Industrial and Domestic Waste Management

Research Article
Volume 5(1), 2025, 1–11
https://doi.org/10.53623/idwm.v5i1.580

Plastic Waste Detection Using Deep Learning: Insights from the WaDaBa Dataset

Suman Kunwar[1]*, Banji Raphael Owabumoye[2], Abayomi Simeon Alade[3]

[1]Faculty of Computer Science, Selinus University of Sciences and Literature, Ragusa, Italy
[2]Department of Geography, Obafemi Awolowo University, Ile-Ife, Nigeria
[3]Department of Physics, University of Ibadan, Oyo, Nigeria

*Correspondence: sumn2u@gmail.com

SUBMITTED: 6 January 2025; REVISED: 19 February 2025; ACCEPTED: 24 February 2025

ABSTRACT: With the increasing use of plastic, the challenges associated with managing plastic waste have become more difficult, emphasizing the need for effective classification and recycling solutions. This study explored the potential of deep learning, focusing on convolutional neural networks (CNNs) and object detection models like YOLO to tackle this issue using the WaDaBa dataset. The results indicated that YOLO-11m achieved the highest accuracy (98.03%) and mAP50 (0.990), while YOLO-11n performed similarly but achieved the highest mAP50 (0.992). Lightweight models like YOLO-10n trained faster but had lower accuracy, whereas MobileNetV2 demonstrated impressive performance (97.12% accuracy) but fell short in object detection. YOLO-11n had the fastest inference time (0.2720s), making it ideal for real-time detection, while YOLO-10m was the slowest (5.9416s). Among CNNs, ResNet50 had the best inference time (1.3260s), whereas MobileNetV2 was the slowest (1.4991s). These findings suggested that by balancing accuracy and computational efficiency, these models could contribute to scalable waste management solutions. The study recommended increasing the dataset size for better generalization, enhancing augmentation techniques, and developing real-time solutions.

KEYWORDS: Waste classification; deep learning; YOLO models; WaDaBa dataset; sustainable waste management

1. Introduction

Plastics have become indispensable in modern society due to their versatility, durability, and low cost. As a result, global plastic production has skyrocketed to over 450 million tons annually [1]. However, the very properties that make plastics so useful—durability and resilience—also contribute to their persistence in the environment, where they can take hundreds of years to degrade. The United Nations Environment Programme estimated that more than 8 million tons of plastic waste enter the oceans annually [2], leading to catastrophic effects on marine life and ecosystems. The most effective strategies to address the global issue of plastic pollution remain uncertain. Borrelle et al. and Lau et al. explored potential solutions and their implications, concluding that substantial reductions in plastic waste generation are achievable over the next few decades if immediate and robust action is taken [3, 4]. However, even under the best circumstances, substantial amounts of plastic are still expected to accumulate in the environment.

Plastic waste management has been a focus of recent research, especially given the limitations of traditional sorting techniques. Studies showed that less than 9% of plastic produced globally is recycled, underscoring inefficiencies in existing waste management frameworks [5, 6]. Traditional waste management systems, reliant on manual sorting, have not been efficient enough to handle the growing volume of plastic waste. To address this issue, the Society of the Plastics Industry introduced Resin Identification Codes (RICs) in 1988. These codes categorized plastics based on polymer content, facilitating recycling and waste management. Common RICs include PET, PE-HD, PVC, PE-LD, PP, and PS, which are often labeled on plastic products to aid in sorting. Artificial intelligence (AI) and machine learning (ML) have emerged as transformative technologies in automating waste management and have introduced various techniques to effectively manage waste [7, 8, 9]. The use of deep learning models such as YOLO (You Only Look Once) has shown significant promise in improving waste classification accuracy, as its ability to detect objects in real-time with high accuracy makes it ideal for waste sorting applications [10, 11].

The study by Redmon and Farhadi showed that the architecture of YOLO allows for simultaneous object detection and classification, making it faster than traditional CNN-based approaches [12]. One study found that the larger YOLO-v5 model outperformed the smaller nano- and medium-sized models in detecting plastic waste along railway lines [13]. The latest version, YOLO-11n, incorporates advanced modules to enhance detection capabilities, particularly for small objects, making it highly effective for waste classification. Furthermore, transfer learning models like MobileNetV2 and ResNet have demonstrated efficiency in recognizing complex waste categories, making them suitable for deployment in resource-constrained environments [14, 15]. Studies have shown that models pre-trained on large datasets like ImageNet can significantly improve performance in specialized tasks such as waste detection [16]. These advancements signal a shift towards AI-driven solutions that hold the potential to make plastic waste classification and recycling much more effective. The aim of this study was to evaluate the effectiveness of YOLO models in identifying plastic waste using the WaDaBa dataset and compare them with MobileNetV2, ResNet-50, and EfficientNet, along with a custom model in various settings. We also implemented various data augmentation techniques to improve dataset diversity and model generalization, as these can enhance accuracy. Finally, the most effective model was embedded in a mobile application.

2. Materials and Methods

This study employed a comprehensive methodology encompassing data preparation, model development and training, and performance evaluation. The models included YOLO variants, a custom deep learning model, and pre-trained transfer learning models—MobileNetV2, ResNet50, and EfficientNet. Each model underwent specific preprocessing and data augmentation techniques to enhance classification accuracy.

2.1. Dataset and preprocessing.

The WaDaBa dataset consists of 4,000 high-resolution RGB images (1920×1277 pixels, 300 dpi) representing five distinct RIC categories: PET, PE-HD, PP, PS, and Others. It has become a benchmark for evaluating deep learning models, as traditional methods struggle to generalize across different plastic types. The images were captured under varying lighting conditions and angles, with key attributes—such as plastic type, color, and deformation level—embedded in the filenames. However, the dataset exhibits a significant class imbalance, with PET comprising 55% of the total images, while the "Others" category accounts for only 1%, as shown in Figure 1.

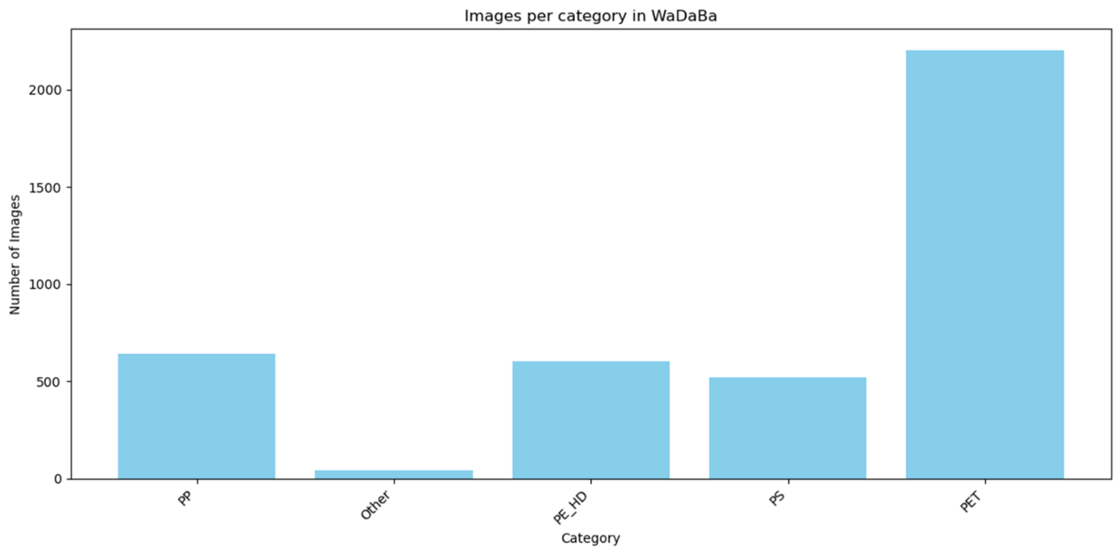

Figure 1. Distribution of the dataset.

To address this imbalance, data oversampling was applied to expand the dataset to 11,000 images. Additionally, data augmentation techniques—including random zoom, rotation, contrast adjustment, and flipping—were implemented to enhance model robustness and generalization across different conditions. The images were annotated using Annotate-Lab, converted into YOLO format, and split into training and testing sets at an 80:20 ratio. Figure 2 illustrates the annotation of an image using Annotate-Lab [17].

Figure 2. Annotation of WaBaDa image using annotate-lab.

2.2. Model architecture and training.

2.2.1. YOLO model training.

The YOLO models—YOLO-10n, YOLO-10m, YOLO-11n, and YOLO-11m—were selected for this study due to their exceptional real-time detection capabilities, making them highly suitable for efficient waste classification. Each model was configured to achieve an optimal balance between precision, recall, and mean Average Precision (mAP) for accurate plastic waste identification. Training was conducted over 20 epochs, utilizing non-max suppression techniques to eliminate overlapping bounding boxes and enhance detection confidence [18]. This approach improved precision while significantly reducing false positives, increasing model reliability in real-world applications. The workflow began with careful dataset annotation and preparation to ensure compatibility with the YOLO framework. The annotated data were then processed and fed into each YOLO model for training. Hyperparameters were fine-tuned to optimize feature extraction, enabling accurate classification of various plastic waste types under diverse conditions. Following model evaluation, the best-performing YOLO configuration was selected based on mAP, accuracy, precision, recall, F1-score, and inference speed. This model was further refined through quantization and optimization to enhance speed and efficiency, making it suitable for real-time waste management deployment. The process flow, illustrated in Figure 3, highlights the seamless integration of data preparation, model training, and deployment, emphasizing the robustness of YOLO models in automated plastic waste detection. Figure 4 illustrates the training and validation loss during the model training of YOLO-10n and YOLO-11n. The overall training and validation box and class loss for YOLO-11n exhibit a smoother trend compared to YOLO-10n.

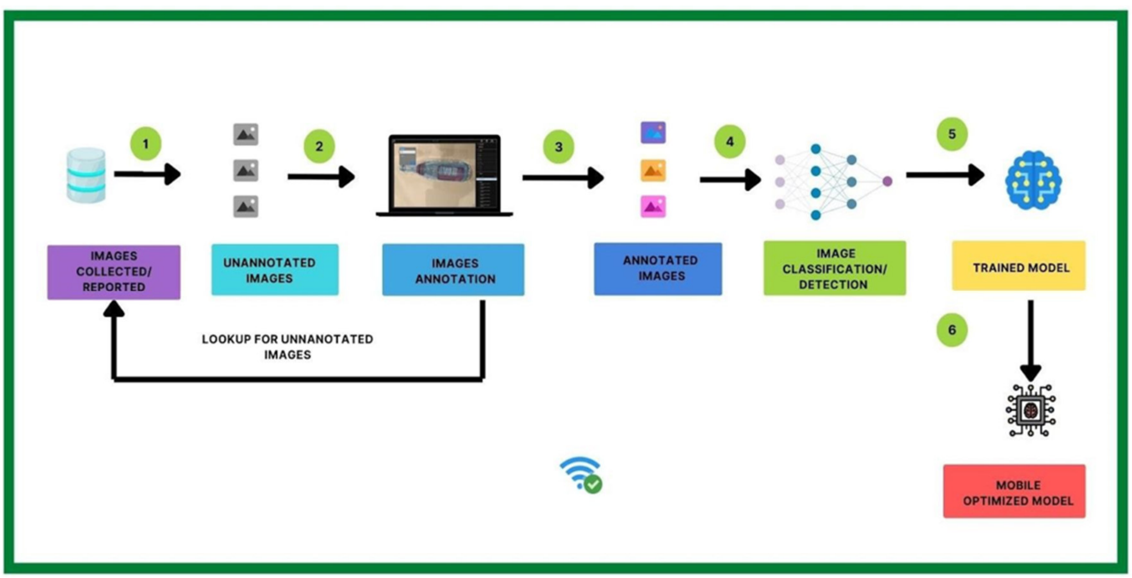

Figure 3. YOLO process flow diagram.

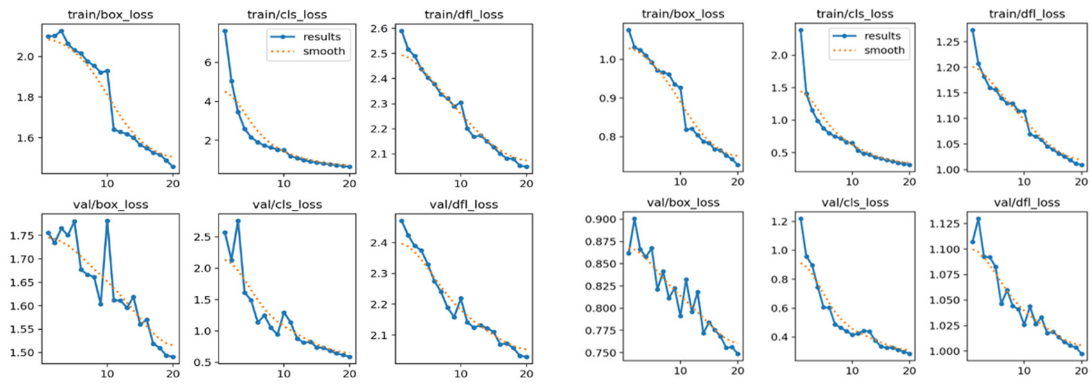

(**a**) (**b**)

Figure 4. Training of YOLO-10n (a); Training of YOLO-11n (b).

2.2.2. Custom model design and training.

The custom model developed in this study was specifically designed to address the dataset's class imbalance while maximizing feature extraction and classification accuracy. Its architecture consisted of three convolutional layers with filter sizes of 16, 32, and 64, each paired with max-pooling and ReLU activation functions to capture intricate data patterns effectively. To prevent overfitting, two dense layers with 256 and 64 neurons were incorporated, along with a 50% dropout rate to enhance generalization. The model was trained over 40 epochs using an 80-20 training-testing split. To balance computational efficiency and accuracy, all images were resized to 180 × 180 pixels, and a batch size of 300 was used. Training was accelerated by utilizing two GPUs. The final output layer employed SoftMax activation to facilitate multi-class classification. Figure 5 illustrates the model's architecture and configuration, highlighting the seamless integration of feature extraction, dimensionality reduction, and classification components. This optimized design ensures adaptability to diverse dataset characteristics while maintaining high classification performance.

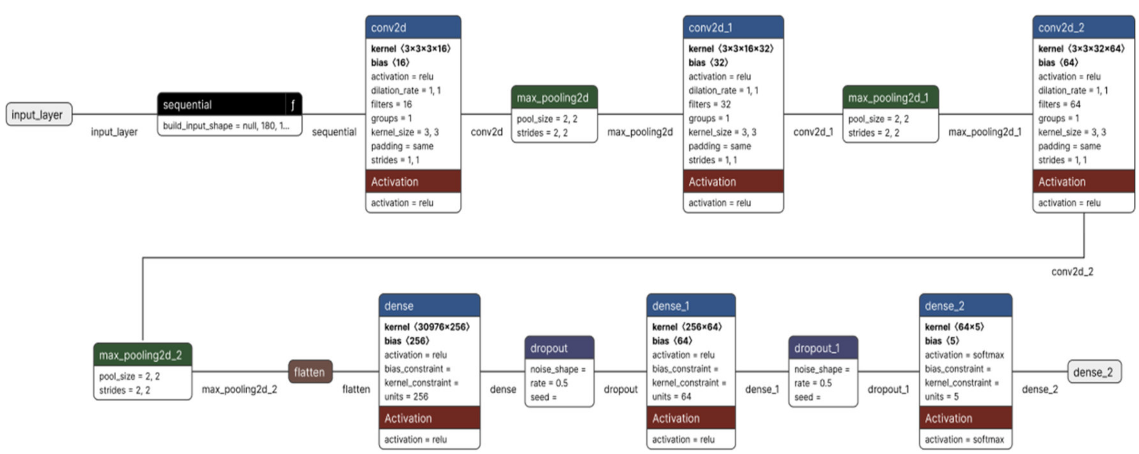

Figure 5. Custom model architecture.

Figure 6 illustrates loss and validation loss and accuracy and validation accuracy during training over epochs and shows MobileNetV2 and Custom model with high accuracy with less loss compared to others. The transfer learning models evaluated include MobileNetV2, ResNet50, and EfficientNet. These models were fine-tuned on the WaDaBa dataset using a batch size of 100 and learning rate of 1e-3. Prior research has demonstrated that transfer learning significantly reduces training time while improving model accuracy.

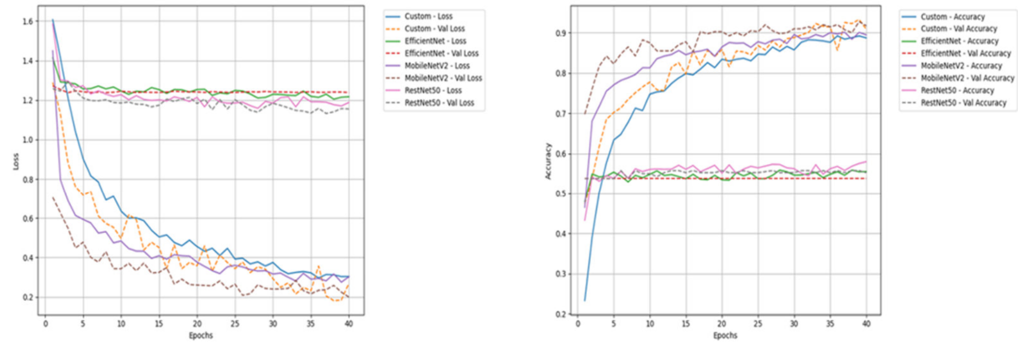

(a) (b)

Figure 6. Loss and validation loss vs. epochs for different models (**a**); Accuracy and validation accuracy vs. epochs for different models (**b**).

3. Results and Discussion

The models were evaluated using multiple performance metrics, including accuracy, precision, recall, F1-score, mAP, and inference time, with each model exhibiting unique strengths across different criteria.

3.1. Result from YOLO models.

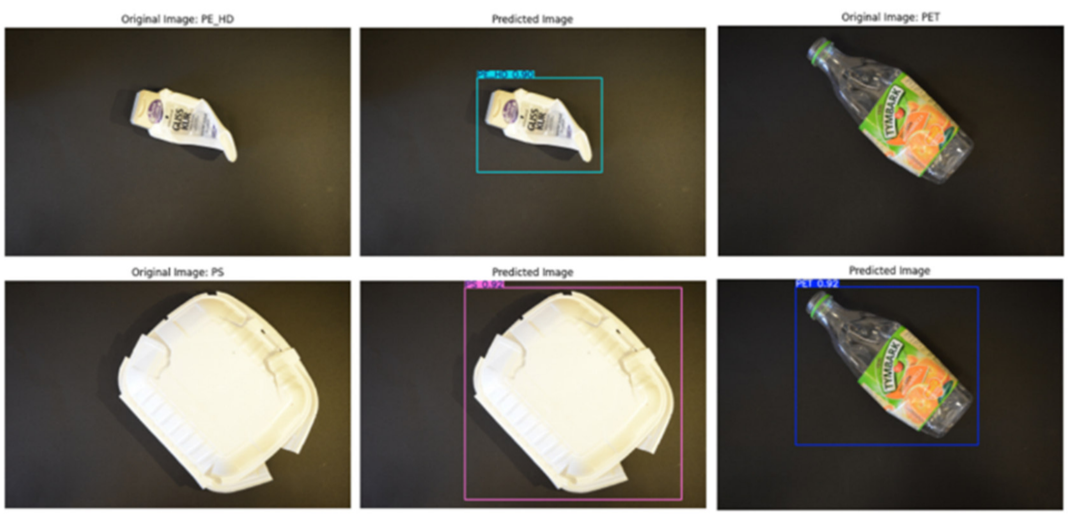

Figure 7. YOLO models' predictions on unseen data.

The YOLO-10n model exhibited well-rounded performance, achieving a precision of 0.9161, recall of 0.9460, F1-score of 0.9842, and mAP50 of 0.984, making it highly suitable for general applications requiring balanced performance across all metrics. Conversely, the YOLO-10m model, optimized for recall, achieved a recall of 0.9374, an F1-score of 0.9563, and an mAP50 of 0.956, making it ideal for tasks prioritizing true positive detection, such as critical identification scenarios. YOLO-11n set a new benchmark for accuracy, attaining the highest accuracy of 0.9553, along with a precision of 0.9803 and an mAP50 of 0.992, making it particularly suitable for accuracy-intensive applications demanding precise classifications.

Meanwhile, YOLO-11m, designed to minimize false positives, achieved a precision of 0.9814 and an mAP50-95 of 0.815, making it the optimal choice for applications where reducing false positives is crucial, such as high-stakes detection environments. Overall, the YOLO-11 models demonstrated superior performance, striking an excellent balance between accuracy and computational efficiency, making them ideal for practical deployment across various applications. The model predictions and their corresponding results are presented in Figure 7.

3.2. Result from custom and transfer learning models.

The custom model demonstrated impressive performance, achieving an accuracy of 93.05%, with precision, recall, and F1-score of 92.96%, 93.06%, and 92.98%, respectively. Trained over 40 epochs, the model completed its training in a swift 8 minutes and 1 second. Remarkably, it successfully classified all samples in an unseen test set of 8 images with complete accuracy. These results highlight the model’s robustness and consistency, underscoring its effectiveness in accurately distinguishing between various plastic waste categories. Such high performance across multiple evaluation metrics indicates that the custom model is well-suited for practical waste classification tasks. The outcomes of its predictions on unseen data are depicted in Figure 8, further validating its generalization capability.

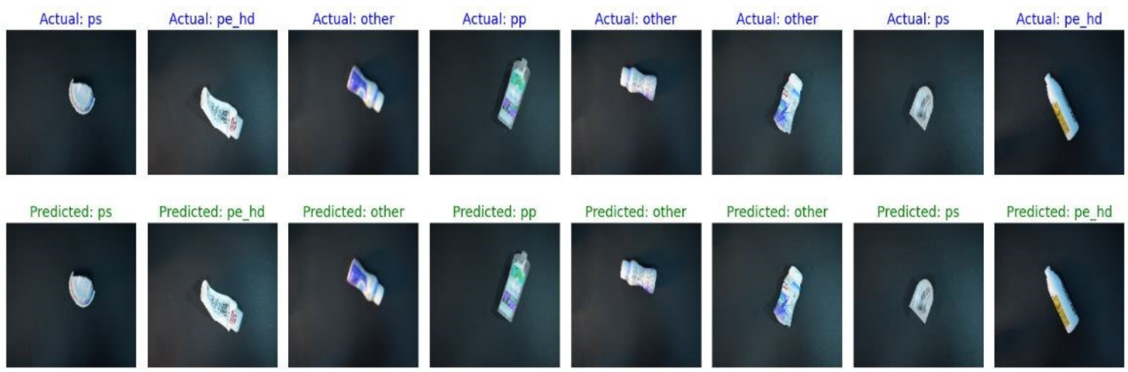

Figure 8. Custom model predictions on unseen data.

MobileNetV2 emerged as the top performer among the transfer learning models, achieving an impressive accuracy of 97.12%. It also demonstrated exceptional precision (96.31%), recall (92.69%), and F1-score (94.26%) after a 40-epoch training period. Notably, MobileNetV2 achieved these results with a remarkably short training time of just 5 minutes and 34 seconds, highlighting both its efficiency and predictive strength. ResNet50 exhibited lower performance in this experiment, achieving an accuracy of 65.20%, precision of 65.43%, recall of 45.89%, and an F1-score of 20.01%. Despite a training time of 8 minutes and 31 seconds over 40 epochs, the model correctly classified only 5 out of the 8 samples in the test set. EfficientNet demonstrated limited performance, achieving an accuracy of 53.25%, precision of 10.65%, recall of 20.00%, and an F1-score of 13.90%. With a training duration of 6 minutes and 35 seconds, the model successfully predicted only 3 out of the 8 test images. This relatively poor performance on ResNet50 and EfficientNet both can be attributed to the challenges posed by the dataset's class imbalance, which hindered the model's ability to generalize effectively, particularly for minority classes.

3.3. Comparative analysis.

The performance of YOLO models along with custom and CNN models is benchmarked using accuracy, precision, recall, F1-score, and the training time. Studies have shown that while mAP50 scores are typically high, mAP50-95 often drops significantly, indicating challenges in fine-grained localization. This decline suggests that YOLO models detect objects well but struggle with precise bounding box placement at higher IoU thresholds, particularly in occlusion-heavy or small object detection scenarios [19]. Furthermore, precision-recall (PR) curve analysis can help assess whether YOLO's higher recall comes at the expense of false positives, which can mislead real-world applications [20]. Comparative studies with Faster R-CNN and EfficientNet have highlighted that YOLO's tradeoff between speed and accuracy often results in more localization errors at stricter IoU thresholds [21]. To strengthen these results, a breakdown of mAP50, mAP75, and mAP50-95, along with detailed PR curves, is necessary to capture performance trends across varying detection strictness [22]. Table 1 summarizes the performance metrics for each model. Results indicate that YOLO models, especially YOLO-11n and YOLO-11m, and MobileNetV2 are the most effective for plastic waste classification.

Table 1. Comparison of performance metrics of various models.

Model	Accuracy	Precision	Recall	F1-score	mAP50	mAP50-95	Inference Speed	Epoch	Training Time
YOLO-10n	0.9161	0.9460	0.9842	0.8706	0.984	0.807	0.2903	20	27m:5s
YOLO-10m	0.8101	0.9374	0.9563	0.7685	0.956	0.78	5.9416	20	43m:12s
YOLO-11n	0.9803	**0.9740**	**0.9921**	**0.9553**	**0.992**	0.813	**0.2720**	20	25m:34s
YOLO-11m	**0.9814**	0.9657	0.9908	0.9483	0.991	**0.815**	1.2238	20	42m:40s
Custom	0.9286	0.9280	0.9285	0.9276	-	-	1.3841	40	8m:01s
MobileNetV2	0.9712	0.9631	0.9269	0.9426	-	-	1.4991	40	**5m:32s**
RestNet50	0.6520	0.6543	0.4589	0.2001	-	-	1.3260	40	8m:31s
EfficientNet	0.5325	0.1065	0.2000	0.1390	-	-	1.4448	40	6m:35s

The inference time was measured to identify the best model for real-world deployment, revealing that YOLO-11n achieved the fastest performance at 0.2720s per inference, followed by YOLO-10n (0.2903s). In contrast, YOLO-10m had the slowest inference time (5.9416s). Among CNN models, ResNet50 performed best (1.3260s), while MobileNetV2 was the slowest (1.4991s). This benchmark provides insights into the performance of these models, highlighting critical factors influencing their suitability for this domain. The custom model, designed specifically to handle the imbalanced nature of the WaDaBa dataset, demonstrated strong classification capabilities. By employing targeted data augmentation strategies—such as random rotation, zoom, contrast adjustments, and flipping—we significantly increased the diversity of training samples. This enhancement improved the robustness of the custom model, leading to greater accuracy and generalization. Additionally, oversampling techniques were applied to balance class distribution, mitigating biases during training and further optimizing model performance. These findings align with recent studies emphasizing the importance of data augmentation and class balancing for improving model accuracy in datasets with skewed distributions. From our experiment, we found that MobileNetV2 emerged as the most effective model, achieving a classification accuracy of 97%. Its lightweight architecture, designed for efficient deployment in mobile and edge computing environments, provided a significant advantage in balancing speed and accuracy. However, it fell behind ResNet50 in inference speed. The use of pre-trained ImageNet weights enabled MobileNetV2 to leverage rich feature representations, allowing it to generalize better to unseen data compared to the custom model. These findings align with previous research on the efficiency of MobileNetV2 in resource-constrained settings. While MobileNetV2 excels in efficiency, it is less robust in complex environments. YOLO models, particularly YOLO-11, outperform MobileNetV2 in occlusion-heavy, low-light, and small-object detection scenarios due to their superior spatial awareness and multi-scale feature extraction [23].

The custom model also showed competitive performance, demonstrating the potential of tailored architectures when designed with domain-specific data augmentation. In contrast, deeper architectures like ResNet50 and EfficientNet did not perform as well, with accuracies of 65% and 53%, respectively. EfficientNet struggled with feature extraction limitations, as its compound scaling assumes that increasing resolution, depth, and width will always improve accuracy. This assumption fails when images have low quality or domain-specific characteristics that do not benefit from higher resolution [24]. ResNet50 suffered from overfitting, as its deep residual blocks tended to memorize patterns in small or imbalanced datasets, leading to poor generalization in real-world noisy data [25]. This observation supports previous research suggesting that deeper models are not always the optimal choice when datasets are limited in size or require domain-specific feature extraction. The suboptimal performance of ResNet50 and EfficientNet underscores the importance of selecting models appropriately scaled to the dataset's requirements, particularly when computational resources are limited. Additionally, our findings align with previous work on the application of deep learning for waste classification. For example, Ren et al. demonstrated the real-time object detection capabilities of YOLO models, which is consistent with the superior performance we observed in the YOLO-11 series during earlier phases of this study. However, our results indicate that in scenarios where computational resources are restricted, leveraging lightweight models like MobileNetV2 or optimized custom architectures may provide a more effective solution than relying on deeper models such as ResNet50 or EfficientNet.

4. Conclusions

This study explored deep learning approaches for detecting plastic waste, focusing on CNNs and YOLO models. The YOLO models (YOLO-11m and YOLO-11n) achieved high accuracy, exceeding 98%, but required longer training times. Meanwhile, CNN models—including the custom model, MobileNetV2, ResNet50, and EfficientNet—trained more quickly but exhibited lower accuracy, particularly ResNet50 and EfficientNet. Among all models, YOLO-11n had the fastest inference speed, making it ideal for real-world deployment, while ResNet50 performed best among non-YOLO models. The study also highlights the potential of combining custom models with pre-trained, lightweight architectures like MobileNetV2 to improve the accuracy of plastic waste classification. The integration of data augmentation and class balancing techniques significantly enhanced performance, particularly for the custom architecture. Future work should focus on further refining and generalizing the WaDaBa dataset, experimenting with more advanced augmentation techniques, and developing real-time solutions for automated waste management systems.

Acknowledgments

The authors confirm that the data used in this experiment are obtained with permission and can be available on request. The experiments can be downloaded from https://github.com/sumn2u/wadaba-analysis.

Author Contribution

The authors confirm their contribution to the paper as follows: study conception and design: Suman Kunwar; experiment setup and analysis: Suman Kunwar and Banji Raphael Owabumoye; drafting introduction and related work section: Suman Kunwar and Abayomi Simeon Alade; draft manuscript preparation: Suman Kunwar; manuscript verification and editing: Suman Kunwar. All authors reviewed the findings and approved the final version of the manuscript.

Competing Interest

The authors declare no conflict of interest.

References

[1] Ren, Y.; Li, Y.; Gao, X. (2024). An MRS-YOLO Model for High-Precision Waste Detection. *Sensors*, *24,* 4339. https://doi.org/10.3390/s24134339.

[2] Our planet is choking on plastic. (accessed on 1 December 2024) Available online: https://www.unep.org/interactives/beat-plastic-pollution/.

[3] Borrelle, S.; Ringma, J.; Law, K.L.; Monnahan, C.C.; Lebreton, L.; McGivern, A.; Murphy, E.; Jambeck, J.; Leonard, G.H. et al. (2020). Predicted growth in plastic waste exceeds efforts to mitigate plastic pollution. *Science*, *369,* 1515–1518. https://doi.org/10.1126/science.aba3656.

[4] Lau, W.W.Y.; Shiran, Y.; Bailey, R.M.; Cook, E.; Stuchtey, M.R.; Koskella, J.; Velis, C.A.; Godfrey, L.; Boucheret, J. et al. (2020). Evaluating scenarios toward zero plastic pollution. *Science*, *369*(6510), 1455–1461. https://doi.org/10.1126/science.aba9475.

[5] Jambeck, J. R.; Geyer, R.; Wilcox, C.; Siegler, T.R.; Perryman, M.; Andrady, A.; Narayan, R.; Law, K.L. (2015). Plastic Waste Inputs from Land into the Ocean. *Science, 347,* 768–771. https://doi.org/10.1126/science.1260352.

[6] Geyer, R.; Jambeck, J.R.; Law, K.L. (2017). Production, Use, and Fate of All Plastics Ever Made. *Science Advances Research Article*, *3*(5). https://doi.org/10.1126/sciadv.1700782.

[7] Kunwar, S.; Alade, A. (2024). Deep learning in waste management: A brief survey. *International Journal of Complexity Applied Science and Technology*, *1*(1), 10068247. https://doi.org/10.1504/IJCAST.2024.10068247.

[8] Abdu, H.; Mohd Noor, M. H. (2022). A survey on waste detection and classification using deep learning. *IEEE Access*, *10*, 128151–128165. https://doi.org/10.1109/ACCESS.2022.3226682.

[9] Olawade, D.B.; Fapohunda, O.; Wada, O.Z.; Usman, S.O.; Ige, A.O.; Ajisafe, O.; Oladapo, B.I. (2024). Smart waste management: A paradigm shift enabled by artificial intelligence. *Waste Management Bulletin*, *2*(2), 244–263. https://doi.org/10.1016/j.wmb.2024.05.001.

[10] Kumar, A.S.; Mathai, A.; Dinesh, R.; Mohan, H.; Deepak, S. (2023). Waste detection and classification using YOLO algorithm and sensors. 16th International Conference on Sensor Technology (ICST), 2023, 1–6. https://doi.org/10.1109/ICST59744.2023.10460825.

[11] Vimal Kumar, M.G.; Kumar, M.; Rao, K.N.; Rao, P.S.; Tirumala, A.; Patnala, E. (2024). Advanced YOLO-based trash classification and recycling assistant for enhanced waste management and sustainability. Second International Conference on Intelligent Cyber-Physical Systems and Internet of Things (ICoICI), 2024, 1238–1246. https://doi.org/10.1109/ICoICI62503.2024.10696214.

[12] Redmon, J.; Farhadi, A. (2018). YOLOv3: An Incremental Improvement. *arXiv preprint*. https://doi.org/10.48550/arXiv.1804.02767.

[13] Liu, L.; Zhou, B.; Liu, G.; Lian, D.; Zhang, R. (2022). YOLO-based multi-model ensemble for plastic waste detection along railway lines. IGARSS 2022 - IEEE International Geoscience and Remote Sensing Symposium, 2022, 7658–7661. https://doi.org/10.1109/IGARSS46834.2022.9883308.

[14] Shetty, T.S. (2021). Identification of Industrial Plastic Waste Using Deep Learning. MSc Thesis, National College of Ireland, Ireland.

[15] Kunwar, S. (2024). Managing household waste through transfer learning. *Industrial and Domestic Waste Management*, *4*(1), 14–22. https://doi.org/10.53623/idwm.v4i1.408.

[16] Chollet, F. (2017). Deep Learning with Python. Manning Publications: New York, United States.

[17] Kunwar, S. (2024). Annotate-Lab: Simplifying Image Annotation. *Journal of Open Source Software*, *9*(103), 7210. https://doi.org/10.21105/joss.07210.

[18] Ge, L.; Dou, L. (2024). Non-maximum suppression for rotated object detection during merging slices of high-resolution images. *IEEE Access*, *12,* 149999–150007. https://doi.org/10.1109/ACCESS.2024.3470815.

[19] Bochkovskiy, A.; Wang, C.Y.; Liao, H.Y.M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint*. https://doi.org/10.48550/arXiv.2004.10934.

[20] Adedeji, O.; Wang, Z. (2019). Urban Waste Classification Using Deep Learning Models. *IEEE Transactions on Automation Science*. https://doi.org/10.1016/j.promfg.2019.05.086.

[21] Tan, M.; Pang, R.; Le, Q.V. (2020). EfficientDet: Scalable and efficient object detection. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, 10781–10790. https://doi.org/10.1109/CVPR42600.2020.01079.

[22] Jocher, G.; Chaurasia, A.; Qiu, J. (2023). YOLO by Ultralytics. GitHub repository. https://github.com/ultralytics/yolov5.

[23] Han, Z.; Yue, Z.; Liu, L. (2025). 3L-YOLO: A Lightweight Low-Light Object Detection Algorithm. *Applied Sciences*, *15*(1), 90. https://doi.org/10.3390/app15010090.

[24] Garay-Maestre, U.; Gallego, A.; Calvo-Zaragoza, J. (2018). Data augmentation via variational auto-encoders. In Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications; Álvarez L., Mejail M., Gómez L., Jacobo J., Eds.; Springer: Berlin, Germany, pp. 231–238. https://doi.org/10.1007/978-3-030-13469-3_21.

[25] He, K.; Zhang, X.; Ren, S.; Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, 770–778. https://doi.org/10.1109/CVPR.2016.90.