

State-of-the-Art Image Recognition for Surface Water Garbage Detection, Classification, and Counting

Prathamesh Garud

prathameshgarud0@gmail.com

Sanjiang University

Research Article

Keywords: water pollution, floating garbage, object detection, YOLOv8, image classification, object counting, deep learning

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11403507/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Surface Water Garbage Detection, Classification and Counting Using Modern Image Recognition

Prathamesh Garud

Sanjivani University, Kopargaon, Ahilyanagar, Maharashtra, India

Abstract—Garbage floating on rivers, lakes and ponds harms water life, blocks drains and spoils the look of public places. Checking it by eye is slow and tiring. This paper builds a simple image-based system that finds floating garbage in a photo, tells what type it is, and counts how many pieces of each type are present. A small YOLOv8 model (YOLOv8n, about 3.0 million parameters) was trained on a public water-garbage image set with 371 training images, 105 validation images and 7 labels. Counting is done by adding up the boxes the model draws for each class. In the training log used for this paper (24 completed epochs), the model reached a best mAP50 of 0.244, with precision of 0.624 and recall of 0.276 at that point, and it took about 155 ms to process one image on a CPU. These scores are modest. They show that the idea works even with very little data and no GPU, and they point clearly to what must improve: more labelled images, cleaner class labels, and a longer training run.

Keywords—water pollution, floating garbage, object detection, YOLOv8, image classification, object counting, deep learning

I. INTRODUCTION

Plastic and other waste reach rivers, lakes and canals every day. Once in the water, it can trap fish and birds, clog drains and break into smaller pieces that are very hard to remove. A well-known estimate by Jambeck et al. [1] showed how large the flow of plastic into the sea is, which makes regular monitoring of inland waters important.

Today most monitoring is done by people walking along the bank or travelling by boat and writing down what they see. This is slow, differs from person to person and cannot be repeated often. A camera on a bridge, a boat or a drone could take the same pictures every day, if a computer could read them. Reading means three jobs: finding each piece of garbage (detection), saying what kind it is (classification) and saying how many there are (counting).

Modern object-detection networks can do the first two jobs in one step, and the third follows naturally once every item has its own box. This study uses one such network, YOLOv8, on a small public data set and asks a practical question: how well can a small model, trained on a normal laptop-class CPU, find and count floating garbage? The aims are (i) to train a detector on water-garbage images, (ii) to measure it with standard detection scores, (iii) to show how counting is obtained from the detections, and (iv) to list honestly what limits the result.

Section II reviews earlier work, Section III explains the data and method, Section IV gives the results, and Section V closes with limits and future work.

II. RELATED WORK

Table I lists the main studies and tools that this work builds on. Detection networks fall into two families. Two-stage models such as Faster R-CNN [3] first suggest regions and then classify them; they are accurate but slower. One-stage models such as YOLO [2] and SSD [4] look at the whole picture once and are quick enough for live video. Fulton et al. [5] compared several of these on underwater plastic litter and showed that deep detectors can work on this kind of scene. Land-based litter has the TACO collection [6], and floating waste on inland water has the FloW benchmark [7]. YOLOv8 [8] is a recent, easy-to-use release of the YOLO family.

TABLE I
SELECTED RELATED WORK

#	Study	Contribution	Open point
1	Jambeck et al. [1]	Estimated plastic entering the sea from land	Needs cheap, repeatable monitoring
2	Redmon et al. [2]	YOLO: one-pass detection	Early version, weak on small items
3	Ren et al. [3]	Faster R-CNN: two-stage detector	Slower for live use
4	Liu et al. [4]	SSD: one-pass, many scales	Not built for waste
5	Fulton et al. [5]	Detectors for underwater litter	Underwater, not surface water
6	Proença and Simões [6]	TACO litter data set	Litter on land
7	Cheng et al. [7]	FloW floating-waste benchmark	Only a few object types
8	Jocher et al. [8]	YOLOv8 software	General purpose, not tuned for waste

Taken together, the earlier work shows that the detectors exist and that litter data sets exist, but few small, low-cost studies report both the type and the number of garbage items on the water surface. This paper tries to fill that gap in a modest way.

III. METHODOLOGY

A. Workflow

Fig. 1 shows the steps. The labelled images are used to fine-tune a YOLOv8n model that was already trained on a large general data set. After each epoch the model is checked on the validation images. Once trained, the model is given a new photo, draws a box and a label around every item it finds, and the boxes are counted per class.




Fig. 1. Workflow of the proposed system.

B. Dataset

The images come from a public Roboflow Universe project called Object-Detection (version 1, licence CC BY 4.0) [9], exported in YOLOv8 format. The split used here has 371 training images and 105 validation images; a separate test folder is also supplied but was not used in the run reported here. The 105 validation images contain 565 labelled objects, which is about five objects per image. The data file lists seven labels: No object detection, Object-Detection, Plastic Bag, Plastic Bottle, Plastic Can, Waste fibre and other waste. The first two look like placeholder names left by the labelling tool rather than real garbage types, and the original file spells some names differently (for example “Bottel” and “faiber”). They are kept as they are so the results match the data file. Per-class scores were not part of the log, so the reported numbers are averages over all classes.

C. Model and Training Setup

YOLOv8n is the smallest model of the YOLOv8 family. It has 130 layers, 3,012,213 parameters and needs 8.2 GFLOPs per image. Training used the Ultralytics library with the settings in Table II. All other settings were left at the library defaults, including the default image augmentation (mosaic, flips and colour changes). Training ran on a CPU, taking about 3.5 minutes per epoch.

TABLE II
TRAINING SETTINGS

Setting	Value
Starting weights	yolov8n.pt (pre-trained)
Image size	640 x 640
Batch size	8
Planned epochs	50
Early stopping patience	10 epochs
Hardware	CPU only

D. Counting

No separate counting network is needed. For one image, the model returns a list of boxes, each with a class and a

confidence value. Boxes with confidence below 0.25 are dropped. The remaining boxes are grouped by class, and the size of each group is the count for that class. For example, three boxes labelled Plastic Bottle give a count of three bottles.

E. Evaluation Measures

A box is counted as correct when it overlaps a true box enough, measured by intersection over union (IoU). Precision is $TP/(TP+FP)$, the share of drawn boxes that are right. Recall is $TP/(TP+FN)$, the share of real objects that were found. Average precision (AP) summarises precision and recall for one class, and mAP is its mean over classes. mAP50 uses an IoU limit of 0.5, while mAP50-95 averages the limits from 0.5 to 0.95 and is stricter [10]. Because counting depends on missed and extra boxes, recall and precision matter directly for count quality.

IV. RESULTS AND DISCUSSION

A. Training Progress

The log used here covers 24 completed epochs; the run was cut off during epoch 25, before the 50 planned epochs or the early-stopping rule ended it. Scores were still rising slowly, so the numbers below are a snapshot and not a final result. Table III gives selected epochs and Fig. 2 plots all of them.

TABLE III
VALIDATION SCORES AT SELECTED EPOCHS

Epoch	Precision	Recall	mAP50
1	0.008	0.268	0.034
5	0.470	0.249	0.108
10	0.648	0.249	0.161
15	0.614	0.256	0.203
18	0.622	0.286	0.235
21	0.577	0.255	0.238
23	0.624	0.276	0.244
24	0.563	0.267	0.231




Fig. 2. Validation scores for each completed epoch.

The highest mAP50 was 0.244 at epoch 23, where precision was 0.624 and recall was 0.276. The stricter mAP50-95 was highest at epoch 21 (0.121); by the library’s default way of ranking epochs (mostly mAP50-95) this is the epoch whose weights would normally be saved as the best model. The training losses fell over the same period: the box loss went from 2.153 at epoch 11 to 1.995 at epoch 24, and the class loss from 2.289 to 1.826.

B. Detection and Counting Example

As a quick check, the trained model was run on one validation image with a confidence limit of 0.25. It found three plastic bottles and the whole step took 154.9 ms on the CPU, plus about 10 ms for reading and cleaning up the result. Counting the boxes therefore gave three bottles for that image. This is one example only and is not a measure of counting accuracy; a proper count test would compare predicted and true counts over many images, for example with mean absolute error.

C. Discussion

The precision of about 0.6 means that roughly six out of ten boxes the model draws are right, while the recall of about 0.27 means it finds only about one in four real objects. In plain terms, the model is careful but misses a lot, so its counts would tend to be too low. The mAP50 of 0.244 is well below what larger models reach on cleaner data sets, and this paper does not claim otherwise.

Several causes are likely. First, 371 training images is small for seven classes. Second, the smallest YOLO model was used, which trades accuracy for speed. Third, two of the seven labels look like placeholders and may confuse the training. Fourth, water scenes are hard: glare, ripples, reflections and partly sunk items change how garbage looks. Finally, training stopped early in the learning curve. Even so, the model is small, runs on an ordinary CPU, and learned useful patterns quickly, which supports using this approach for low-cost monitoring.

The study has clear limits. The result is a snapshot from an unfinished run, there are no per-class scores, the test folder was not used, only one example image was used to show counting, and no comparison with other detectors was made.

V. CONCLUSION

This paper built a simple pipeline that detects, names and counts floating garbage using a small YOLOv8 model. In the logged training run, the model reached an mAP50 of 0.244 (precision 0.624, recall 0.276) after 24 epochs on 371 training images, using only a CPU. The results are modest but show that the pipeline works end to end and is cheap to run.

Future work includes finishing the full training run and reporting results on the test folder, cleaning the class list, adding many more images from different weather and light conditions, trying larger YOLOv8 models and other detectors, reporting per-class scores, and checking counts against hand-counted images. Running the model on a drone or boat camera would be the final real-world test.

REFERENCES

- [1] J. R. Jambeck et al., "Plastic waste inputs from land into the ocean," *Science*, vol. 347, no. 6223, pp. 768–771, 2015.
- [2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [4] W. Liu et al., "SSD: Single shot multibox detector," in *Proc. European Conf. Computer Vision (ECCV)*, 2016.
- [5] M. S. Fulton, J. Hong, M. J. Islam, and J. Sattar, "Robotic detection of marine litter using deep visual detection and classification models," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2019.
- [6] P. F. Proença and P. Simões, "TACO: Trash annotations in context for litter detection," *arXiv:2003.06975*, 2020.
- [7] Y. Cheng et al., "FloW: A dataset and benchmark for floating waste detection in inland waters," in *Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV)*, 2021.
- [8] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," software, 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [9] Roboflow Universe, "Object-Detection, version 1," workspace small-intestine-cancer, CC BY 4.0. [Online]. Available: <https://universe.roboflow.com/small-intestine-cancer/object-detection-zffem/dataset/1>
- [10] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. European Conf. Computer Vision (ECCV)*, 2014.