



# Optimized Custom Dataset for Efficient Detection of Underwater Trash

Jaskaran Singh Walia<sup>1</sup>(✉) and Karthik Seemakurthy<sup>2</sup>

<sup>1</sup> Vellore Institute of Technology, Chennai, India  
karanwalia2k3@gmail.com

<sup>2</sup> University of Lincoln, Lincoln, UK  
kseemakurthy@lincoln.ac.uk

**Abstract.** Accurately quantifying and removing submerged underwater waste plays a crucial role in safeguarding marine life and preserving the environment. While detecting floating and surface debris is relatively straightforward, quantifying submerged waste presents significant challenges due to factors like light refraction, absorption, suspended particles, and color distortion. This paper addresses these challenges by proposing the development of a custom dataset and an efficient detection approach for submerged marine debris. The dataset encompasses diverse underwater environments and incorporates annotations for precise labeling of debris instances. Ultimately, the primary objective of this custom dataset is to enhance the diversity of litter instances and improve their detection accuracy in deep submerged environments by leveraging state-of-the-art deep learning architectures. The source code to replicate the results in this paper can be found at [GitHub](#).

**Keywords:** Deep Learning · Computer Vision · Visual Object Detection · Artificial Intelligence · Robotics · Marine Debris · Trash Detection · Dataset Generation · Data Pre-processing

## 1 Introduction

Over the past few years, the increase in underwater debris due to poor waste management practices, littering, and international industry expansion has resulted in numerous environmental issues, such as water pollution and harm to aquatic life [1, 2]. The debris, which remains in the epipelagic and mesopelagic zones (Fig. 1) for years after it is dumped into the water, not only pollutes the water but also harms aquatic animals [3]. However, removing debris from beneath the aquatic surface is challenging and expensive. Thus, there is a need for a cost-effective solution that can operate effectively and efficiently in a wide range of environments.

Recent advances in robotics, artificial intelligence, and automated driving [5, 6] have made it feasible to use intelligent robots for underwater debris removal. Nevertheless, existing approaches are costly, and computationally demanding. Additionally, publicly available datasets are environment-specific, which limits their ability to produce a generalized and robust model. Therefore, we propose a




Fig. 1. Zones of the Oceans [4].

new dataset where the main focus is to increase the diversity of litter instances and enhance the generalization ability of state-of-the-art object detectors.

Autonomous underwater vehicles (AUVs) are a crucial component of a successful strategy for removing debris from maritime ecosystems. Therefore, the primary requirement for Autonomous Underwater Vehicles (AUVs) is the detection of underwater debris, specifically plastic debris. To address this challenge, we evaluated the dataset using advanced computer vision techniques to establish a baseline for litter detection. Our goal is to replace resource-intensive, time consuming algorithms with more efficient ones that will aid in real-time underwater debris detection. In this regard, we explore various deep learning based visual object detection networks that are trained and tested using the proposed dataset. The effectiveness of these detectors is evaluated using multiple metrics to validate their performance accurately.

The following are the main contributions of this paper:

- Proposed a new dataset with a focus to increase the diversity of litter instances under different environments.
- Trash, Rover and Bio are the three classes in the proposed dataset.
- Benchmarked the litter detection by using various deep learning-based object detectors.

## 2 Related Work

The literature in underwater robotics has focused on the development of multi-robot systems for surface and deep water natural aquatic environment applications such as marine monitoring using learning-based recognition of underwater biological occurrences by the National Oceanic and Atmospheric Administration [7]. Underwater robots have also been utilized for environmental monitoring, environmental mapping [8], maritime robotics for guidance and localization [9–12].

Underwater debris, particularly plastic waste, has become a significant environmental concern due to its detrimental effects on marine ecosystems. Plastic

debris can persist in the marine environment for long periods, posing threats to marine organisms through entanglement, ingestion, and habitat destruction [13,14]. It can also disrupt marine food webs and alter the biodiversity of marine ecosystems.

Efforts have been made to address the issue of underwater debris removal. Various methods have been employed, including manual clean-up operations, the use of remotely operated underwater vehicles (ROVs) equipped with gripping arms to physically capture debris, and the development of autonomous robotic systems specifically designed for marine debris removal. However, these approaches often face challenges in terms of efficiency, cost-effectiveness, and the vast scale of the problem. Researchers and organizations continue to explore innovative strategies and technologies to effectively tackle underwater debris and minimize its impact on the marine environment [15,16].

Recently, underwater robotics (ROVs) is considered as a popular alternative over the harmful manual methods to remove the marine debris. The vision system of a robot will aid in localising the debris and provide appropriate feedback to physically control a gripper limb to capture the objects of interest. A non-profit group for environmental protection and cleaning, Clear Blue Sea [17], has proposed the FRED (Floating Robot for Eliminating Debris). However, the FRED platform is not autonomous. In order to find garbage in marine habitats, another nonprofit, the Rozalia project, has employed underwater ROVs fitted with multibeam and side-scan sonars [3]. Autonomous garbage identification and collection for terrestrial settings have also been studied, such as in the case of Kulkarni et al. [18], who employed ultrasonic devices and applied them to interior garbage. However, a vision-based system can also be envisioned.

In a study on at-sea tracking of marine detritus by Mare [19], various tactics and technological possibilities were addressed. Following the 2011 tsunami off the shore in Japan, researchers have examined the removal of detritus from the ocean's top [20] using advanced Deep Visual Detection Models. In the study by M. Bernstein [21], LIDAR was used to locate and record beach garbage.

In recent research by Valdenegro-Toro[22], it was shown that a deep convolutional neural network (CNN) trained on forward-looking sonar (FLS) images could identify underwater debris with an accuracy of about 80%. This study made use of a custom made dataset created by capturing FLS images of items frequently discovered with marine debris in a water tank. Data from water tanks were also used in the assessment.

As mentioned above, the majority of the literature which dealt with debris detection used either sonar or lidar sensors. However, visual sensors have superior resolution over sensors such as sonars or lidar. A sizable, labeled collection of underwater detritus is required to allow visual identification of underwater litter using a deep learning-based model. This collection needs to include information gathered from a broad variety of underwater habitats to accurately capture the various looks across wide-ranging geographic areas. There are very few datasets that are publicly available and majority of them are unlabeled. The Monterey Bay Aquarium Research Institute (MBARI) has amassed a dataset

over 22 years to survey trash strewn across the ocean floor off the western coast of the United States of America [23], specifically plastic and metal inside and around the undersea Monterey Canyon, which traps and transports the debris in the deep oceans. The Global Oceanographic Data Center, a division of the Japan Agency for Marine-Earth Science and Technology (JAMSTEC), is another example of a publicly available large dataset. As part of the larger J-EDI (JAMSTEC E-Library of Deep-sea Images) collection, JAMSTEC has made a dataset of deep-sea detritus available online [11]. This dataset provides type-specific debris data in the form of short video clips and images dating back to 1982. The annotated data was utilized to construct deep learning-based models for the work discussed in this paper.

In summary, various studies have been conducted on the use of autonomous robots for underwater monitoring and debris collection. The development of multi-robot systems for environmental monitoring, environmental mapping, maritime robotics, and other applications have utilized undersea robots. Researchers have also explored learning-based recognition of underwater biological occurrences for marine monitoring. Additionally, the use of remotely operated underwater vehicles (ROVs) and autonomous garbage identification and collection for terrestrial settings has been studied. A significant labeled collection of underwater trash is necessary for accurate identification using deep learning-based models.

## 3 Dataset

### 3.1 Existing Datasets

Although several publicly available datasets, including JAMSTEC, J-ED, Trash-CAN 1.0, and Trash-ICRA19, exist for automatic waste detection, they are highly domain-specific and restricted to very limited environmental variations [24]. Table 1 and 2 shows the statistics of existing detection and classification datasets, respectively. This limits the generalisation ability for using the vision based debris detectors across wide variety of water bodies. Also, lack of diversity in the existing datasets can induce bias into the object detection networks. The main aim of the proposed dataset is to increase the diversity of images in identifying three classes, namely Trash, Rover, and Bio, which are most useful for classifying submerged debris.

The Bio class provides an aspect of marine life in the environment and how much trash has affected it relative to nearby environments which can be used to further prioritise the trash cleaning. The Rover class helps differentiate the rover from being misclassified as trash in some input imagery. Finally, the Trash class helps to detect and quantify the amount of trash present in the input image/video. The dataset was curated by collecting inputs from various open-source datasets and videos across different oceans and water bodies with varying conditions and environments. We manually annotated marine debris in frames of images, focusing on selecting images with tricky object detection conditions

**Table 1.** Comparison of existing litter detection datasets.

Type	Dataset	# classes	# images	Annotation type	Environment
Detection	Wade-AI	1	1396	instance masks	outdoor
Detection	Extended TACO	7	4562	bounding box	outdoor
Detection	UAVWaste	1	772	instance masks	outdoor
Detection	TrashCan	8	7212	instance masks	underwater
Detection	Trash-ICRA	7	7668	bounding box	underwater
Detection	Drinking waste	4	4810	bounding box	indoor
Detection	MJU-Waste	1	2475	instance masks	indoor
Detection	MJU-Waste	1	2475	instance masks	indoor

**Table 2.** Comparison of existing litter classification datasets.

Type	Dataset	Images	Classes	Instances	Environment	Annotation type
Classification	Open Litter Map	> 100k	> 100	> 100k	outdoor	multilabels
Classification	Waste pictures	23633	34	23633	outdoor	labels
Classification	TrashNet	2194	5	2194	indoor	labels

like occlusion, noise, and illumination. We used an annotation tool [25] to create the final dataset, which comprises of 9625 images.

### 3.2 Data Preparation

The first step to create this dataset is to collect inputs from various open-source datasets and videos with varying ocean environments from different countries. We manually annotated the marine debris in frames of images, focusing on selecting images with difficult object detection conditions such as occlusion, noise, and illumination. The annotations were done using a free annotation tool [25], resulting in 9,625 images in the dataset. A few of the sample images from our dataset are shown in Fig. 2. It can be seen that the diversity of objects and the environments that were considered in this paper.

**Fig. 2.** Representative images from proposed dataset.

A Deep learning analysis was also performed on the pre-existing datasets, which can be viewed on our paper’s source repository. While the models performed well on training, they failed to accurately detect classes when tested on unseen data from a slightly varying environment. Our dataset comprises of bounding box labels and image annotations and is available in more than ten different formats, making it readily importable for use with different algorithms. The dataset was prepared using the following steps.

1. **Data collection:** The input images were selectively picked manually, comprising of varying environments across different regions of the world.
2. **Annotation:** The unlabelled raw images were annotated and the annotations of labelled images were merged and renamed into three final categories, Trash, Rov and Bio which stands for underwater debris, rover (autonomous vehicle) and biological marine life respectively.
3. **Pre-processing:** These images were then rescaled to  $416 \times 416$ . A total of 26 classes were dropped and mapped into the final three classes. Clear water images that comprised of no annotations were also added to make the model more robust towards different environments. The dataset was further improved by randomly distorting the brightness and saturation of the images using PyTorch’s built-in Transforms augmentation tool. This was done in order to mitigate the effects of spurious correlations on the model and to replicate variable underwater conditions such as illumination, occlusion, and coloring.

The total dataset consisted of 9625 images which were split into approximately 7300 for training, 1800 for validation and 473 for test. The Labels of the dataset were as follows:

- **Trash:** All sorts of marine debris (plastics, metals, etc.).
- **Bio:** All naturally occurring biological material, marine life, plants, etc.
- **Rover:** Parts of the rover such as a robotic arm, sensors, or any part of the AUV to avoid misclassification.

The main objective behind choosing these three particular classes is that the trash class will contain all forms of trash, this increases the model’s robustness when encountering unseen/new form of trash. The Bio class provides an aspect of current marine life in the environment and how much trash has affected it relative to nearby environments which can be used to prioritise the trash cleaning based on the quality of the marine life present. The Rover class helps the rover’s components from being misclassified as trash in some input imagery.

## 4 Benchmarking

This section presents the latest trash detection and classification models, followed by benchmarks for the proposed dataset and statistical evaluation of the training metrics.

## 4.1 Object Detection Algorithms

The various architectures selected for this project were chosen from the most recent, efficient, and successful object detection networks currently in use. Each has its advantages and disadvantages, with different levels of accuracy, execution speeds, and other metrics. We utilized several state-of-the-art neural network architectures, including YOLOv7, YOLOv6s, YOLOv5s, and YOLOv4 Darknet, using their respective repositories. We also trained a custom FasterR-CNN and Mask R-CNN.

## 4.2 GPU Hardware

In this project, we utilized an Nvidia K80 GPU with a memory of 12 GB and a memory clock rate of 0.82GHz. This GPU was released in 2014 and has two CPU cores with 358 GB of disk space.

## 4.3 Models

In this section, we discuss the latest models used and the results produced.

**You Only Look Once (YOLO).** You Only Look Once, or YOLO, is a popular object detection technique that can recognize multiple items in a real-time video or image. In one evaluation, it utilizes a single neural network to predict bounding boxes and class probabilities straight from the complete image. Due to this approach, YOLO is faster and more accurate than other object detection systems and therefore it can provide fast inference speeds for the real-time application of this research.

- **YOLOv7 tiny** [26]: The YOLOv7 algorithm outperforms its older versions in terms of speed and accuracy. It requires significantly less hardware than conventional neural networks and can be trained much more quickly on small datasets with no pre-learned weights.
- **YOLOv5s small** [27] and **YOLOv6s (small)** [28]: Both of these algorithms have similar performances and results.

**Faster R-CNN and Mask R-CNN.** Faster R-CNN and Mask R-CNN are two popular region-based object detection models that use a two-stage approach for object detection. The first stage generates region proposals, and the second stage predicts the class and refines the bounding box of each proposal.

**Mask R-CNN.** [28]: Mask R-CNN extends Faster R-CNN by adding a branch to predict segmentation masks for each object. This allows the model to also segment the detected objects in addition to predicting their bounding boxes and class probabilities.

#### 4.4 Evaluation Metrics

After the model has been trained, we utilize the testing and validation datasets, which comprise images that are mutually exclusive from the training dataset, as input to analyze the network’s accuracy. The model generates a bounding box around correctly recognized items with a confidence value of .50 or higher. The amount of true positive bounding boxes drawn around marine plastic waste and true negatives serves as the basis for evaluation.

The following performance metrics were used to validate and compare the performance of the detectors used:

- **True positive and True negative values.**
- **Precision and Recall:** Reflects whether the model predicted debris in the input image.

$$Recall = \frac{TP}{TP + FN}, \quad Precision = \frac{TP}{TP + FP} \quad (1)$$

- **Mean Average Precision:** - Determines how frequently the network can correctly identify plastic. After gathering the true and false positive data, use the Intersection over Union (IoU) formula to build a precision-recall curve:

$$IOU = \frac{BBox_{pred} \cap BBox_{GroundTruth}}{BBox_{pred} \cup BBox_{GroundTruth}} \quad (2)$$

where  $BBox_{Pred}$  and  $BBox_{GroundTruth}$  are the expected areas under the curve for predicted and ground truth bounding boxes. To maximize accuracy, a high threshold for confidence and IoU must be specified, with a correct prediction, indicated by the threshold being exceeded. After that, the mAP can be calculated by integrating the precision-recall curve. obtained by integrating the precision-recall curve [29]:

$$mAP = \int_0^1 p(x) dx \quad (3)$$

## 5 Results

The results obtained for debris localization on our custom-curated dataset outperform previous models that used individual datasets for training. In this study, we tested the individual components of two frameworks by conducting exhaustive research on publicly available waste data in various contexts, including clean waters, natural or man-made lakes/ponds, and ocean beds. The broad range of baseline results for different contexts and diverse object dimensions will assist future researchers in this field. The tested models exhibit high average precision, mAP, and F1 scores compared to their inference speed.

The outcomes of a comprehensive study comparing several architectural networks are presented in Table 3. These trade-offs suggest that the results reported



in this study better reflect long-term performance in a wider range of marine conditions, enabling a more comprehensive evaluation of the object identification model’s performance in the field. Our findings suggest that YOLOv5-Small and YOLOv6s both achieve strong debris localization metrics in the real-time detection of epipelagic plastic. However, YOLOv7 yields a notably higher F1 score despite a slight reduction in inference performance. The results of a comprehensive research comparing several architectural networks is shown in the table below.

**Table 3.** Comparison between various algorithms for the purpose of benchmarking.


<u>Network</u>	<u>mAP 0.5</u>	<u>Precision</u>	<u>Recall</u>	<u>Epochs</u>
YOLOv7	0.96	0.96	0.93	120
YOLOv5s	0.96	0.95	0.93	110
YOLOv6s	0.90	0.94	0.92	180
Faster R-CNN	0.81	0.88	–	100
Mask R-CNN	0.83	0.85	–	100

The trade-offs observed in our study demonstrate that the reported outcomes reflect the long-term performance of the object identification model in a wider range of marine conditions, thereby facilitating a more comprehensive evaluation of the model’s performance in the field. Our findings suggest that YOLOv5-Small and YOLOv6s achieve excellent debris localization metrics in real-time detection of epipelagic plastic. However, YOLOv7 achieves a significantly higher F1 score despite a slight decrease in inference performance.

After evaluating multiple advanced algorithms within the same family, including YOLOv5x, v7E6E, and v8x, it was determined that the nano/small/tiny network architecture demonstrated the highest performance in evaluations, had a smaller parameter count, and required less computational power. As a result, this architecture was selected for the study. These algorithms outperformed classic Faster-RCNN and Mask-RCNN algorithms in terms of both speed and F1 score.

The performance of the model in real-world scenarios was found to be consistent with the evaluation results presented in Table 3, with only slight variations observed in a near-real-time setting. These results demonstrate the model’s efficacy in identifying and categorizing underwater debris in practical applications. Furthermore, the proposed research can serve as a crucial baseline and benchmark for future investigations focused on the identification and classification of marine debris.

$$ImgCord_k = BoxScore_i^j * Width \quad (4)$$



**Fig. 3.** Quantitative Analysis. (a) YOLOv5 and (b) YOLOv8. First row: Precision curves. Second row: Recall curves.

where  $k$  belongs to the top, down, right and left corners,  $i$  is the box index,  $j \in 0, 1, 2, 3$ , and  $Width$  is the image width. In the test photos, these image coordinates were utilized to illustrate the results of predicted bounding boxes.

## 6 Conclusion

In this research, our objective was to improve object detection models by reducing dependence on environment-specific datasets. By employing our mixed, curated dataset and the latest state-of-the-art computer vision models, we were able to evaluate the feasibility of monitoring submerged marine debris in near-real-time for debris quantification. Through the use of robotic arms within Autonomous Underwater Vehicles (AUVs), our rapid inference speeds achieved a high level of performance, making real-time object detection of marine plastic litter in the ocean's epipelagic and mesopelagic layer possible, as well as the automatic detection, classification, and sorting of various submerged objects, including the collection of debris in locations such as sea-beds that are inaccessible to humans due to high pressure and other environmental factors. This application has the potential to automate trash recycling in the extreme aquatic environment with the help of deep learning. Furthermore, our proposed research serves as a fundamental baseline and benchmark for future research involving the identification and classification of underwater debris.

## References

1. Coyle, R., Hardiman, G., Driscoll, K.O.: Microplastics in the marine environment: A review of their sources, distribution processes, uptake and exchange in ecosystems. *Case Stud. Chem. Environ. Eng.* **2**, 100010 (2020). <https://doi.org/10.1016/j.csee.2020.100010>, <https://www.sciencedirect.com/science/article/pii/S2666016420300086>
2. Derraik, J.G.B.: The pollution of the marine environment by plastic debris: a review. *Marine Pollution Bull.* **44**(9), 842–52 (2002)
3. Honingh, D., van Emmerik, T., Uijttewaal, W., Kardhana, H., Hoes, O., van de Giesen, N.: Urban river water level increase through plastic waste accumulation at a rack structure. *Front. Earth Sci.* **8**. <https://doi.org/10.3389/feart.2020.00028>
4. Laves of the oceans. <https://www.sas.upenn.edu/msheila/biolumevolution.html>
5. Xiao, Y., et al.: A review of object detection based on deep learning. *Multimedia Tools Appl.* **79**(33–34), 23729–23791 (2020). <https://doi.org/10.1007/s11042-020-08976-6>
6. Andreu-Perez, J., Deligianni, F., Ravi, D., Yang, G.-Z.: Artificial Intelligence and Robotics (2018). <https://doi.org/10.48550/ARXIV.1803.10813>. <https://arxiv.org/abs/1803.10813>
7. N Oceanic, A Administration: What is marine debris?. <https://oceanservice.noaa.gov/facts/marinedebris.html>
8. Yuan, X., Martínez-Ortega, J.-F., Fernández, J.A.S., Eckert, M.: Aekf-slam: A new algorithm for robotic underwater navigation. *Sensors* **17**(5), 1174 (2017)
9. Torrey, L.A., Shavlik, J.W.: Chapter 11 transfer learning (2009)
10. Drever, M.C., Provencher, J.F., O'Hara, P.D., Wilson, L., Bowes, V., Bergman, C.M.: Are ocean conditions and plastic debris resulting in a ‘double whammy’ for marine birds? *Marine Pollution Bull.* **133**, 684–692 (2018). <https://doi.org/10.1016/j.marpolbul.2018.06.028>, <https://www.sciencedirect.com/science/article/pii/S0025326X18304259>
11. Tensorflow “tensorflow object detection zoo. <https://github.com/tensorflow/models/blob/master/research/objectdetection/g3doc/detectionmodelzoo.md>
12. Derraik, J.G.: The pollution of the marine environment by plastic debris: a review. *Marine Pollution Bull.* **44**(9), 842–852 (2002). [https://doi.org/10.1016/S0025-326X\(02\)00220-5](https://doi.org/10.1016/S0025-326X(02)00220-5), <https://www.sciencedirect.com/science/article/pii/S0025326X02002205>
13. Thompson, R., et al.: Lost at sea: where is all the plastic? *Science (New York)* **304**, 838 (2004). <https://doi.org/10.1126/science.1094559>
14. Jambeck, J., et al.: Marine pollution. plastic waste inputs from land into the ocean. *Science (New York)* **347**, 768–771 (2015) . <https://doi.org/10.1126/science.1260352>
15. Jia, T., et al.: Deep learning for detecting macroplastic litter in water bodies: A review. *Water Res.* **119632** (2023)
16. Zocco, F., Lin, T.-C., Huang, C.-I., Wang, H.-C., Khyam, M.O., Van, M.: Towards more efficient efficientdets and real-time marine debris detection. *IEEE Robotics Autom. Lett.* **8**(4), 2134–2141 (2023)
17. Fred cars. [https://www.sandiego.edu/news/detail.php?\\_focus=72984](https://www.sandiego.edu/news/detail.php?_focus=72984)
18. Kulkarni, S., Junghare, S.: Robot based indoor autonomous trash detection algorithm using ultrasonic sensors, pp. 1–5 (2013). <https://doi.org/10.1109/CARE.2013.6733698>
19. Girdhar, Y., et al.: MARE: marine autonomous robotic explorer. In: Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011), San Francisco, USA, pp. 5048–5053 (2011)

20. Fulton, M., Hong, J., Islam, M., Sattar, J.: Robotic detection of marine litter using deep visual detection models
21. Bernstein, M., Graham, R., Cline, D., Dolan, J.M., Rajan, K.: Learning-Based Event Response For Marine Robotics, pp. 3362–3367 (2013)
22. Singh, D., Valdenegro-Toro, M.: The marine debris dataset for forward-looking sonar semantic segmentation, [arXiv: 2108.06800](https://arxiv.org/abs/2108.06800)
23. Stutters, L., Liu, H., Tiltman, C., Brown, D.J.: Navigation technologies for autonomous underwater vehicles. *IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.)* **38**(4), 581–589 (2008). <https://doi.org/10.1109/TSMCC.2008.919147>
24. Majchrowska, S., et al.: Learning-based waste detection in natural and urban environments. *Waste Manag.* **138**, 274–284 (2022). <https://doi.org/10.1016/j.wasman.2021.12.001>, <https://www.sciencedirect.com/science/article/pii/S0956053X21006474>
25. Alexandrova, S., Tatlock, Z., Cakmak, M.: Roboflow: a flow-based visual programming language for mobile manipulation tasks In: 2015 IEEE International Conference on Robotics and Automation (ICRA), pp. 5537–5544 (2015). <https://doi.org/10.1109/ICRA.2015.7139973>
26. Wang, C.-Y., Bochkovskiy, A., Liao, H.-Y.M.: Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors (2022). <https://doi.org/10.48550/ARXIV.2207.02696>, <https://arxiv.org/abs/2207.02696>
27. Jocher, G. et al.: Ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation (Nov 2022). <https://doi.org/10.5281/zenodo.7347926>
28. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015). <https://doi.org/10.48550/ARXIV.1512.03385>, <https://arxiv.org/abs/1512.03385>
29. Boyd, K., Eng, K.H., Page, C.D.: Area under the precision-recall curve: point estimates and confidence intervals. In: Blockeel, H., Kersting, K., Nijssen, S., Železný, F. (eds.) *ECML PKDD 2013. LNCS (LNAI)*, vol. 8190, pp. 451–466. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-40994-3\\_29](https://doi.org/10.1007/978-3-642-40994-3_29)