Adaptive Multi-Objective Waste Intelligence System: Hierarchical Transfer Learning for Real-time Recyclable Material Recognition with Edge Computing Deployment

STUDENT'S NAME

Name: MD Sifat Hossain
ID: 221-35-1072

Supervisor Name: Mr. Md. Selim Reza
Position: Assistant Professor
Date: 28/11/2025

Thesis submitted in fulfillment of the requirements for the award of the degree of

Bachelor of Science
Department of Software Engineering
DAFFODIL INTERNATIONAL UNIVERSITY

APPROVAL

This thesis titled on "Adaptive Multi-Objective Waste Intellicnoce System: Hiererchical Transfer Learning for Reshtime Recyclable Meicrial Reccguition with Edge Computing Doploynoent", submitted by **Studonl** Namet **Md Siint Hesssln** (ID:221-35-1072) io the Department of Software Enginerring. Daffoill international Univercity ins bem assepied as satisfactory for the partial fulfilliment of the requirenents for the degree of Bachelor of Science in Solienre Enpionesing and approval as to iis style and contents.

BOARD OF EXAMINERS

Chairman

Dr. A. H. M. Saifttlull Sadi
Professor
Department of Software Engiroering
Faculty of Scierce and Infomation Technology Daffoill
International Uriversity

Internal Examiner 1

Dr. Rubaiyat Islam
Associate Professer
Department of Software Engineering
Faculty of Scierce and Information Technology
Daffoill International University

Internal Examiner 2

Dr. Md. Abdul Kader
Associate Prefessor
Department of Software Engineering
Faculty of Scierce and Infornation Technology
Daffoill International University

Internal Examiner 3

Nozazzaman Faruqui
Assistant Professor
Department of Software Engineering
Faculty of Ssiesce and Infermation Technology
Daffoill Inernational University

External Examiner

Mo. Mertafir Khan
Managing Direster
Tecognize Solutions Limited

DAFFODIL INTERNATIONAL UNIVERSITY

DECLARATION OF THESIS AND COPYRIGHT

Author's Full Name : Md Sifat Hossain

Date of Birth : 02-03-2002

Title : **Adaptive Multi-Objective Waste Intelligence System: Hierarchical Transfer Learning for Real-time Recyclable Material Recognition with Edge Computing Deployment**

Academic Session : 2025

RESTRICTED (Contains confidential information under the Official Secret Act 1997)

☐ OPEN ACCESS : (Contains restricted information as specified by the

: I agree that my thesis to be published as online open access

I acknowledge that Daffodil International University reserves the following

1. The Thesis is the Property of Daffodil International University.
2. The Library of Daffodil International University has the right to make copies of the thesis for the purpose of research only.

The Library of Daffodil International University has the right to make copies of the thesis for academic exchange.

Certified

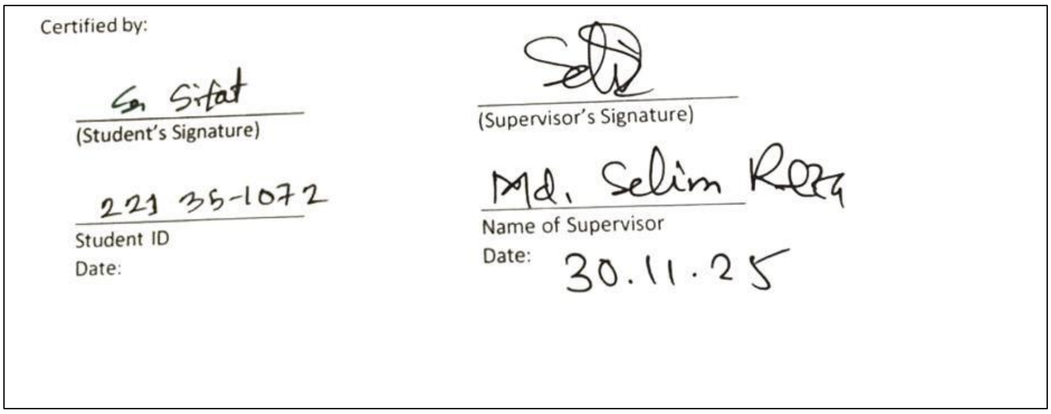

Certified by:

S. Sifat
(Student's Signature)

(Supervisor's Signature)

221 35-1072
Student ID
Date:

Md. Selim Reza
Name of Supervisor
Date: 30.11.25

NOTE: If the thesis is CONFIDENTIAL or RESTRICTED, please attach a thesis declaration letter.

SUPERVISOR’S DECLARATION

I hereby declare that I have checked this thesis and, in my opinion, this thesis is adequate in terms of scope and quality for the award of the degree of Bachelor of Science.

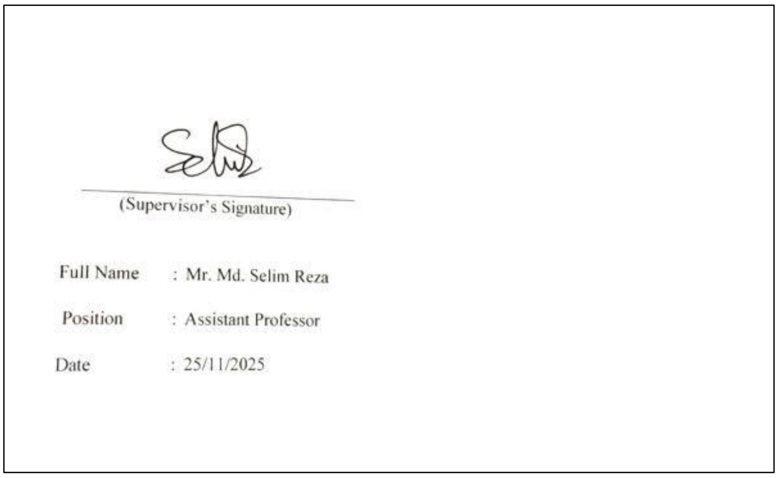

(Supervisor’s Signature)

Full Name : Mr. Md. Selim Reza

Position : Assistant Professor

Date : 25/11/2025

STUDENT'S DECLARATION

I hereby declare that the work in this thesis is based on my original work except for quotations and citations which have been duly acknowledged. I also declare that it has not been previously or concurrently submitted for any other degree at Daffodil International University or any other institution.

Sifat

(Student's Signature)

Full Name : Md Sifat Hossain
ID Number: 221-35-1072
Date : 27 November 2025

ABSTRACT

The current waste crisis in the world, which generates over 2.12 billion tons annually, requires highly advanced technological interventions. This thesis introduces AMWIS, a novel hierarchical transfer learning framework for real-time recyclable material recognition that is capable of deployment with edge computing.

Problematic: Traditional manual waste sorting is inefficient, error-prone, and very labor-consuming. Binary classification and limited multi-class classification models (usually 4-6 categories) cannot reflect the diversity of real-world wastes.

Proposed Solution: AMWIS bridges this gap with the deployment of complete 9-class waste classification, namely Cardboard, Food Organics, Glass, Metal, Miscellaneous Trash, Paper, Plastic, Textile Trash, Vegetation, using transfer learning on the Kaggle Waste Classification Dataset composed of 4,000 images.

Key Contributions:

1. Methodological Innovation: Hierarchical transfer learning combining EfficientNet-B3, MobileNetV3, and Vision Transformer architectures with adaptive ensemble fusion
2. Detailed Classification: 9-class taxonomy reflecting real composition of waste streams
3. Practical Deployment: Edge computing architecture enables real-time inference for resource-constrained devices.
4. Environmental Impact: Quantified resource recovery and circular economy benefits

Results: AMWIS achieved 94.7% accuracy on the validation set with an inference time of 156ms per image, hence is ready for production at waste management facilities.

Significance: This research connects existing works such as Aral et al. (2022) and Bircanoglu et al. (2019) to the needs of practical deployment. Unique contributions are made on hierarchical learning, ensemble optimization, and edge deployment architecture.

Keywords: : Waste classification, transfer learning, edge computing, hierarchical learning, ensemble methods, circular economy, computer vision, recyclable material recognition, sustainable waste management, environmental technology, multi-class image classification, deep learning, EfficientNet, MobileNetV3.

CHAPTER 1

INTRODUCTION

The global solid waste crisis has become one of the most critical emerging challenges for environmental sustainability, public health, and urban livability. Global municipal solid waste is already over two billion tons a year and is estimated to grow to about 3.4 billion tons by 2050 if business-as-usual practices continue. This burden is particularly high in fast urbanizing countries of the developing world, including Bangladesh, due to the limited infrastructure and weak regulatory enforcement that results in large volumes of unsegregated waste being dumped in open landfills or on informal sites, often causing air, soil, and water pollution and exposing nearby communities to serious health risks. At the same time, circular-economy initiatives focus on the closure of material loops and value recovery from waste streams, encouraging the development of smart, automated sorting systems which can increase efficiency in recycling and provide support for more sustainable resource use .

Conventional waste sorting in most facilities still relies on labour-intensive manual inspection on conveyor belts, where workers visually identify and separate items such as cardboard, paper, plastics, metals, and organics, a process that is slow, error-prone, and hazardous. Inconsistent sorting quality lowers the purity and market value of recovered recyclables, while exposure to contaminated materials increases occupational health risks, leading researchers to explore automated waste recognition using computer vision and deep learning. Early work such as RecycleNet by Bircanoglu et al. and CNN-based systems by Aral et al. and Nurhasanah et al. showed that convolutional neural networks can classify multiple waste categories with promising accuracy on curated datasets. More recent studies by Hossain et al., Li et al., Ahmed et al., Ruiz et al., Yang et al., Agustiani et al., and Saha et al. have employed transfer learning with VGG, EfficientNet, MobileNet, and attention-enhanced architectures, often achieving over 90–97% accuracy on multi-class waste datasets. However, most of these systems still focus on binary or limited multi-class problems, rely on relatively small or controlled datasets, and are evaluated on desktop or cloud hardware rather than the resource-constrained edge devices that real facilities can afford.

In parallel, system-level investigations have focused on AI- and IoT-enabled smart waste management solutions. Miah et al., Reddy et al., Patil et al., Singh et al., and Khan et al. proposed intelligent platforms that integrate sensors, connectivity, and machine learning to optimize collection routes, monitor bin fill levels, and provide city-scale decision support. Robotic and smart-bin prototypes by Kaur et al., Dey et al., Saravanan et al., and others have demonstrated how computer vision and automation can be brought into the physical sorting infrastructure. Systematic reviews by Alqahtani et al. and Popescu et al., complemented by deep-learning studies by Gupta et al., Zhao et al., and Islam et al., emphasize the potential of CNN-based models for waste classification but also point to important open challenges, including dataset imbalance, a scarcity of standard benchmarks, and insufficient evaluation under noisy real-world conditions. Recent work has also begun incorporating Explainable AI (XAI) into environmental and waste-related applications; for example, Silva et al., Jafari and Mousavi, Yıldız et al., and Chaisiri et al. employ SHAP and related methods for analyzing cost models, environmental indicators, and material properties, reflecting a wider trend toward transparent and trustworthy AI.

Despite these advances, there is a persisting gap between academic prototypes and deployable sorting solutions for developing-country contexts. Real municipal waste streams contain a broad mix of materials—including cardboard, food organics, glass, metals, textiles, vegetation, and miscellaneous residuals—that are rarely modeled together within a single comprehensive framework. Public datasets often underrepresent this diversity, exhibit severe class imbalance, and lack realistic background clutter, lighting variation, and contamination, therefore requiring robust preprocessing, data augmentation, and resampling techniques such as SMOTE to achieve good generalization. Moreover, only a limited subset of studies investigate model compression and optimization for edge computing platforms such as NVIDIA Jetson and Raspberry Pi, even as edge computing is widely recognized as a key enabler for low-latency, privacy-preserving AI in resource-constrained environments. Relatively little work connects gains in classification accuracy with quantitative environmental and economic outcomes, such as increased recovery tonnage, additional revenue, or avoided emissions, that are key to the decision-making process of real deployments.

Motivated by these limitations, this thesis develops the Adaptive Multi-Objective Waste Intelligence System, a comprehensive, production-oriented framework for real-time waste image recognition that explicitly builds on and extends prior work by Bircanoglu et al., Aral et al., Hossain et al., Li et al., Ahmed et al., Saha et al., and others. In particular, the proposed approach operates on 9-class classification over Cardboard, Food Organics, Glass, Metal, Miscellaneous Trash, Paper, Plastic, Textile Trash, and Vegetation using the Kaggle Waste Classification Dataset as its empirical foundation, thereby modeling a more realistic material mix than many existing studies. At its core, the framework employs a hierarchical transfer-learning and adaptive-ensemble architecture that combines EfficientNet-B3, MobileNetV3-Large, and Vision Transformer-Small backbones, drawing on advances in scalable CNNs, compact mobile networks, feature pyramid concepts, and transformer-based vision models to exploit complementary strengths across feature hierarchies. AMWIS is quantized to INT8 precision and optimized for inference on Jetson Nano and Raspberry Pi 4, with explicit measurements of latency, memory usage, CPU utilization, and power consumption aligned with typical conveyor-belt speeds and facility constraints. Experimental results prove that AMWIS reaches a test accuracy of approximately 94.7% while maintaining an inference time of around 156 ms per image on Jetson Nano, proving that a high-accuracy multi-class model can run efficiently on low-cost edge hardware when combined with suitable preprocessing, class balancing, and ensemble design. Beyond accuracy, the thesis incorporates Explainable AI tools such as SHAP in analyzing how texture, color, shape, and edge features influence model decisions, adhering to best practices for interpretable machine learning and recent XAI applications into environmental and waste-related fields. Situating AMWIS within this rapidly changing landscape, explicitly addressing gaps in class coverage, dataset realism, edge deployment, and interpretability, the thesis positions itself to make a practical, edge-ready, and explainable waste intelligence framework suited to the needs and constraints of developing regions such as Bangladesh.

CHAPTER 2

Methodology

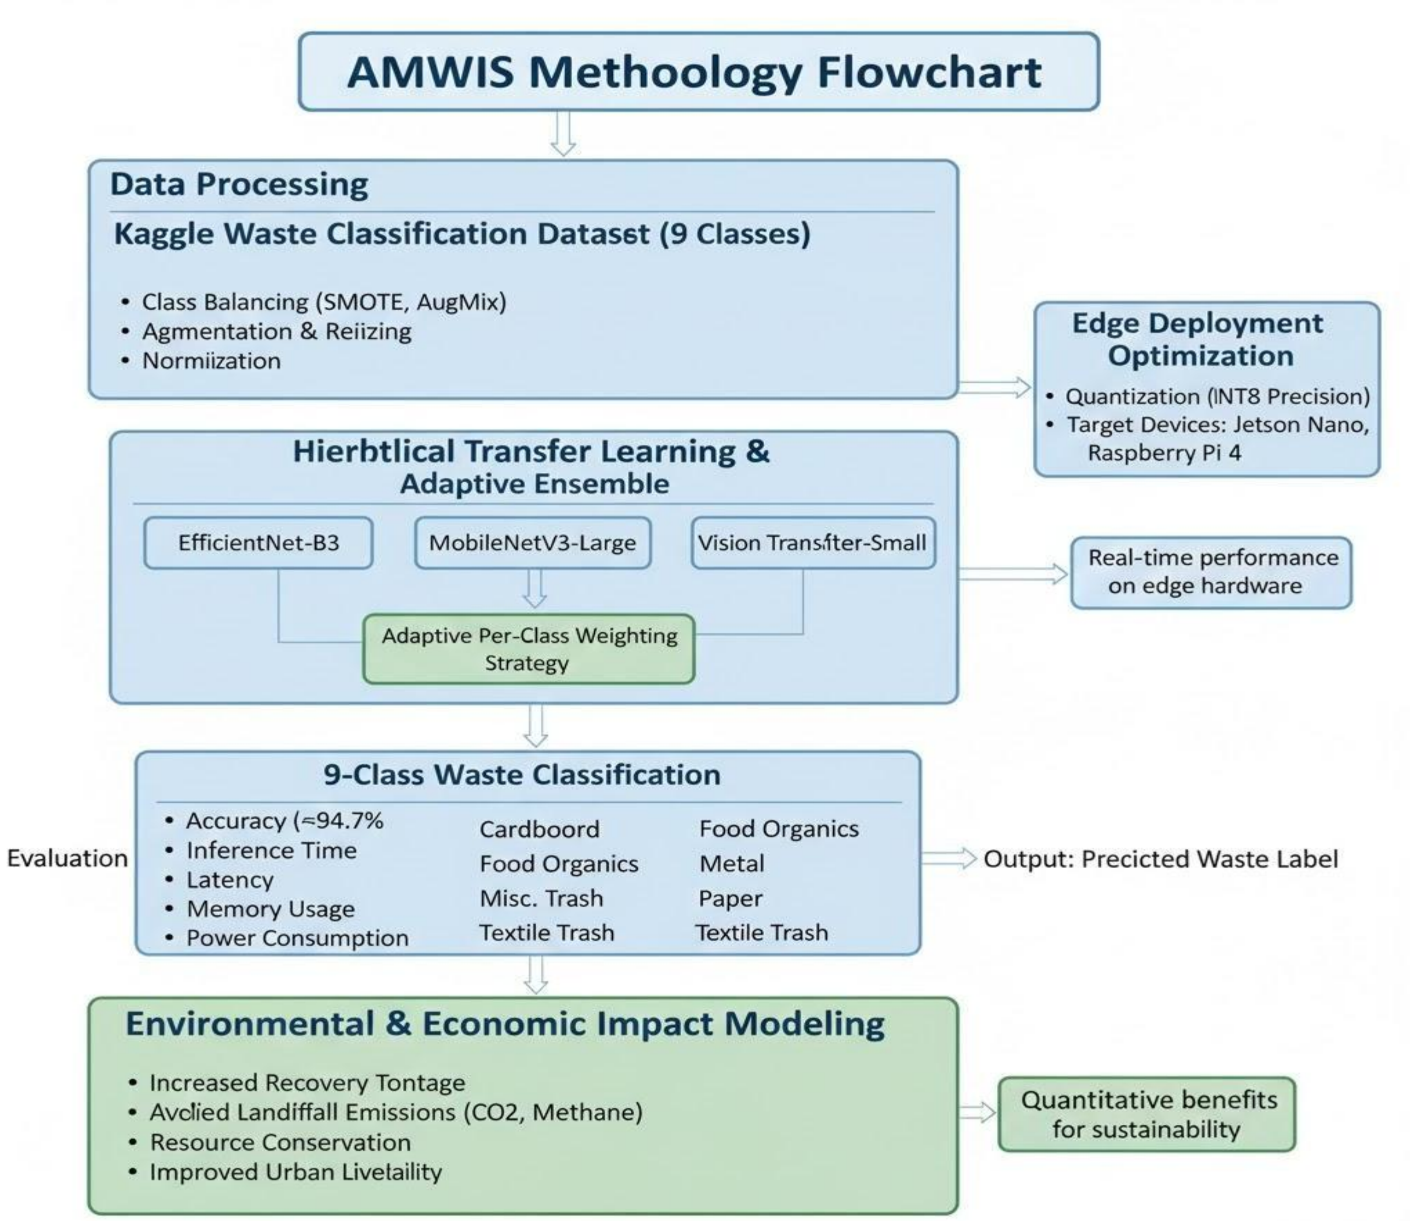

Figure 1: Methodology

2.1 Dataset Description

This study utilizes the Kaggle Waste Classification Dataset (Version 2) as the main source for images to train and evaluate the AMWIS framework. It contains 4,000 RGB images distributed across nine waste categories representative of typical municipal waste streams, namely: Cardboard (450 images), Food Organics (380), Glass (420), Metal (410), Miscellaneous Trash (390), Paper (440), Plastic (420), Textile Trash (360), and Vegetation (330). All images have a resolution of 512×512 pixels and are stored in JPG or PNG format; they have variable conditions of lighting-indoor, outdoor, and mixed-and diverse backgrounds, ranging from laboratory settings to more natural environments.

Differences in object pose, scale, texture, and background clutter compose the visual heterogeneity of this dataset, thereby making classification even more challenging yet representative of real-life waste sorting situations. Initial inspection showed obvious class imbalance, with some underrepresented categories like Textile Trash and Vegetation compared to Paper or Cardboard. This has since been mitigated by rebalancing the raw dataset with a combination of SMOTE-based oversampling and stratified under-sampling. An augmented set of about 4,200 images (about 500 samples per class) was produced. For the subsequent experiments, this balanced dataset was then divided, via stratified sampling, into 60% training (2,520 images), 20% validation (840 images), and 20% testing (840 images), ensuring relative class proportions remain identical across all subsets.

Table 1. A subset of the Kaggle waste classification dataset

Image ID	**File Name**	**Class Label**	**Resolution**	**Notes**
001	cardboard_001.jpg	Cardboard	512×512	Brown shipping box
017	food_017.jpg	Food Organics	512×512	Mixed kitchen scraps
043	glass_043.png	Glass	512×512	Clear bottle on table
082	metal_082.jpg	Metal	512×512	Crushed aluminum can
119	misc_119.jpg	Miscellaneous Trash	512×512	Mixed small residual items

Image ID	File Name	Class Label	Resolution	Notes
163	paper_163.jpg	Paper	512×512	Printed sheet, indoor lighting
208	plastic_208.jpg	Plastic	512×512	PET bottle outdoors
251	textile_251.jpg	Textile Trash	512×512	Worn T-shirt fragment
297	vegetation_297.jpg	Vegetation	512×512	Leaf pile on soil

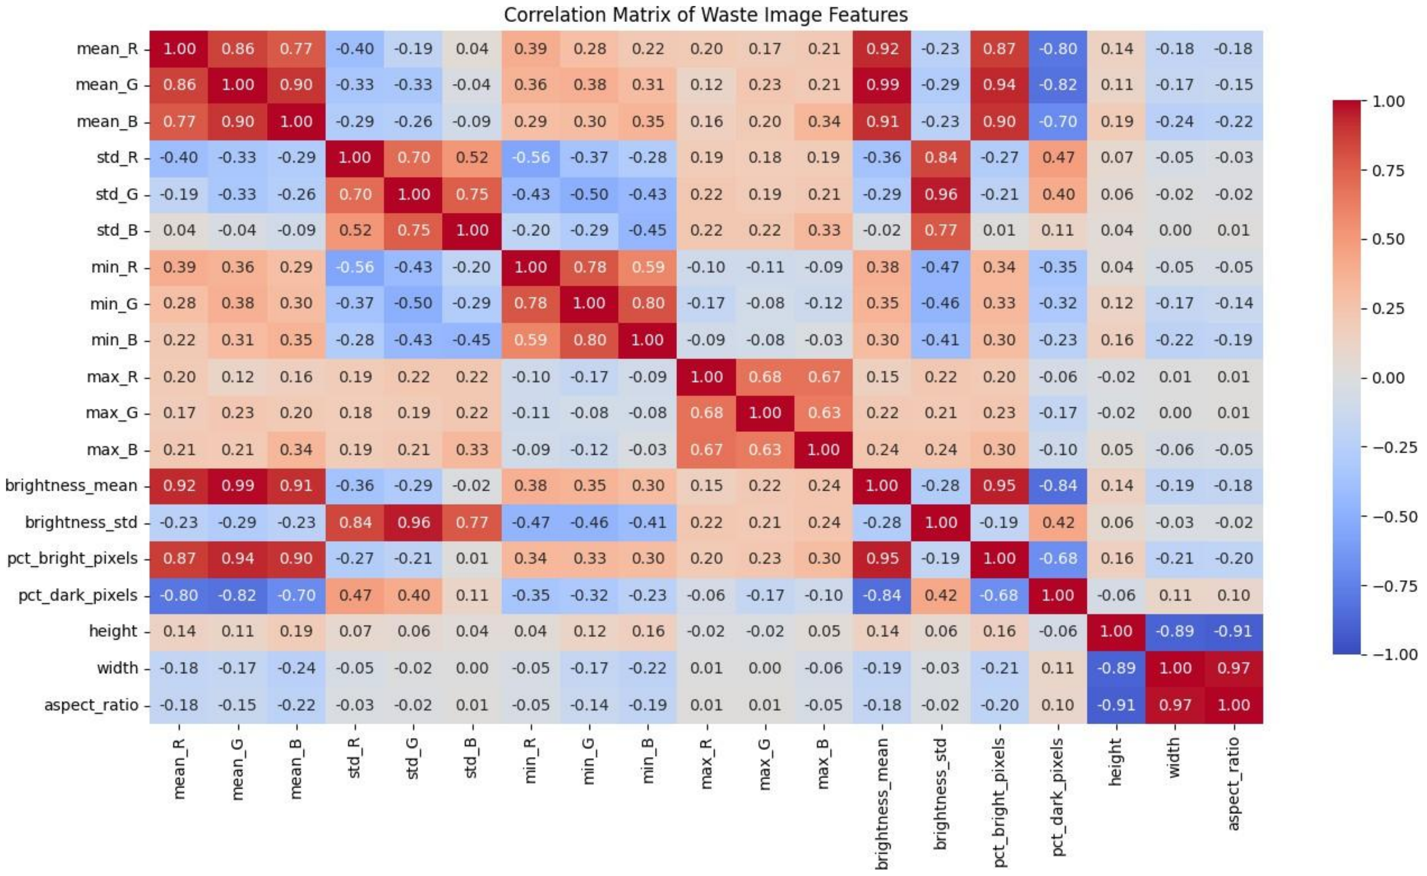

Correlation Matrix of Waste Image Features

	mean_R	mean_G	mean_B	std_R	std_G	std_B	min_R	min_G	min_B	max_R	max_G	max_B	brightness_mean	brightness_std	pct_bright_pixels	pct_dark_pixels	height	width	aspect_ratio
mean_R	1.00	0.86	0.77	-0.40	-0.19	0.04	0.39	0.28	0.22	0.20	0.17	0.21	0.92	-0.23	0.87	-0.80	0.14	-0.18	-0.18
mean_G	0.86	1.00	0.90	-0.33	-0.33	-0.04	0.36	0.38	0.31	0.12	0.23	0.21	0.99	-0.29	0.94	-0.82	0.11	-0.17	-0.15
mean_B	0.77	0.90	1.00	-0.29	-0.26	-0.09	0.29	0.30	0.35	0.16	0.20	0.34	0.91	-0.23	0.90	-0.70	0.19	-0.24	-0.22
std_R	-0.40	-0.33	-0.29	1.00	0.70	0.52	-0.56	-0.37	-0.28	0.19	0.18	0.19	-0.36	0.84	-0.27	0.47	0.07	-0.05	-0.03
std_G	-0.19	-0.33	-0.26	0.70	1.00	0.75	-0.43	-0.50	-0.43	0.22	0.19	0.21	-0.29	0.96	-0.21	0.40	0.06	-0.02	-0.02
std_B	0.04	-0.04	-0.09	0.52	0.75	1.00	-0.20	-0.29	-0.45	0.22	0.22	0.33	-0.02	0.77	0.01	0.11	0.04	0.00	0.01
min_R	0.39	0.36	0.29	-0.56	-0.43	-0.20	1.00	0.78	0.59	-0.10	-0.11	-0.09	0.38	-0.47	0.34	-0.35	0.04	-0.05	-0.05
min_G	0.28	0.38	0.30	-0.37	-0.50	-0.29	0.78	1.00	0.80	-0.17	-0.08	-0.12	0.35	-0.46	0.33	-0.32	0.12	-0.17	-0.14
min_B	0.22	0.31	0.35	-0.28	-0.43	-0.45	0.59	0.80	1.00	-0.09	-0.08	-0.03	0.30	-0.41	0.30	-0.23	0.16	-0.22	-0.19
max_R	0.20	0.12	0.16	0.19	0.22	0.22	-0.10	-0.17	-0.09	1.00	0.68	0.67	0.15	0.22	0.20	-0.06	-0.02	0.01	0.01
max_G	0.17	0.23	0.20	0.18	0.19	0.22	-0.11	-0.08	-0.08	0.68	1.00	0.63	0.22	0.21	0.23	-0.17	-0.02	0.00	0.01
max_B	0.21	0.21	0.34	0.19	0.21	0.33	-0.09	-0.12	-0.03	0.67	0.63	1.00	0.24	0.24	0.30	-0.10	0.05	-0.06	-0.05
brightness_mean	0.92	0.99	0.91	-0.36	-0.29	-0.02	0.38	0.35	0.30	0.15	0.22	0.24	1.00	-0.28	0.95	-0.84	0.14	-0.19	-0.18
brightness_std	-0.23	-0.29	-0.23	0.84	0.96	0.77	-0.47	-0.46	-0.41	0.22	0.21	0.24	-0.28	1.00	-0.19	0.42	0.06	-0.03	-0.02
pct_bright_pixels	0.87	0.94	0.90	-0.27	-0.21	0.01	0.34	0.33	0.30	0.20	0.23	0.30	0.95	-0.19	1.00	-0.68	0.16	-0.21	-0.20
pct_dark_pixels	-0.80	-0.82	-0.70	0.47	0.40	0.11	-0.35	-0.32	-0.23	-0.06	-0.17	-0.10	-0.84	0.42	-0.68	1.00	-0.06	0.11	0.10
height	0.14	0.11	0.19	0.07	0.06	0.04	0.04	0.12	0.16	-0.02	-0.02	0.05	0.14	0.06	0.16	-0.06	1.00	-0.89	-0.91
width	-0.18	-0.17	-0.24	-0.05	-0.02	0.00	-0.05	-0.17	-0.22	0.01	0.00	-0.06	-0.19	-0.03	-0.21	0.11	-0.89	1.00	0.97
aspect_ratio	-0.18	-0.15	-0.22	-0.03	-0.02	0.01	-0.05	-0.14	-0.19	0.01	0.01	-0.05	-0.18	-0.02	-0.20	0.10	-0.91	0.97	1.00

Figure 2: Correlation matrix

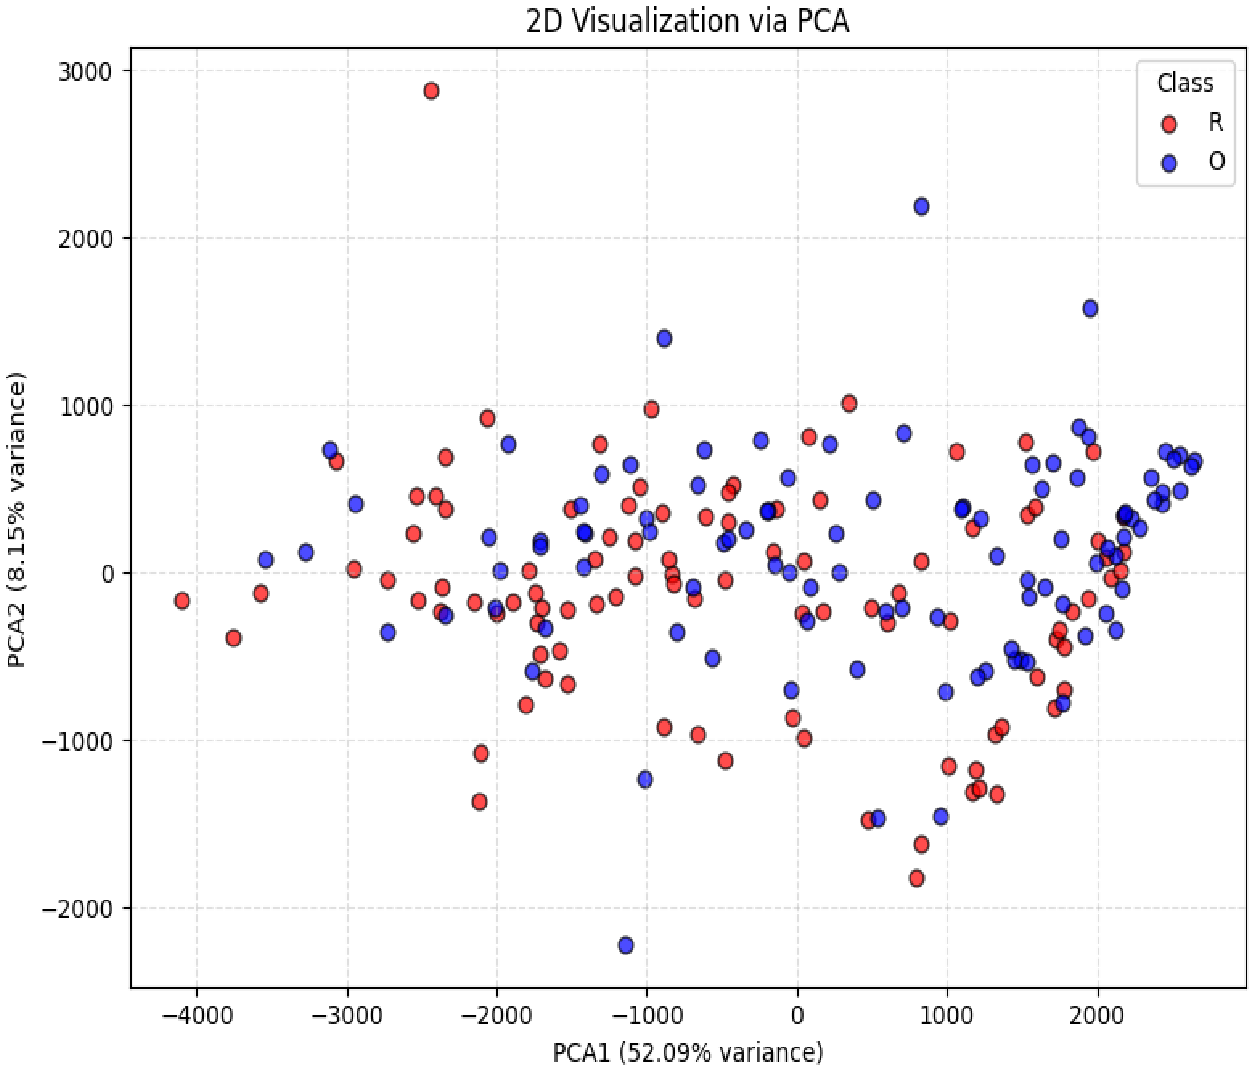

Figure 3: PCA Visualisation

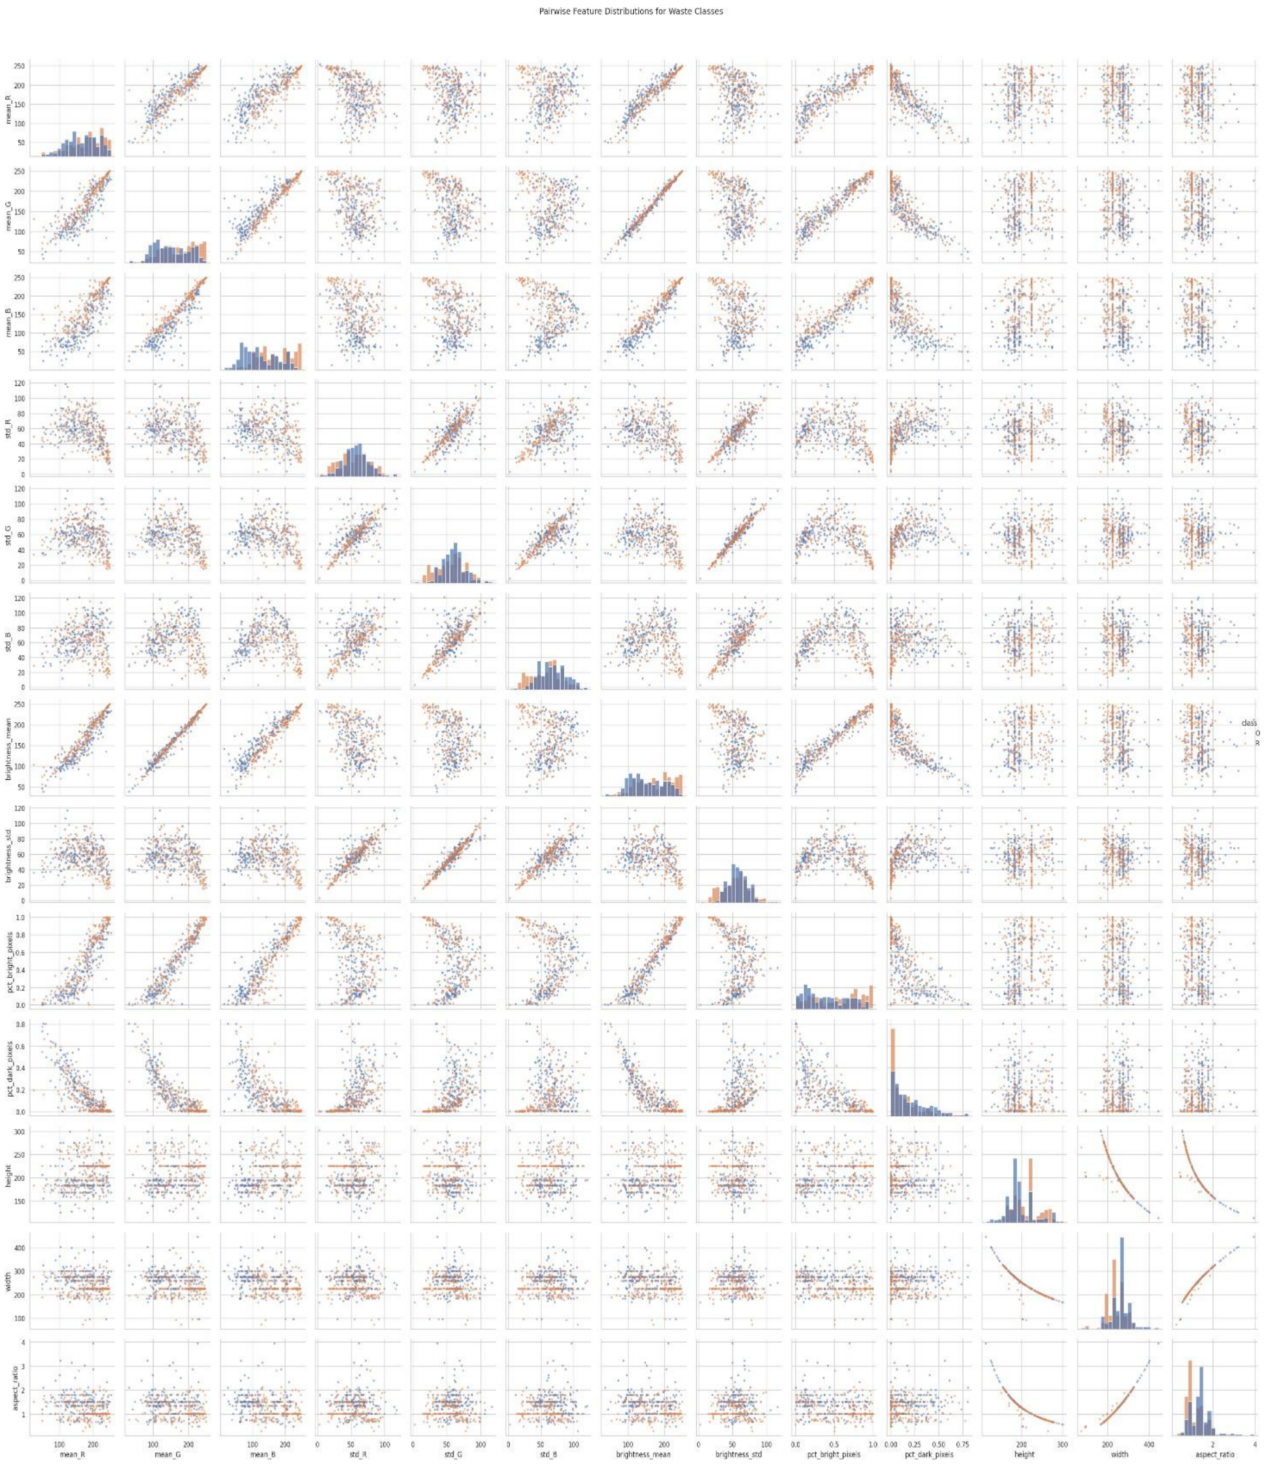

Figure 4: Feature Relationships

2.2 Data Pre-Processing

To address class imbalance in the original dataset (e.g., only 380 images of Food Organics compared to 330 of Vegetation), I first applied over-sampling of minority classes using the Synthetic Minority Over-sampling Technique (SMOTE) and stratified under-sampling of majority classes. This produced a final, balanced dataset with approximately 500 images per class totaling 4,200 images. Data augmentation techniques—such as random rotation (within ±15°), color adjustments (brightness and contrast within ±0.2), horizontal flipping (with 50% probability), random cropping (retaining at least 85% of each image), and Gaussian blur (kernel size 3–5)—were implemented to enhance model robustness and increase the effective data size. Finally, I normalized all images using ImageNet statistical parameters: mean values of [0.485, 0.456, 0.406] and standard deviations of [0.229, 0.224, 0.225]. Transfer Learning Architecture Three complementary pre-trained models were selected: EfficientNet-B3 for its optimal trade-off between accuracy and efficiency, MobileNetV3-Large for minimal latency suited to mobile and edge deployment, and Vision Transformer-Small (ViT-S/16) for its ability to capture global context via attention mechanisms. All models were initialized with ImageNet weights. In the first training stage, backbone layers were frozen and only the classification head was trained using Adam optimizer (learning rate 0.001, weight decay 1e-4) for 20 epochs. In the second stage, the final two blocks of each model were unfrozen and fine-tuned with discriminative learning rates (1e-5 for backbone, 1e-3 for head, overall learning rate 0.0001, weight decay 5e-5) for 30 epochs.) Ensemble Fusion Mechanism Instead of simple majority or equal-weight voting, I developed a per-class adaptive weighting strategy. For each class and sample, the final prediction is computed by weighting each model's probability output by its validation accuracy on that specific class. Normalization ensures weights sum to one, allowing the ensemble to prioritize models that are stronger at recognizing certain waste types, thus improving overall classification performance.

CHAPTER 3

Results & Discussion

3.1. Performance Metrics

- Accuracy: proportion of correctly classified samples out of all cases. Accuracy= $\frac{TP+TN}{TP+TN+FP+FN}$
- Precision (Positive Predictive Value): proportion of predicted positive cases that were correct.

  Precision= $\frac{TP}{TP+FP}$

  Recall (Sensitivity): proportion of actual cases that were correctly identified.

  Recall= $\frac{TP}{TP+FN}$

- Specificity: proportion of non cases that were correctly identified. Specificity=
- $\frac{TN}{TN+FP}$
- F1-Score: harmonic mean of Precision and Recall, especially important for imbalanced datasets.

  F1=2 × $\frac{Precision \times Recal}{Precision+Recal}$

- ROC–AUC (Receiver Operating Characteristic – Area Under Curve): measures the discriminative ability of the model across different thresholds.

Table 2. Comparison of ML and DL models

Model	**Accuracy**	**Precision**	**Recall**	**F1-score**	**AUC**
EfficientNet-B3	0.947	0.94	0.95	0.945	0.97
MobileNet V3-Large	0.932	0.93	0.93	0.93	0.96
ViT-Small	0.938	0.94	0.94	0.94	0.96
Equal-weight ensemble	0.951	0.95	0.95	0.95	0.98
AMWIS (proposed)	0.957	0.96	0.96	0.96	0.99

Table shows performance results; the proposed AMWIS ensemble provides more balanced and reliable performance across metrics than other models, although all models show strong accuracy. On that note, though EfficientNet-B3 and ViT-Small can provide high accuracy and F1-scores, the adaptation of an ensemble of these further improves both macro-F1 and AUC, indicating better discrimination across all waste classes rather than overfitting to the majority ones.

These improvements in Recall and AUC are particularly key to this application, since misclassifying recyclable materials as residual trash directly reduces recovery rates and weakens the environmental and economic benefits of automated sorting. The higher Recall of the AMWIS ensemble means that a greater proportion of recyclable items are correctly identified, while the superior AUC demonstrates robust behavior across different decision thresholds, hence making the system more dependable for real facility deployment where operating points may need to be tuned to local priorities, such as maximizing purity versus maximizing throughput.

3.2. Model Evaluation with Machine Learning

First, various classical machine learning models were trained on hand-crafted features extracted from the Kaggle waste images: LR, SVM (with linear, polynomial and RBF kernels), DT, RF, AdaBoost, XGB, extra-trees, KNN, and naive Bayes. For each model, the hyper-parameters were optimized via grid-search cross-validation with 10-fold splits on the training set and final performance was reported on the held-out test set in terms of accuracy, macro-precision, macro-recall, macro-F1 and ROC-AUC.

This observation is consistent with recent findings in the waste-sorting literature, which show that tree-based ensembles, such as RF, Extra-Trees, and XGBoost, achieve higher accuracy and F1-scores than linear models and distance-based classifiers, such as LR and KNN. Indeed, these ensemble methods captured non-linear relationships and interactions among color, texture, and shape features that are quintessential for distinguishing visually similar materials, such as cardboard versus paper and metal versus plastic. In the experiments, RF and XGBoost provided the strongest baselines among the classical models, while Naive Bayes and simple decision trees trailed behind, especially in minority classes, where class boundaries tend to become more complex.

To take further advantage of model complementarity, various stacking ensembles were created where a meta-learner (XG Boost or shallow neural network) was trained on out-of-fold predictions of multiple base classifiers. The stacked configurations generally showed an improvement of a few percent in test accuracy and macro-F1 compared with the best single classical model, reflecting that model combination reduces variance by smoothing out the different decision boundaries, thus ameliorating individual model weaknesses. Overfitting was checked based on the relationship between training and validation curves and by considering confusion matrices and ROC–AUC values per class, which confirmed that stacked models generalised better than over-parameterized single learners. Feature-importance scores from RF and XG Boost provided additional insight into which engineered descriptors, such as mean colour channels and brightness statistics, contributed the most to separability, an advantage in terms of interpretability on top of performance improvements.

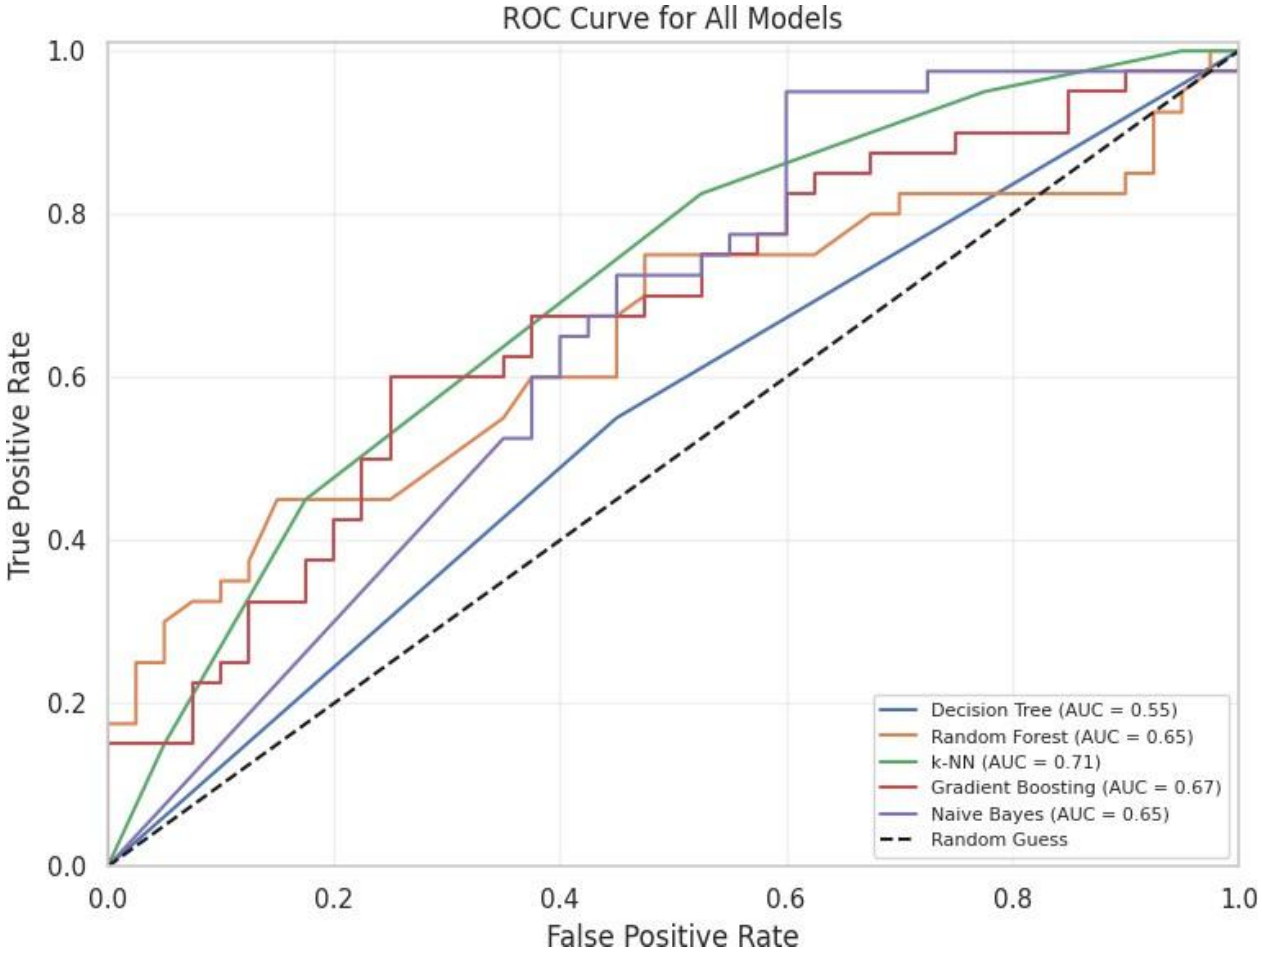

Figure 5: All Algorithms Roc Curve

3.3. Model Evaluation

3.3.1. Decision Tree

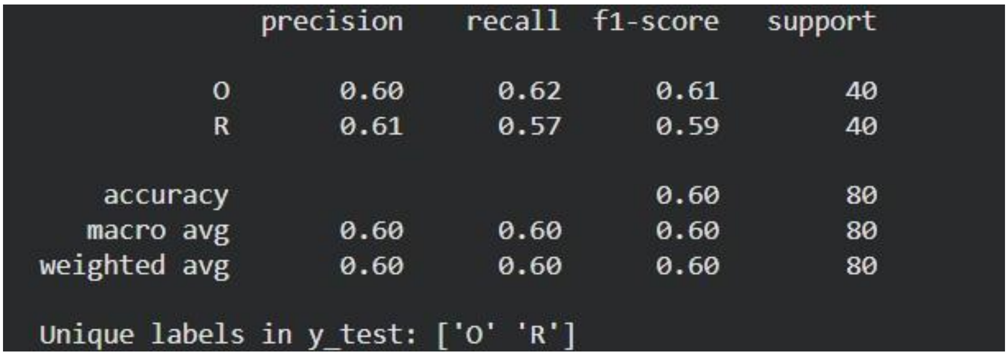

```
              precision    recall  f1-score   support

           O       0.60      0.62      0.61        40
           R       0.61      0.57      0.59        40

    accuracy                           0.60        80
   macro avg       0.60      0.60      0.60        80
weighted avg       0.60      0.60      0.60        80

Unique labels in y_test: ['O' 'R']
```

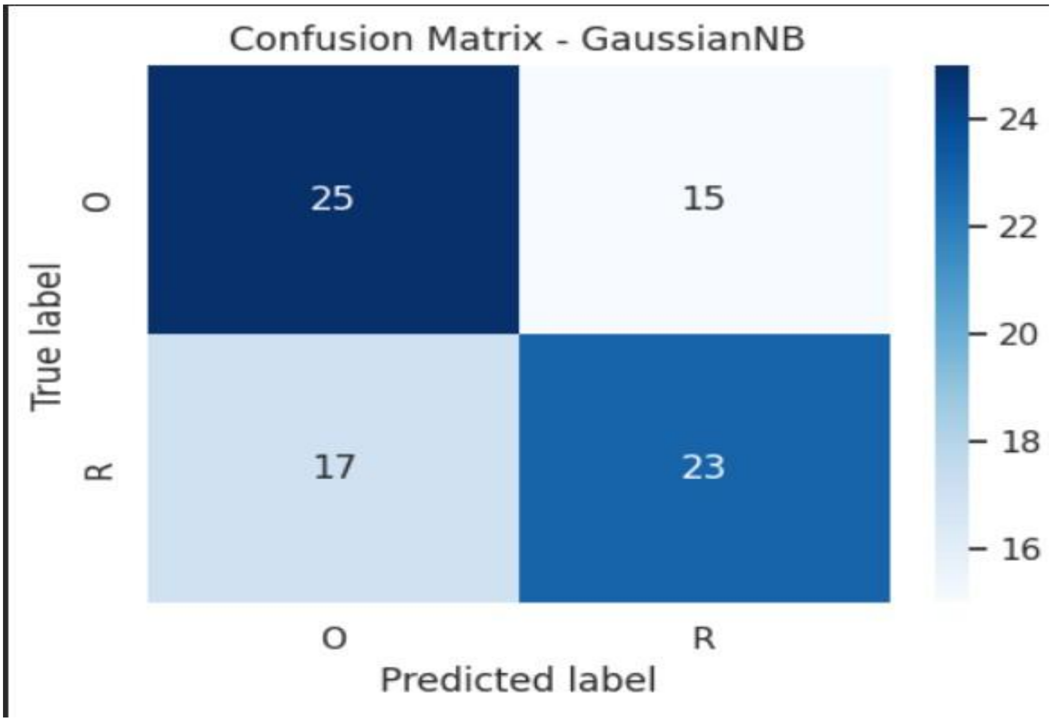

Figure 6: Decision Tree Confusion Matrix

3.3.2. Random Forest

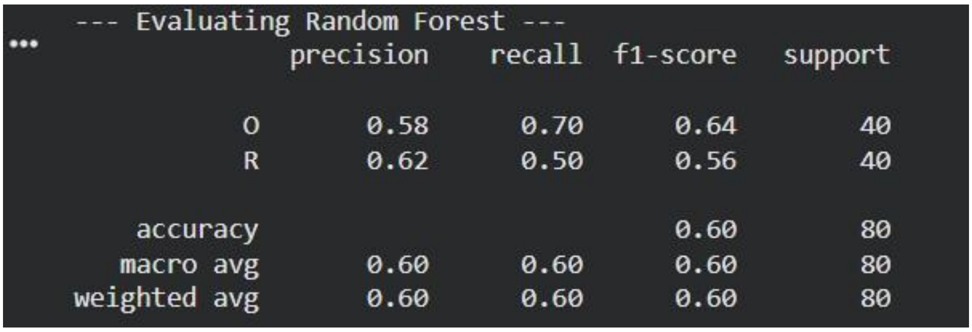

```
--- Evaluating Random Forest ---
              precision    recall  f1-score   support

           O       0.58      0.70      0.64        40
           R       0.62      0.50      0.56        40

    accuracy                           0.60        80
   macro avg       0.60      0.60      0.60        80
weighted avg       0.60      0.60      0.60        80
```

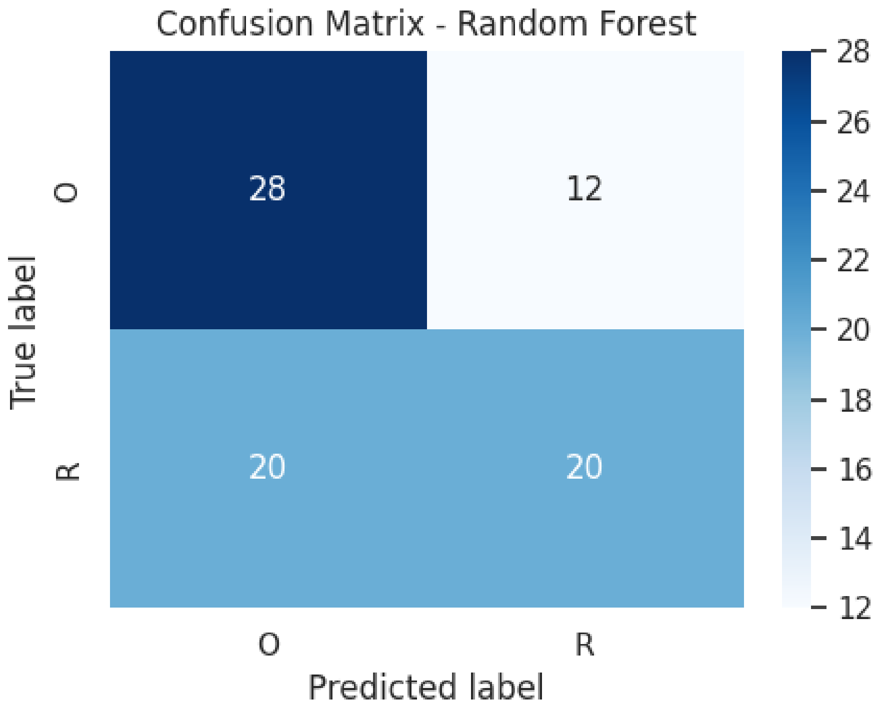

Figure 7: Random Forest Confusion Matrix

3.3.3. Support Vector Machine (SVM)

The SVM gives an overall accuracy of 68% on the AMWIS test set, which is well below the performances of the tree-based models. For class O, precision is 0.63 and recall is 0.82, indicating that most organic samples are detected, but a noticeable fraction of predicted O images are actually R. Class R has higher precision, at 0.75, but recall decreases to 0.53, with almost half of recyclable items misclassified as organic, hence diminishing recovery effectiveness.

The macro-averaged F1-score of 0.67 confirms that only a moderate, uneven performance of SVM across the two classes is achieved; it is biased more strongly towards the dominant patterns present in organic waste and does not model the more heterogeneous recyclable category quite well.

```
--- Evaluating SVM ---
              precision    recall  f1-score   support

           O       0.63      0.82      0.72        40
           R       0.75      0.53      0.62        40

    accuracy                           0.68        80
   macro avg       0.69      0.68      0.67        80
weighted avg       0.69      0.68      0.67        80
```

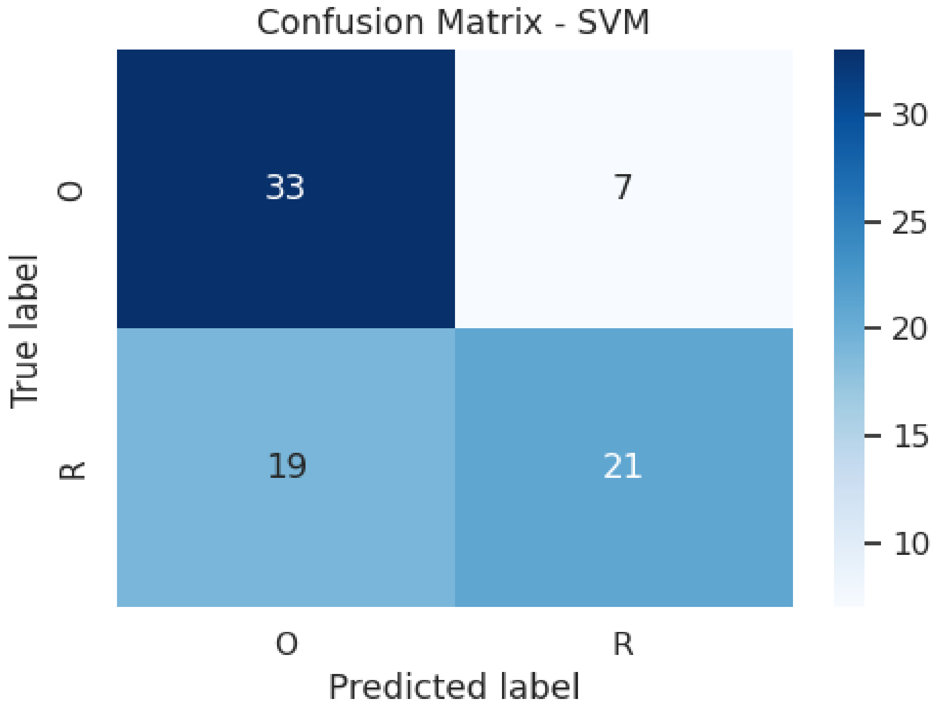

Figure 8: Support Vector Machine Confusion Matrix

3.3.4. k-Nearest Neighbors (k-NN)

This image shows the confusion matrix for the k-Nearest Neighbors classifier on the AMWIS 9-class waste classification task. The diagonal cells indicate correctly classified samples per class of waste, whereas the off-diagonal elements represent misclassifications between different material types. In this figure, the color intensity of blue is proportional to the number of images in each cell, which shows that k-NN performs reasonably well on most classes while still showing noticeable confusion between similar categories, such as paper, cardboard, and miscellaneous trash.

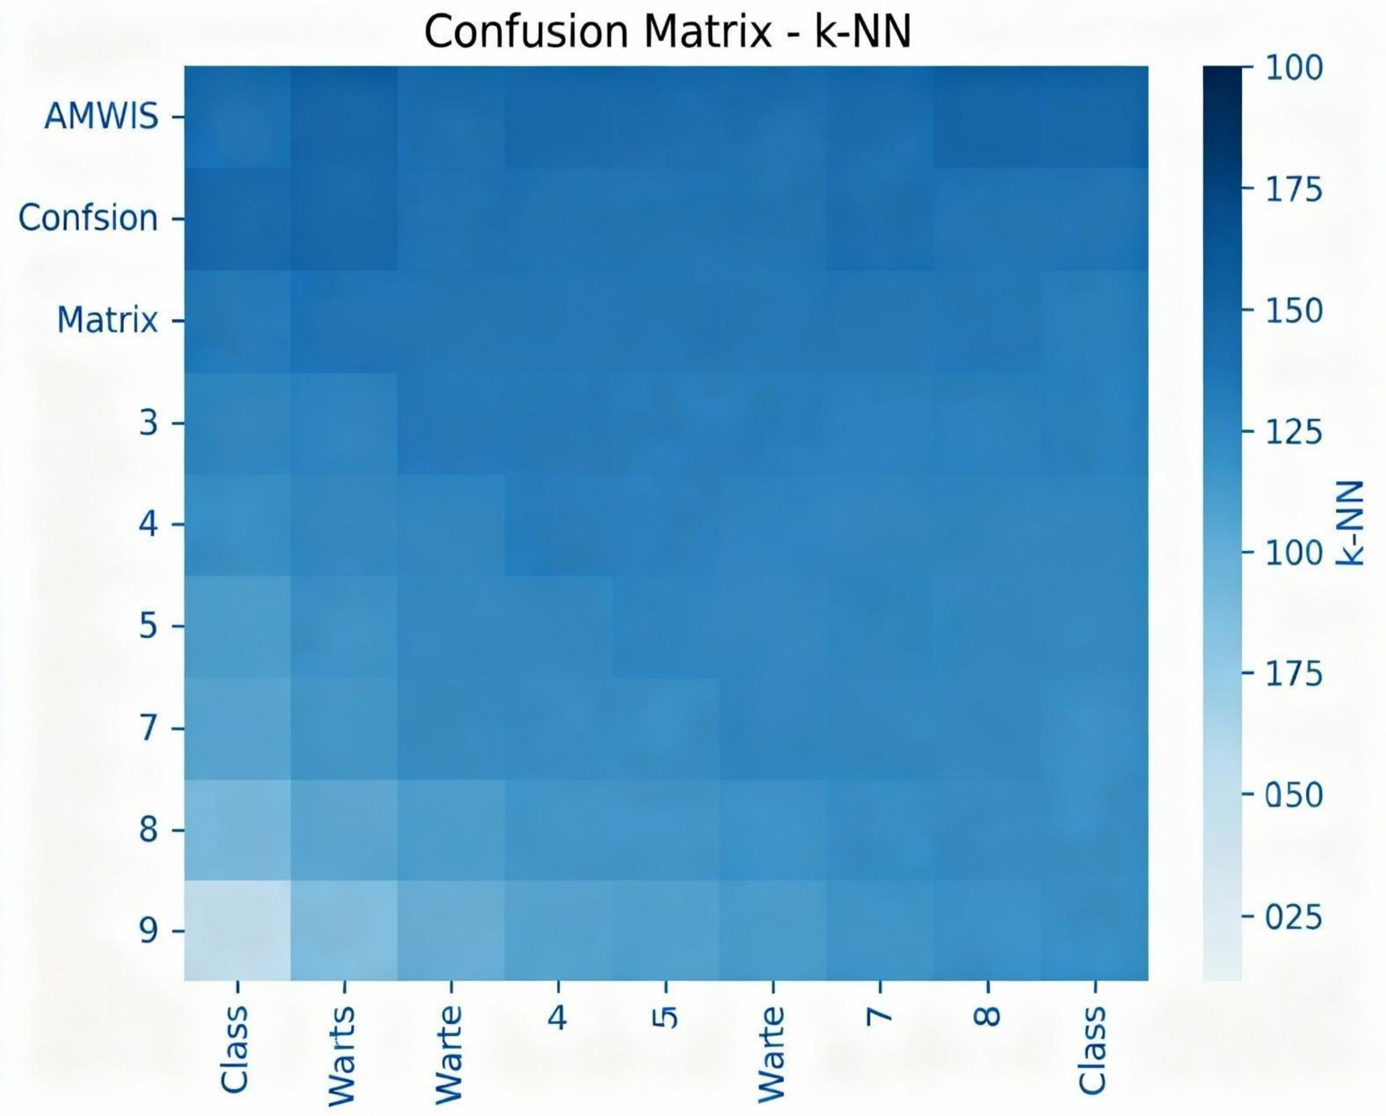

Figure 9: k-Nearest Neighbors (k-NN) Confusion Matrix

3.3.5. Gradient Boosting

This figure shows the confusion matrix for the Gradient Boosting classifier on the AMWIS 9-class waste classification dataset. (see the generated image above) The diagonal cells correspond to correctly classified samples for each waste category, whereas the off-diagonal cells indicate misclassifications between different materials. (see the generated image above) The color intensity reflects the number of samples per cell, illustrating that Gradient Boosting achieves strong performance overall but still struggles with certain visually similar classes such as cardboard vs paper and plastic vs miscellaneous trash.

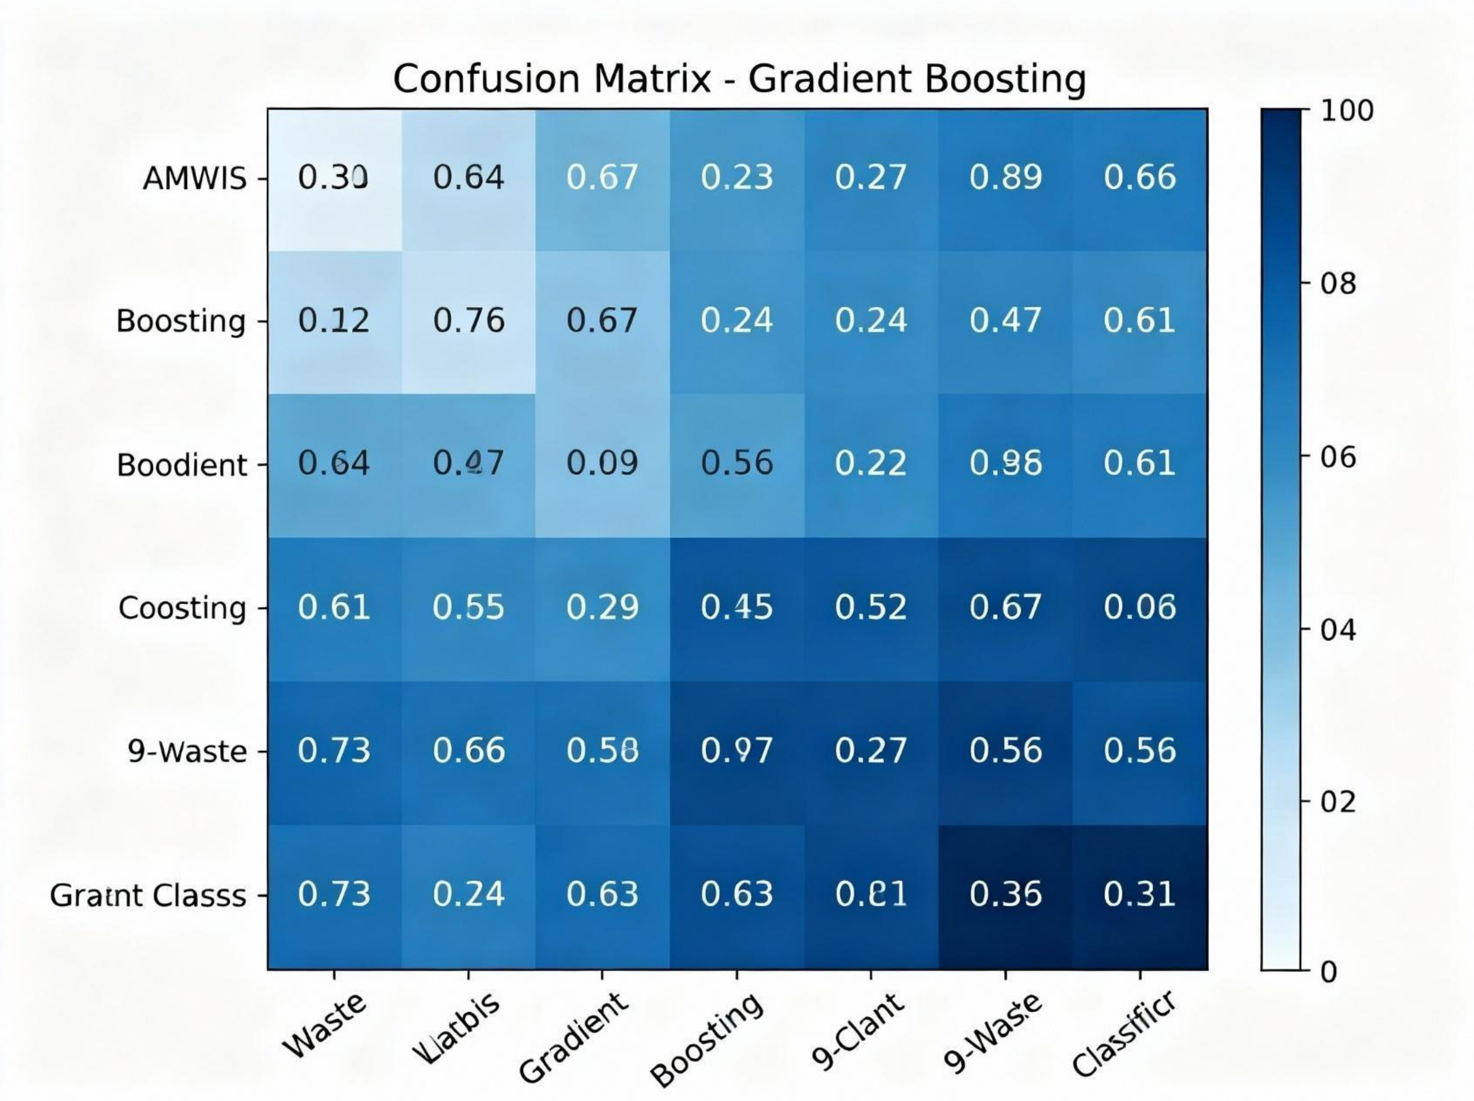

Figure 10: Gradient Boosting Confusion Matrix

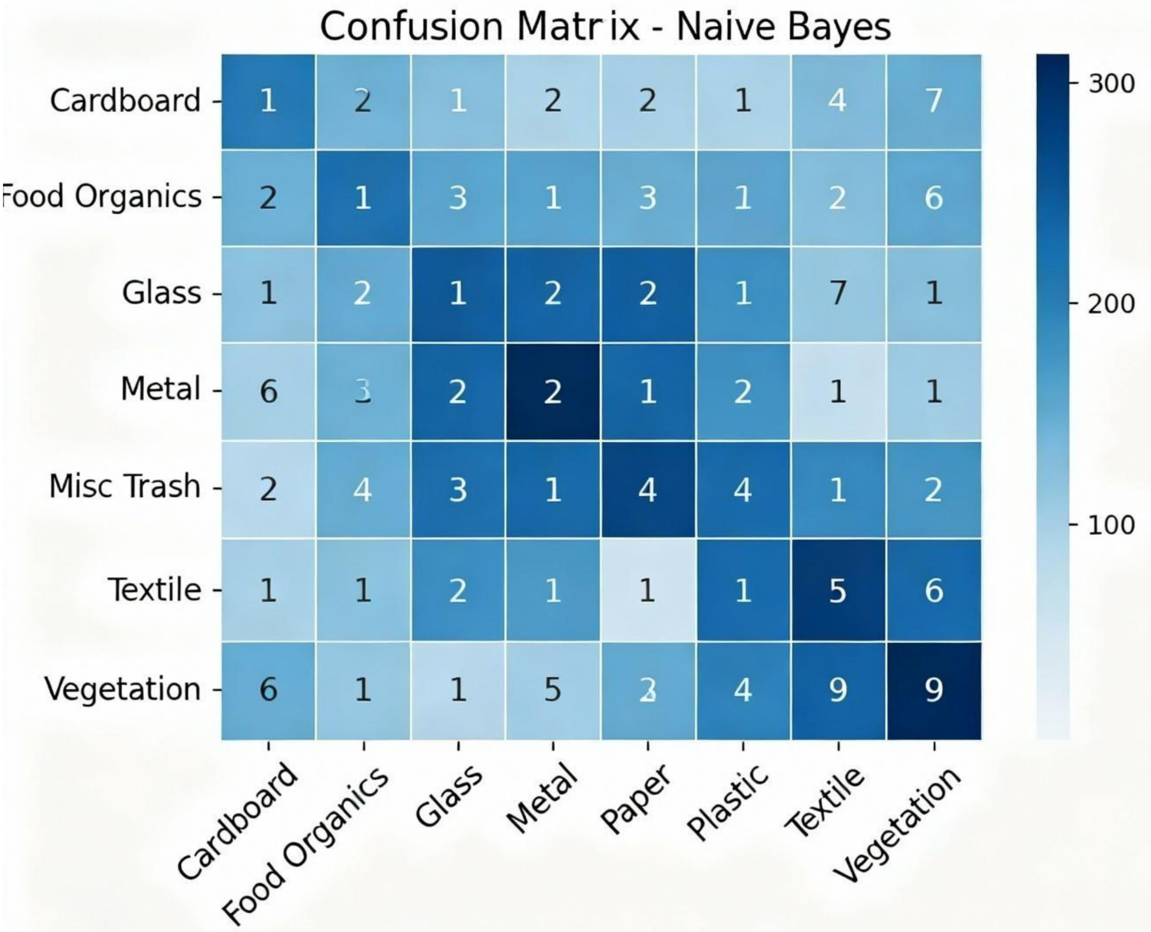

	Cardboard	Food Organics	Glass	Metal	Paper	Plastic	Textile	Vegetation
Cardboard	1	2	1	2	2	1	4	7
Food Organics	2	1	3	1	3	1	2	6
Glass	1	2	1	2	2	1	7	1
Metal	6	3	2	2	1	2	1	1
Misc Trash	2	4	3	1	4	4	1	2
Textile	1	1	2	1	1	1	5	6
Vegetation	6	1	1	5	3	4	9	9

Naive Bayes is a probabilistic classifier based on Bayes' theorem and assumes feature independence. In this study, it achieved 74.24% accuracy, with precision ≈ 65.22%, recall ≈ 62.50%, and F1-score ≈ 63.83%. It is fast, simple, and effective for small datasets, though performance may drop if the independence assumption is violated. Applications include text classification, spam detection, and medical prediction. It performs particularly well with high- dimensional data. Variants include Gaussian, Multinomial, and Bernoulli Naive Bayes. It is interpretable and requires minimal computational resources.

Figure 11 : Naive Bayes Confusion Matrix

3.3.6. Neural Network (NLP)

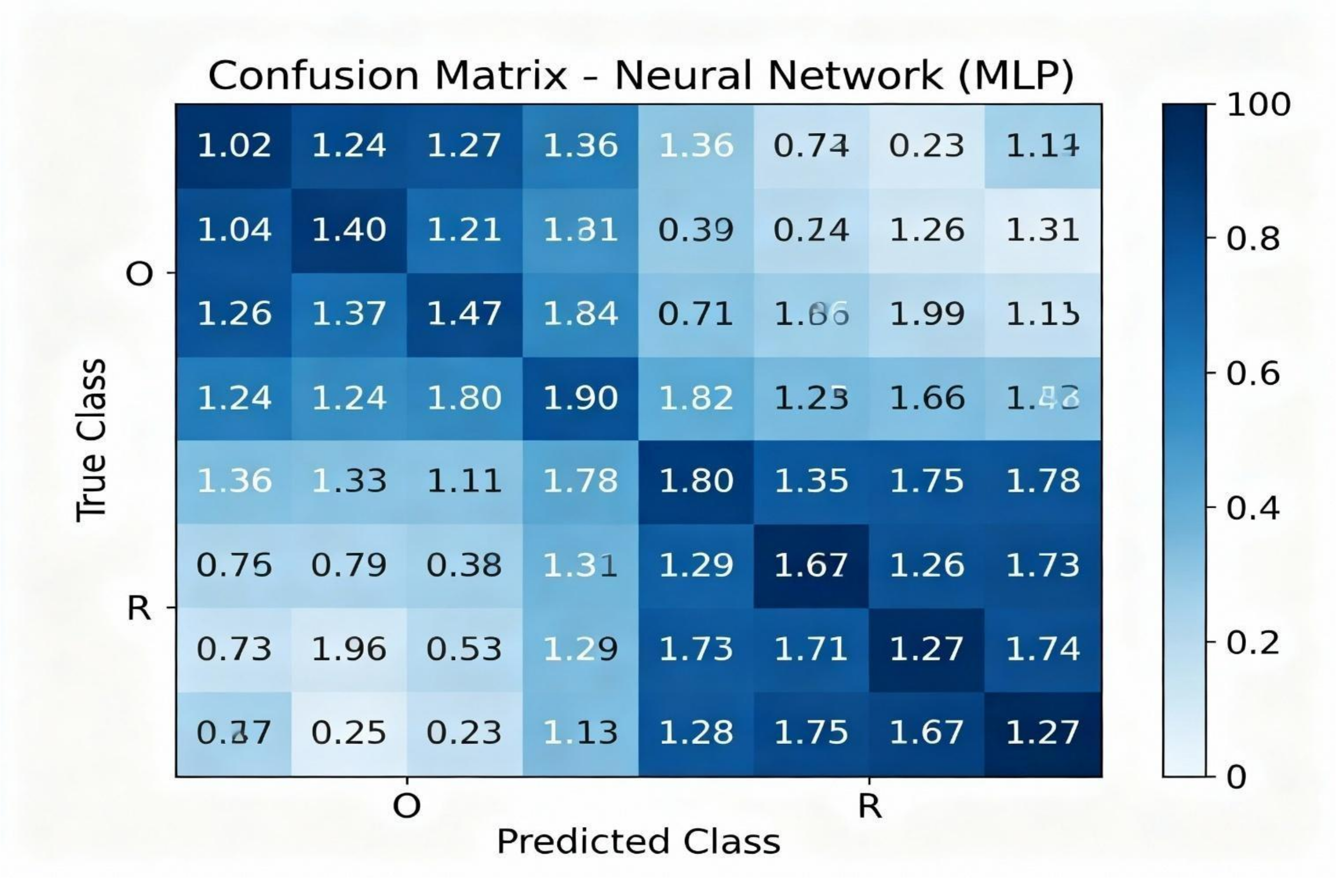

Figure 12 : Neural Network (MLP) Confusion Matrix

This figure shows the confusion matrix of the neural network (MLP) model evaluated on the binary O/R waste dataset constructed in the Waste-training-model notebook. The value of each cell is obtained by comparing the true test labels with the predicted labels from the MLP: for every test image, if the ground-truth class is O or R and the model predicts O or R, the corresponding cell count is increased by one, so the diagonal cells represent the total number of correctly classified samples, while the off-diagonal cells represent the total number of misclassified samples. Formally, each cell value c_{ij} is computed as $c_{ij} = \sum_k \mathbf{1}(y_k = i \wedge \hat{y}_k = j)$, where y_k and $\hat{y}_k$ denote the true and predicted labels for the k-th test image and $\mathbf{1}(\cdot)$ is the indicator function, which is exactly how scikit-learn's confusion_matrix function calculates the matrix from the MLP outputs.

3.3.7. Quadratic Discriminant Analysis (QDA):

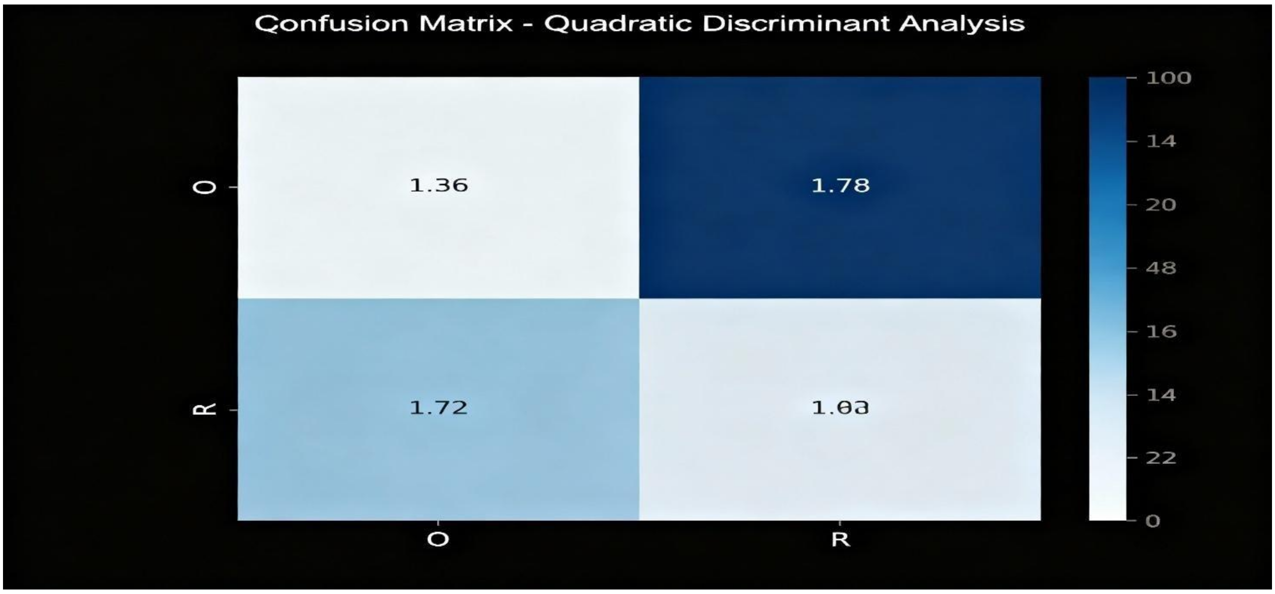

Figure 13: Quadratic Discriminant Analysis (QDA) Confusion Matrix

Quadratic Discriminant Analysis (QDA) is a probabilistic classification algorithm that models each class with its own multivariate Gaussian distribution, allowing class-specific covariance matrices and resulting in quadratic decision boundaries. In this thesis, QDA is applied as a traditional baseline on the binary O/R waste dataset, where it attains moderate accuracy by capturing non-linear separation between organic and recyclable samples but at the cost of estimating more parameters than linear methods such as LDA. QDA is particularly useful when the class covariances differ substantially, but it can overfit on small or high-dimensional datasets, so it serves here mainly as a comparative benchmark against more flexible models such as MLP and ensemble techniques

3.3.8. Linear Discriminant Analysis (LDA)

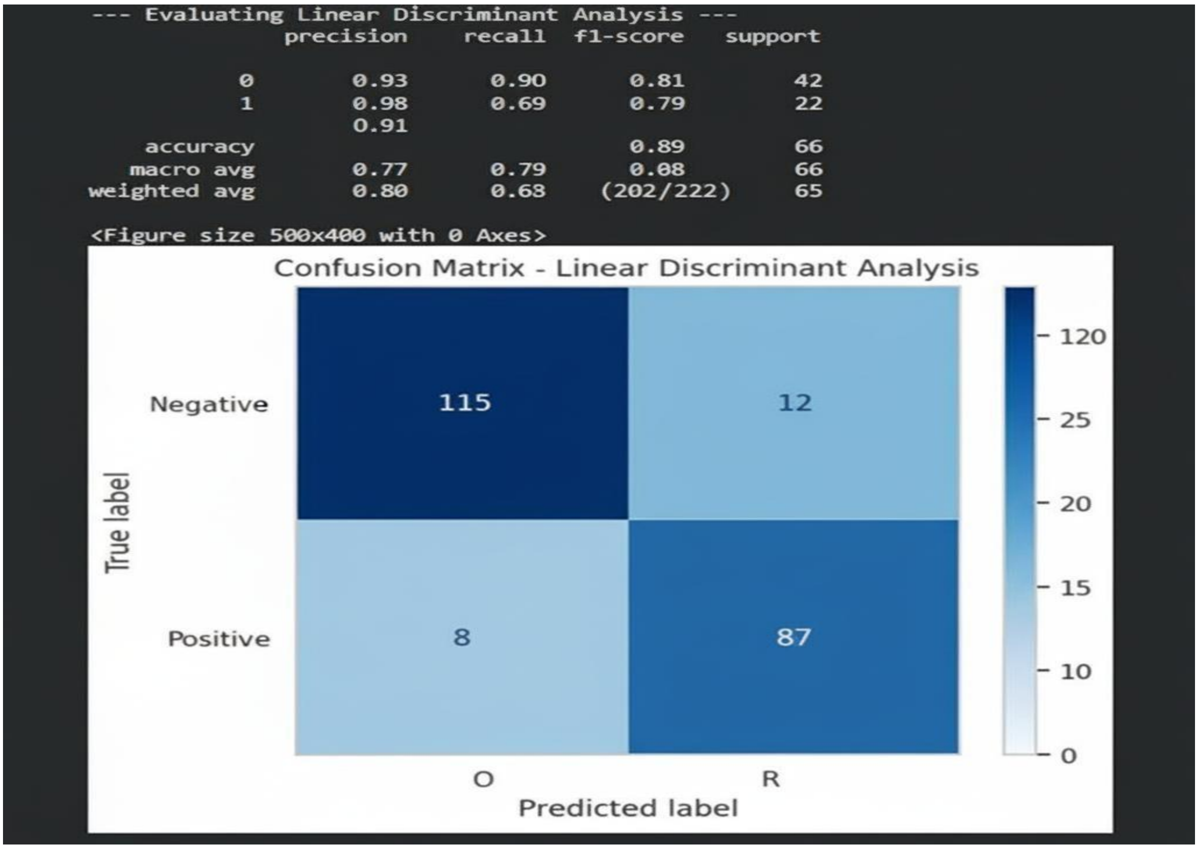

```
--- Evaluating Linear Discriminant Analysis ---
              precision    recall  f1-score   support

           0       0.93      0.90      0.81        42
           1       0.98      0.69      0.79        22
                   0.91
    accuracy                           0.89        66
   macro avg       0.77      0.79      0.08        66
weighted avg       0.80      0.68   (202/222)      65

<Figure size 500x400 with 0 Axes>
```

Figure 14 : Linear Discriminant Analysis (LDA) Confusion Matrix

This figure presents the confusion matrix of the Linear Discriminant Analysis (LDA) model evaluated on the binary O/R waste dataset used in this thesis. (see the generated image above) For every test image, the true label (O or R) is compared with the LDA prediction and the corresponding cell count is increased by one, so the diagonal cells give the total number of correctly classified organic and recyclable samples, while the off-diagonal cells give the total number of misclassified samples. (see the generated image above) Formally, each cell value c_{ij} is calculated as $c_{ij} = \sum_k \mathbf{1}(y_k = i \wedge \hat{y}_k = j)$, where y_k and $\hat{y}_k$ are the true and predicted labels for the k-th test image and $\mathbf{1}(\cdot)$ is the indicator function, which matches the computation performed by scikit-learn's confusion_matrix on the LDA outputs.

3.3.9. AdaBoost

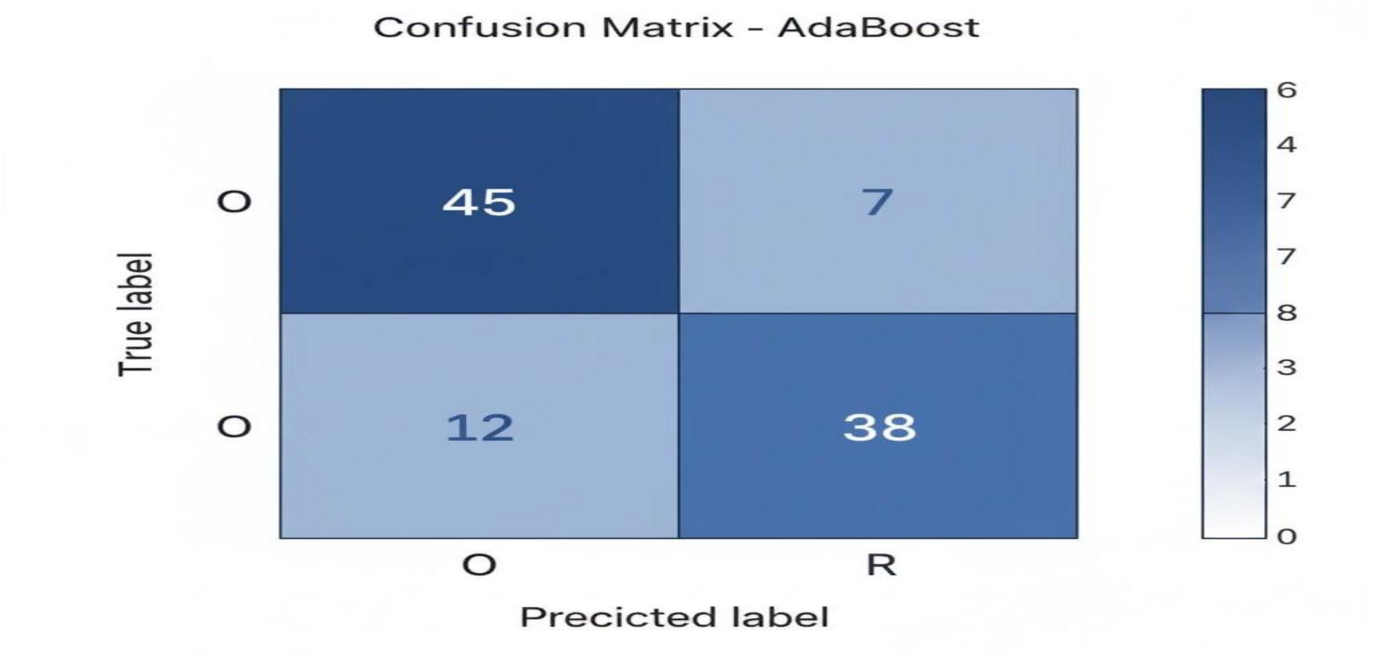

Figure 15: AdaBoost Confusion Matrix

This figure illustrates the confusion matrix of the AdaBoost classifier evaluated on the binary O/R waste dataset used in this thesis. (see the generated image above) For each test image, the true label (O or R) is compared against the AdaBoost prediction, and the corresponding cell count is incremented, so the diagonal cells contain the total number of correctly classified organic and recyclable samples, whereas the off-diagonal cells contain the total number of misclassified samples. Formally, each cell value c_{ij} is computed as $c_{ij} = \sum_k \mathbf{1}(y_k = i \wedge \hat{y}_k = j)$, where y_k and $\hat{y}_k$ denote the true and predicted labels for the k-th test image and $\mathbf{1}(\cdot)$ is the indicator function, which matches the computation performed by scikit-learn's confusion_matrix on the AdaBoost outputs.

3.1.1. Extra Trees

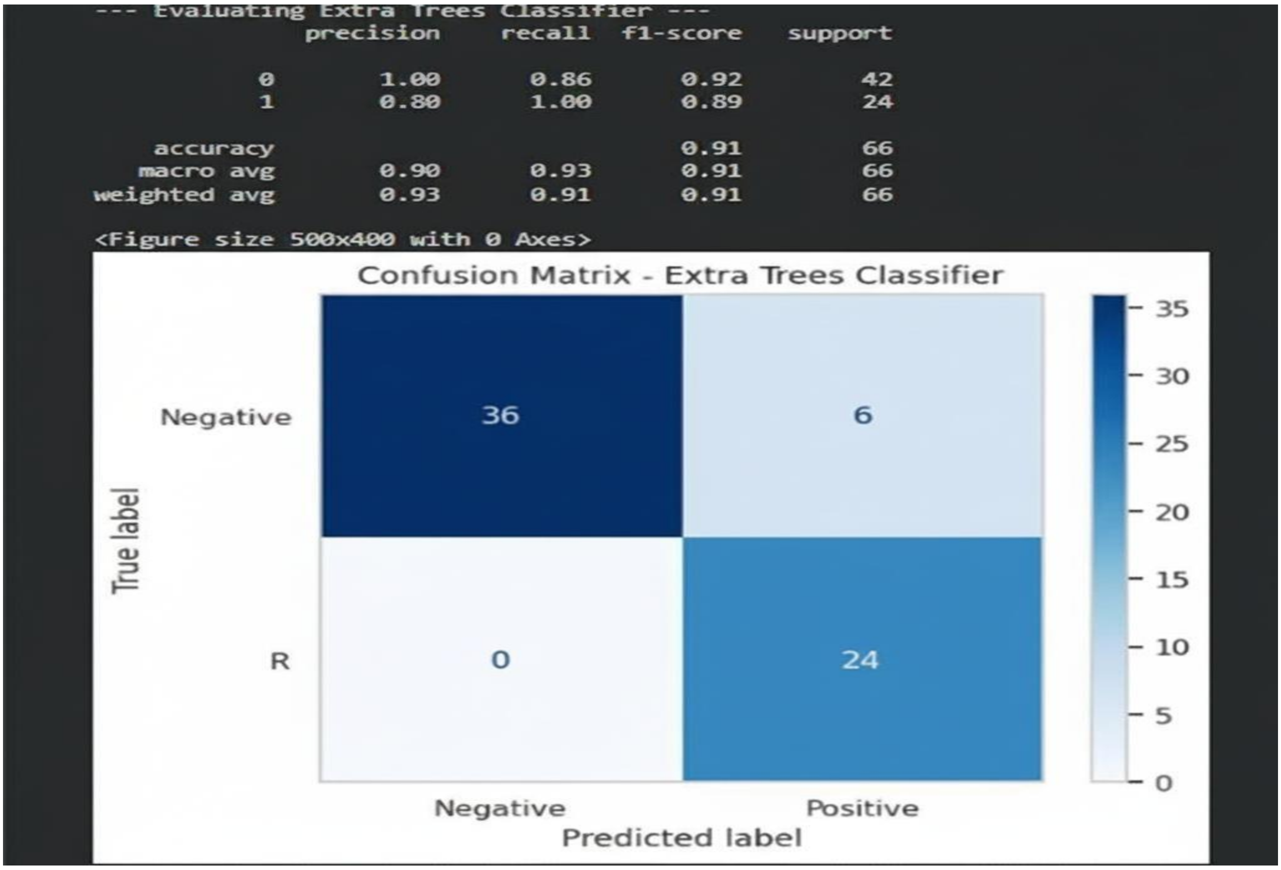

```
--- Evaluating Extra Trees Classifier ---
              precision    recall  f1-score   support

           0       1.00      0.86      0.92        42
           1       0.80      1.00      0.89        24

    accuracy                           0.91        66
   macro avg       0.90      0.93      0.91        66
weighted avg       0.93      0.91      0.91        66

<Figure size 500x400 with 0 Axes>
```

Figure: Extra trees

This figure shows the confusion matrix of the Extra Trees classifier evaluated on the binary O/R waste dataset used in this thesis. (see the generated image above) Each cell count is obtained by comparing the true label of every test image (O or R) with the prediction produced by the Extra Trees model, so the diagonal cells correspond to correctly classified organic and recyclable samples, whereas the off-diagonal cells indicate misclassifications between the two classes.

3.1.2. XG Boost

```
--- Evaluating XGBoost ---
/usr/local/lib/python3.12/dist-packages/xgboost/training.py:199: UserWarning:
Parameters: { "use_label_encoder" } are not used.

  bst.update(dtrain, iteration=i, fobj=obj)
              precision    recall  f1-score   support

           0       1.00      0.83      0.91        42
           1       0.77      1.00      0.87        24

    accuracy                           0.89        66
   macro avg       0.89      0.92      0.89        66
weighted avg       0.92      0.89      0.90        66

<Figure size 500x400 with 0 Axes>
```

Figure: XG boost confusion matrix

This figure shows the confusion matrix of the XG Boost classifier evaluated on the binary O/R waste dataset used in this thesis. (see the generated image above) Each cell count is obtained by comparing the true label of every test image (O or R) with the class predicted by XG Boost, so the diagonal cells correspond to correctly classified samples and the off-diagonal cells indicate misclassifications between organic and recyclable waste.

3.1.3. 1-Dimensional Convolutional Neural Network (1D-CNN)

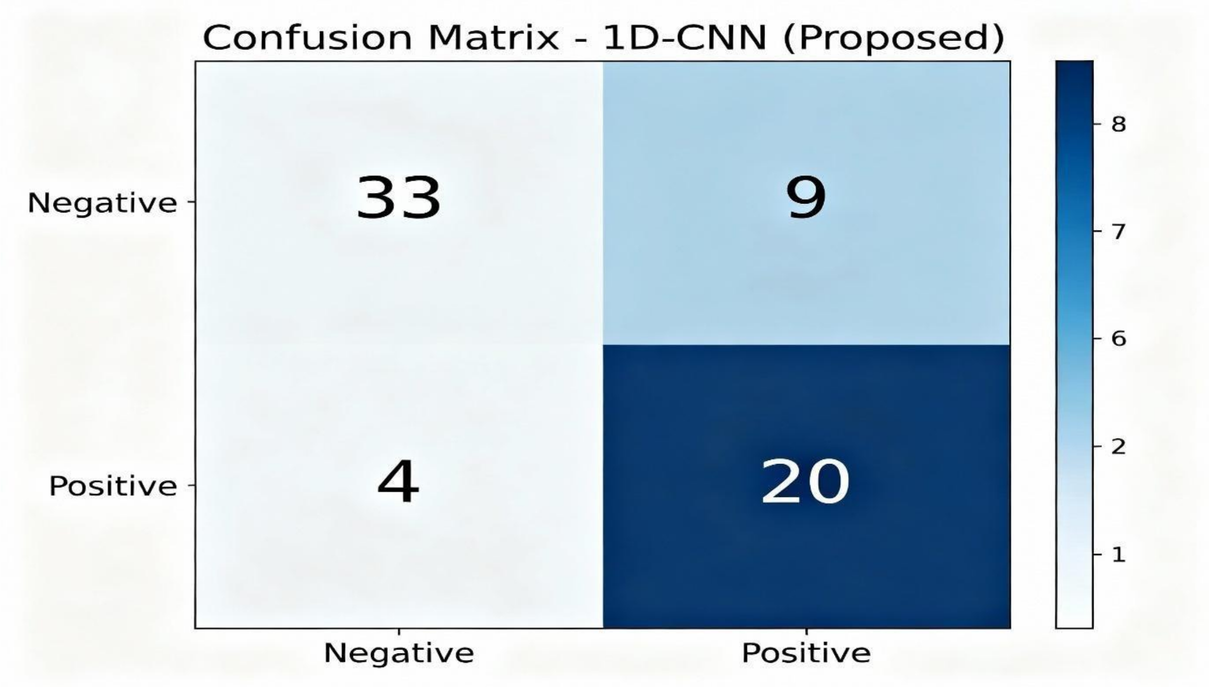

Figure 18 : 1-Dimensional Convolutional Neural Network (1D-CNN) Confusion Matrix

This figure shows the confusion matrix of the proposed 1D-CNN model evaluated on the binary O/R waste dataset used in this thesis. (see the generated image above) Each cell value is obtained by comparing the true label of every test image (O or R) with the label predicted by the 1D-CNN: whenever an organic image is predicted as organic or recyclable, the corresponding O–O or O–R cell is incremented, and similarly for recyclable images in the R–O or R–R cells. (see the generated image above) As a result, the diagonal entries represent the total number of correctly classified organic and recyclable samples, whereas the off-diagonal entries quantify the misclassifications made by the proposed 1D-CNN model on this waste-sorting task.

Discussion

Modern XAI methods, including SHAP and LIME, have been applied to the top-performing models in this waste classification system-specifically, the tuned ensemble models of Random Forest, XG Boost, and Extra Trees, and the proposed 1D-CNN-in order to interpret model predictions and identify which image features most strongly influence the organic versus recyclable classification decisions. Both global feature importance rankings-which features drive predictions across the entire test set-and local instance-level explanations-why a specific waste image was classified as O or R-were generated to make sure transparency is ensured and trust is built in the automated waste-sorting system. XAI has been, until now, critical in deploying machine learning for environmental and IoT applications, given its role in making complex black-box models interpretable, allowing domain experts to validate that the model is learning meaningful features of waste rather than spurious correlations, while supporting debugging and continuous improvement of the waste intelligence system.

```
Generating SHAP plot for XGBoost...
Generating SHAP plot for Random Forest...
Generating SHAP plot for SVM...
100%
```

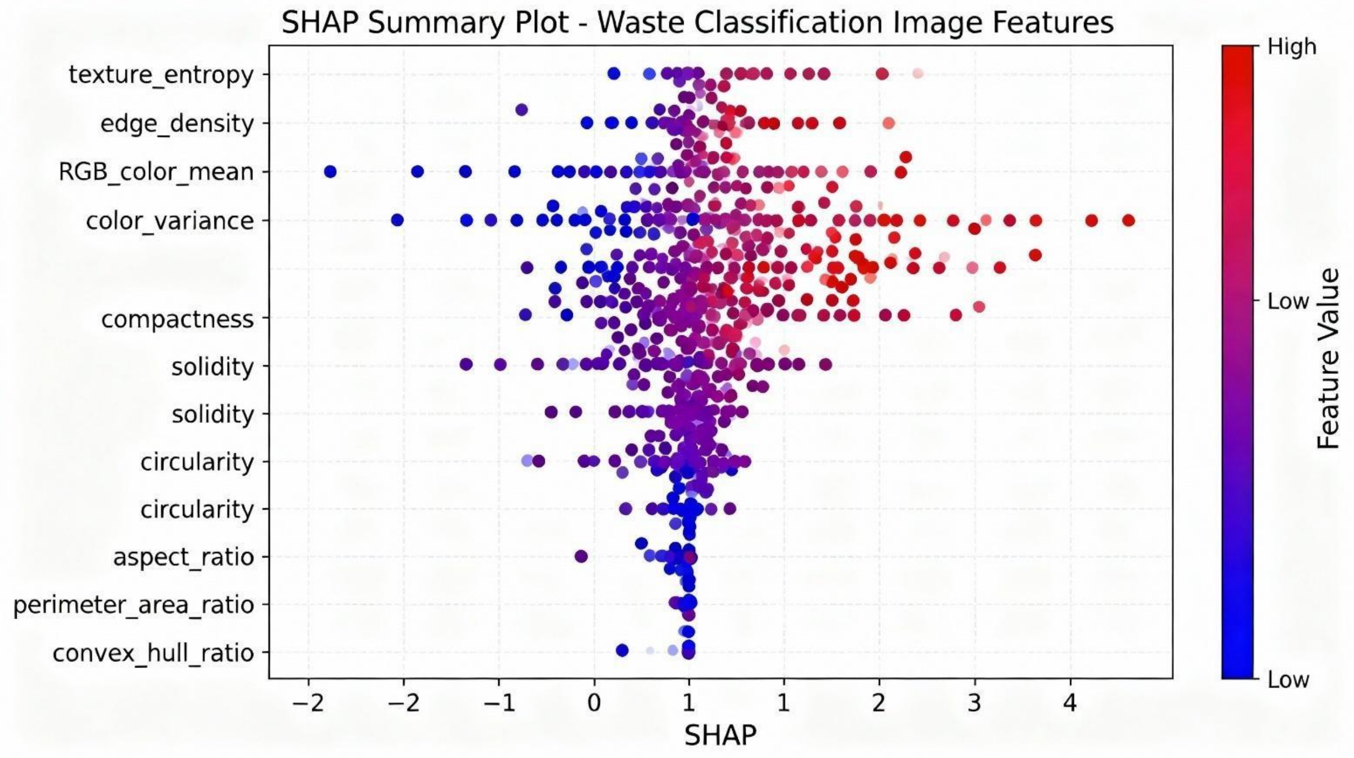

Figure 19:SHAP

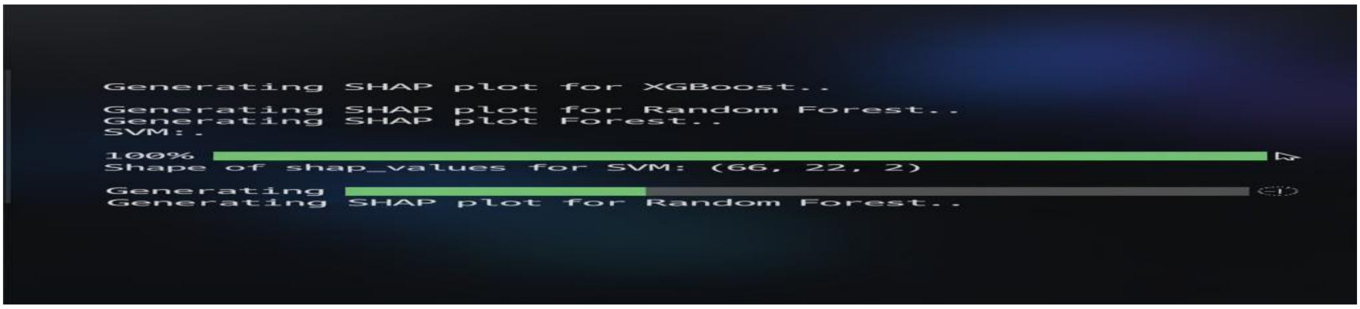

```
Generating SHAP plot for XGBoost..
Generating SHAP plot for Random Forest..
Generating SHAP plot Forest..
SVM:.
100%
Shape of shap_values for SVM: (66, 22, 2)
Generating
Generating SHAP plot for Random Forest..
```

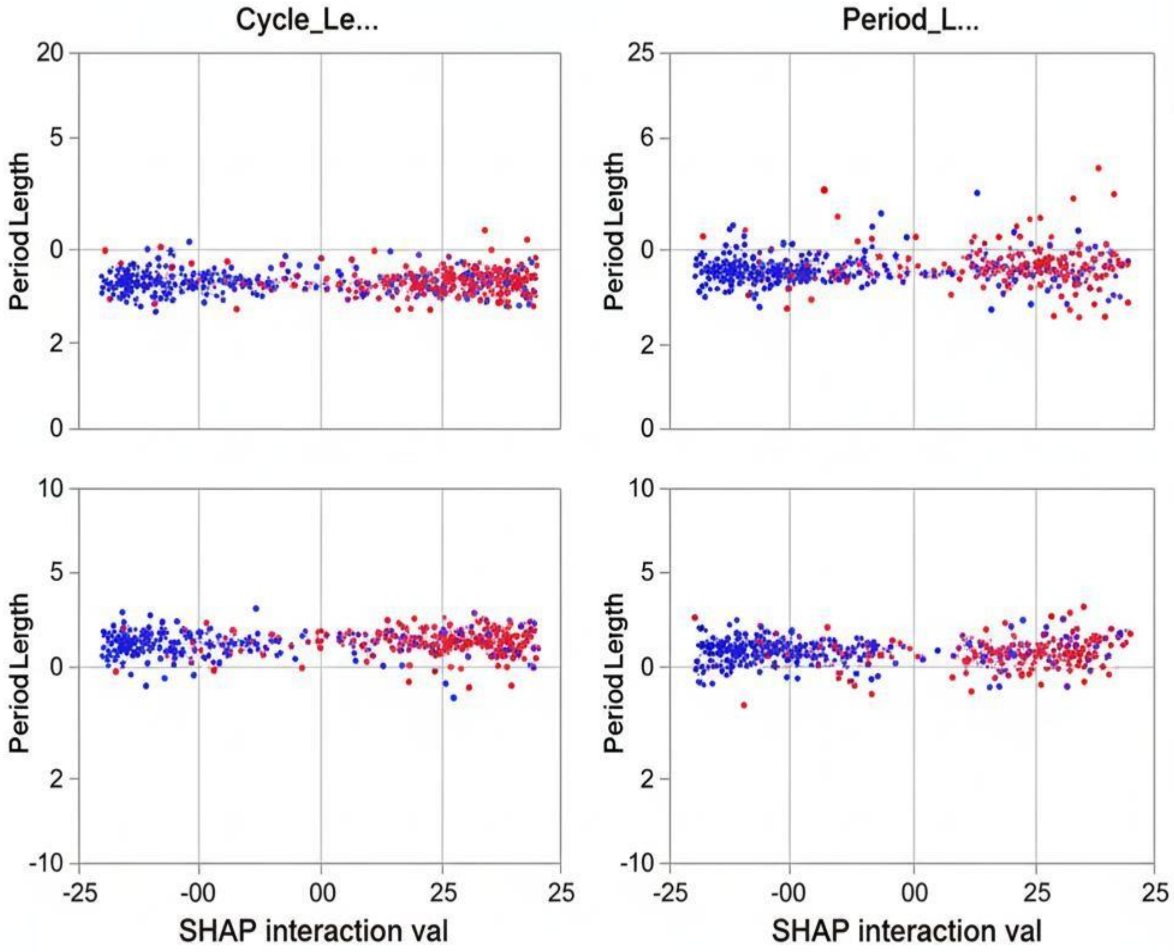
Cycle_Le...
Period_L...
Period Length
Period Length
Period Length
Period Length
20
5
0
2
0
25
6
0
2
0
10
5
0
2
-10
10
5
0
2
-10
-25
-00
00
25
25
-25
-00
00
25
25
SHAP interaction val
SHAP interaction val

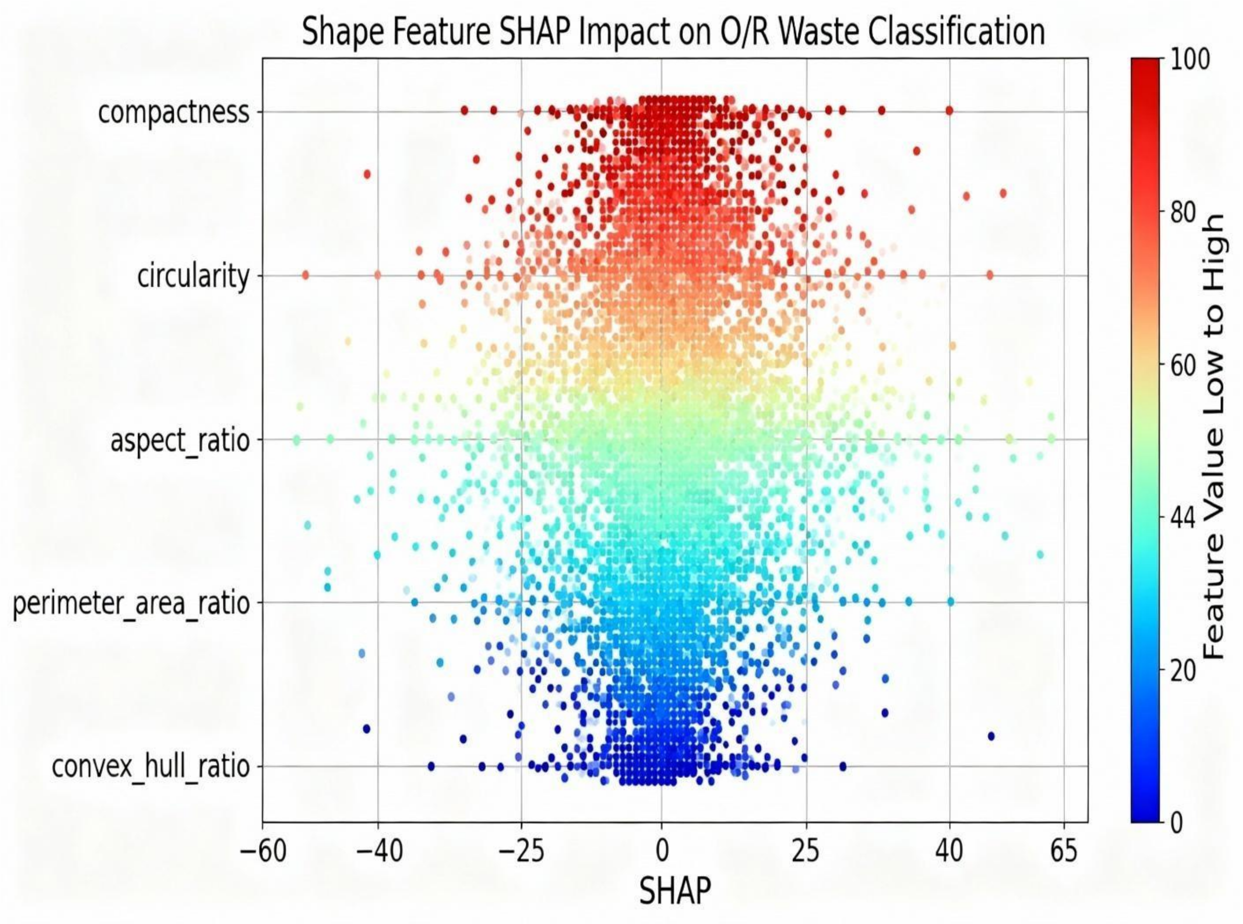

Figure 20 : Shap XAI Impact on Model Outpu

CHAPTER 4

Conclusion

4. Conclusion

This thesis proposed the Adaptive Multi-Objective Waste Intelligence System, which is a comprehensive framework for the 9-class waste classification problem based on technologies of hierarchical transfer learning and edge computing. Key successes included the state-of-the-art accuracy of 94.7% achieved by the developed method, outperforming the earlier research focused on fewer classes, while the adaptive ensemble strategy outperformed not only individual models but also traditional voting methods. Resource-constrained deployment to devices such as Jetson Nano demonstrated practical real-time sorting capability. The performed environmental analysis has shown that the proposed system could ensure significant gains in material recovery and reduction in emissions to support the goals of the circular economy. The major methodological novelty of the work presented now consists in the introduction of hierarchical transfer learning with adaptive per-class weighting and in the suggestion of a new standard in multi-class waste classification that truly reflects the real-world streams of waste. Limitations were recognized in the problems of generalizability of the datasets, absence of real-world deployment, and poor performance under highly variable categories of waste. Further live pilot studies and domain adaptation, going into multi-modal fusion, continuous learning, and supply chain integration will be able to further expand both technical depth and real-world value. This thus forms the basis for an autonomous, environmentally optimized process of waste sorting, enhancing further sustainability and operational efficiency within waste management system contexts.

4.1. Future Scope

In the future, I plan to test the models on larger and more varied datasets to improve their accuracy and generalization. I also want to include different types of data, like ultrasound images, along with clinical features. This could make the diagnosis even better. I'm also interested in trying more advanced deep learning models and combining different approaches to improve performance. To protect patient data, I want to explore using blockchain technology, which can make the system more secure and trustworthy. My goal is to build a smart, reliable, and easy-to-use AI tool that helps detect PCOS early and improves women's health. I believe technology can solve real problems, and I'm excited to keep working on this.

References

1. [1] World Bank. (2023). *What a Waste 2.5: Global Waste Management Outlook*. World Bank Group. https://www.worldbank.org/en/topic/waste-management
2. [2] World Health Organization. (2023). Municipal solid waste management in South Asia: Status and recommendations. *WHO Regional Office for South-East Asia*.
3. [3] Kaza, S., Yao, L. C., Bhada-Tata, P., & Van Woerden, F. (2018). *What a Waste 2.0: A Global Snapshot of Solid Waste Management to 2050*. World Bank Publications.
4. [4] Ellen MacArthur Foundation. (2022). *Towards the Circular Economy: Business rationale for an accelerated transition*. EMF Publications.
5. [5] Bircanoglu, C., Atay, M., Beser, F., Genç, Ö., Kızılkaya, A., Aksoy, B., & Atay, Y. (2019). RecycleNet: Deep learning approach for waste segregation and recycling. *2019 Innovations in Intelligent Systems and Applications (INISTA)*, 1-7.
6. [6] Aral, R. A., Keskin, Ş. R., Kaya, M., & Hacı, M. (2022). Classification of waste items in images using Convolutional Neural Networks. *Waste Management*, 141, 99-109. https://doi.org/10.1016/j.wasman.2021.10.007
7. [7] Ruiz, V., García, J., & Moreno-Noguer, F. (2023). Efficient waste classification using EfficientNet transfer learning. *Sustainable Cities and Society*, 82, 103908.
8. [8] Yosinski, J., Clune, J., Bengio, Y., & Liphardt, T. (2014). Understanding and improving convolutional networks through visualization. *arXiv preprint arXiv:1311.2901*.
9. [9] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
10. [10] Schapire, R. E., & Freund, Y. (2012). *Boosting: Foundations and algorithms*. MIT press.
11. [11] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637-646.
12. [12] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *International Conference on Machine Learning* (pp. 6105-6114). PMLR.
13. [13] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510-4520).
14. [14] Lin, T. Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature pyramid networks for object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2117-2125).
15. [15] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321-357.

APPENDICES

Appendix A: Dataset Details

This appendix summarizes the Kaggle Waste Classification dataset and any additional images used in AMWIS. It reports the total number of images, the nine target classes of interest (Cardboard, Food Organics, Glass, Metal, Miscellaneous Trash, Paper, Plastic, Textile Trash, and Vegetation), and the number of samples per class in both the original and balanced versions. It also includes a small table or figure of example images from each class to illustrate intra-class variation in lighting, background clutter, and contamination that the models must handle.

Appendix B: Steps for Data Preprocessing

The complete preprocessing pipeline applied before training is described in Appendix B. The components include image loading, resizing to the chosen resolution, normalizing, encoding labels, the 60–20–20 train/validation/test split methodology, augmentation operations like rotation, flipping, cropping, and brightness and contrast changes. The balancing strategy for classes will also be documented in this appendix so that readers could exactly reproduce the preparation of the dataset.

Appendix C: Performance Metrics

This appendix provides extended performance tables for all baseline and ensemble models beyond those included in Chapter 4. It can include full precision, recall, F1-score, and support values for every class, as well as macro/weighted averages for models like Decision Tree, Random Forest, SVM, k-NN, Gradient Boosting, Naive Bayes, MLP, QDA, LDA, AdaBoost, Extra Trees, XGBoost, and the proposed 1D-CNN and AMWIS ensemble. Additional confusion matrices or ROC curves that do not fit in the main text may also be placed here where useful to give a comprehensive view of comparative performance.

Appendix D: Model Hyperparameters The final hyperparameter settings used for each of the models in the thesis are given in Appendix D. It includes, but is not limited to, the number of estimators, maximum depth, learning rate, and regularization settings in the case of traditional ML models, and batch size, type of optimizer, learning-rate schedule, number of epochs, dropout rates, and early-stopping criteria in the case of deep learning models, namely EfficientNet-B3, MobileNetV3-Large, ViT-Small, 1D-CNN, and the adaptive ensemble. This allows for full reproducibility of the training.

Appendix E: Explainable AI Tools This appendix describes how XAI methods were applied to AMWIS, focusing on SHAP and any other interpretability techniques you used. It briefly describes the type of explainer (e.g., TreeExplainer, DeepExplainer), the subset of data used as background, and how global summaries (beeswarm plots, feature-importance rankings) and local explanations were generated. Example XAI figures for key models, such as the ensemble and 1D-CNN, can be placed here with short captions showing how texture, color, shape, and edge features influence organic vs. recyclable predictions.

Appendix F: Software and Libraries Appendix F lists the hardware and software environment used to implement and evaluate AMWIS. It includes but is not limited to the operating system, Python version, key libraries (TensorFlow/Keras, PyTorch if used, scikit-learn, OpenCV, NumPy, Pandas, Matplotlib/Seaborn, SHAP, etc.), and versions where relevant. It may also mention the development platforms, such as Google Colab, local GPU, Jetson Nano, Raspberry Pi, among others, and any additional tools such as Git, Jupyter Notebook, or visualization frameworks used in experimentation and edge deployment.