HIGH SCHOOL EDITION
Journal of Student Research

Volume 13 Issue 2 (2024)

Trashnet: An Object Detection Model That Classifies Images of Trash in Real-Time

Veerrohit Veeravadivel Santhanalakshmi[1] and Hieu Nguyen[#]

[1]The Quarry Lane School, USA
[#]Advisor

ABSTRACT

Global warming and pollution are huge problems today. One of the main factors behind these two problems is plastic pollution. As plastic takes millennia to decompose, an estimated 270,000 tons of plastic are floating around our oceans, which is too much to be considered "safe." This problem is mainly caused by the fact that people improperly dispose of plastic, whether that be through littering or putting it into the trash instead of a recycling bin. To better identify and correct plastic that was inappropriately disposed of, a deep-learning model (YOLOv5) that uses object detection and classification was implemented to detect which bin someone's trash should go in. The YOLOv5 uses the PyTorch framework for the object detection model. Using this model helped solve this problem, as a custom object-detection model would have needed to be developed, which would not have been efficient. The model was tested by trying to run the model on pieces of trash placed on a tabletop and analyzing the code output on which trash bin the waste will have to be thrown away. After conducting multiple tests, the model exhibited a commendable accuracy rate of 90%, which is noteworthy given the substantial amount of data leveraged. To further improve its efficacy and real-world value, future research could explore augmenting the training data, refining the object detection model for greater precision, and expanding the dataset to encompass a broader range of use cases.

Introduction

Proper recycling of materials positively affects our planet, but throwing away trash in the wrong bin causes many problems for this world's environment, according to a paper from the University of Washington's School of Environmental and Marine Affairs (Baechler, 2018). The process of sorting recyclable materials and landfills can go smoothly if the proper items are thrown away in the appropriate bins at different places. Suppose this is not done at all places, for example. If food is thrown into the recycling bin, recyclable items will be denied because of contamination in the recycled material shipment. This will cause a profit loss for the companies that buy the repurposing materials and result in logistic issues. An article from mic.com says that sometimes, when the wrong items are put into recycling bins, around $1 million worth of time is lost due to improper recycling in the United States(Ocampo, 2018). This same study also cites an increase in the economy, as an improved recycling program could create almost 2.3 million jobs in the recycling sector.

Along with logistics issues, many environmental issues come with not throwing away the proper type of waste in the respective trash bin (Vogel, 2019). For example, if someone throws away a plastic water bottle or a paper plate in a landfill bin, this will reduce the number of items that can be recycled that were initially meant to be recycled (EPA) (Rowe, 2022). This makes landfills fill up faster (Dan, 2022), reducing the number of items that can be recycled and causing more pollution and adverse environmental effects (Brenner, 2019). This includes climate change, wildlife loss, and other harmful issues.

Over the past decade, the field of machine learning ("ML") and artificial intelligence ("AI") has been growing. Existing studies suggest that multiple breakthroughs in AI have made this technology more accessible to others (Kothari, 2023; Vishwanath, 2021). For example, open-source Python projects facilitate the rapid development of AI and ML models, such as Tensorflow(Multiple et al.) and Pytorch(Multiple Authors, 2019). Data availability has also increased rapidly these past few years, with services like Kaggle offering access to thousands of different datasets at a click to train AI and ML programs.

Previously, another project was carried out to generate an ML algorithm to classify trash (Alsubaei, 2022). This project uses a different method for identifying and classifying garbage (DLSODC-GWM technique). This technique revolves around creating a new model from scratch, which is not the main focus of this paper. Our approach revolves around a real-time implementation of a trash-detection model instead of a custom-built model that can only work with images being uploaded to it. Both approaches are suitable for different tasks, as the DSLODC-GWM technique is suited more for labeling and classifying static images. In contrast, the approach used in this paper (YOLOv5) is suited for real-time object detection and classification.

This paper aims to provide a solution to the issue of pollution by creating a deep-learning model that identifies a piece of trash in *real-time* to help reduce environmental harm and damage. Here, we used the YOLOv5, an object-detection model trained with a custom dataset from Kaggle(cchangs, 2018). The name of this dataset is “Garbage Classification” and is authored by the user “cchangcs.” As this project evolves with more research and development, accuracy will likely exceed the 90% yield over the course of this project. This highly accurate trash classification and integration of this methodology into our trash management system will have a positive environmental impact by reducing the amount of recyclables in landfills.

Problem Statement

This project was done for two main reasons: The first reason is to be able to detect trash effectively, and the second reason is to contribute to the field of AI and the open-source community. “Open source” refers to projects, code, and datasets anyone can use.

Using AI will make this daunting task much easier, as the backbone of AI rests on learning from data rather than brute force every possibility. To complete this project, the YOLOv5 model was selected. The YOLOv5 model, authored by Glenn Jocher, is an object detection and classification model based on the Convolutional Neural Network (“CNN”) architecture (Ultralytics, n.d.). Many corporations, such as the EV car manufacturer Rivian, also use the YOLOv5 model for their products. A dataset on Kaggle, an open-source data website with different types of trash was selected for use in this project. To format the data to be trained using YOLOv5, border boxes have to be added around all the objects in the images. The Roboflow image annotation tool was utilized to label the 2,200 images manually. The dataset had six classes (cardboard, glass, trash, plastic, paper, and metal), and each image was labeled with a respective class. The dataset is split into 70% of the images being used for training, 20% for testing, and 10% for validation of the model during the training process.

Figure 1. Sample images of various trash types. These represent all classes of trash available in the training data.

HIGH SCHOOL EDITION
Journal of Student Research
Volume 13 Issue 2 (2024)

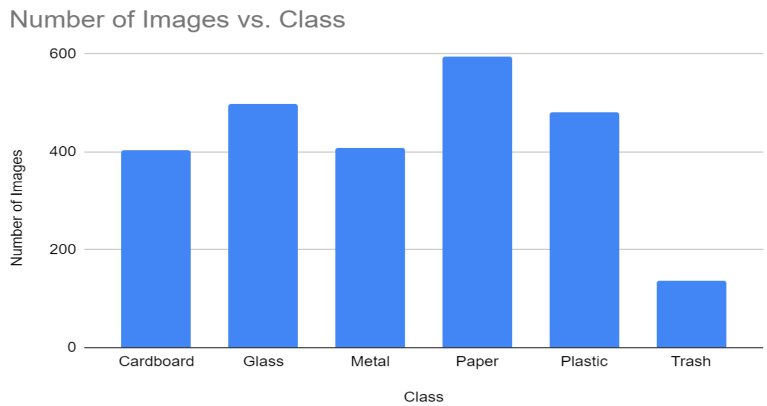

Figure 2. Number of images per class in the data. The number of images per class in the dataset (including training set, validation set, and testing set) Cardboard: 403, Glass: 499, Metal: 409, Paper: 594, Plastic: 482, Trash: 137.

Methods

With the data labeling phase of the project complete, a suitable ML model is required to complete the task at hand. As mentioned, YOLOv5 was selected because of its flexibility with different use cases and a relatively lower training time without significantly impacting the model accuracy. The model was trained using the YOLOv5s weights for the fastest training on a local device with average specifications. The trained model was then used with a user interface (UI) to interact with the model. The model then detects objects using the device's webcam and outputs a boundary box with the detected class in a built-in live webcam streamer provided with the YOLOv5 model.

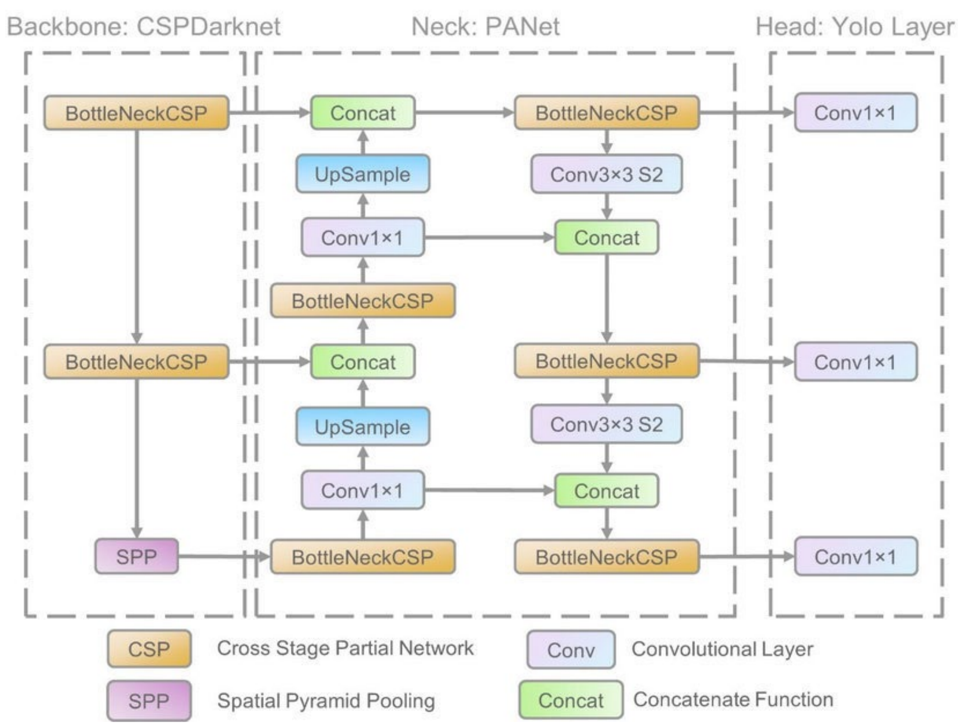

HIGH SCHOOL EDITION
Journal of Student Research
Volume 13 Issue 2 (2024)

Figure 3. YOLOv5 architecture flowchart (Ding, 2021)

To deploy this ML model, a Python-based web server was used to host a UI that would let anyone interact with the trained model. There would be buttons to open up a live webcam service (provided by Ultralytics, the maintainers of YOLOv5) that runs real-time object detection. If a webcam is unavailable, another button allows a user to upload an image, and the model will run inference on it and provide a new image with the bounding box labeled on the image with the class. To use this in commercial services, a webcam can be hooked up to any form of mini-computer with a monitor, and the monitor can display this live webcam with the output of the model being processed. This processed output from the model can be outputted through a speaker to tell the user which trash bin to put the piece of garbage in.

Results

The following charts provide us with a deep insight into how the model performed on the testing set of data. A testing data set is a subset of randomly picked images from the original dataset that tests the accuracy of the model after the training process. The charts also provide insight into the training process and how the model "learned" from the data. The confusion matrix provides us insight into the percentage of images from each class that were predicted accurately. The F_1 curve gives us insight into the proper confidence threshold to get the best performance out of our model. The final image gives us an insight into how the model learned during training.

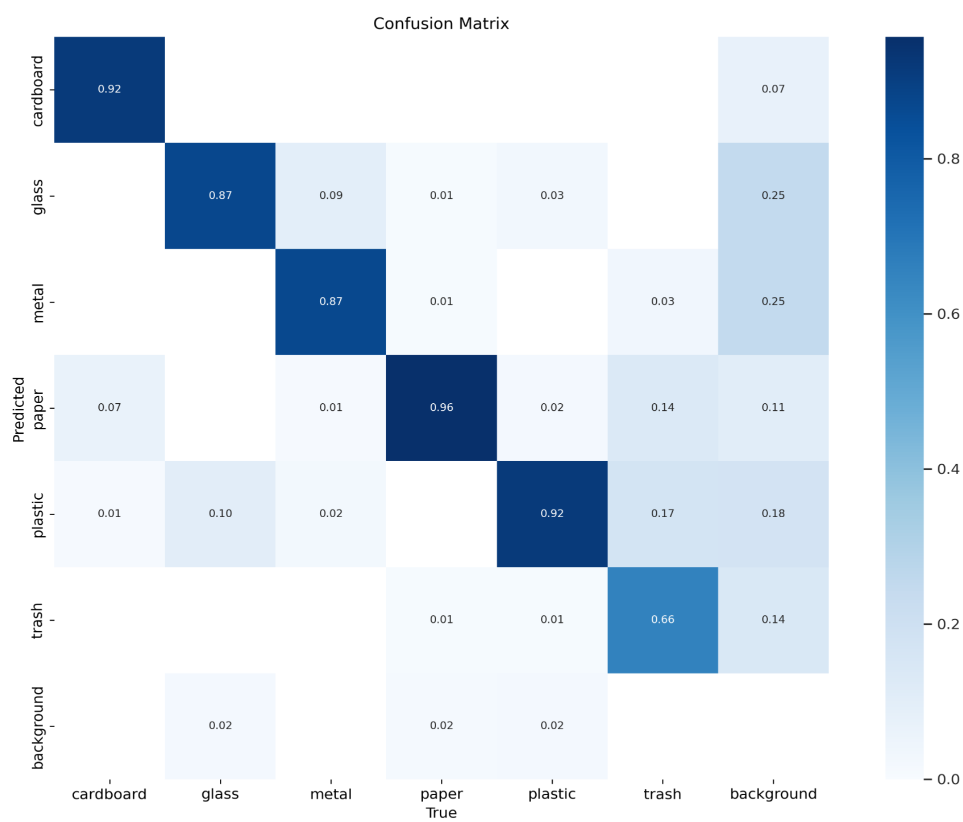

Figure 4. Confusion Matrix. This figure shows the percentage of images in each class that were predicted accurately. Each axis represents the predicted class and the actual class of the image. The top-left bottom-right diagonal represents all the accurately labeled images, as the predicted and actual class match. The darker the

HIGH SCHOOL EDITION
Journal of Student Research

blue squares, the higher the percentages of images that were accurately labeled, as shown in the bar on the right of the image (92% of the cardboard images, 87% of the glass and metal images, 96% of the paper, 92% of the plastic, and 66% of the trash images were accurately labeled).

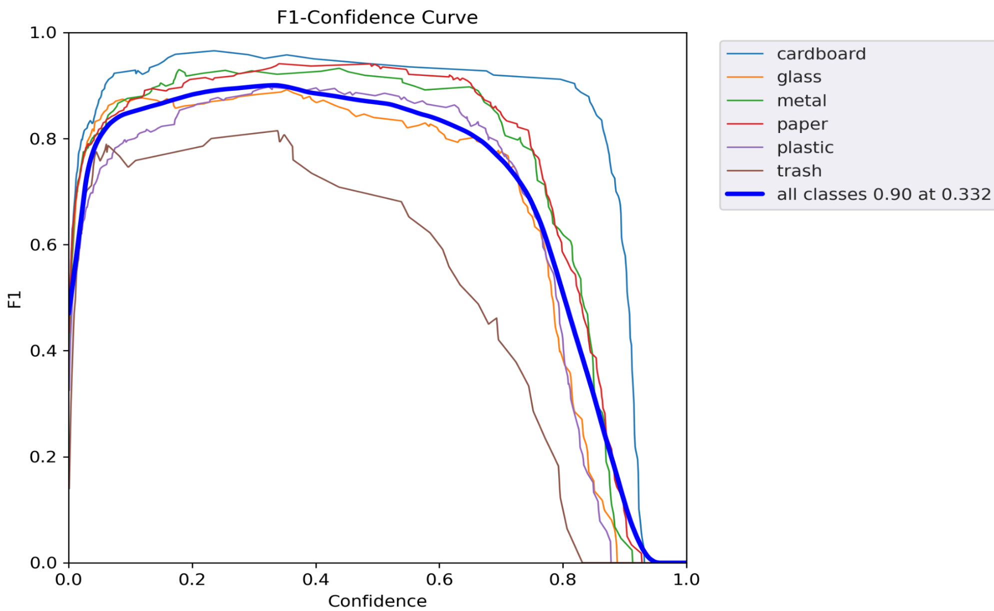

Figure 5. F_1 Confidence curve: The figure above plots the model's accuracy given multiple confidence thresholds. The legend shows the different accuracies for each class in different colors for different confidence thresholds. The overall model confidence is plotted as the thick, dark-blue line. The legend also tells us the optimal confidence threshold for our model, which is 0.332, where the model has a 90.0% accuracy.

The F_1 statistic explains how accurate the model is given many different confidence values (see Appendix for details about the F_1 statistic). Figure 5 provides insight into our optimal value for the confidence threshold so that our model performs at its best. The confidence threshold represents how confident the model is that it detects a class in an image. The bounding box will be outlined in the processed image if the confidence value is above a specific value.

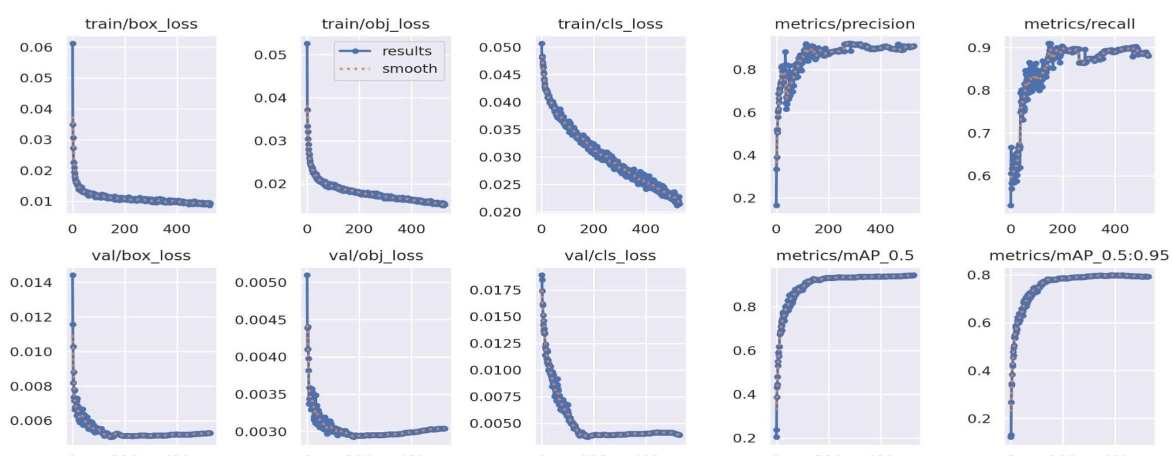

Figure 6. Overview of all model statistics. This image shows the model's stats. The most important stat for the YOLOv5 model is the metrics/mAP_0.5 graph. mAP[9](Yohanandan, 2020), also known as Mean Average Precision, measures how many predictions the model made were correct. The mAP score considers each run's precision, recall, and IoU scores during training. mAP0.5 refers to the mAP score (precision and recall) with an IoU threshold of 0.5. The charts show that the mAP0.5 score is 0.95, proving that the model performs well on the validation dataset.

Discussion

This experiment aimed to build a real-time platform that helps people near trash bins identify which bin a piece of trash is supposed to go into. This project was successful after all the research and testing, as it can predict what type of trash a given piece of junk is. Respectively, it can predict what bin the garbage should go into, given what kind of trash it is, with an accuracy of 90.0% on the validation and testing datasets.

Figure 7. The model runs on some images of recyclable and non-recyclable material in real time (screenshots provided).

For further research on this paper, I could create a mobile app allowing photo-based or real-time trash detection and classification. This would make this project more accessible to others. Further research includes using more images and backgrounds to prevent the model from overfitting to one type of image and background. Improving the website UI and making the model more accessible so that it can be used in many different places without having a huge energy footprint is also some other research that can be done.

Finally, this is the second paper published about identifying and classifying trash. This paper will inspire others to contribute to this niche object detection field.

Conclusion

In conclusion, this project was an overall success, as the YOLOv5 model allowed me to develop a solution that identifies the type of waste and helps people dispose of it properly in real time. The model's accuracy of 90% provided me with a promising start to further research and development in the future. Furthermore, this project raised awareness about proper waste management and protecting the environment. The user-friendly interface of the solution will make it possible to make the application accessible by many others, including those who are not tech-savvy. Along with the environmental impacts, this project highlighted the potential of AI and ML and its capability of solving real-world problems.

Limitations

One significant limitation of this project is that if the background of the live stream is not white in the real-time platform, the model will assume that the entire material or object is paper.

Figure 8. A piece of trash is classified as paper and trash simultaneously. The piece of trash shown here is an empty wrapper for some Indian biscuits, which the model has never seen before.

This issue can be traced back to the variety in the dataset. Since all the images in the dataset had a white background (photographed with a plain, white background), the live stream should have a white background to counteract this issue. The model will not work as intended if there is no white background.

Acknowledgments

I would like to acknowledge my mentor, Mr. Hieu Nguyen, a Ph.D. student from the University of Connecticut, for providing me with extensive knowledge and support through the process of writing this research paper. I was able to write a paper involving very advanced topics with his support and guidance. Without my mentor, I would not have written this research paper, nor would I have done the project as well. Thank you very much for your guidance and support.

References

Alsubaei, F. S., et al. (2022). Deep Learning-Based Small Object Detection and Classification Model for Garbage Waste Management in Smart Cities and IOT Environment. MDPI. https://www.mdpi.com/2076-3417/12/5/2281

Baechler, N. (2018, February 12). *'WISHFUL RECYCLING': MORE HARM THAN GOOD*. School of Marine and Environmental Affairs. https://smea.uw.edu/currents/wishful-recycling-more-harm-than-good/

Brenner, L. (2019, March 2). The Effects of Not Recycling. Sciencing. https://sciencing.com/effects-not-recycling-4827369.html

Cchangcs. (2018, November 24). Garbage Classification. Kaggle. https://www.kaggle.com/datasets/asdasdasasdas/garbage-classification

Dan. (2022, January 25). The Compounding Effects of Not Recycling. Napa Recycling and Waste Services. https://naparecycling.com/the-compounding-effects-of-not-recycling/

Environmental Protection Agency. (n.d.). The U.S. Recycling System. EPA. https://www.epa.gov/circulareconomy/us-recycling-system

Kothari, A. (2023, May 31). Top 11 Frameworks in the AI World. Geekflare. https://geekflare.com/ai-frameworks/

Lebiedzinski, P. (2021, June 24). A Single Number Metric for Evaluating Object Detection Models. Medium, Towards Data Science. https://towardsdatascience.com/a-single-number-metric-for-evaluating-object-detection-models-c97f4a98616d

Multiple Authors (n.d.). *TensorFlow: A system for large-scale machine learning*. Google Research. https://research.google/pubs/pub45381/

Multiple Authors (2019, December 3). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. ArXiv. https://arxiv.org/abs/1912.01703

Ocampo, J. (2018, August 29). Americans Are Terrible at Recycling - This Is What Happens When You Put Something in the Wrong Bin. Mic. https://www.mic.com/articles/190974/americans-are-terrible-at-recycling-this-is-what-happens-when-you-put-something-in-the-wrong-bin

ResearchGate. (n.d.). YOLO v5 Architecture Overview. ResearchGate. https://www.researchgate.net/figure/YOLO-v5-Architecture-Overview-3_fig2_351355008

Rowe, K. (2022, March 22). 5 Effects of Not Recycling Properly. CivilDigital. https://civildigital.com/5-effects-of-not-recycling-properly/

S, V. V. (2021, April 22). Top 15 Machine Learning (ML) Frameworks to Know in 2023. Mindmajix Technologies Inc. https://mindmajix.com/machine-learning-frameworks

Ultralytics (n.d.). *Ultralytics YOLOv5 Architecture*. Ultralytics YOLO. https://docs.ultralytics.com/yolov5/tutorials/architecture_description/

Veeravadivel, V., & Nguyen, H. (2023). TrashNet: A Set of Annotated Images of Trash That Can Be Used for Object Detection Object Detection Dataset and Pre-Trained Model by Polygence Project. Roboflow. https://universe.roboflow.com/polygence-project/trashnet-a-set-of-annotated-images-of-trash-that-can-be-used-for-object-detection

Vogel Disposal Services. (2019, July 15). What Happens When The Wrong Things Go Into Recycling? https://www.vogeldisposal.com/vogel-disposal-news/what-happens-when-the-wrong-things-go-into-recycling-4022

Yohanandan, S. (2020, June 9). Map (Mean Average Precision) Might Confuse You! Medium, Towards Data Science. https://towardsdatascience.com/map-mean-average-precision-might-confuse-you-5956f1bfa9e2