. Each framework has unique displayed behaviour when applied to varying datasets (*Li and Luo*). Consequently, a framework that performs exceptionally well with a dataset might perform poorly with dataset. The variability necessitates the development of multiple frameworks and versions (*"Machine Learning Frameworks"*).

For instance, during the aftermath of the meltdown of nuclear reactors in the Fukushima Daiichi Nuclear Powerplant (*"Fukushima Daiichi Accident"*). There was an urgent need for rescue robots to (*Westcott*) assess the damage, search for survivors and conduct repairs. The site was also experiencing harsh weather conditions such as heavy rain and strong winds. This made it hard for the available robots to navigate or perform tasks as they were unable to navigate through dust, weather and debris (*Li et al.*).

Due to the collection of environmental data through hardware such as cameras, there is little to no validation of the data received which could help differentiate and eliminate visuals which can't be interpreted by the model, thus raising a problem statement.

1.2 Research Question

In the recent advances in Artificial Intelligence; (AI) You Only Look Once (YOLO), which is a widely used framework and influential method for object detection (OD) has been brought to significant attention, and is often used for Real-time object detection tasks. Redmon et al introduced YOLO in 2015 since then, scholars have published updated iterations of the concept.

This investigation aims to examine algorithms YOLOv8, YOLOv9 and YOLOv10. Recent versions of YOLO frameworks such as YOLOv10 released in May 2024 and YOLOv9 released in February 2024 lack extensive documentation, particularly in comparison with YOLOv8. Notably, these three frameworks have yet to be compared against each other regarding the robustness of their performance in varying weather conditions.

This investigation will conduct an extensive evaluation of the performances of frameworks; YOLOv8, YOLOv9 and YOLOv10 evaluating these concerning the mean average precision, recall, and precision. Therefore, the research question proposed is: ***“How do the real-time object detection frameworks; YOLOv8, YOLOv9, and YOLOv10, comparatively perform in terms of model efficiency in detecting various objects in adverse weather conditions?”*** Although this investigation does not rely on multimodal data inputs nor videos as datasets, it emphasises the

detection of objects in images compromised by weather conditions— providing more insight into object detection in still images. With the recognition that video data or multimodal datasets could further enhance the quality of the research, the study aims to maintain simplicity while offering a foundation for future development and research.

1.3 Related work

Though there have been numerous studies on CV there seems to be a lack of documentation with regards to YOLOv10 and YOLOv9. A journal published in ISPRS examines the necessity of specialised enhancements in object detection models for adverse weather conditions but there has been a lack of specificity in the analysis of YOLO (*Toma et al.*). Similarly, another paper released in 2019 shares a valuable understanding of how YOLO faces challenges in adverse weather conditions and further suggests the use of multimodal approaches— this paper however lacks in-depth research using the latest YOLO models (*Wang et al.*).

2 Background

2.1 Neural Networks

2.1.1 Definitions

In addition to this, it is recommended to hover through Appendix A and Appendix B to gain a profound understanding of Neural networks and Convolutional neural networks.

| Terminology | Definition |
|---------------------|--|
| Backbone | Extracts crucial features from input dataset through the use of CNNs |
| Neck | Connects the backbone to the head. Helps aggregate features provided by backbone |
| Head | Uses features extracted through the neck and backbone to make predictions. |
| Convolutional Layer | Appliance of convolutional operations to input data to capture spatial hierarchies in images by extracting features. (<i>LeCun et al.</i>) |
| Convolution | The mathematical operation is applied in convolutional layers to filter data by sliding a kernel/filter across input aimed at producing a feature map. (<i>LeCun et al.</i>) |
| Feature maps | Represents various features example, edges and textures. (<i>LeCun, Bengio, and Hinton</i>) |
| Kernal size | Dimension of filter used in a convolutional layer expressed as height multiplied by width. (<i>LeCun et al.</i>) |
| Channel count | Number of channels in an image or features (RGB or feature representations). (<i>Jain</i>) |
| Post-Processing | Techniques applied after a model has done learning. (" <i>References for Papers</i> ") |

| | |
|-----------------------|--|
| Activation Function | Mathematical structures applied to a neural network's output to introduce non-linearity, helping the network to model complex patterns. (<i>Jain</i>) |
| Inference | Using a trained machine learning model to make predictions based on new data. (<i>Jain</i>) |
| Inference cost | Computational resources are required to generate predictions for inference. (<i>Chollet</i>) |
| Bottleneck | Series of layers where the dimensionality of feature representation is reduced to compress the information or improve computational efficiency. (<i>Chollet</i>) |
| Anchor free detection | Methods that don't rely on predefined anchor boxes, rather predicting object locations directly. (<i>Tian et al.</i>) |
| Lightweight model | A network designed to be resource-efficient. (<i>Redmon and Farhadi</i>) |
| Up-sample | Process of increasing the spatial resolution of feature maps. (<i>Odena et al.</i>) |
| Downsample | Process of reducing spatial resolution of feature maps. (<i>LeCun et al.</i>) |
| Decoupling | Separation of different aspects of processing to allow more efficient architectural designs. |
| Spatial reduction | Decreasing spatial dimensions of feature maps. (<i>LeCun et al.</i>) |
| Pointwise convolution | Convolutional operation with kernel (<i>Redmon and Farhadi</i>) |

| | |
|---------------------------------|--|
| Cross-stage partial connections | A technique where feature maps are partially propagated through different stages of the network– enhances gradient flow. (<i>Wang et al.</i>) |
| Spatial Pyramid Pooling Faster | Pools feature maps at multiple scales to improve OD. (<i>Hirschfeld and Ziemann</i>) |
| Non-Maximum Suppression (NMS) | Post-processing technique which eliminates redundant bounding box predictions and selects the ones with high confidence. |
| Non-uniform matching | Approach to OD where matching between predicted bounding boxes and ground truth boxes are adjusted dynamically or irregularly. |
| NUMS post-processing | Post-processing technique which leverages non-uniform matching strategies during. |
| One-to-one head | Each anchor point is assigned a single object label, facilitating simpler matching. (<i>He et al.</i>) |
| One-to-one matching predictions | Predicted objects are matched with one ground truth object during training. (<i>Zhou et al.</i>) |
| Auxiliary framework | Additional components or branches are added to the main model to support the learning process by supplying gradient information and enhancing performance. (<i>Szegedy et al.</i>) |
| Gradient path planning | A technique used to optimise the flow of gradients when backpropagation occurs. |

| | |
|--------------------|---|
| Transformer Model | Relies on self-attention mechanisms to process input in parallel. Efficiently handles sequential data. |
| Vision transformer | NN architecture that uses transformer models. |
| Python SDK | Software development kit with a compilation of kits, documentation and libraries. |

Table 1: Important Definitions and Terminologies

2.2 You Only Look Once (YOLO)

You Only Look Once (YOLO) algorithm is a family of object detectors that have iterated throughout the years since its initial release by Joseph Redmon in 2015 (*Redmon et al.*). Each iteration advanced its predecessors by including a novel method. YOLO is notoriously known for its ability to proceed with real-time object detection by dividing the input images into a grid matrix and predicting the bounding boxes along with its class probabilities in parallel (*Redmon et al.*).




Figure 1: Process of bounding box prediction; centered coordinate prediction; use of sigmoid function; use of the location of the grid cell.

2.2.1 YOLOv8

Compared to its predecessors YOLOv8 includes a newer backbone that consists of a CSPDarknet architecture (*Touvron et al.*) which holds 53 convolutional layers and as well as equips itself with cross-stage partial connections, along with this, YOLOv8 uses the SiLU activation function which mitigates the vanishing gradient problem (*Chen et al.*). This intensifies information transmission in deep neural networks. The C2f module combines features on a high level with context to enhance detection accuracy. Spatial pyramid pooling faster (SPPF) (*Simonyan and Zisserman*) is another module and the other convolutional layers process the features in variable scales.

In YOLOv8 the head is detachable hence it handles classification, object scores and regression work independently— due to this the overall accuracy is increased. Up sample (U) layers help increase the resolution of provided feature maps. Convolutional layers are included in the head to analyse these feature maps. Overall, the head is designed to optimise its speed and accuracy and hence consideration is given to kernel sizes and channel count of each layer.

YOLOv8 uses anchor-free detection to speed up post-processing (Non-Maximum Suppression) additionally YOLOv8 also includes a newer convolutional layer aimed to detect features using learnable filters. The input layers are detected in variable resolutions and sizes to allow the network to be versatile (*Fischer et al.*).

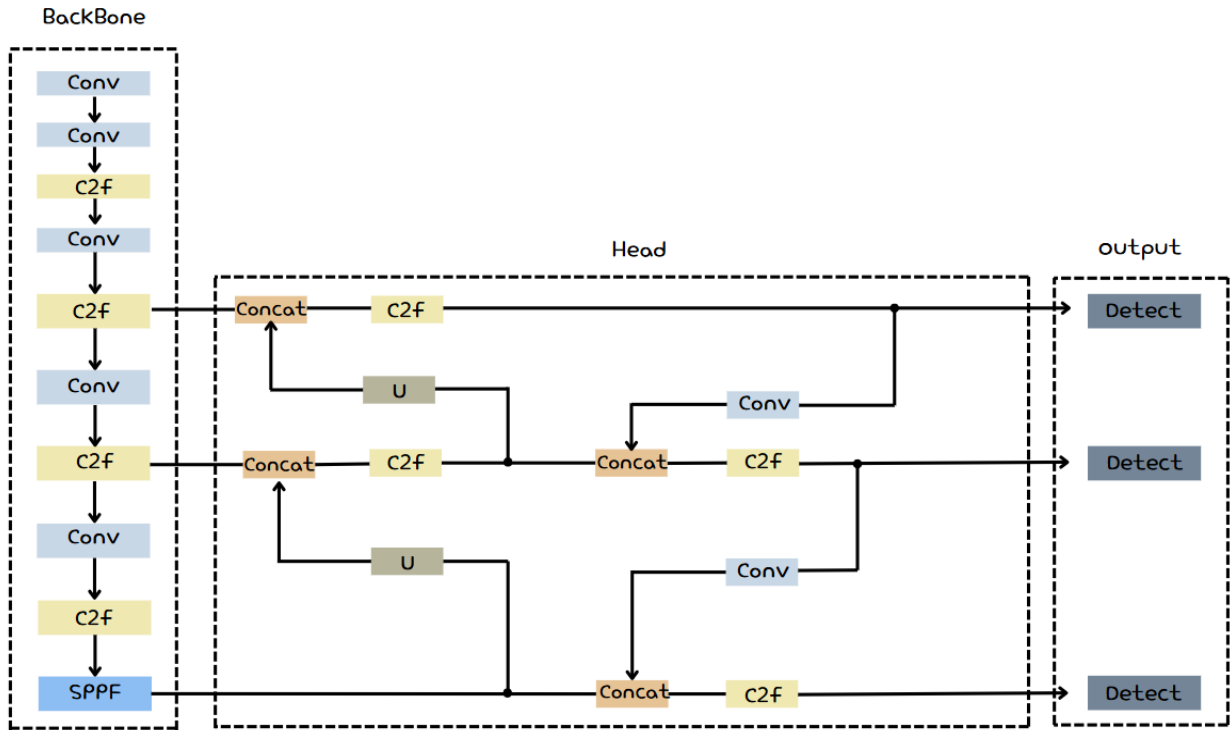


Figure 2: YOLOv8 Architecture

Predecessors of YOLOv8 had equipped a C3 convolutional layer which YOLOv8 replaces using a C2f layer. The C2f layers help utilise all the bottlenecks. Splitting YOLOv8 achieves parallel processing. Additionally, the kernel size has been increased in YOLOv8.

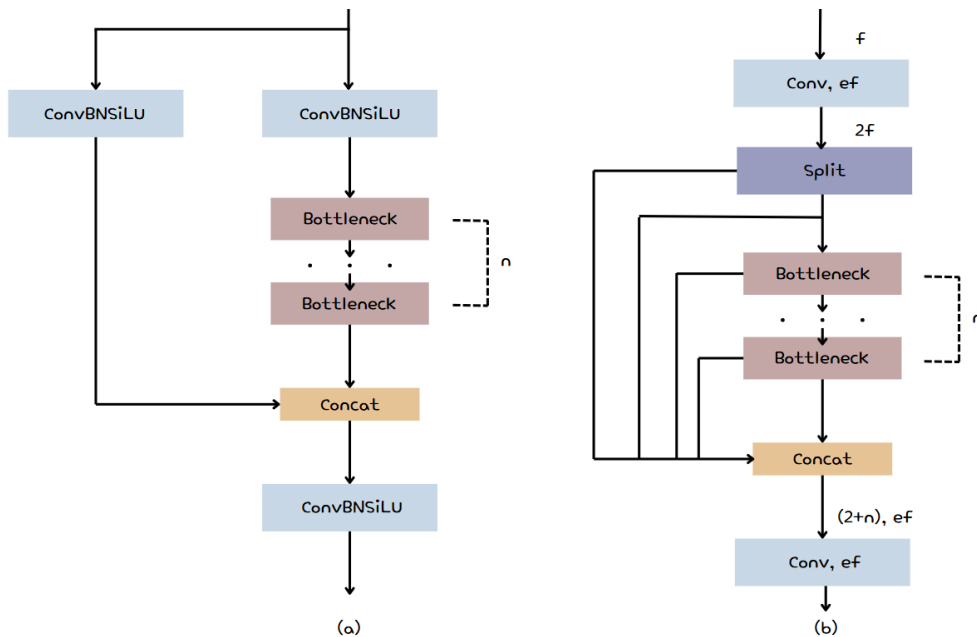


Figure 3: (a) defines bottlenecks in YOLOv5 defined as n (b) defines bottleneck layers in YOLOv8 defined as n

In addition to these architectural features, YOLOv8 employs user accessibility by integrating Python SDK and Command Line Interface (CLI) which further supports programmers to be able to use this model (*Ahmed et al.*).

2.2.2 YOLOv9

During bottleneck (*Kingma and Ba*) conditions, it can be especially hard for ML models to be able to interpret data. Traditional methods of Masked modelling (*Wang et al.*) and reversible architectures (*Nijkamp et al.*) had to be iterated due to their drawbacks. To address these issues especially the bottleneck situation YOLOv9 has proposed Programmable gradient information (PGI). PGI generates reliable gradient information for model network weight updates. Generalised ELAN (*Jiang et al.*) (GELAN) simultaneously takes multiple parameters such as accuracy, computational efficiency, and speed into account which allows this design to allow programmers to make decisions upon choosing appropriate computational blocks for different devices. The combination of PGI and GELAN gave rise to YOLOv9. YOLOv9 proposes that increasing the model size accumulating more parameters and adding enhanced data transformers can help information retainment– however, this does not address the issue completely.

The introduced auxiliary framework PGI is divided into three components. As shown in the figure 4. PGI is dependent on the main branch and hence it does not require inference costs, the remainder is utilised to solve and precisely investigate important issues in learning methods. Auxiliary reversible branch tries to deal with information bottleneck which occurs when the network is deepened which will cause the loss function to be incapable of providing reliable gradients.

Finally, multi-level auxiliary information handles error accumulation problems which are generally caused due to deep supervision.




Figure 4: Three divisions of PGI

GELAN is a new network architecture made from the combination of CSPNet (Ramesh *et al.*) and ELAN (Jiang *et al.*) which are innately made with gradient path planning. YOLOv9’s GELAN takes into consideration the weight of the model, inference speed, and accuracy. GELAN validates PGI, especially in lightweight models.

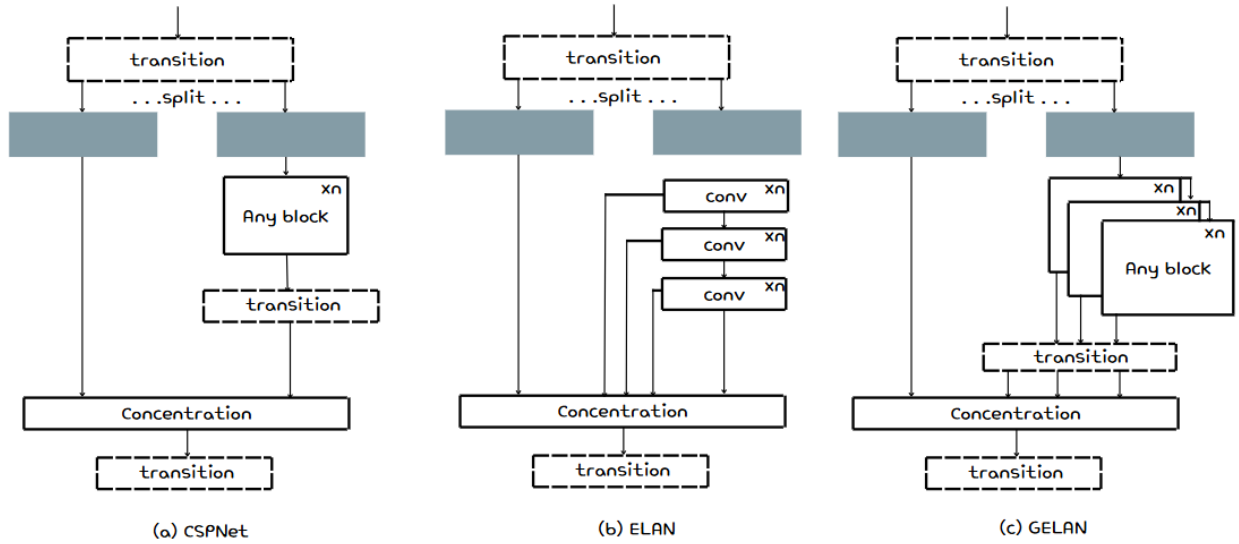


Figure 5: (a) CSPNet (b) ELAN and (c) GELAN architectures

2.2.3 YOLOv10

End-to-end object detection is a shift of methodology from the past where traditional pipelines (*Srivastava et al.*) were used. Baidu's RT-DETR (*Raffel et al.*) is used which is a vision transformer architecture. Hungarian loss is also occupied by it to reach one-to-one matching predictions (*Viana*) and hence it eliminates post-processing.

YOLOs rely on NUMS post-processing to allocate positive samples for each instance to leverage TAL (*Vaswani et al.*) during training— this is detrimental to the model's inference efficiency. YOLOv10 provides an NMS-free training strategy through the use of dual labels and consistent matching metrics (*"YOLOv10 Documentation"*).

As shown in Figure 6 the incorporation of one-to-one head is seen. The optimisation objective remains the same as the original branch of one-to-many but this leverages to obtain label assignments (*Xu et al.*).

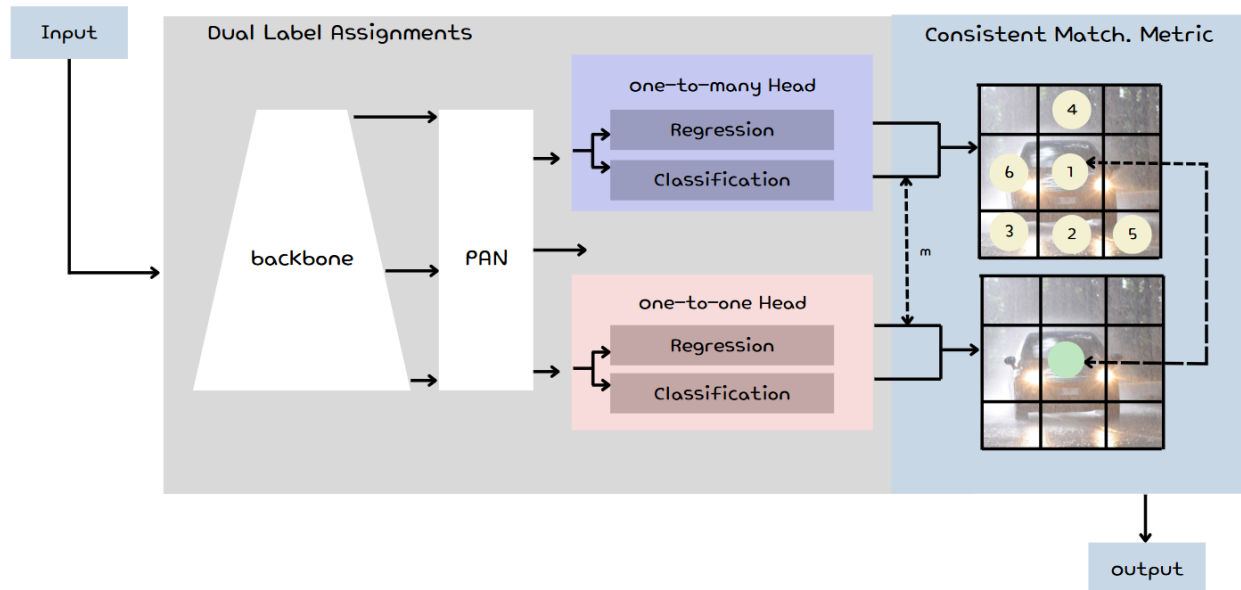


Figure 6: Architecture of YOLOv10

The conjoining of two heads allows the backbone to have optimal supervision (*"YOLOv10: An Introduction"*). While inference is employed to avoid inference cost, a one-to-many head is discarded and a one-to-one head is used, which beats Hungarian matching by utilising lesser training time (*Xu et al.*).

The final regression head takes the significance performance of YOLO compared to the classification head. We can reduce the overhead of the classification head without hurting the performance; therefore, a lightweight architecture is employed (*"YOLOv10: Everything You Need to Know"*).

YOLOs use standard convolution stride to achieve spatial downsampling and channel transformation simultaneously. This raises computational costs and parameter counts. YOLOv10 decouples the spatial reduction and channel increase operations. Pointwise convolution is leveraged to modulate the dimension of the channel and further utilize depthwise convolution to achieve spatial downsampling. This maximises the information retained during downsampling as well as computational cost (*Xu et al.*).

Rank-guided block design scheme decreases the complexity of redundant stages; essentially making a compact model. Compact inverted block (CIB) adopts depthwise convolution for spatial mixing and pointwise convolutions. Rank-guided block allocation helps achieve maximum efficiency while also maintaining good capacity (*Xu et al.*).

3 Methodology

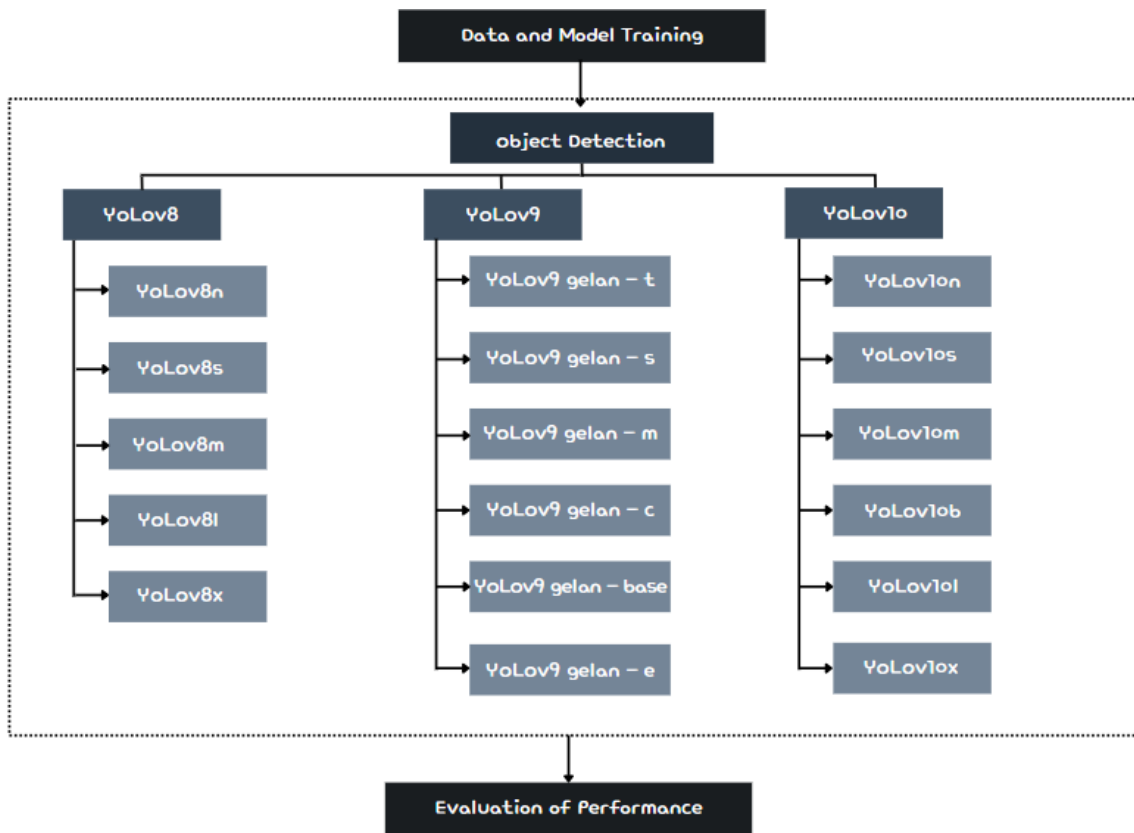


Figure 7: Flow chart of investigation to be proceeded

3.1 Dataset

Roboflow is an online directory of CV model resources, including datasets. It allows the download of these data sets through multiple version formats and these datasets can be easily integrated into Google Colaboratory through the use of APIs. BDD100K is another larger dataset comprising videos and still images of objects in adverse weather conditions. However, due to limitations in hardware along with Roboflow's ease of integrating datasets into Google Colaboratory and its service of cloud-based data storage, Roboflow was utilised for this investigation.

The selected public dataset “O.D IN BAD WEATHER “(*"Object Detection in Bad Weather Dataset"*) contains images of labelled subjects including Bikes, Buses, Cars, Motors, Persons and Riders. The dataset has innately been pre-processed using auto-orientation and resized to 640x640 pixels by its creator (*"Foreign Object Aerodromes"*). Furthermore, it is split into training (815 images) validation (218 images), and test (117 images) subsets.

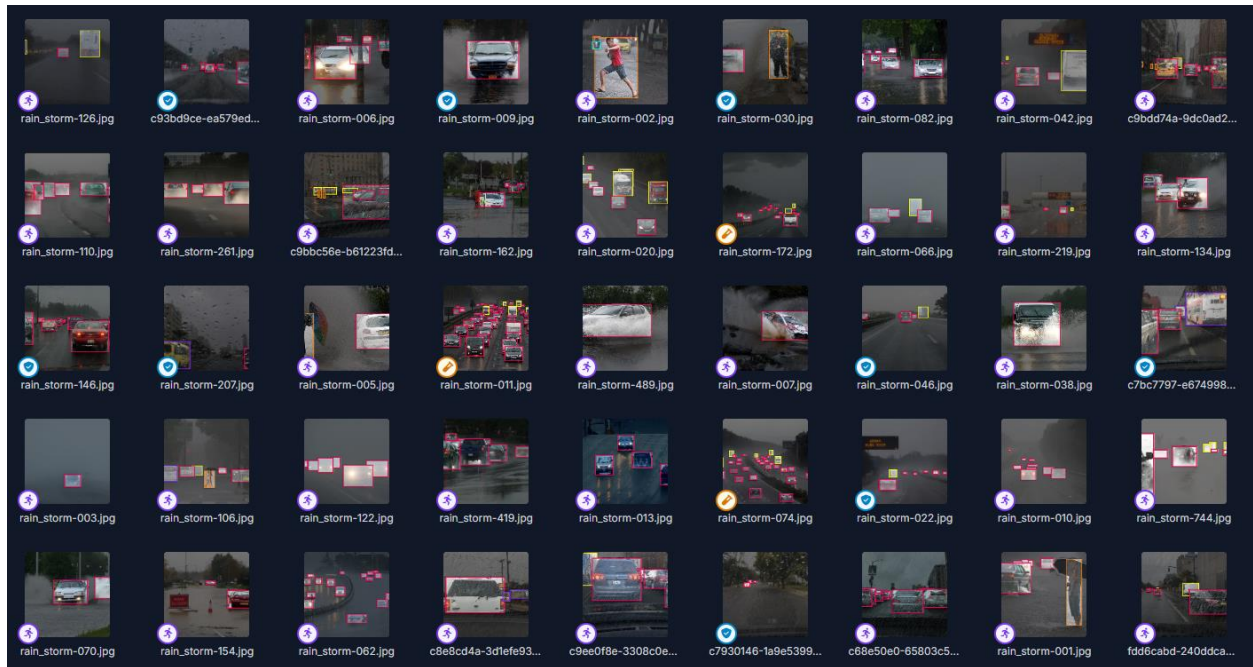


Figure 8: Visual view of the images compiled in the public dataset (*"Object Detection in Bad Weather Dataset"*)

The dataset is integrated from Roboflow by the following snippet:

```
!pip install roboflow

from roboflow import Roboflow
rf = Roboflow(api_key="XXXXXX")
project = rf.workspace("XXXXXX").project("XXXXXX")
version = project.version(x)
dataset = version.download("XX")
```

Figure 9: Integration of Roboflow API

3.2 Variables

During training, the following parameters were kept constant across each framework:

| Parameter | Value |
|-----------|-------|
| Epoch | 25 |
| Batch | 16 |
| img/imgz | 640 |
| Plots | true |

Table 2: Showcases the parameters and chosen value

The dataset used remained constant throughout the investigation, with only the dataset version differing, which does not affect the final result; these variables act independently in this investigation.

The dependent variables are the obtained results from each framework post-training. These results are retrieved by visualising the trained models. The results are stored in the training results directory in the workspace and the visualization is accomplished using `IPython.display`. These raw results are then evaluated and compared against each other.

3.3 Evaluation Metrics

All models were evaluated using post-training results which are accessible in each model's results directory. All the results were measured using a constant evaluation metric. The following metrics were visualized:

1. Confusion Matrix
2. F1-Confidence Curve
3. Precision-Recall Curve
4. Precision-Confidence Curve
5. Recall-Confidence Curve

Additionally, overall model results were also visualized, including mean average precision (mAP), losses in validation, test and training datasets, precision and recall.

$$Precision = \frac{True\ Positive\ (TP)}{True\ Positive\ (TP) + False\ Positive\ (FP)} \quad [1]$$

$$Recall = \frac{True\ Positive\ (TP)}{True\ Positive\ (TP) + False\ Negative\ (FN)} \quad [2]$$

$$F1\ score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad [3]$$

$$mAP = \frac{\sum_{j+1}^k AP_i}{k} \quad [4]$$

Precision helps evaluate the robustness of the object detected whereas recall indicates the model's ability to be able to detect instances of concern in the data input. The **F1 score** helps us measure both precision and recall in a balanced manner. Finally, **mAP** helps compare the growth truth bounding box and the bounding box detected by the model. The **Confusion matrix** summarizes the performance of an ML model on a set data set and includes specifics such as classes also known as the error matrix.

A **confusion matrix** is usually structured through the use of:

1. True Positive (TP): Number of instances where the model had correctly predicted the positive classes
2. False Positive (FP): Number of instances where the model had incorrectly predicted positive classes. (Predicted positive, but the real class was negative.)
3. True Negative (TN): The number of instances where the model had correctly predicted the negative classes.
4. False Negative (FN): The number of instances where the model had incorrectly predicted negative classes. (Predicted negative, but the real class value was positive.)

Higher TP and TN values showcase the model's performance as well subsequently higher FP values suggest the model is making multiple mistakes.

3.4 Model Training

To proceed with training the frameworks; Cloud GPU provided by Google Colaboratory was utilized due to GPU limitations. The following details specify the GPU used:

| | | | | | | | |
|-----------------------|----------|---------------|------------------|----------------------------|----------------------|--------------------|--------|
| NVIDIA-SMI 535.104.05 | | | | Driver Version: 535.104.05 | | CUDA Version: 12.2 | |
| GPU | Name | Persistence-M | Bus-Id | Disp.A | Volatile Uncorr. ECC | | |
| Fan | Temp | Perf | Pwr:Usage/Cap | Memory-Usage | GPU-Util | Compute M. | MIG M. |
| 0 | Tesla T4 | Off | 00000000:00:04.0 | Off | 0 | | |
| N/A | 47C | P8 | 9W / 70W | 0MiB / 15360MiB | 0% | Default | N/A |

Figure 10 : details of GPU used for investigation

Each framework was run through constant parameters and visualised post-training.

3.4.1 YOLOv8

YOLOv8 was installed directly in the workspace using a pip function

```
!pip install ultralytics==8.0.196

from IPython import display
display.clear_output()

import ultralytics
ultralytics.checks()
```

Figure 11: Installation of YOLOv8

Subsequently, the dataset was imported to the workspace using RoboFlow. The dataset version used for this model was yolov8 and the training was conducted using the set parameters, as shown in the code snippet below.

```
!yolo task=detect mode=train model=yolov8s.pt
data=/content/datasets/O.D-IN-BAD-WEATHER-1/data.yaml epochs=25 imgsz=640
batch=16 plots=True
```

Figure 12: Data training on YOLOv8

Post the training, the model was visualised by outputting the training results.

3.4.2 YOLOv9

Despite the similarities with YOLOv8 while initializing the training process, due to the recent release of YOLOv9 it yet cannot be imported through a pip function, therefore the framework had to be cloned from its official GitHub repository.

```
!git clone https://github.com/SkalskiP/yolov9.git
%cd yolov9
!pip install -r requirements.txt -q
```

Figure 13: Installation of YOLOv9

Additionally, the YOLOv9 model at the current stage cannot download the weights directly, so they were manually downloaded from GitHub and stored separately in the workspace.

```
# YOLOv9 can not automatically download the weights so they are manually
downloaded
!wget -P {HOME}/weights -q
https://github.com/WongKinYiu/yolov9/releases/download/v0.1/yolov9-c.pt
!wget -P {HOME}/weights -q
https://github.com/WongKinYiu/yolov9/releases/download/v0.1/yolov9-e.pt
!wget -P {HOME}/weights -q
https://github.com/WongKinYiu/yolov9/releases/download/v0.1/gelan-c.pt
!wget -P {HOME}/weights -q
https://github.com/WongKinYiu/yolov9/releases/download/v0.1/gelan-e.pt
```

Figure 14: Installation of YOLOv9 weights

The parameters were kept constant for YOLOv8, and the training results were extracted and visualized post-training.

```
%cd {HOME}/yolov9

!python train.py \
--batch 16 --epochs 25 --img 640 --device 0 \
--data {dataset.location}/data.yaml \
--weights {HOME}/weights/gelan-c.pt \
--cfg models/detect/gelan-c.yaml \
--hyp hyp.scratch-high.yaml
```

Figure 15: Data training on YOLOv9

3.4.3 YOLOv10

Similar to YOLOv9, the recent release of YOLOv10 means it cannot be downloaded using a pip function and hence it was to be imported by cloning its GitHub repository.

```
!pip install -q git+https://github.com/THU-MIG/yolov10.git
```

Figure 16: Installation of YOLOv10

Following this, the weights were manually downloaded and stored in the workspace separately.

```
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10n.pt
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10s.pt
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10m.pt
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10b.pt
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10x.pt
!wget -P {HOME}/weights -q
https://github.com/THU-MIG/yolov10/releases/download/v1.1/yolov10l.pt
```

Figure 17: Installation of YOLOv10 weights

The dataset was downloaded in the same manner similar to YOLOv9 and YOLOv8. However, the model version remained yolov9 due to the unavailability of a yolo10 data type version. This change did not affect the performance. YOLOv10 was trained using constant parameters and visualized post-training results.

```
!yolo task=detect mode=train epochs=25 batch=16 imgsz=640 plots=True \
model={HOME}/weights/yolov10n.pt \
data=/content/datasets/O.D-IN-BAD-WEATHER-1/data.yaml
```

Figure 18: Data training on YOLOv10

3.5 Hypothesis

This study hypothesises that YOLOv10 will outperform the other models in precision, especially for well-defined object classes like cars and motors. NMS-free strategy and dual-head architecture employed by YOLO will aid this. However, YOLOv8 may have a higher recall because of its CSPDarknet backbone and finally YOLOv9 will exhibit higher precision in detecting small or occluded objects due to the introduction of PGI and GELAN.

4 Results and Analysis

4.1 Tabular Representation

| | Precision | | | Recall | | | mAP50 | | |
|--------|-----------|--------|---------|--------|--------|---------|--------|--------|---------|
| Class | YOLOv8 | YOLOv9 | YOLOv10 | YOLOv8 | YOLOv9 | YOLOv10 | YOLOv8 | YOLOv9 | YOLOv10 |
| All | 0.677 | 0.475 | 0.658 | 0.263 | 0.377 | 0.192 | 0.286 | 0.378 | 0.203 |
| Bike | 0.603 | 0.334 | 1 | 0.04 | 0.12 | 0 | 0.7 | 0.115 | 0 |
| Bus | 0.523 | 0.424 | 0.418 | 0.425 | 0.525 | 0.275 | 0.412 | 0.469 | 0.288 |
| Car | 0.704 | 0.705 | 0.603 | 0.709 | 0.749 | 0.631 | 0.741 | 0.777 | 0.655 |
| Motor | 1 | 0.575 | 1 | 0 | 0.145 | 0 | 0.02 | 0.209 | 0.7 |
| Person | 0.498 | 0.454 | 0.301 | 0.261 | 0.355 | 0.198 | 0.285 | 0.348 | 0.171 |
| Rider | 1 | 0.372 | 1 | 0 | 0.211 | 0 | 0.081 | 0.286 | 0 |
| Truck | 0.414 | 0.461 | 0.288 | 0.403 | 0.532 | 0.242 | 0.369 | 0.445 | 0.203 |

Table 3: Comparison of YOLOv8, YOLOv9 and YOLOv10's performances regarding Precision, Recall and mAP50

Precision: YOLOv8 provides the highest Precision overall (0.677), compared to YOLOv9 (0.475) and YOLOv10 (0.658). YOLOv9 has the lowest overall precision, suggesting that it may produce more false positives, lastly, YOLOv10 has slightly lower precision than YOLOv8 but higher than YOLOv9, indicating YOLOv10 holding a balance.

Recall: YOLOv8 employs the lowest Recall (0.263) hence it misses more objects in the input; This could be a significant drawback in addressing adverse weather conditions. YOLOv9 has a higher Recall (0.377) than YOLOv8 but still is behind YOLOv10, indicating a better ability to detect in challenging environments. YOLOv10 has the highest recall (0.658), making it able to detect most objects which is crucial in adverse weather.

mAP50: YOLOv8 shows a middle-group performance in mAP50 (0.286) but is however outperformed by YOLOv9 (0.378), indicating YOLOv9 has better accuracy. Lastly, YOLOv10 has the lowest mAP (0.203), indicating that while it detects multiple objects, the accuracy of these predictions is lower compared to YOLOv9 and YOLOv8.

4.2 Graphical Representation


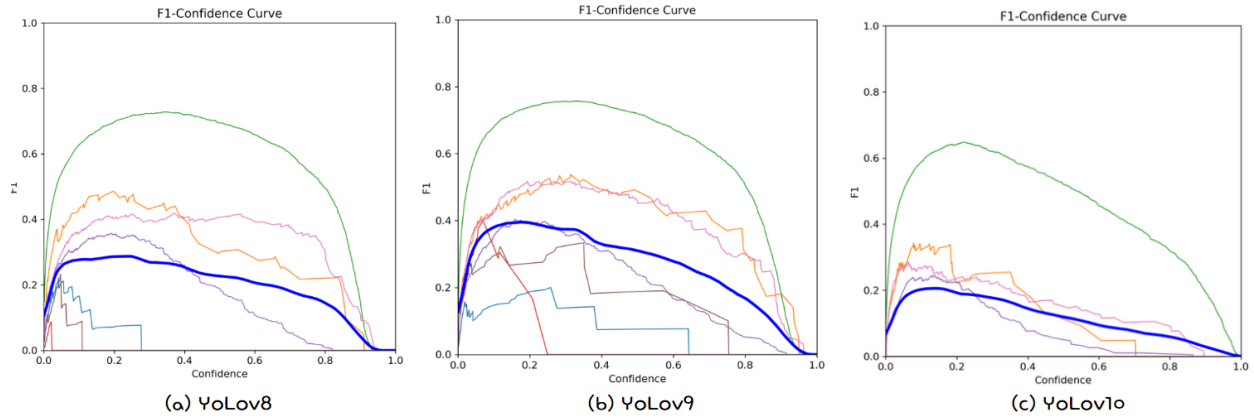


Figure 19: Colour coding for object classes

The image above showcases the colour assigned to each object class. The results graph each object class result on each evaluation metric.

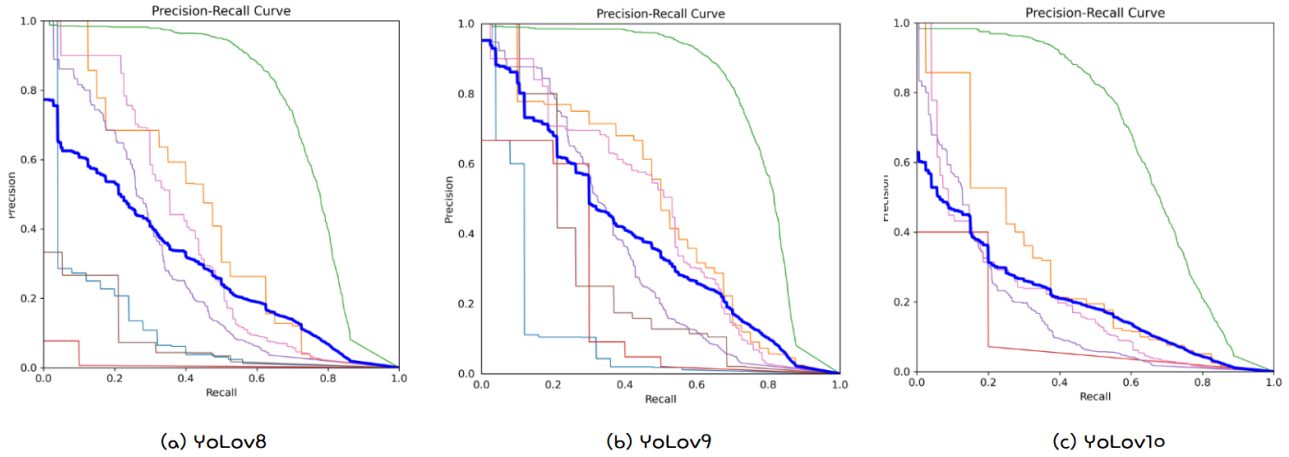
4.2.1 F1-Confidence



Graph 1: Graphical Result of YOLOv8, YOLOv9 and YOLOv10 on F1-Confidence Curve

YOLOv9 stands out to be the most balanced model across all the classes, with a higher and more stable F1 score over a range of confidence levels. YOLOv10's tabular results reflect its generally lower F1 scores. YOLOv8 lastly, provides a more moderate performance, with less consistent F1 scores which makes this framework more consistent than YOLOv10 and less favourable compared to YOLOv9.

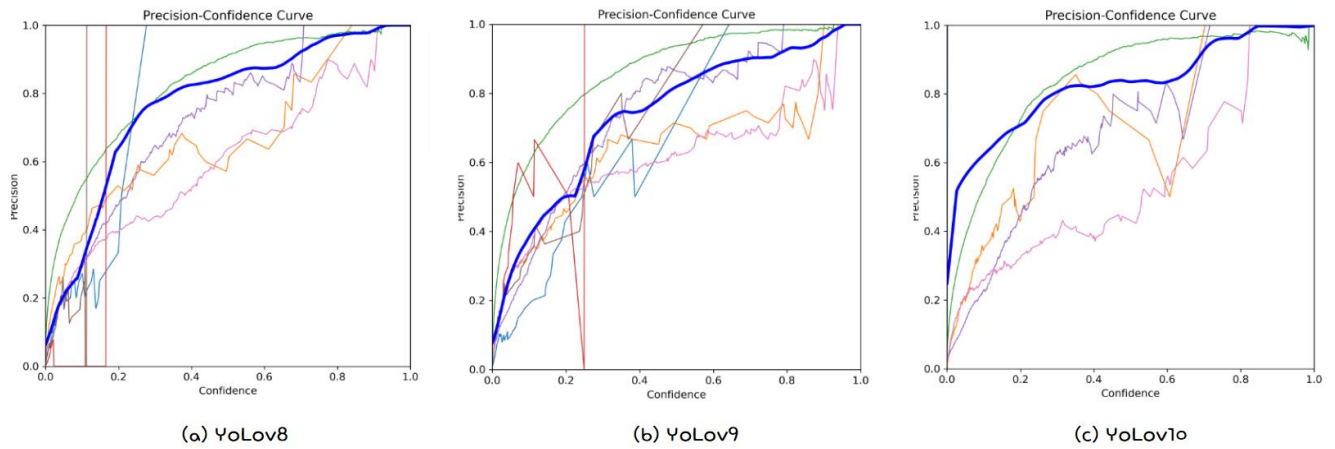
4.2.2 Precision on Recall



Graph 2: Graphical Result of YOLOv8, YOLOv9 and YOLOv10 on Precision/Recall Curve

YOLOv9 exhibits its balance by offering the best trade-off between precision and recall across most object classes consequently YOLOv10 shows strong precision for certain classes but tends to lose precision rapidly as recall increases, particularly for smaller or more complex objects. Lastly, YOLOv8 is less consistent than YOLOv9 with a more significant drop in precision as recall increases across several classes.

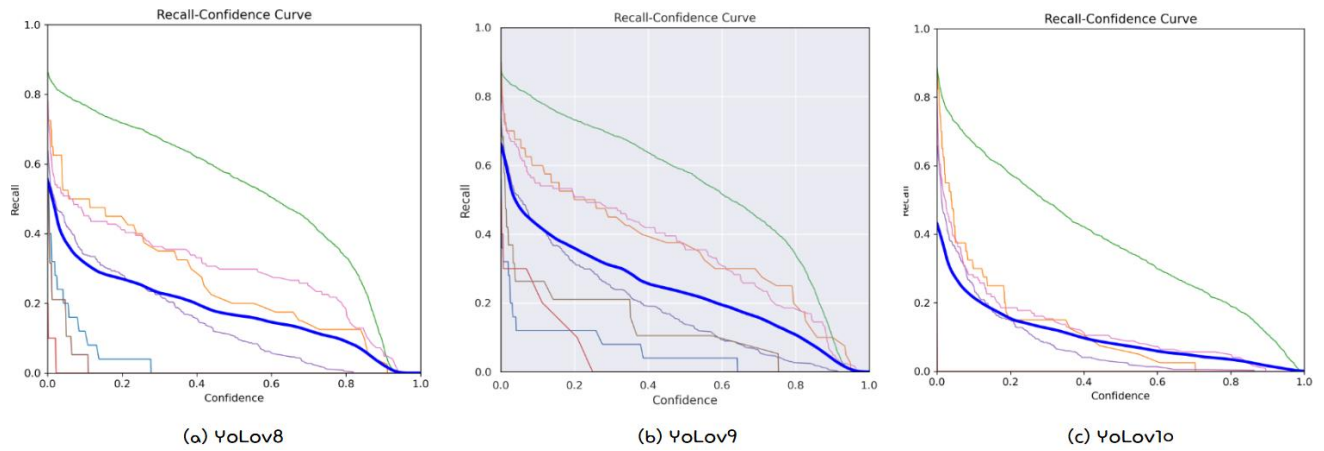
4.2.3 Precision on Confidence



Graph 3: Graphical Result of YOLOv8, YOLOv9 and YOLOv10 on Precision/Confidence Curve

YOLOv9 is presented as the most balanced model, by showcasing consistent and dependable accuracy thresholds for the majority of object classes. YOLOv10 achieves excellent precision but with greater variability between classes; overall, its performance may be less consistent and more class-dependent. Lastly, while YOLOv8 performs reasonably well, YOLOv9 often outperforms it, particularly when it comes to consistency across various object classes.


4.2.4 Recall on Confidence




Graph 4: Graphical Result of YOLOv8, YOLOv9 and YOLOv10 on Recall/Confidence Curve

YOLOv9 once again is exhibited to be the most balanced, by maintaining higher recall across a range of confidence levels. YOLOv10 rapidly declines in recall as confidence increases, indicating a strong preference for high precision at the expense of recall. YOLOv8 generally struggles to maintain recall as the confidence is increased especially compared against YOLOv9

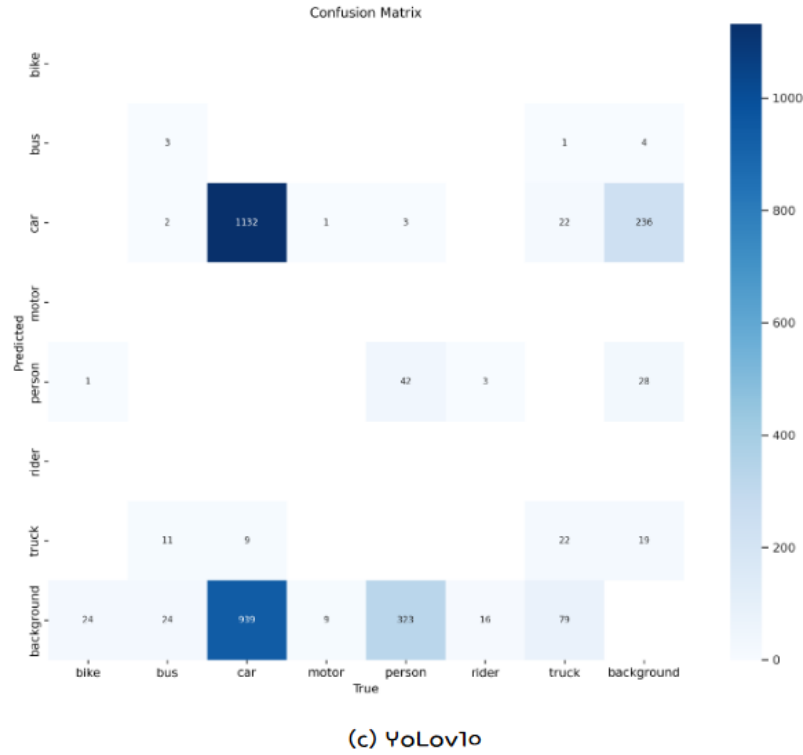
4.2.5 Confusion Matrix



Graph 5: Confusion Matrices of YOLOv8



Graph 6: Confusion Matrices of YOLOv8, YOLOv9 and YOLOv10



Graph 7: Confusion Matrices of YOLOv8, YOLOv9 and YOLOv10

YOLOv8 excels in predicting the class car but then showcases confusion with the background class and often misclassifies the background as car. YOLOv9 however improves this by offering a more balanced performance with normalized data regardless it does struggle with misclassification. Lastly, YOLOv10 outperforms the other models by maintaining overall accuracy and not consuming background with other classes, making it the most effective of the three of them.

4.3 Analysis

YOLOv8 has a high precision for classes such as Motor and Car and low precision for classes Bus and Truck this can be because of the CSPDarknet backbone as it is optimized for well-designed objects like cars and motors. Larger scale objects such as trucks could introduce challenges due to variability of size especially when distorted by adverse weather conditions.

Furthermore, reliance on convolutional layers, which are effective for high-contrast features may be the cause of low recall as high-contrast features are diminished by adverse conditions. Additionally, YOLOv8's anchor-free detection can make the model less reliable regardless of increasing the speed of detections.

The C2f model in YOLOv8 combines high-level features with context which can be effective in environments which are high in contrast and clear. However, given a dataset with distortions for this investigation, YOLOv8 is unable to maintain a balance in mAP50. It becomes low in classes where these distortions are more pronounced.

The use of anchor-free detection may have caused YOLOv8 to struggle with objects whose position and size vary unpredictably this directly contributes to the steep in the recall furthermore, the confusion matrix showcases misclassification between objects and the background, especially for Truck, this can be justified because of the use of CSPDarknet which is generally excellent for extracting complex features and hence there is a chance that noise and other elements could have confused the architecture to pick it up as well and in return why the model confuses objects such as trucks with the background.

YOLOv9 shows variability in precision-recall curves, and includes sharp declines in classes like Bike and Person— though YOLOv9 can maintain high precision it struggles with recall, especially in challenging detection scenarios. This is possible due to the introduction of PGI— which can affect gradient stability and make the model more sensitive to confidence thresholds.

YOLOv9 exhibits more fluctuation in the precision/confidence curve, especially at lower confidence levels— GELAN's impact could have led to affect gesture aggregation, which may not always provide constant gesture quality across different classes, in return, this could lead to the precision and confidence variability. Similarly, the recall/confidence curve also showcases a noticeable sharp drop in YOLOv9 in recall as confidence increases— This suggests the model is cautious in its detections when operating under high confidence thresholds.

Achieving high true positive rates for Cars and trucks but also showing confusion in the background category, particularly in bike and motor can be justified due to the model's sensitivity to background features whilst using GELAN— which aggregates features from different scales but doesn't effectively separate foreground objects from the background. Higher cross-entropy losses for these classes define the model's struggle with class separability, particularly under the noise of adverse conditions.

YOLOv10 exhibits a conservative precision-recall trade-off. This is evident in the sharp decline of recall as precision increases— the NMS-free strategy employed can likely be the cause of this, which optimizes precision by eliminating overlapping detections but this costs the recall.

The weighting on precision may cause the model to prioritise certain high-confidence predictions which can lead to a low recall, especially for objects that are harder to detect—for example, objects smaller in size or covered by fog or adverse factors.

The dual head architecture could reason for the model's high precision at higher confidence and variability of precision at lower confidence—the model might over-prune detection from one-to-many heads during inference. Consequently, the recall/confidence curve showcases the decline of YOLOv10—this is because the model sacrifices recall to ensure higher precision. Decoupling of spatial and channel operations may also contribute to this making the model less sensitive to details which is necessary for detecting smaller objects.

The Confusion matrix suggests that YOLOv10 shows significant confusion in the background particularly with smaller classes such as Person and Motor. YOLOv10 reduces post-processing through NMS-free strategies this can contribute to filtering out lower confidence detection that would otherwise be true positives. The higher rates of false positives specifically in the background can be justified by the confidence-weighted prediction strategy where the model focuses on precision during inference and could lead to overconfidence in background misclassifications.

5 Conclusion

From this investigation, it was interpreted that in the evolution of frameworks from YOLOv8 to YOLOv10, there is a trend towards refining precision and recall. YOLOv9 serves substantial improvements compared to YOLOv8 making it robust for a wide range of OD tasks, especially in challenging environments. Technically, even though YOLOv10 is more advanced, it does not justify its performance. Given the evidence from the results, generally, YOLOv9 is suggested for OD in bad weather due to its balance and reliability compared to other models used in this investigation. However, in conclusion after the result analysis, the best model does not exist and can only be generalised. To gain maximum benefit, it is recommended to consider the performance of these frameworks to be case-dependent. YOLOv10 can be considered for specialized scenarios such as situations where the need for high precision is required, high confidence environments and specialised OD, similarly, YOLOv8 can be considered for lightweight applications and environments with low complexity.

6 Further Scope

According to this investigation; resulted in a sequence of best-performing to least-performing frameworks. There is scope to implement multiple framework models in a singular workspace aimed at one task and leverage the importance of each model's output to be case-dependent, similar to how weights work in a neural network— this approach could help the rescue robot not be

dependent on a singular framework rather multiple working frameworks which are enhanced in a particular field such as precision.

7 Limitations and Challenges

The investigation would have been carried out to be more sophisticated perhaps if the dataset was not too small. It is hard to generalise and train a model using a small dataset. The BDDK100 dataset was not used to due poor computational power and hardware limitations. Furthermore, the usage of other modes of data types such as videos and other multimodal options would have further enhanced the integrity of this research but it was not proceeded with due to yet again limitations of resources, hence the investigation improvised to providing investigation upon a smaller data set with only still-images which gives rise for further studies in the future.

It took plentiful hours to find the dataset which matched this investigation's aim in addition to more time spent playing around with Google Colaboratory and the framework implementation and numerous inferences and model training to be done before this investigation.

The lack of proper documentation which YOLOv10 held made it immensely hard to be able to understand the architecture of the model. However, through the use of official documents, the GitHub repository and a few research journals this was made possible. Highlighting the extensive time taken to evaluate each research before referencing through their acceptance in terms of peer review and citations.

8 Bibliography

Flynn, Helen, et al. "Machine Learning Applied to Object Recognition in Robot Search and Rescue Systems." *ResearchGate*, 20, https://www.researchgate.net/profile/Helen-Flynn/publication/224773218_Machine_Learning_Applied_to_Object_Recognition_in_Robot_Search_and_Rescue_Systems/links/5440edd50cf2ebb036905a88/Machine-Learning-Applied-to-Object-Recognition-in-Robot-Search-and-Rescue-Systems.pdf. Accessed 1 April. 2024.

"Computer Vision." *IBM*, <https://www.ibm.com/topics/computer-vision>. Accessed 4 April. 2024.

Li, H., and L. S. Luo. "A Critical Review of Multiphase Flow Simulation Using the Lattice Boltzmann Method." *Archives of Computational Methods in Engineering*, 2019 <https://link.springer.com/article/.07/s11831-018-09312-w>. Accessed 1 June. 2024.

"Machine Learning Frameworks." *GeeksforGeeks*, 23 May 2023, <https://www.geeksforgeeks.org/machine-learning-frameworks/>. Accessed 3 May 2024.

"Fukushima Daiichi Accident." *World Nuclear Association*, June 2023, <https://world-nuclear.org/information-library/safety-and-security/safety-of-plants/fukushima-daiichi-accident>. Accessed 2 June. 2024.

Westcott, Lucy. "Robots Sent into Fukushima Have 'Died.'" *Newsweek*, Mar. 2016, <https://www.newsweek.com/robots-sent-fukushima-have-died-435332>. Accessed 5 June. 2024.

Li, Qiang, et al. "An Overview of the Cuckoo Search Algorithm and Its Applications." *Proceedings of the IEEE*, 2011, <https://ieeexplore.ieee.org/document/5876227>. Accessed 5 June. 2024.

Toma, David, et al. "Comprehensive Survey of Machine Learning-Based Prediction Models for the Solar Energy System and Evaluations of Solar Irradiance." *Neurocomputing*, 2022 <https://www.sciencedirect.com/science/article/pii/S0924271622003367>. Accessed 5 June. 2024.

Wang, Cheng, et al. "A Review of Machine Learning Applications in Wireless Networks: Algorithms, Datasets, and Results." *arXiv*, 19 Jul. 2019, <https://arxiv.org/abs/1907.09408>. Accessed 5 May 2024.

LeCun, Yann, et al. "Gradient-Based Learning Applied to Document Recognition." *Proceedings of the IEEE*, ResearchGate, https://www.researchgate.net/publication/2985446_Gradient-Based_Learning_Applied_to_Document_Recognition. Accessed 25 June 2024.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep Learning." *Nature*, 2015, <https://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf>. Accessed 2 April 2024.

Jain, Avijit. *MIT Deep Learning Book PDF*. GitHub, <https://github.com/janishar/mit-deep-learning-book-pdf>. Accessed 25 June 2024.

"References for Papers." *Scientific Research Publishing*, <https://www.scirp.org/reference/ReferencesPapers?ReferenceID=1308330>. Accessed 25 June 2024.

Chollet, François. "Xception: Deep Learning with Depthwise Separable Convolutions." *IEEE Xplore*, 2017, <https://ieeexplore.ieee.org/document/7780459>. Accessed 7 January 2024.

Tian, Zhi, et al. "FCOS: Fully Convolutional One-Stage Object Detection." *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, https://openaccess.thecvf.com/content_ICCV_2019/html/Tian_FCOS_Fully_Convolutional_One-Stage_Object_Detection_ICCV_2019_paper.html. Accessed 15 July 2024.

Redmon, Joseph, and Ali Farhadi. "YOLO9000: Better, Faster, Stronger." *arXiv*, 2017, <https://arxiv.org/abs/1704.04861>. Accessed 15 July 2024.

Odena, Augustus, et al. "Deconvolution and Checkerboard Artifacts." *Distill*, 2016, <https://distill.pub/2016/deconv-checkerboard/>. Accessed 15 July 2024.

Wang, Chien-Yao, et al. "CSPNet: A New Backbone that Can Enhance Learning Capability of CNN." *NYCU Scholar*, <https://scholar.nycu.edu.tw/en/publications/cspnet-a-new-backbone-that-can-enhance-learning-capability-of-cnn>. Accessed 15 July 2024.

Hirschfeld, Gerrit, and Ulf Ziemann. "Neuroplasticity and Functional Recovery After Stroke: Evidence from TMS and fMRI." *PubMed*, 2015, <https://pubmed.ncbi.nlm.nih.gov/26353135/>. Accessed 15 July 2024.

He, Tong, et al. "Bag of Tricks for Image Classification with Convolutional Neural Networks." *arXiv*, 2017, <https://arxiv.org/abs/1711.06897>. Accessed 15 July 2024.

Zhou, Zheng, et al. "Multi-Scale Feature Attention for Person Re-Identification." *Lecture Notes in Computer Science (LNCS)*, vol. 12335, Springer, 2020
https://link.springer.com/chapter/10.1007/978-3-030-58452-8_13. Accessed 15 July 2024.

Szegedy, Christian, et al. "Going Deeper with Convolutions." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, https://www.cv-foundation.org/openaccess/content_cvpr_2015/html/Szegedy_Going_Deepier_With_2015_CVPR_paper.html. Accessed 15 July 2024.

Sanderson, Grant. "Neural Networks." *BlueBrown*,
<https://www.3blue1brown.com/lessons/neural-networks>. Accessed 7 January 2024.

Redmon, Joseph, et al. "You Only Look Once: Unified, Real-Time Object Detection." *arXiv*, 8 Jun. 2015, <https://arxiv.org/abs/1506.02640>. Accessed 7 June 2024.

Chen, Zhenyu, et al. "An Extensive Survey on Vision Transformers: Transforming Vision for Better Future." *arXiv*, 5 Jul. 2024, <https://arxiv.org/html/2407.02988v1#S3>. Accessed 7 June 2024.

Kingma, Diederik P., and Jimmy Ba. "Adam: A Method for Stochastic Optimization." *arXiv*, 3 Mar. 2015, <https://arxiv.org/abs/1503.02406>. Accessed 7 June 2024.

Wang, Alex, et al. "GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding." *arXiv*, 22 Oct. 2020, <https://arxiv.org/abs/20.11929>. Accessed 8 June 2024.

Nijkamp, Erik, et al. "A Survey on Visual Prompt Learning." *arXiv*, 23 Dec. 2022, <https://arxiv.org/abs/2212.11696>. Accessed 8 June 2024.

Jiang, Zhengkai, et al. "Text2Human: Text-Driven Controllable Human Image Generation." *arXiv*, 9 Nov. 2022, <https://arxiv.org/abs/2211.04800>. Accessed 7 June 2024.

Zhou, Xiangning, et al. "OpenXGPT: An Open-Source Instruction-Following Language Model." *arXiv*, 27 Feb. 2024, <https://arxiv.org/pdf/2402.13616>. Accessed 7 June 2024.

Ramesh, Aditya, et al. "Zero-Shot Text-to-Image Generation." *arXiv*, May 2021, <https://arxiv.org/abs/25.04206>. Accessed 7 June 2024.

Jiang, Zhengkai, et al. "Text2Human: Text-Driven Controllable Human Image Generation." *arXiv*, 9 Nov. 2022, <https://arxiv.org/abs/2211.04800>. Accessed 7 June 2024.

Srivastava, Rupesh Kumar, et al. "Highway Networks." *arXiv*, 15 Jun. 2015, <https://arxiv.org/abs/1506.04878>. Accessed 8 June 2024.

Raffel, Colin, et al. "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer." *arXiv*, 26 May 2020, <https://arxiv.org/abs/2005.12872>. Accessed 8 June 2024.

Viana, André Lucas. "Hungarian-Loss: A Python Library for the Hungarian Algorithm." *PyPI*, <https://pypi.org/project/hungarian-loss/>. Accessed 9 June. 2024.

Vaswani, Ashish, et al. "Scaling Vision Transformers." *arXiv*, 17 Aug. 2021, <https://arxiv.org/abs/2106.04551>. Accessed 9 June. 2024.

"YOLOv10 Documentation." *Ultralytics*, <https://docs.ultralytics.com/models/yolov10/>. Accessed 15 June 2024.

Xu, Ming, et al. "Emerging Trends in Large Language Models: A Survey." *arXiv*, 25 May 2024, <https://arxiv.org/pdf/2405.14458>. Accessed 15 June . 2024.

"YOLOv10: An Introduction to the Latest Version of YOLO." *LearnOpenCV*, <https://learnopencv.com/yolov10/>. Accessed 15 June . 2024.

"YOLOv10: Everything You Need to Know." *Roboflow Blog*, 25 July 2023, <https://blog.roboflow.com/what-is-yolov10/>. Accessed 15 June 2024.

"Object Detection in Bad Weather Dataset." *Roboflow Universe*, <https://universe.roboflow.com/foreignobjectaerodromes/o.d-in-bad-weather/dataset/1>. Accessed 15 June 2024.

"Foreign Object Aerodromes." *Roboflow Universe*, <https://universe.roboflow.com/foreignobjectaerodromes>. Accessed 15 June. 2024.

Hinton, Geoffrey E., et al. "Improving Neural Networks by Preventing Co-adaptation of Feature Detectors." *arXiv*, 24 Jun. 2012, <https://arxiv.org/abs/1206.5538>. Accessed 19 June 2024.

Karn, Ujjwal. "Ujjwal Karn's Blog." *Ujjwalkarn.me*, <https://ujjwalkarn.me/>. Accessed 19 June 2024.

Liu, Zhuang, et al. "Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows." *arXiv*, 8 Jul. 2021, <https://arxiv.org/abs/2104.02111>. Accessed 19 June 2024.

Devlin, Jacob, et al. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." *arXiv*, 8 Sep. 2018, <https://arxiv.org/abs/1809.03193>. Accessed 19 June 2024.

Simonyan, Karen, and Andrew Zisserman. "Very Deep Convolutional Networks for Large-Scale Image Recognition." *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2015 <https://ieeexplore.ieee.org/document/7005506>. Accessed 19 June 2024.

Touvron, Hugo, et al. "Training Data-Efficient Image Transformers and Distillation Through Attention." *arXiv*, 23 Apr. 2020, <https://arxiv.org/abs/2004.934>. Accessed 19 June 2024.

Ahmed, Sameer, et al. "A Review on YOLOv8 and Its Advancements." *ResearchGate*, 2024, https://www.researchgate.net/publication/377216968_A_Review_on_YOLOv8_and_Its_Advancements. Accessed 25 June 2024.

9 Appendices

9.1 Appendix A: Terminology

| Terminology | Definition |
|--------------------------|---|
| Multimodal | Integration of multiple data types. |
| Grid Matrix | Division of an image into separate grid cells. |
| Bounding boxes | Boxes which are used to localize and locate objects in an image are defined by coordinates. |
| Deep Neural Networks | Neural network with multiple hidden layers between input and output. |
| Computation | Process of performing calculations through technology |
| Information retainment | The ability of a NN to retain information to utilize later |
| Cloud-based data storage | Storing data on remote servers which are accessed via the internet |
| API keys | Identifiers which are used to authenticate access to a program |

Table 1: Additional basic terminology

9.2 Appendix B: Background

9.2.1 Artificial Neural Networks (NNs)

A neuron contains data and each neuron contains different activation values in essence the activation value decides upon the significance of the neuron (*Sanderson*). The more significant a neuron it is likely to increase significance of the neuron subsequently to the neuron forward, the less significant the neuron is cancelled out. Assigning weights to neuron connections between

layers helps classify and manually tweak the significance of each neuron to be specific about results (*Sanderson*).

$$w_1a_1 + w_2a_2 + w_3a_3 + \dots + w_na_n \quad [1]$$


The above equation describes the addition of weight (*Sanderson*). The weighted sum would give the specific feature which is looked up for. Each neuron consists of biases which is a scalar value added to the weighted sum before passing the result through an activation function which is used to adjust the output of each neuron along with the weighted inputs.

9.2.2 Convolutional Neural Networks (CNNs)

Computer vision involves the computer interpreting stimuli from its environment-- to achieve this a CNN is used. CNN helps interpret visual data on a sophisticated basis. There are multiple ways to interpret data such as; Object recognition, detection or segmentation.



Classification



Classification & Localization



Multiple Detection

Figure 1: Classification, Localization and Segmentation Visualization

CNN is implemented through the use of " blocks " of convolution, it is essential to define Kernels/filters as wanted to be able to feature extract effectively (*Hinton et al.*). Kernels and Filters

describe small matrices used for convolutional operations. They help detect specific features in input data.

$$q_i^l = f(b_i^l + \sum_{j=0}^{d-1} w_{ij} \times x_j + j) \quad [1]$$

$$q_{ij}^l = f\left(b_{ij} + \sum_{k=0}^{d_1} \sum_{m=0}^{d_2} w_{(i+k)(j+m)}^{x_{(i+k)(j+m)}}\right) \quad [2]$$

Equation [1] The equation calculates the output q_{ij}^l of a neuron by applying intended activation f to the weighted sum of its inputs plus the bias, with an additional term j included. q_{ij}^l represents the output of a neuron i in layer l . b_{ij}^l represents the bias which has been added to the weighted sum of inputs. Equation [2] describes a neuron in a CNN where the output is calculated by applying the activation function to the sum of weighted inputs like [1] but over a local receptive field. The double summation reflects the process of convolving a filter across the input space to compute the neuron's output.

The output undergoes pooling to aggregate an emphasis on key features, achieved by reducing the spatial dimensions. Various pooling techniques are available such as average pooling, sum pooling, and max pooling (*Karn*).




Figure 2: General structure of a CNN

(Liu *et al.*) Single-stage detectors identify objects in a single pass and hence remove the necessity for a separate region proposal step (like R-CNN). By employing multiple convolutional feature maps and varying scales for bounding box predictions, these detectors can effectively identify objects of different sizes and shapes in a single forward pass. Examples of single-shot detectors include the YOLO framework (Devlin *et al.*)




Figure 3: Abstract architecture of a single-stage object detector

9.3 Appendix D: Investigation Code

Training YOLOv8

GPU access verification

```
In [1]: !nvidia-smi
```

Tue Aug 6 15:47:43 2024

```

+-----+
| NVIDIA-SMI 535.104.05                Driver Version: 535.104.05   CUDA Version: 12.2     |
+-----+-----+-----+-----+-----+
| GPU Name                               Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap       |              |           |         |
| MIG M. |                                         |              |           |         |
+-----+-----+-----+-----+-----+
| 0 Tesla T4                            Off          | 00000000:00:04.0 Off |                    |
| N/A   47C   P8             9W / 70W    | 0MiB / 15360MiB |         0%      D |
| Default                                |                 |           |         |
| N/A                                     |                 |           |         |
+-----+-----+-----+-----+-----+

+-----+
| Processes:
| GPU   GI   CI        PID   Type   Process name                      GPU
| Memory |   ID   ID              |              |             |       |
| Usage  |       |       |              |              |             |       |
+-----+-----+-----+-----+-----+
| No running processes found
+-----+

```

```
In [2]: import os
HOME = os.getcwd()
print(HOME)
```

/content

Installation of YOLOv8 (pip)

```
In [3]: !pip install ultralytics==8.0.196

from IPython import display
display.clear_output()

import ultralytics
ultralytics.checks()
```

Ultralytics YOLOv8.0.196 🚀 Python-3.10.12 torch-2.3.1+cu121 CUDA:0 (Tesla T4, 15102MiB)
Setup complete ✅ (2 CPUs, 12.7 GB RAM, 33.6/78.2 GB disk)

```
In [4]: from ultralytics import YOLO

        from IPython.display import display, Image
```

Dataset Upload

```
In [5]: !mkdir {HOME}/datasets
        %cd {HOME}/datasets
```

/content/datasets

```
In [6]: !pip install roboflow

        from roboflow import Roboflow
        rf = Roboflow(api_key="jPCXLMbZJU137MRBek9F")
        project = rf.workspace("foreignobjectaerodromes").project("o.d-in-bad-weather")
        version = project.version(1)
        dataset = version.download("yolov8")
```

Collecting roboflow

Downloading roboflow-1.1.37-py3-none-any.whl.metadata (9.4 kB)

Requirement already satisfied: certifi in /usr/local/lib/python3.10/dist-packages (from roboflow) (2024.7.4)

Collecting chardet==4.0.0 (from roboflow)

Downloading chardet-4.0.0-py2.py3-none-any.whl.metadata (3.5 kB)

Requirement already satisfied: idna==3.7 in /usr/local/lib/python3.10/dist-packages (from roboflow) (3.7)

Requirement already satisfied: cyclur in /usr/local/lib/python3.10/dist-packages (from roboflow) (0.12.1)

Requirement already satisfied: kiwisolver>=1.3.1 in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.4.5)

Requirement already satisfied: matplotlib in /usr/local/lib/python3.10/dist-packages (from roboflow) (3.7.1)

Requirement already satisfied: numpy>=1.18.5 in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.26.4)

Requirement already satisfied: opencv-python-headless==4.10.0.84 in /usr/local/lib/python3.10/dist-packages (from roboflow) (4.10.0.84)

Requirement already satisfied: Pillow>=7.1.2 in /usr/local/lib/python3.10/dist-packages (from roboflow) (9.4.0)

Requirement already satisfied: python-dateutil in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.8.2)

Collecting python-dotenv (from roboflow)

Downloading python_dotenv-1.0.1-py3-none-any.whl.metadata (23 kB)

Requirement already satisfied: requests in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.31.0)

Requirement already satisfied: six in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.16.0)

Requirement already satisfied: urllib3>=1.26.6 in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.0.7)

Requirement already satisfied: tqdm>=4.41.0 in /usr/local/lib/python3.10/dist-packages (from roboflow) (4.66.4)

Requirement already satisfied: PyYAML>=5.3.1 in /usr/local/lib/python3.10/dist-packages (from roboflow) (6.0.1)

Collecting requests-toolbelt (from roboflow)

Downloading requests_toolbelt-1.0.0-py2.py3-none-any.whl.metadata (14 kB)

Collecting filetype (from roboflow)

Downloading filetype-1.2.0-py2.py3-none-any.whl.metadata (6.5 kB)

Requirement already satisfied: contourpy>=1.0.1 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (1.2.1)

Requirement already satisfied: fonttools>=4.22.0 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (4.22.0)


```

t-packages (from matplotlib->roboflow) (4.53.1)
Requirement already satisfied: packaging>=20.0 in /usr/local/lib/python3.10/dist-
packages (from matplotlib->roboflow) (24.1)
Requirement already satisfied: pyparsing>=2.3.1 in /usr/local/lib/python3.10/dist-
packages (from matplotlib->roboflow) (3.1.2)
Requirement already satisfied: charset-normalizer<4,>=2 in /usr/local/lib/python
3.10/dist-packages (from requests->roboflow) (3.3.2)
Downloading roboflow-1.1.37-py3-none-any.whl (76 kB)
76.9/76.9 kB 3.9 MB/s eta 0:00:00
Downloading chardet-4.0.0-py2.py3-none-any.whl (178 kB)
178.7/178.7 kB 11.9 MB/s eta 0:00:00
Downloading filetype-1.2.0-py2.py3-none-any.whl (19 kB)
Downloading python_dotenv-1.0.1-py3-none-any.whl (19 kB)
Downloading requests_toolbelt-1.0.0-py2.py3-none-any.whl (54 kB)
54.5/54.5 kB 4.5 MB/s eta 0:00:00
Installing collected packages: filetype, python-dotenv, chardet, requests-toolbel
t, roboflow
  Attempting uninstall: chardet
    Found existing installation: chardet 5.2.0
    Uninstalling chardet-5.2.0:
      Successfully uninstalled chardet-5.2.0
Successfully installed chardet-4.0.0 filetype-1.2.0 python-dotenv-1.0.1 requests-
toolbelt-1.0.0 roboflow-1.1.37
loading Roboflow workspace...
loading Roboflow project...
Downloading Dataset Version Zip in O.D-IN-BAD-WEATHER-1 to yolov8:: 100%|
75164/75164 [00:01<00:00, 42494.88it/s]
Extracting Dataset Version Zip to O.D-IN-BAD-WEATHER-1 in yolov8:: 100%|
2312/2312 [00:00<00:00, 4936.68it/s]

Training with Dataset

```

```

In [5]: %cd {HOME}

!yolo task=detect mode=train model=yolov8s.pt data=/content/datasets/O.D-IN-BAD
/content
New https://pypi.org/project/ultralytics/8.2.74 available 🤖 Update with 'pip in
stall -U ultralytics'
Ultralytics YOLOv8.0.196 🚀 Python-3.10.12 torch-2.3.1+cu121 CUDA:0 (Tesla T4, 1
5102MiB)
engine/trainer: task=detect, mode=train, model=yolov8s.pt, data=/content/dataset
s/O.D-IN-BAD-WEATHER-1/data.yaml, epochs=25, patience=50, batch=16, imgsz=640, sa
ve=True, save_period=-1, cache=False, device=None, workers=8, project=None, name=
None, exist_ok=False, pretrained=True, optimizer=auto, verbose=True, seed=0, dete
rministic=True, single_cls=False, rect=False, cos_lr=False, close_mosaic=10, resu
me=False, amp=True, fraction=1.0, profile=False, freeze=None, overlap_mask=True,
mask_ratio=4, dropout=0.0, val=True, split=val, save_json=False, save_hybrid=Fals
e, conf=None, iou=0.7, max_det=300, half=False, dnn=False, plots=True, source=Non
e, show=False, save_txt=False, save_conf=False, save_crop=False, show_labels=Tru
e, show_conf=True, vid_stride=1, stream_buffer=False, line_width=None, visualize=
False, augment=False, agnostic_nms=False, classes=None, retina_masks=False, boxes
=True, format=torchscript, keras=False, optimize=False, int8=False, dynamic=Fals
e, simplify=False, opset=None, workspace=4, nms=False, lr=0.01, lrf=0.01, moment
um=0.937, weight_decay=0.0005, warmup_epochs=3.0, warmup_momentum=0.8, warmup_bia
s_lr=0.1, box=7.5, cls=0.5, dfl=1.5, pose=12.0, kobj=1.0, label_smoothing=0.0, nb
s=64, hsv_h=0.015, hsv_s=0.7, hsv_v=0.4, degrees=0.0, translate=0.1, scale=0.5, s
hear=0.0, perspective=0.0, flipud=0.0, fliplr=0.5, mosaic=1.0, mixup=0.0, copy_pa
ste=0.0, cfg=None, tracker=botsort.yaml, save_dir=runs/detect/train3
Downloading https://ultralytics.com/assets/Arial.ttf to /root/.config/Ultralytic
s/Arial.ttf...
100% 755k/755k [00:00<00:00, 21.4MB/s]
2024-08-06 15:51:53.010289: E external/local_xla/xla/stream_executor/cuda/cuda_ff
t.cc:485] Unable to register cuFFT factory: Attempting to register factory for pl
ugin cuFFT when one has already been registered

```

```

2024-08-06 15:51:53.244881: E external/local_xla/xla/stream_executor/cuda/cuda_dnn
n.cc:8454] Unable to register cuDNN factory: Attempting to register factory for p
lugin cuDNN when one has already been registered
2024-08-06 15:51:53.311034: E external/local_xla/xla/stream_executor/cuda/cuda_bla
s.cc:1452] Unable to register cuBLAS factory: Attempting to register factory for
plugin cuBLAS when one has already been registered
Overriding model.yaml nc=80 with nc=7

      from  n  params  module
arguments
  0         -1  1      928  ultralytics.nn.modules.conv.Conv
[3, 32, 3, 2]
  1         -1  1     18560  ultralytics.nn.modules.conv.Conv
[32, 64, 3, 2]
  2         -1  1     29056  ultralytics.nn.modules.block.C2f
[64, 64, 1, True]
  3         -1  1     73984  ultralytics.nn.modules.conv.Conv
[64, 128, 3, 2]
  4         -1  2     197632  ultralytics.nn.modules.block.C2f
[128, 128, 2, True]
  5         -1  1     295424  ultralytics.nn.modules.conv.Conv
[128, 256, 3, 2]
  6         -1  2     788480  ultralytics.nn.modules.block.C2f
[256, 256, 2, True]
  7         -1  1    1180672  ultralytics.nn.modules.conv.Conv
[256, 512, 3, 2]
  8         -1  1    1838080  ultralytics.nn.modules.block.C2f
[512, 512, 1, True]
  9         -1  1     656896  ultralytics.nn.modules.block.SPPF
[512, 512, 5]
 10         -1  1         0  torch.nn.modules.upsampling.Upsample
[None, 2, 'nearest']
 11         [-1, 6]  1         0  ultralytics.nn.modules.conv.Concat
[1]
 12         -1  1     591360  ultralytics.nn.modules.block.C2f
[768, 256, 1]
 13         -1  1         0  torch.nn.modules.upsampling.Upsample
[None, 2, 'nearest']
 14         [-1, 4]  1         0  ultralytics.nn.modules.conv.Concat
[1]
 15         -1  1     148224  ultralytics.nn.modules.block.C2f
[384, 128, 1]
 16         -1  1     147712  ultralytics.nn.modules.conv.Conv
[128, 128, 3, 2]
 17         [-1, 12]  1         0  ultralytics.nn.modules.conv.Concat
[1]
 18         -1  1     493056  ultralytics.nn.modules.block.C2f
[384, 256, 1]
 19         -1  1     590336  ultralytics.nn.modules.conv.Conv
[256, 256, 3, 2]
 20         [-1, 9]  1         0  ultralytics.nn.modules.conv.Concat
[1]
 21         -1  1    1969152  ultralytics.nn.modules.block.C2f
[768, 512, 1]
 22         [15, 18, 21]  1    2118757  ultralytics.nn.modules.head.Detect
[7, [128, 256, 512]]
Model summary: 225 layers, 11138309 parameters, 11138293 gradients, 28.7 GFLOPs

Transferred 349/355 items from pretrained weights
TensorBoard: Start with 'tensorboard --logdir runs/detect/train3', view at htt
p://localhost:6006/
Freezing layer 'model.22.dfl.conv.weight'
AMP: running Automatic Mixed Precision (AMP) checks with YOLOv8n...
Downloading https://github.com/ultralytics/assets/releases/download/v0.0.0/yolov8
n.pt to 'yolov8n.pt'...
100% 6.23M/6.23M [00:00<00:00, 135MB/s]
WARNING ⚠ NMS time limit 0.550s exceeded

```

```

AMP: checks passed ✓
train: Scanning /content/datasets/O.D-IN-BAD-WEATHER-1/train/labels... 815 images, 5 backgrounds, 0 corrupt: 100% 815/815 [00:00<00:00, 2077.86it/s]
train: New cache created: /content/datasets/O.D-IN-BAD-WEATHER-1/train/labels.cache
INFO:albumentations.check_version:A new version of Albumentations is available: 1.4.13 (you have 1.4.12). Upgrade using: pip install -U albumentations. To disable automatic update checks, set the environment variable NO_ALBUMENTATIONS_UPDATE to 1.
/usr/local/lib/python3.10/dist-packages/albumentations/core/composition.py:161: UserWarning: Got processor for bboxes, but no transform to process it.
  self._set_keys()
albumentations: Blur(p=0.01, blur_limit=(3, 7)), MedianBlur(p=0.01, blur_limit=(3, 7)), ToGray(p=0.01), CLAHE(p=0.01, clip_limit=(1, 4.0), tile_grid_size=(8, 8))
/usr/lib/python3.10/multiprocessing/popen_fork.py:66: RuntimeWarning: os.fork() was called. os.fork() is incompatible with multithreaded code, and JAX is multithreaded, so this will likely lead to a deadlock.
  self.pid = os.fork()
val: Scanning /content/datasets/O.D-IN-BAD-WEATHER-1/valid/labels... 218 images, 2 backgrounds, 0 corrupt: 100% 218/218 [00:00<00:00, 1589.12it/s]
val: New cache created: /content/datasets/O.D-IN-BAD-WEATHER-1/valid/labels.cache
Plotting labels to runs/detect/train3/labels.jpg...
optimizer: 'optimizer=auto' found, ignoring 'lr=0.01' and 'momentum=0.937' and determining best 'optimizer', 'lr' and 'momentum' automatically...
optimizer: AdamW(lr=0.000909, momentum=0.9) with parameter groups 57 weight(decay=0.0), 64 weight(decay=0.0005), 63 bias(decay=0.0)
Image sizes 640 train, 640 val
Using 2 dataloader workers
Logging results to runs/detect/train3
Starting training for 25 epochs...

```

| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 1/25 | 4.34G | 1.577 | 2.238 | 1.275 | 282 | 640: 10 |
| 0% 51/51 [00:35<00:00, 1.45it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:04<00:00, 1.74it/s] | | | | | | |
| | all | 218 | 2666 | 0.696 | 0.169 | 0.169 |
| 0.096 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
| 2/25 | 4.26G | 1.444 | 1.273 | 1.203 | 239 | 640: 10 |
| 0% 51/51 [00:19<00:00, 2.63it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 2.02it/s] | | | | | | |
| | all | 218 | 2666 | 0.616 | 0.166 | 0.144 |
| 0.0758 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
| 3/25 | 4.23G | 1.483 | 1.297 | 1.241 | 249 | 640: 10 |
| 0% 51/51 [00:21<00:00, 2.43it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.43it/s] | | | | | | |
| | all | 218 | 2666 | 0.631 | 0.134 | 0.144 |
| 0.0781 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
| 4/25 | 4.29G | 1.511 | 1.267 | 1.239 | 159 | 640: 10 |
| 0% 51/51 [00:19<00:00, 2.60it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.59it/s] | | | | | | |
| | all | 218 | 2666 | 0.626 | 0.162 | 0.159 |
| 0.0888 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
| 5/25 | 4.14G | 1.443 | 1.184 | 1.193 | 194 | 640: 10 |
| 0% 51/51 [00:20<00:00, 2.54it/s] | | | | | | |

| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
|----------------------------------|-----------------------------|----------|-----------|----------|-----------|---------|----|
| P50-95): 100% | 7/7 [00:04<00:00, 1.68it/s] | | | | | | |
| | all | 218 | 2666 | 0.654 | 0.176 | 0.165 | |
| 0.0961 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 6/25 | 4.31G | 1.427 | 1.161 | 1.196 | 271 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.56it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.50it/s] | | | | | | |
| | all | 218 | 2666 | 0.665 | 0.165 | 0.178 | |
| 0.1 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 7/25 | 4.23G | 1.412 | 1.126 | 1.198 | 226 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.64it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:04<00:00, 1.49it/s] | | | | | | |
| | all | 218 | 2666 | 0.631 | 0.214 | 0.2 | |
| 0.11 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 8/25 | 4.24G | 1.379 | 1.082 | 1.17 | 186 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.57it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.48it/s] | | | | | | |
| | all | 218 | 2666 | 0.721 | 0.201 | 0.219 | |
| 0.12 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 9/25 | 4.31G | 1.376 | 1.046 | 1.169 | 323 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.58it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:04<00:00, 1.74it/s] | | | | | | |
| | all | 218 | 2666 | 0.671 | 0.215 | 0.212 | |
| 0.12 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 10/25 | 4.27G | 1.354 | 1.023 | 1.165 | 285 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.57it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.38it/s] | | | | | | |
| | all | 218 | 2666 | 0.714 | 0.225 | 0.237 | |
| 0.13 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 11/25 | 4.47G | 1.357 | 1.021 | 1.159 | 292 | 640: 10 | |
| 0% 51/51 [00:20<00:00, 2.54it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.44it/s] | | | | | | |
| | all | 218 | 2666 | 0.718 | 0.226 | 0.232 | |
| 0.132 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 12/25 | 4.24G | 1.318 | 0.9773 | 1.137 | 210 | 640: 10 | |
| 0% 51/51 [00:21<00:00, 2.38it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:03<00:00, 2.02it/s] | | | | | | |
| | all | 218 | 2666 | 0.703 | 0.228 | 0.234 | |
| 0.133 | | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size | |
| 13/25 | 4.21G | 1.309 | 0.9812 | 1.139 | 316 | 640: 10 | |
| 0% 51/51 [00:19<00:00, 2.57it/s] | | | | | | | |
| | Class | Images | Instances | Box(P) | R | mAP50 | mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.44it/s] | | | | | | |
| | all | 218 | 2666 | 0.759 | 0.225 | 0.258 | |

0.145

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 14/25 | 4.38G | 1.28 | 0.9235 | 1.119 | 163 | 640: 10 |
| 0% 51/51 [00:21<00:00, 2.42it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:04<00:00, 1.67it/s] | | | | | | |
| | all | 218 | 2666 | 0.777 | 0.226 | 0.261 |

0.140

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 15/25 | 4.25G | 1.269 | 0.9259 | 1.114 | 232 | 640: 10 |
| 0% 51/51 [00:19<00:00, 2.64it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.44it/s] | | | | | | |
| | all | 218 | 2666 | 0.725 | 0.23 | 0.24 |

0.139

Closing dataloader mosaic

```
/usr/local/lib/python3.10/dist-packages/albumentations/core/composition.py:161: UserWarning: Got processor for bboxes, but no transform to process it.
```

self._set_keys()

```
albumentations: Blur(p=0.01, blur_limit=(3, 7)), MedianBlur(p=0.01, blur_limit=(3, 7)), ToGray(p=0.01), CLAHE(p=0.01, clip_limit=(1, 4.0), tile_grid_size=(8, 8))
```

```
/usr/lib/python3.10/multiprocessing/popen_fork.py:66: RuntimeWarning: os.fork() was called. os.fork() is incompatible with multithreaded code, and JAX is multithreaded, so this will likely lead to a deadlock.
```

self.pid = os.fork()

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 16/25 | 4.23G | 1.314 | 0.903 | 1.123 | 180 | 640: 10 |
| 0% 51/51 [00:29<00:00, 1.75it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.59it/s] | | | | | | |
| | all | 218 | 2666 | 0.696 | 0.264 | 0.232 |

0.133

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 17/25 | 4.26G | 1.289 | 0.8594 | 1.112 | 131 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.72it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.76it/s] | | | | | | |
| | all | 218 | 2666 | 0.557 | 0.245 | 0.242 |

0.139

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 18/25 | 4.23G | 1.268 | 0.8262 | 1.103 | 158 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.73it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 2.18it/s] | | | | | | |
| | all | 218 | 2666 | 0.75 | 0.254 | 0.28 |

0.161

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 19/25 | 4.24G | 1.262 | 0.8251 | 1.106 | 172 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.83it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:04<00:00, 1.74it/s] | | | | | | |
| | all | 218 | 2666 | 0.782 | 0.232 | 0.271 |

0.153

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 20/25 | 4.23G | 1.232 | 0.7854 | 1.088 | 192 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.81it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.62it/s] | | | | | | |
| | all | 218 | 2666 | 0.64 | 0.262 | 0.274 |

0.158

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 21/25 | 4.25G | 1.214 | 0.7618 | 1.076 | 156 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.76it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.55it/s] | | | | | | |
| | all | 218 | 2666 | 0.508 | 0.266 | 0.273 |

0.156

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 22/25 | 4.23G | 1.203 | 0.747 | 1.067 | 140 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.80it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 1.78it/s] | | | | | | |
| | all | 218 | 2666 | 0.677 | 0.263 | 0.288 |

0.162

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 23/25 | 4.21G | 1.19 | 0.7252 | 1.053 | 161 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.79it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 1.88it/s] | | | | | | |
| | all | 218 | 2666 | 0.66 | 0.262 | 0.276 |

0.156

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 24/25 | 4.23G | 1.168 | 0.7198 | 1.051 | 175 | 640: 10 |
| 0% 51/51 [00:17<00:00, 2.89it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.61it/s] | | | | | | |
| | all | 218 | 2666 | 0.546 | 0.255 | 0.281 |

0.154

| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
|---|---------|----------|-----------|----------|-----------|----------|
| 25/25 | 4.2G | 1.172 | 0.7061 | 1.052 | 169 | 640: 10 |
| 0% 51/51 [00:18<00:00, 2.72it/s] | | | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.40it/s] | | | | | | |
| | all | 218 | 2666 | 0.647 | 0.25 | 0.278 |

0.157

25 epochs completed in 0.172 hours.

Optimizer stripped from runs/detect/train3/weights/last.pt, 22.5MB

Optimizer stripped from runs/detect/train3/weights/best.pt, 22.5MB

Validating runs/detect/train3/weights/best.pt...

Ultralytics YOLOv8.0.196 Python-3.10.12 torch-2.3.1+cu121 CUDA:0 (Tesla T4, 15102MiB)

Model summary (fused): 168 layers, 11128293 parameters, 0 gradients, 28.5 GFLOPs

| Class | Images | Instances | Box(P | R | mAP50 mA |
|---|--------|-----------|-------|-------|----------|
| P50-95): 100% 7/7 [00:12<00:00, 1.78s/it] | | | | | |
| all | 218 | 2666 | 0.677 | 0.263 | 0.286 |
| bike | 218 | 25 | 0.603 | 0.04 | 0.107 |
| bus | 218 | 40 | 0.523 | 0.425 | 0.412 |
| car | 218 | 2080 | 0.704 | 0.709 | 0.741 |
| motor | 218 | 10 | 1 | 0 | 0.0102 |
| person | 218 | 368 | 0.498 | 0.261 | 0.285 |
| rider | 218 | 19 | 1 | 0 | 0.0809 |
| truck | 218 | 124 | 0.414 | 0.403 | 0.369 |

0.258
 Speed: 0.3ms preprocess, 5.4ms inference, 0.0ms loss, 5.2ms postprocess per image
 Results saved to runs/detect/train3
 Learn more at <https://docs.ultralytics.com/modes/train>

In [6]: `!ls {HOME}/runs/detect/train/`

args.yaml weights


Training Results

Confusion Matrix

In [10]: `from IPython.display import Image`

In [11]: `%cd {HOME}`
`Image(filename='/content/runs/detect/train3/confusion_matrix.png', width=1000)`


/content




Results

In [13]: `%cd {HOME}`
`Image(filename=f'/content/runs/detect/train3/results.png', width=1000)`


/content





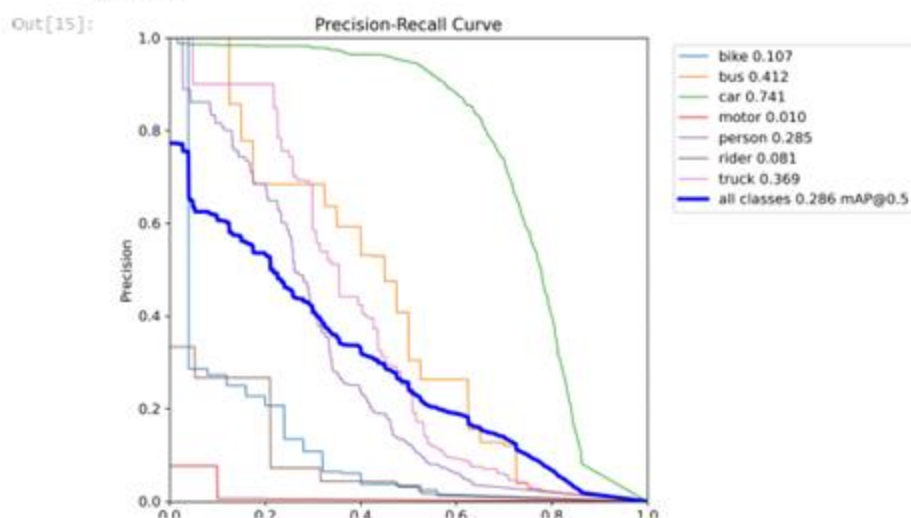
```
In [14]: %cd {HOME}
         Image(filename=f'/content/runs/detect/train3/F1_curve.png', width=1000)
```

/content




```
In [15]: %cd {HOME}
         Image(filename=f'/content/runs/detect/train3/PR_curve.png', width=1000)
```

/content




```
In [16]: %cd {HOME}
         Image(filename=f'/content/runs/detect/train3/P_curve.png', width=1000)
```

/content



```
In [17]: %cd {HOME}
         Image(filename=f'/content/runs/detect/train3/R_curve.png', width=1000)
```

Training YOLOv9

GPU access verification

In [1]: `!nvidia-smi`

```
Tue Aug 6 16:09:27 2024
+-----+
+-----+
| NVIDIA-SMI 535.104.05                Driver Version: 535.104.05   CUDA Version: 12.2
+-----+-----+
+-----+
| GPU  Name                       Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Comp.  M. |
| MIG M. |
+-----+-----+
| 0  Tesla T4                       Off | 00000000:00:04.0 Off |
| N/A   42C    P8             9W / 70W |  0MiB / 15360MiB |      0%    D
| default |
| N/A   |
+-----+-----+
+-----+
+-----+
| Processes:
| GPU  GI    CI          PID    Type    Process name                  GPU
Memory |
|      ID    ID                                   |          Usag
e      |
+-----+
| No running processes found
+-----+
+-----+
```

In [2]: `import os`
`HOME = os.getcwd()`
`print(HOME)`

/content

Installation of YOLOv9 (cloning because not distributed through pip packages)

In [3]: `!git clone https://github.com/SkalskiP/yolov9.git`
`%cd yolov9`
`!pip install -r requirements.txt -q`

```
Cloning into 'yolov9'...
remote: Enumerating objects: 325, done.
remote: Counting objects: 100% (218/218), done.
remote: Compressing objects: 100% (62/62), done.
```

```
remote: Total 325 (delta 159), reused 156 (delta 156), pack-reused 107
Receiving objects: 100% (325/325), 2.23 MiB | 13.38 MiB/s, done.
Resolving deltas: 100% (165/165), done.
/content/yolov9
_____ 207.3/207.3 kB 7.0 MB/s eta 0:00:00
_____ 62.7/62.7 kB 70.2 kB/s eta 0:00:00
```

```
In [4]: # YOLOv9 can not automatically download the weights so they are manually downlo
!wget -P {HOME}/weights -q https://github.com/WongKinYiu/yolov9/releases/downlo
!wget -P {HOME}/weights -q https://github.com/WongKinYiu/yolov9/releases/downlo
!wget -P {HOME}/weights -q https://github.com/WongKinYiu/yolov9/releases/downlo
!wget -P {HOME}/weights -q https://github.com/WongKinYiu/yolov9/releases/downlo
```

```
In [5]: !ls -la {HOME}/weights

total 402444
drwxr-xr-x 2 root root 4096 Aug 6 16:10 .
drwxr-xr-x 1 root root 4096 Aug 6 16:10 ..
-rw-r--r-- 1 root root 51508261 Feb 18 12:36 gelan-c.pt
-rw-r--r-- 1 root root 117203713 Feb 18 12:36 gelan-e.pt
-rw-r--r-- 1 root root 103153312 Feb 18 12:36 yolov9-c.pt
-rw-r--r-- 1 root root 140217688 Feb 18 12:36 yolov9-e.pt
```

Dataset Upload

```
In [6]: %cd {HOME}/yolov9
```

```
/content/yolov9
```

```
In [7]: !pip install roboflow

from roboflow import Roboflow
rf = Roboflow(api_key="JPCXLMZJU137MRBek9F")
project = rf.workspace("foreignobjectaerodromes").project("o.d-in-bad-weather")
version = project.version(1)
dataset = version.download("yolov9")
```

```
Collecting roboflow
  Downloading roboflow-1.1.37-py3-none-any.whl.metadata (9.4 kB)
Requirement already satisfied: certifi in /usr/local/lib/python3.10/dist-packages
(from roboflow) (2024.7.4)
Collecting chardet==4.0.0 (from roboflow)
  Downloading chardet-4.0.0-py2.py3-none-any.whl.metadata (3.5 kB)
Requirement already satisfied: idna==3.7 in /usr/local/lib/python3.10/dist-packag
es (from roboflow) (3.7)
Requirement already satisfied: cyclor in /usr/local/lib/python3.10/dist-packages
(from roboflow) (0.12.1)
Requirement already satisfied: kiwisolver>=1.3.1 in /usr/local/lib/python3.10/dis
t-packages (from roboflow) (1.4.5)
Requirement already satisfied: matplotlib in /usr/local/lib/python3.10/dist-packa
ges (from roboflow) (3.7.1)
Requirement already satisfied: numpy>=1.18.5 in /usr/local/lib/python3.10/dist-pa
ckages (from roboflow) (1.26.4)
Requirement already satisfied: opencv-python-headless==4.10.0.84 in /usr/local/li
b/python3.10/dist-packages (from roboflow) (4.10.0.84)
Requirement already satisfied: Pillow>=7.1.2 in /usr/local/lib/python3.10/dist-pa
ckages (from roboflow) (9.4.0)
Requirement already satisfied: python-dateutil in /usr/local/lib/python3.10/dist-
packages (from roboflow) (2.8.2)
Collecting python-dotenv (from roboflow)
  Downloading python_dotenv-1.0.1-py3-none-any.whl.metadata (23 kB)
Requirement already satisfied: requests in /usr/local/lib/python3.10/dist-package
s (from roboflow) (2.31.0)
```

```

Requirement already satisfied: six in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.16.0)
Requirement already satisfied: urllib3>=1.26.6 in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.0.7)
Requirement already satisfied: tqdm>=4.41.0 in /usr/local/lib/python3.10/dist-packages (from roboflow) (4.66.4)
Requirement already satisfied: PyYAML>=5.3.1 in /usr/local/lib/python3.10/dist-packages (from roboflow) (6.0.1)
Collecting requests-toolbelt (from roboflow)
  Downloading requests_toolbelt-1.0.0-py2.py3-none-any.whl.metadata (14 kB)
Collecting filetype (from roboflow)
  Downloading filetype-1.2.0-py2.py3-none-any.whl.metadata (6.5 kB)
Requirement already satisfied: contourpy>=1.0.1 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (1.2.1)
Requirement already satisfied: fonttools>=4.22.0 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (4.53.1)
Requirement already satisfied: packaging>=20.0 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (24.1)
Requirement already satisfied: pyparsing>=2.3.1 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (3.1.2)
Requirement already satisfied: charset-normalizer<4,>=2 in /usr/local/lib/python3.10/dist-packages (from requests->roboflow) (3.3.2)
Downloading roboflow-1.1.37-py3-none-any.whl (76 kB)
----- 76.9/76.9 kB 4.9 MB/s eta 0:00:00
Downloading chardet-4.0.0-py2.py3-none-any.whl (178 kB)
----- 178.7/178.7 kB 12.4 MB/s eta 0:00:00
Downloading filetype-1.2.0-py2.py3-none-any.whl (19 kB)
Downloading python_dotenv-1.0.1-py3-none-any.whl (19 kB)
Downloading requests_toolbelt-1.0.0-py2.py3-none-any.whl (54 kB)
----- 54.5/54.5 kB 4.4 MB/s eta 0:00:00
Installing collected packages: filetype, python-dotenv, chardet, requests-toolbelt, roboflow
  Attempting uninstall: chardet
    Found existing installation: chardet 5.2.0
    Uninstalling chardet-5.2.0:
      Successfully uninstalled chardet-5.2.0
Successfully installed chardet-4.0.0 filetype-1.2.0 python-dotenv-1.0.1 requests-toolbelt-1.0.0 roboflow-1.1.37
loading Roboflow workspace...
loading Roboflow project...
Downloading Dataset Version Zip in O.D-IN-BAD-WEATHER-1 to yolov9:: 100%|██████████| 75164/75164 [00:03<00:00, 23377.79it/s]
Extracting Dataset Version Zip to O.D-IN-BAD-WEATHER-1 in yolov9:: 100%|██████████| 2312/2312 [00:00<00:00, 5743.00it/s]

```

Training with Dataset

```

In [8]: %cd {HOME}/yolov9

!python train.py \
  --batch 16 --epochs 25 --img 640 --device 0 \
  --data {dataset.location}/data.yaml \
  --weights {HOME}/weights/gelan-c.pt \
  --cfg models/detect/gelan-c.yaml \
  --hyp hyp.scratch-high.yaml

```

```

/content/yolov9
2024-08-06 16:11:39.547702: E external/local_xla/xla/stream_executor/cuda/cuda_fft.cc:485] Unable to register cuFFT factory: Attempting to register factory for plugin cuFFT when one has already been registered
2024-08-06 16:11:39.841425: E external/local_xla/xla/stream_executor/cuda/cuda_dnn.cc:8454] Unable to register cuDNN factory: Attempting to register factory for plugin cuDNN when one has already been registered
2024-08-06 16:11:39.926061: E external/local_xla/xla/stream_executor/cuda/cuda_blas.cc:1452] Unable to register cuBLAS factory: Attempting to register factory for plugin cuBLAS when one has already been registered

```



```

256]
17      [-1, 12] 1      0 models.common.Concat [1]
18      -1 1 2988544 models.common.RepNCSPELAN4 [768,
512, 512, 256, 1]
19      -1 1 656384 models.common.ADown [512,
512]
20      [-1, 9] 1      0 models.common.Concat [1]
21      -1 1 3119616 models.common.RepNCSPELAN4 [102
4, 512, 512, 256, 1]
22      [15, 18, 21] 1 5496037 models.yolo.DDetect [7,
[256, 512, 512]]
gelan-c summary: 621 layers, 25442469 parameters, 25442453 gradients, 103.2 GFLOP
s

```

Transferred 931/937 items from /content/weights/gelan-c.pt

AMP: checks passed 

optimizer: SGD(lr=0.01) with parameter groups 154 weight(decay=0.0), 161 weight(decay=0.0005), 160 bias

INFO: albumentations.check_version: A new version of Albumentations is available: 1.4.13 (you have 1.4.12). Upgrade using: pip install -U albumentations. To disable automatic update checks, set the environment variable NO_ALBUMENTATIONS_UPDATE to 1.

albumentations: Blur(p=0.01, blur_limit=(3, 7)), MedianBlur(p=0.01, blur_limit=(3, 7)), ToGray(p=0.01), CLAHE(p=0.01, clip_limit=(1, 4.0), tile_grid_size=(8, 8))

/usr/lib/python3.10/multiprocessing/popen_fork.py:66: RuntimeWarning: os.fork() was called. os.fork() is incompatible with multithreaded code, and JAX is multithreaded, so this will likely lead to a deadlock.

self.pid = os.fork()

train: Scanning /content/yolov9/O.D-IN-BAD-WEATHER-1/train/labels... 815 images, 5 backgrounds, 0 corrupt: 100% 815/815 [00:00<00:00, 1690.07it/s]

train: New cache created: /content/yolov9/O.D-IN-BAD-WEATHER-1/train/labels.cache

val: Scanning /content/yolov9/O.D-IN-BAD-WEATHER-1/valid/labels... 218 images, 2 backgrounds, 0 corrupt: 100% 218/218 [00:00<00:00, 650.85it/s]

val: New cache created: /content/yolov9/O.D-IN-BAD-WEATHER-1/valid/labels.cache

Plotting labels to runs/train/exp/labels.jpg...

Image sizes 640 train, 640 val

Using 2 dataloader workers

Logging results to runs/train/exp

Starting training for 25 epochs...

| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
|---|---------|------------------------|----------|----------|-----------|---------|
| 0/24 | 14.5G | 1.568 | 2.077 | 1.285 | 321 | 640: 10 |
| 0% 51/51 [00:56<00:00, 1.10s/it] | | | | | | |
| | | Class Images Instances | | P | R | mAP50 n |
| AP50-95: 100% 7/7 [00:07<00:00, 1.06s/it] | | | | | | |
| | | all 218 2666 | | 0.719 | 0.229 | 0.235 |
| 0.141 | | | | | | |

| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
|---|---------|------------------------|----------|----------|-----------|---------|
| 1/24 | 14.5G | 1.39 | 1.229 | 1.175 | 435 | 640: 10 |
| 0% 51/51 [00:41<00:00, 1.24it/s] | | | | | | |
| | | Class Images Instances | | P | R | mAP50 n |
| AP50-95: 100% 7/7 [00:05<00:00, 1.27it/s] | | | | | | |
| | | all 218 2666 | | 0.7 | 0.254 | 0.236 |
| 0.146 | | | | | | |

| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
|---|---------|------------------------|----------|----------|-----------|---------|
| 2/24 | 12.5G | 1.417 | 1.285 | 1.198 | 337 | 640: 10 |
| 0% 51/51 [00:42<00:00, 1.21it/s] | | | | | | |
| | | Class Images Instances | | P | R | mAP50 n |
| AP50-95: 100% 7/7 [00:05<00:00, 1.29it/s] | | | | | | |
| | | all 218 2666 | | 0.626 | 0.246 | 0.236 |
| 0.139 | | | | | | |

| Epoch | GPU_mem | box_loss | cls_loss | dfl_loss | Instances | Size |
|-------|---------|----------|----------|----------|-----------|---------|
| 3/24 | 13G | 1.461 | 1.272 | 1.221 | 358 | 640: 10 |

```

0% 51/51 [00:43<00:00, 1.18it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:06<00:00, 1.15it/s]
      all        218      2666      0.566      0.237      0.22
0.129

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      4/24         13G         1.474         1.2         1.215         355         640: 10
0% 51/51 [00:42<00:00, 1.19it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:06<00:00, 1.16it/s]
      all        218      2666      0.676      0.204      0.213
0.103

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      5/24         13G         1.454         1.136         1.207         306         640: 10
0% 51/51 [00:41<00:00, 1.22it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.17it/s]
      all        218      2666      0.66      0.222      0.231
0.133

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      6/24         13G         1.452         1.11         1.205         354         640: 10
0% 51/51 [00:42<00:00, 1.21it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.24it/s]
      all        218      2666      0.593      0.235      0.232
0.128

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      7/24         13G         1.458         1.129         1.231         346         640: 10
0% 51/51 [00:42<00:00, 1.20it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.31it/s]
      all        218      2666      0.619      0.212      0.217
0.124

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      8/24         13G         1.447         1.071         1.197         310         640: 10
0% 51/51 [00:42<00:00, 1.20it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.30it/s]
      all        218      2666      0.751      0.209      0.231
0.132

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      9/24         13G         1.411         1.061         1.21         317         640: 10
0% 51/51 [00:41<00:00, 1.22it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.18it/s]
      all        218      2666      0.644      0.235      0.245
0.137

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      10/24         13G         1.408         1.045         1.203         355         640: 10
0% 51/51 [00:41<00:00, 1.21it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.18it/s]
      all        218      2666      0.721      0.244      0.255
0.142

      Epoch      GPU_mem      box_loss      cls_loss      dfl_loss      Instances      Size
      11/24         13G         1.394         1.019         1.191         533         640: 10
0% 51/51 [00:41<00:00, 1.22it/s]
      Class      Images Instances      P      R      mAP50  n
AP50-95: 100% 7/7 [00:05<00:00, 1.26it/s]

```

| | | | | | | |
|---------------------------------|---------|-----------|----------|----------|-----------|---------|
| 0.142 | all | 218 | 2666 | 0.602 | 0.218 | 0.268 |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 12/24 | 13G | 1.373 | 0.9734 | 1.187 | 234 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.31it/s] | | | | |
| all | 218 | 2666 | 0.582 | 0.275 | 0.275 | |
| 0.152 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 13/24 | 13G | 1.365 | 0.9574 | 1.176 | 200 | 640: 10 |
| 0% 51/51 [00:41<00:00, | | 1.23it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.17it/s] | | | | |
| all | 218 | 2666 | 0.465 | 0.3 | 0.279 | |
| 0.155 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 14/24 | 13G | 1.354 | 0.9434 | 1.153 | 361 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.23it/s] | | | | |
| all | 218 | 2666 | 0.567 | 0.253 | 0.254 | |
| 0.146 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 15/24 | 13G | 1.33 | 0.919 | 1.138 | 296 | 640: 10 |
| 0% 51/51 [00:41<00:00, | | 1.22it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.25it/s] | | | | |
| all | 218 | 2666 | 0.477 | 0.344 | 0.305 | |
| 0.175 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 16/24 | 13G | 1.315 | 0.9253 | 1.146 | 265 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.19it/s] | | | | |
| all | 218 | 2666 | 0.386 | 0.273 | 0.296 | |
| 0.178 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 17/24 | 13G | 1.306 | 0.8933 | 1.129 | 302 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:06<00:00, | | 1.16it/s] | | | | |
| all | 218 | 2666 | 0.379 | 0.326 | 0.3 | |
| 0.175 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 18/24 | 13G | 1.302 | 0.8924 | 1.146 | 408 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:05<00:00, | | 1.20it/s] | | | | |
| all | 218 | 2666 | 0.474 | 0.276 | 0.307 | |
| 0.176 | | | | | | |
| Epoch | GPU_mem | box_loss | cls_loss | df1_loss | Instances | Size |
| 19/24 | 13G | 1.283 | 0.8528 | 1.128 | 390 | 640: 10 |
| 0% 51/51 [00:42<00:00, | | 1.21it/s] | | | | |
| Class | Images | Instances | P | R | mAP50 | n |
| AP50-95: 100% 7/7 [00:06<00:00, | | 1.15it/s] | | | | |
| all | 218 | 2666 | 0.452 | 0.328 | 0.328 | |
| 0.19 | | | | | | |


```

Epoch GPU_mem box_loss cls_loss dfl_loss Instances Size
20/24 13G 1.284 0.8476 1.132 278 640: 10
0% 51/51 [00:42<00:00, 1.20it/s]
Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:05<00:00, 1.19it/s]
all 218 2666 0.488 0.318 0.339
0.191

```

```

Epoch GPU_mem box_loss cls_loss dfl_loss Instances Size
21/24 13G 1.275 0.8364 1.13 176 640: 10
0% 51/51 [00:41<00:00, 1.22it/s]
Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:05<00:00, 1.19it/s]
all 218 2666 0.575 0.335 0.337
0.19

```

```

Epoch GPU_mem box_loss cls_loss dfl_loss Instances Size
22/24 13G 1.24 0.7961 1.114 296 640: 10
0% 51/51 [00:42<00:00, 1.20it/s]
Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:05<00:00, 1.27it/s]
all 218 2666 0.625 0.319 0.347
0.207

```

```

Epoch GPU_mem box_loss cls_loss dfl_loss Instances Size
23/24 13G 1.248 0.7864 1.105 332 640: 10
0% 51/51 [00:42<00:00, 1.19it/s]
Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:05<00:00, 1.31it/s]
all 218 2666 0.632 0.315 0.357
0.214

```

```

Epoch GPU_mem box_loss cls_loss dfl_loss Instances Size
24/24 13G 1.233 0.7746 1.101 337 640: 10
0% 51/51 [00:42<00:00, 1.20it/s]
Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:05<00:00, 1.31it/s]
all 218 2666 0.475 0.377 0.379
0.218

```

25 epochs completed in 0.359 hours.

Optimizer stripped from runs/train/exp/weights/last.pt, saved as runs/train/exp/weights/last_stripped.pt, 51.5MB

Optimizer stripped from runs/train/exp/weights/best.pt, saved as runs/train/exp/weights/best_stripped.pt, 51.5MB

Validating runs/train/exp/weights/best.pt...

Fusing layers...

gelan-c summary: 467 layers, 25416357 parameters, 0 gradients, 102.5 GFLOPs

```

Class Images Instances P R mAP50 n
AP50-95: 100% 7/7 [00:13<00:00, 2.00s/it]
all 218 2666 0.475 0.377 0.378
0.218
bike 218 25 0.334 0.12 0.115
0.0521
bus 218 40 0.424 0.525 0.469
0.325
car 218 2080 0.705 0.749 0.777
0.494
motor 218 10 0.575 0.145 0.209
0.131
person 218 368 0.454 0.355 0.348
0.149
rider 218 19 0.372 0.211 0.286
0.0765
truck 218 124 0.461 0.532 0.445
0.297

```

Results saved to runs/train/exp

```
In [9]: !ls {HOME}/yolov9/runs/train/exp/
```

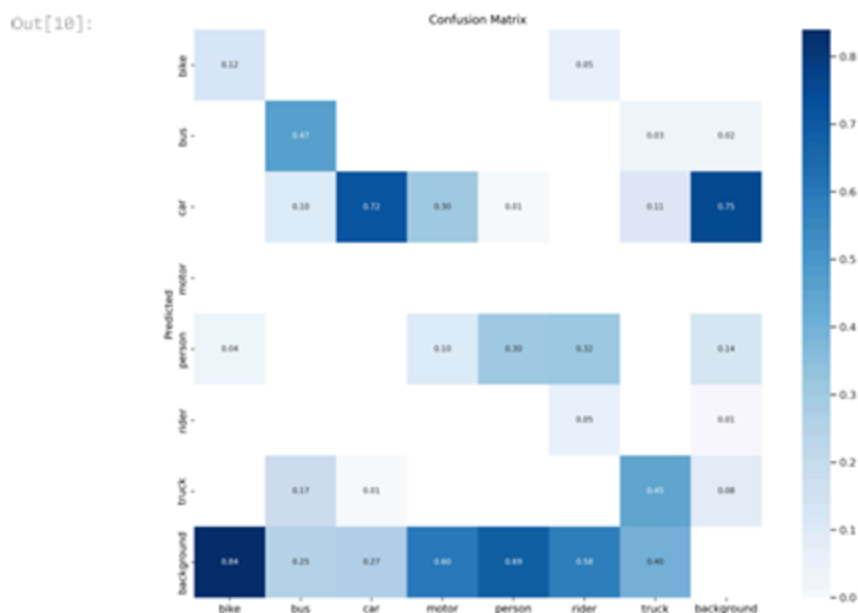
```
confusion_matrix.png      PR_curve.png             val_ba
tch0_pred.jpg             R_curve.png             val_ba
events.out.tfevents.1722960702.5ee47e966f1a.1251.0  results.csv             val_ba
tch1_labels.jpg           results.png              val_ba
F1_curve.png              train_batch0.jpg         val_ba
tch1_pred.jpg             train_batch1.jpg         weight
hyp.yaml                  train_batch2.jpg
tch2_labels.jpg           val_batch0_labels.jpg
labels_correlogram.jpg
tch2_pred.jpg
labels.jpg
s
opt.yaml
P_curve.png
```

Training Results

Confusion Matrix

```
In [10]: from IPython.display import Image

Image(filename=f"{HOME}/yolov9/runs/train/exp/confusion_matrix.png", width=1000)
```




Results

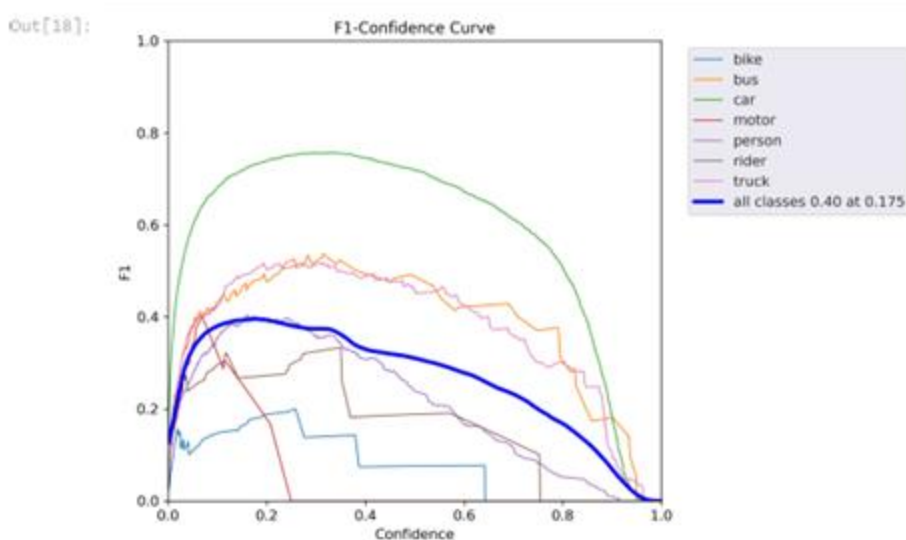
```
In [11]: from IPython.display import Image

Image(filename=f"{HOME}/yolov9/runs/train/exp/results.png", width=1000)
```

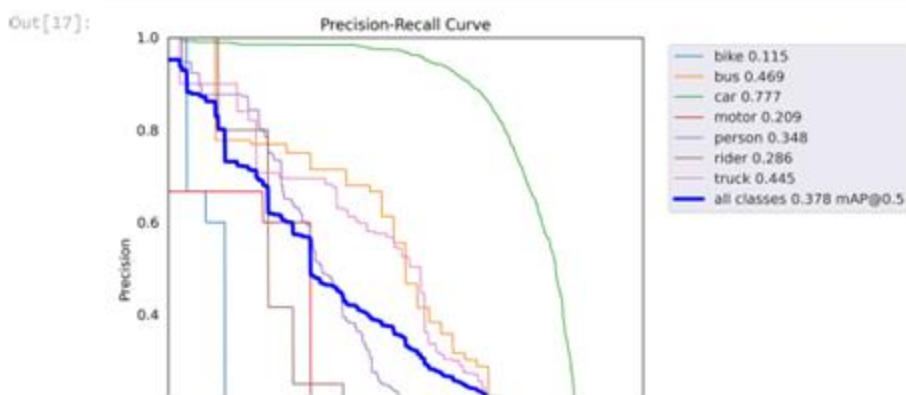





```
In [18]: from IPython.display import Image
Image(filename=f'/content/yolov9/runs/train/exp/F1_curve.png', width=1000)
```

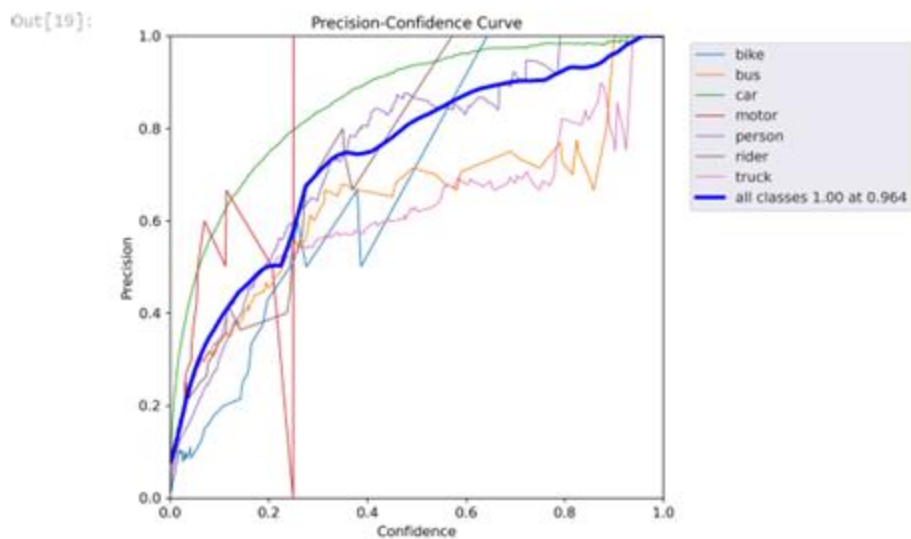


```
In [17]: from IPython.display import Image
Image(filename=f'/content/yolov9/runs/train/exp/PR_curve.png', width=1000)
```





```
In [19]: from IPython.display import Image
          Image(filename=f'/content/yolov9/runs/train/exp/P_curve.png', width=1000)
```



Training YOLOv10

GPU access verification

In [2]: `!nvidia-smi`

```
Tue Aug  6 16:39:56 2024
+-----+
+-----+
| NVIDIA-SMI 535.104.05                 Driver Version: 535.104.05   CUDA Version: 12.2     |
+-----+-----+
| GPU Name                               Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap       |      Memory-Usage | GPU-Util  Comp.  M. |
| MIG M. |
+-----+-----+
| 0 Tesla T4                             Off | 00000000:00:04.0 Off |                    |
| N/A   55C    P8              10W / 70W   |  0MiB / 15360MiB |      0%    D     |
| default |
| N/A   |
+-----+-----+

+-----+
+-----+
| Processes:
|
| GPU   GI   CI        PID   Type   Process name                      GPU
| Memory |
| ID   ID
+-----+-----+
| No running processes found
|
+-----+
+-----+
```

In [1]: `import os
HOME = os.getcwd()
print(HOME)`

/content

Installation of YOLOv10

In [3]: `!pip install -q git+https://github.com/THU-MIG/yolov10.git`

```
Installing build dependencies ... done
Getting requirements to build wheel ... done
Preparing metadata (pyproject.toml) ... done
Building wheel for ultralytics (pyproject.toml) ... done
```

In [4]: `!pip install -a supervision roboflow`

```

===== 0.0/135.7 kB ? eta -:--:--
===== 135.7/135.7 kB 4.9 MB/s eta 0:00:00
===== 0.0/76.9 kB ? eta -:--:--
===== 76.9/76.9 kB 6.5 MB/s eta 0:00:00
===== 178.7/178.7 kB 15.9 MB/s eta 0:00:00
===== 54.5/54.5 kB 5.0 MB/s eta 0:00:00

```

```

In [5]: !mkdir -p {HOME}/weights
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
!wget -P {HOME}/weights -q https://github.com/THU-MIG/yolov10/releases/download
ls -lh {HOME}/weights

```

```

total 408M
-rw-r--r-- 1 root root 80M May 26 15:53 yolov10b.pt
-rw-r--r-- 1 root root 100M May 26 15:53 yolov10l.pt
-rw-r--r-- 1 root root 64M May 26 15:54 yolov10m.pt
-rw-r--r-- 1 root root 11M May 26 15:54 yolov10n.pt
-rw-r--r-- 1 root root 32M May 26 15:54 yolov10s.pt
-rw-r--r-- 1 root root 123M May 26 15:54 yolov10x.pt

```

Dataset Upload

```

In [6]: !mkdir {HOME}/datasets
%cd {HOME}/datasets
!pip install -q roboflow
from google.colab import userdata
from roboflow import Roboflow

!pip install roboflow

from roboflow import Roboflow
rf = Roboflow(api_key="jPCXLMbZJU137MRBek9F")
project = rf.workspace("foreignobjectaerodromes").project("o.d-in-bad-weather")
version = project.version(1)
dataset = version.download("yolov9")

```

/content/datasets

```

Requirement already satisfied: roboflow in /usr/local/lib/python3.10/dist-packages
(1.1.37)
Requirement already satisfied: certifi in /usr/local/lib/python3.10/dist-packages
(from roboflow) (2024.7.4)
Requirement already satisfied: chardet==4.0.0 in /usr/local/lib/python3.10/dist-pa
ckages (from roboflow) (4.0.0)
Requirement already satisfied: idna==3.7 in /usr/local/lib/python3.10/dist-packag
es (from roboflow) (3.7)
Requirement already satisfied: cyclur in /usr/local/lib/python3.10/dist-packages
(from roboflow) (0.12.1)
Requirement already satisfied: kiwisolver>=1.3.1 in /usr/local/lib/python3.10/dis
t-packages (from roboflow) (1.4.5)
Requirement already satisfied: matplotlib in /usr/local/lib/python3.10/dist-packa
ges (from roboflow) (3.7.1)
Requirement already satisfied: numpy>=1.18.5 in /usr/local/lib/python3.10/dist-pa
ckages (from roboflow) (1.26.4)
Requirement already satisfied: opencv-python-headless==4.10.0.84 in /usr/local/li
b/python3.10/dist-packages (from roboflow) (4.10.0.84)
Requirement already satisfied: Pillow>=7.1.2 in /usr/local/lib/python3.10/dist-pa
ckages (from roboflow) (9.4.0)
Requirement already satisfied: python-dateutil in /usr/local/lib/python3.10/dist-
packages (from roboflow) (2.8.2)

```

```

Requirement already satisfied: python-dotenv in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.0.1)
Requirement already satisfied: requests in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.31.0)
Requirement already satisfied: six in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.16.0)
Requirement already satisfied: urllib3>=1.26.6 in /usr/local/lib/python3.10/dist-packages (from roboflow) (2.0.7)
Requirement already satisfied: tqdm>=4.41.0 in /usr/local/lib/python3.10/dist-packages (from roboflow) (4.66.4)
Requirement already satisfied: PyYAML>=5.3.1 in /usr/local/lib/python3.10/dist-packages (from roboflow) (6.0.1)
Requirement already satisfied: requests-toolbelt in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.0.0)
Requirement already satisfied: filetype in /usr/local/lib/python3.10/dist-packages (from roboflow) (1.2.0)
Requirement already satisfied: contourpy>=1.0.1 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (1.2.1)
Requirement already satisfied: fonttools>=4.22.0 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (4.53.1)
Requirement already satisfied: packaging>=20.0 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (24.1)
Requirement already satisfied: pyparsing>=2.3.1 in /usr/local/lib/python3.10/dist-packages (from matplotlib->roboflow) (3.1.2)
Requirement already satisfied: charset-normalizer<4,>=2 in /usr/local/lib/python3.10/dist-packages (from requests->roboflow) (3.3.2)
loading Roboflow workspace...
loading Roboflow project...

Downloading Dataset Version Zip in O.D-IN-BAD-WEATHER-1 to yolov9:: 100%|██████████| 75164/75164 [00:02<00:00, 25864.88it/s]

Extracting Dataset Version Zip to O.D-IN-BAD-WEATHER-1 in yolov9:: 100%|██████████| 2312/2312 [00:00<00:00, 4765.98it/s]

```

Training with Dataset

In [9]: %cd {HOME}

```

!yolo task=detect mode=train epochs=25 batch=16 imgsz=640 plots=True \
model={HOME}/weights/yolov10n.pt \
data=/content/datasets/O.D-IN-BAD-WEATHER-1/data.yaml

```

/content

New <https://pypi.org/project/ultralytics/8.2.74> available 🤖 Update with 'pip install -U ultralytics'

Ultralytics YOLOv8.1.34 🚀 Python-3.10.12 torch-2.3.1+cu121 CUDA:0 (Tesla T4, 15102MiB)

```

engine/trainer: task=detect, mode=train, model=/content/weights/yolov10n.pt, data=/content/datasets/O.D-IN-BAD-WEATHER-1/data.yaml, epochs=25, time=None, patience=100, batch=16, imgsz=640, save=True, save_period=-1, val_period=1, cache=False, device=None, workers=8, project=None, name=train3, exist_ok=False, pretrained=True, optimizer=auto, verbose=True, seed=0, deterministic=True, single_cls=False, rect=False, cos_lr=False, close_mosaic=10, resume=False, amp=True, fraction=1.0, profile=False, freeze=None, multi_scale=False, overlap_mask=True, mask_ratio=4, dropout=0.0, val=True, split=val, save_json=False, save_hybrid=False, conf=None, iou=0.7, max_det=300, half=False, dnn=False, plots=True, source=None, vid_stride=1, stream_buffer=False, visualize=False, augment=False, agnostic_nms=False, classes=None, retina_masks=False, embed=None, show=False, save_frames=False, save_txt=False, save_conf=False, save_crop=False, show_labels=True, show_conf=True, show_boxes=True, line_width=None, format=torchscript, keras=False, optimize=False, int8=False, dynamic=False, simplify=False, opset=None, workspace=4, nms=False, lr0=0.01, lrf=0.01, momentum=0.937, weight_decay=0.0005, warmup_epochs=3.0, warmup_momentum=0.8, warmup_bias_lr=0.1, box=7.5, cls=0.5, dfl=1.5, pose=12.0, kobj=1.0, label_smoothing=0.0, nbs=64, hsv_h=0.015, hsv_s=0.7, hsv_v=0.4, degrees=0.0, translate=0.1, scale=0.5, shear=0.0, perspective=0.0, flipud=0.0, fliplr=0.5, bgr=0.0, mosaic=1.0, mixup=0.0, copy_paste=0.0, auto_augment=randaug, erasing=0.4, crop_fr

```

```

action=1.0, cfg=None, tracker=botsort.yaml, save_dir=runs/detect/trains
Downloading https://ultralytics.com/assets/Arial.ttf to '/root/.config/yolov10/Arial.ttf'...
100% 755k/755k [00:00<00:00, 14.7MB/s]
2024-08-06 16:47:24.676931: E external/local_xla/xla/stream_executor/cuda/cuda_fft.cc:485] Unable to register cuFFT factory: Attempting to register factory for plugin cuFFT when one has already been registered
2024-08-06 16:47:24.935453: E external/local_xla/xla/stream_executor/cuda/cuda_dnn.cc:8454] Unable to register cuDNN factory: Attempting to register factory for plugin cuDNN when one has already been registered
2024-08-06 16:47:25.013554: E external/local_xla/xla/stream_executor/cuda/cuda_blas.cc:1452] Unable to register cuBLAS factory: Attempting to register factory for plugin cuBLAS when one has already been registered
Overriding model.yaml nc=80 with nc=7

```

| arguments | from | n | params | module |
|---------------------------|------|---|--------|---------------------------------------|
| 0 | -1 | 1 | 464 | ultralytics.nn.modules.conv.Conv |
| [3, 16, 3, 2] | -1 | 1 | 4672 | ultralytics.nn.modules.conv.Conv |
| 1 | -1 | 1 | 4672 | ultralytics.nn.modules.conv.Conv |
| [16, 32, 3, 2] | -1 | 1 | 7360 | ultralytics.nn.modules.block.C2f |
| 2 | -1 | 1 | 7360 | ultralytics.nn.modules.block.C2f |
| [32, 32, 1, True] | -1 | 1 | 18560 | ultralytics.nn.modules.conv.Conv |
| 3 | -1 | 1 | 18560 | ultralytics.nn.modules.conv.Conv |
| [32, 64, 3, 2] | -1 | 2 | 49664 | ultralytics.nn.modules.block.C2f |
| 4 | -1 | 2 | 49664 | ultralytics.nn.modules.block.C2f |
| [64, 64, 2, True] | -1 | 1 | 9856 | ultralytics.nn.modules.block.SCDOWN |
| 5 | -1 | 1 | 9856 | ultralytics.nn.modules.block.SCDOWN |
| [64, 128, 3, 2] | -1 | 2 | 197632 | ultralytics.nn.modules.block.C2f |
| 6 | -1 | 2 | 197632 | ultralytics.nn.modules.block.C2f |
| [128, 128, 2, True] | -1 | 1 | 36096 | ultralytics.nn.modules.block.SCDOWN |
| 7 | -1 | 1 | 36096 | ultralytics.nn.modules.block.SCDOWN |
| [128, 256, 3, 2] | -1 | 1 | 460288 | ultralytics.nn.modules.block.C2f |
| 8 | -1 | 1 | 460288 | ultralytics.nn.modules.block.C2f |
| [256, 256, 1, True] | -1 | 1 | 164608 | ultralytics.nn.modules.block.SPPF |
| 9 | -1 | 1 | 164608 | ultralytics.nn.modules.block.SPPF |
| [256, 256, 5] | -1 | 1 | 249728 | ultralytics.nn.modules.block.PSA |
| 10 | -1 | 1 | 249728 | ultralytics.nn.modules.block.PSA |
| [256, 256] | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| 11 | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| [None, 2, 'nearest'] | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| 12 | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| [1] | -1 | 1 | 148224 | ultralytics.nn.modules.block.C2f |
| 13 | -1 | 1 | 148224 | ultralytics.nn.modules.block.C2f |
| [384, 128, 1] | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| 14 | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| [None, 2, 'nearest'] | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| 15 | -1 | 1 | 0 | torch.nn.modules.upsampling.Upsample |
| [1] | -1 | 1 | 37248 | ultralytics.nn.modules.block.C2f |
| 16 | -1 | 1 | 37248 | ultralytics.nn.modules.block.C2f |
| [192, 64, 1] | -1 | 1 | 36992 | ultralytics.nn.modules.conv.Conv |
| 17 | -1 | 1 | 36992 | ultralytics.nn.modules.conv.Conv |
| [64, 64, 3, 2] | -1 | 1 | 0 | ultralytics.nn.modules.conv.Concat |
| 18 | -1 | 1 | 0 | ultralytics.nn.modules.conv.Concat |
| [1] | -1 | 1 | 123648 | ultralytics.nn.modules.block.C2f |
| 19 | -1 | 1 | 123648 | ultralytics.nn.modules.block.C2f |
| [192, 128, 1] | -1 | 1 | 18048 | ultralytics.nn.modules.block.SCDOWN |
| 20 | -1 | 1 | 18048 | ultralytics.nn.modules.block.SCDOWN |
| [128, 128, 3, 2] | -1 | 1 | 0 | ultralytics.nn.modules.conv.Concat |
| 21 | -1 | 1 | 0 | ultralytics.nn.modules.conv.Concat |
| [1] | -1 | 1 | 282624 | ultralytics.nn.modules.block.C2fCIB |
| 22 | -1 | 1 | 282624 | ultralytics.nn.modules.block.C2fCIB |
| [384, 256, 1, True, True] | -1 | 1 | 864058 | ultralytics.nn.modules.head.v10Detect |
| 23 | -1 | 1 | 864058 | ultralytics.nn.modules.head.v10Detect |
| [7, [64, 128, 256]] | -1 | 1 | 864058 | ultralytics.nn.modules.head.v10Detect |

YOLOv10n summary: 385 layers, 2709770 parameters, 2709754 gradients, 8.4 GFLOPs

Transferred 402/505 items from pretrained weights


```

Transferred 422/522 items from pretrained weights
TensorBoard: Start with 'tensorboard --logdir runs/detect/train3', view at http://localhost:6006/
Freezing layer 'model.23.dfl.conv.weight'
AMP: running Automatic Mixed Precision (AMP) checks with YOLOv8n...
Downloading https://github.com/ultralytics/assets/releases/download/v8.1.0/yolov8n.pt to 'yolov8n.pt'...
100% 6.23M/6.23M [00:00<00:00, 72.4MB/s]
AMP: checks passed ✓
train: Scanning /content/datasets/O.D-IN-BAD-WEATHER-1/train/labels... 815 images, 5 backgrounds, 0 corrupt: 100% 815/815 [00:00<00:00, 1851.90it/s]
train: New cache created: /content/datasets/O.D-IN-BAD-WEATHER-1/train/labels.cache
INFO:albumentations.check_version:A new version of Albumentations is available: 1.4.13 (you have 1.4.12). Upgrade using: pip install -U albumentations. To disable automatic update checks, set the environment variable NO_ALBUMENTATIONS_UPDATE to 1.
/usr/local/lib/python3.10/dist-packages/albumentations/core/composition.py:161: UserWarning: Got processor for bboxes, but no transform to process it.
  self._set_keys()
albumentations: Blur(p=0.01, blur_limit=(3, 7)), MedianBlur(p=0.01, blur_limit=(3, 7)), ToGray(p=0.01), CLAHE(p=0.01, clip_limit=(1, 4.0), tile_grid_size=(8, 8))
/usr/lib/python3.10/multiprocessing/popen_fork.py:66: RuntimeWarning: os.fork() was called. os.fork() is incompatible with multithreaded code, and JAX is multithreaded, so this will likely lead to a deadlock.
  self.pid = os.fork()
val: Scanning /content/datasets/O.D-IN-BAD-WEATHER-1/valid/labels... 218 images, 2 backgrounds, 0 corrupt: 100% 218/218 [00:00<00:00, 1466.02it/s]
val: New cache created: /content/datasets/O.D-IN-BAD-WEATHER-1/valid/labels.cache
Plotting labels to runs/detect/train3/labels.jpg...
optimizer: 'optimizer=auto' found, ignoring 'lr0=0.01' and 'momentum=0.937' and determining best 'optimizer', 'lr0' and 'momentum' automatically...
optimizer: AdamW(lr=0.000909, momentum=0.9) with parameter groups 95 weight(decay=0.0), 108 weight(decay=0.0005), 107 bias(decay=0.0)
TensorBoard: model graph visualization added ✓
Image sizes 640 train, 640 val
Using 2 dataloader workers
Logging results to runs/detect/train3
Starting training for 25 epochs...

```

| Epoch | GPU_mem | box_om | cls_om | dfl_om | box_oo | cls_oo |
|---------------|-----------------------------|-----------|-------------------------------|---------|--------|----------|
| 1/25 | 3.65G | 1.758 | 3.537 | 1.292 | 1.95 | 4.249 |
| 1.187 | 352 | 640: 100% | 51/51 [00:32<00:00, 1.59it/s] | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% | 7/7 [00:03<00:00, 1.79it/s] | | | | | |
| | all | 218 | 2666 | 0.00568 | 0.177 | 0.0311 |
| 0.0209 | | | | | | |

| Epoch | GPU_mem | box_om | cls_om | dfl_om | box_oo | cls_oo |
|---------------|-----------------------------|-----------|-------------------------------|--------|--------|----------|
| 2/25 | 3.38G | 1.732 | 2.098 | 1.281 | 1.908 | 2.981 |
| 1.175 | 203 | 640: 100% | 51/51 [00:23<00:00, 2.22it/s] | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% | 7/7 [00:04<00:00, 1.46it/s] | | | | | |
| | all | 218 | 2666 | 0.935 | 0.0315 | 0.0509 |
| 0.0311 | | | | | | |

| Epoch | GPU_mem | box_om | cls_om | dfl_om | box_oo | cls_oo |
|---------------|-----------------------------|-----------|-------------------------------|--------|--------|----------|
| 3/25 | 3.57G | 1.663 | 1.679 | 1.28 | 1.947 | 2.676 |
| 1.205 | 211 | 640: 100% | 51/51 [00:22<00:00, 2.25it/s] | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% | 7/7 [00:02<00:00, 2.38it/s] | | | | | |
| | all | 218 | 2666 | 0.711 | 0.0928 | 0.0813 |
| 0.0417 | | | | | | |

| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
|---|-----------|---|-----------|--------|--------|----------|
| df1_oo | Instances | Size | | | | |
| 4/25 | 3.77G | 1.619 | 1.579 | 1.264 | 1.922 | 2.491 |
| 1.198 | 141 | 640: 100% 51/51 [00:25<00:00, 2.02it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.42it/s] | | | | | | |
| | all | 218 | 2666 | 0.715 | 0.102 | 0.097 |
| 0.0526 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 5/25 | 3.39G | 1.551 | 1.524 | 1.246 | 1.86 | 2.419 |
| 1.192 | 139 | 640: 100% 51/51 [00:22<00:00, 2.28it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 2.22it/s] | | | | | | |
| | all | 218 | 2666 | 0.56 | 0.12 | 0.104 |
| 0.057 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 6/25 | 3.31G | 1.54 | 1.481 | 1.236 | 1.857 | 2.283 |
| 1.191 | 244 | 640: 100% 51/51 [00:23<00:00, 2.17it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.61it/s] | | | | | | |
| | all | 218 | 2666 | 0.565 | 0.122 | 0.11 |
| 0.06 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 7/25 | 3.53G | 1.527 | 1.39 | 1.223 | 1.855 | 2.132 |
| 1.177 | 229 | 640: 100% 51/51 [00:22<00:00, 2.23it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:04<00:00, 1.50it/s] | | | | | | |
| | all | 218 | 2666 | 0.594 | 0.149 | 0.129 |
| 0.0729 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 8/25 | 3.37G | 1.474 | 1.368 | 1.213 | 1.818 | 2.077 |
| 1.17 | 223 | 640: 100% 51/51 [00:21<00:00, 2.34it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 2.27it/s] | | | | | | |
| | all | 218 | 2666 | 0.596 | 0.119 | 0.123 |
| 0.0697 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 9/25 | 3.28G | 1.484 | 1.322 | 1.199 | 1.829 | 1.976 |
| 1.159 | 333 | 640: 100% 51/51 [00:24<00:00, 2.06it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:03<00:00, 1.92it/s] | | | | | | |
| | all | 218 | 2666 | 0.607 | 0.149 | 0.124 |
| 0.068 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 10/25 | 3.37G | 1.457 | 1.302 | 1.204 | 1.788 | 1.931 |
| 1.169 | 256 | 640: 100% 51/51 [00:22<00:00, 2.29it/s] | | | | |
| | Class | Images | Instances | Box(P | R | mAP50 mA |
| P50-95): 100% 7/7 [00:02<00:00, 2.52it/s] | | | | | | |
| | all | 218 | 2666 | 0.647 | 0.14 | 0.131 |
| 0.075 | | | | | | |
| Epoch | GPU_mem | box_om | cls_om | df1_om | box_oo | cls_oo |
| df1_oo | Instances | Size | | | | |
| 11/25 | 3.69G | 1.451 | 1.273 | 1.192 | 1.809 | 1.879 |
| 1.166 | 214 | 640: 100% 51/51 [00:23<00:00, 2.17it/s] | | | | |

```

Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:02<00:00, 2.63it/s]
      all      218      2666      0.613      0.15      0.141
0.0814

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
12/25      3.7G      1.43      1.232      1.176      1.771      1.795
1.148      202      640: 100% 51/51 [00:21<00:00, 2.35it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:03<00:00, 2.27it/s]
      all      218      2666      0.611      0.177      0.145
0.0799

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
13/25      3.29G      1.434      1.242      1.18      1.767      1.822
1.153      142      640: 100% 51/51 [00:22<00:00, 2.24it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:02<00:00, 2.45it/s]
      all      218      2666      0.631      0.151      0.144
0.085

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
14/25      3.42G      1.392      1.177      1.156      1.743      1.717
1.132      207      640: 100% 51/51 [00:21<00:00, 2.33it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:04<00:00, 1.75it/s]
      all      218      2666      0.597      0.164      0.142
0.0816

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
15/25      3.54G      1.394      1.168      1.15      1.759      1.702
1.13      167      640: 100% 51/51 [00:22<00:00, 2.25it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:02<00:00, 2.55it/s]
      all      218      2666      0.585      0.168      0.139
0.0821
Closing dataloader mosaic
/usr/local/lib/python3.10/dist-packages/albumentations/core/composition.py:161: UserWarning: Got processor for bboxes, but no transform to process it.
  self._set_keys()
albumentations: Blur(p=0.01, blur_limit=(3, 7)), MedianBlur(p=0.01, blur_limit=(3, 7)), ToGray(p=0.01), CLAHE(p=0.01, clip_limit=(1, 4.0), tile_grid_size=(8, 8))
/usr/lib/python3.10/multiprocessing/popen_fork.py:66: RuntimeWarning: os.fork() was called. os.fork() is incompatible with multithreaded code, and JAX is multithreaded, so this will likely lead to a deadlock.
  self.pid = os.fork()

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
16/25      3.56G      1.409      1.16      1.159      1.748      1.709
1.141      178      640: 100% 51/51 [00:27<00:00, 1.85it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:02<00:00, 2.34it/s]
      all      218      2666      0.601      0.16      0.14
0.0774

Epoch      GPU_mem      box_om      cls_om      dfl_om      box_oo      cls_oo
dfl_oo Instances      Size
17/25      3.14G      1.389      1.13      1.159      1.711      1.664
1.14      134      640: 100% 51/51 [00:20<00:00, 2.52it/s]
      Class      Images  Instances  Box(P      R      mAP50  mA
P50-95): 100% 7/7 [00:02<00:00, 2.73it/s]

```

```

                                all          218          2666          0.624          0.174          0.161
0.0968

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
18/25 3.15G 1.362 1.082 1.137 1.684 1.583
1.119 154 640: 100% 51/51 [00:21<00:00, 2.37it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:03<00:00, 1.77it/s]
                                all          218          2666          0.61          0.16          0.164
0.0942

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
19/25 3.14G 1.365 1.078 1.14 1.666 1.59
1.125 176 640: 100% 51/51 [00:19<00:00, 2.55it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:02<00:00, 2.56it/s]
                                all          218          2666          0.645          0.169          0.18
0.104

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
20/25 3.16G 1.346 1.035 1.13 1.672 1.537
1.12 190 640: 100% 51/51 [00:21<00:00, 2.32it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:03<00:00, 1.79it/s]
                                all          218          2666          0.668          0.191          0.188
0.11

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
21/25 3.14G 1.329 1.02 1.114 1.658 1.519
1.104 167 640: 100% 51/51 [00:20<00:00, 2.44it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:02<00:00, 2.58it/s]
                                all          218          2666          0.668          0.184          0.18
0.107

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
22/25 3.16G 1.319 1 1.115 1.641 1.483
1.107 142 640: 100% 51/51 [00:20<00:00, 2.48it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:02<00:00, 2.35it/s]
                                all          218          2666          0.659          0.19          0.203
0.117

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
23/25 3.17G 1.301 0.9793 1.108 1.628 1.468
1.1 166 640: 100% 51/51 [00:21<00:00, 2.33it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:03<00:00, 2.23it/s]
                                all          218          2666          0.642          0.207          0.202
0.112

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo
dfl_oo Instances Size
24/25 3.09G 1.294 0.9693 1.102 1.613 1.433
1.1 170 640: 100% 51/51 [00:20<00:00, 2.47it/s]
Class Images Instances Box(P R mAP50 mA
P50-95): 100% 7/7 [00:02<00:00, 2.80it/s]
                                all          218          2666          0.659          0.202          0.193
0.118

Epoch GPU_mem box_om cls_om dfl_om box_oo cls_oo

```

```

df1_oo Instances      Size
      25/25      3.1G      1.277      0.9493      1.09      1.597      1.437
1.091      174      640: 100% 51/51 [00:20<00:00, 2.48it/s]
      Class      Images      Instances      Box(P      R      mAP50      mA
P50-95): 100% 7/7 [00:04<00:00, 1.64it/s]
      all      218      2666      0.672      0.196      0.197
0.116

```

25 epochs completed in 0.194 hours.

Optimizer stripped from runs/detect/train3/weights/last.pt, 5.8MB

Optimizer stripped from runs/detect/train3/weights/best.pt, 5.8MB

Validating runs/detect/train3/weights/best.pt...

Ultralytics YOLOv8.1.34 Python-3.10.12 torch-2.3.1+cu121 CUDA:0 (Tesla T4, 15102MiB)

YOLOv10n summary (fused): 285 layers, 2697146 parameters, 0 gradients, 8.2 GFLOPs

```

      Class      Images      Instances      Box(P      R      mAP50      mA
P50-95): 100% 7/7 [00:09<00:00, 1.31s/it]
      all      218      2666      0.658      0.192      0.203
0.117

```

```

      bike      218      25      1      0      0
0

```

```

      bus      218      40      0.418      0.275      0.288
0.201

```

```

      car      218      2080      0.603      0.631      0.655
0.395

```

```

      motor      218      10      1      0      0.107
0.0249

```

```

      person      218      368      0.301      0.198      0.171
0.0692

```

```

      rider      218      19      1      0      0
0

```

```

      truck      218      124      0.288      0.242      0.203
0.128

```

Speed: 0.2ms preprocess, 5.6ms inference, 0.0ms loss, 0.0ms postprocess per image

Results saved to runs/detect/train3

Learn more at <https://docs.ultralytics.com/modes/train>

```
In [10]: !ls {HOME}/runs/detect/train/
```

args.yaml weights

Training Results

```
In [12]: from IPython.display import Image
```

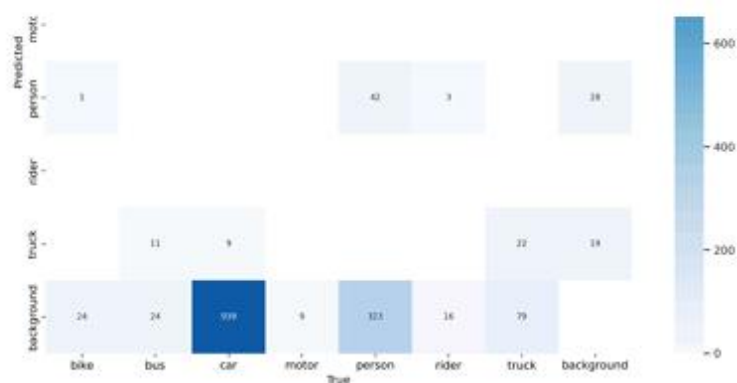
Confusion Matrix

```
In [14]: %cd {HOME}
Image(filename=f'/content/runs/detect/train3/confusion_matrix.png', width=1000)
```

/content

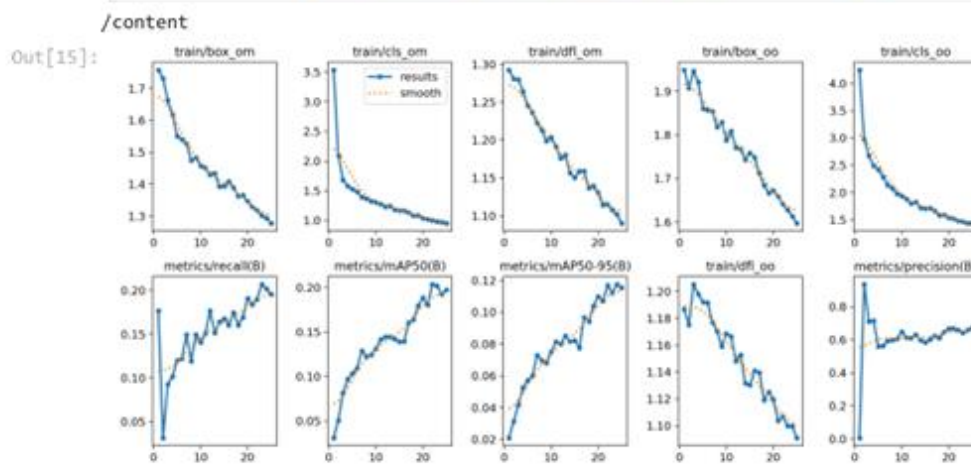
Out[14]:







Results

```
In [15]: %cd {HOME}
          Image(filename=f'/content/runs/detect/train3/results.png', width=1000)
```




```
In [16]: %cd {HOME}
          Image(filename=f'/content/runs/detect/train3/F1_curve.png', width=1000)
```





```
In [17]: %cd {HOME}
          Image(filename=f'/content/runs/detect/train3/PR_curve.png', width=1000)
```

/content



```
In [18]: %cd {HOME}
          Image(filename=f'/content/runs/detect/train3/P_curve.png', width=1000)
```

/content

