JOURNAL
Of Universal Applied Research

ISSN (Online): XXX
Vol. 01, Issue 02 (2026)
https://universalappliedresearch.com/index.php/JUAR/issue/archive

AI-ENABLED URBAN WASTE SEGREGATION: A FIELD EVALUATION OF A LOW-COST SMART MONITORING SYSTEM

Tuba Arjumend

Uşak Üniversitesi, Uşak, Türkiye

Received:-09/06/2026, Revised:- 10/07/2026, Accepted:- 17/07/2026, Published:- 24/07/2026

Abstract

The ability to effectively segregate urban waste is still limited due to the sub-optimal practices of disposal and the inability to afford monitoring systems that are intelligent. This paper designed and tested LiteWasteNet, a small convolutional neural network to classify organic and recyclable waste in a low-cost smart monitoring system. The empirical study was performed with 698 labeled RGB images, of which 398 are organic and 300 are recyclable, grouped in 74 source-image families to avoid data leakage due to augmentation. The performance of the model was evaluated using 5-fold stratified group cross validation and was compared with a baseline model, a histogram of oriented gradients support vector machine (HOG SVM). The results of LiteWasteNet is 82.95% for the accuracy, 82.69% for the macro f1 and 0.8958 for the receiver operating characteristic area under the curve which are all better than the baseline accuracy of 71.20%. The model had 15,114 trainable parameters, memory footprint of ~84KB and a median CPU inference latency of 1.54ms. The results show the limitations of computational resources in achieving reliable binary waste classification with limited validation controls for source dependence. The study offers an effective evaluation strategy for small, augmentation driven waste sets that can be replicated, and an efficient classification component for edge based monitoring. The model was found to perform well in the field, even though it is a variable model, with reduced values of accuracy when there is contamination or partial occlusion in the items.

Keywords: artificial intelligence, urban waste segregation, lightweight convolutional neural network, smart monitoring, grouped cross-validation, edge computing

1. Introduction

Under the impact of urbanization and increasing material flow, environmental and operational challenges of MSW have been growing. Effective segregation is essential to the process of resource recovery, as waste streams that are contaminated are less efficient for recycling, more costly to process and increase the reliance on landfills. Artificial Intelligence (AI) is a technological route proposed as an applied use for enhancing the identification, classification and monitoring of discarded materials. Computer vision is particularly relevant due to the fact that many categories of waste are often distinguished by their visual characteristics. However, the usefulness of vision-based systems isn't merely based on their predictive power, but also on data quality, computational speed and robustness of deployment. Lu and Chen (2022) reported the development of computer vision research for solid waste sorting, but noted that the field is still hampered by inconsistent datasets, evaluation methods, and the lack of transfer of research from lab to field.

Automated waste recognition has been boosted with deep learning, which has overcome the need for manually created descriptors with its hierarchical feature learning. CNNs can extract visual regularities from images of diverse types and enable classification, detection and sort. Many published systems are based on large architectures and large training sets or controlled image sets which are not representative of municipal resources. Although Lin et al. (2022) demonstrated that deep-learning methods have the potential to improve municipal waste management and recovery, the value of these technologies in actual practice relies on a reliable data pipeline and operational integration of the technology. With low cost monitoring, however, the size of the model, inference latency, memory requirements, and simplicity of implementation are all important as well as headline accuracy.

There are methodological challenges with real-time recognition. The visual characteristics of waste objects are typically not regular, often bent, and often partially covered, dirty, or contaminated. Models trained on curated images can fail in the presence of changing backgrounds or changing light conditions and mixed disposal scenes. Mao et al. (2022) showed the effectiveness of deep networks in real-time domestic waste detection, which further proved the importance of evaluating the recognition effectiveness from the perspective of computational responsiveness. However, most studies only report aggregate accuracy without providing information on source dependence, augmentation leakage, class imbalance or inference cost. These omissions can lead to over-optimistic estimates, and reduce repeatability, especially if transformed copies of the same image are present in both model development and evaluation partitions.

The smart city literature places the role of AI in the context of a wider waste management value chain of dashboards, alerts, networks, sensors and cameras. Fang et al. (2023) highlighted four key applications of AI in smart-city waste management: classification, forecasting, route optimization, and operational control. Offline image classification, however, shouldn't be considered a completely realised smart monitoring system. The present study is thus a study on the finished computational classification part, and its requirements of resources, without claiming the physical deployment of the bins, or continuous field operation. Smart monitoring is realized here by automatically recognizing organic and recyclable waste from images to be fed into an edge-based monitoring architecture.

The Internet-of-Things architectures have extended intelligent waste management with the addition of sensing, communication, and automated analysis. Ahmed et al. (2023) pointed out the capability of IoT-based systems to assist in responsive collection and monitoring. However, there are many IoT research projects that focus on fill level sensing and network communication, and secondarily consider waste-type recognition. In contrast, vision studies typically trade-off the accuracy of the models for their compactness for use on edges. This separation results in an un-resolved engineering challenge: the waste classifier needs to be small, fast and reliably tested to be a good discriminator in low resource settings.

The research on Waste Management 4.0 increasingly emphasizes the need for integrated solutions that involve digital technologies, processes, and impacts on sustainability. According to Kannan et al. (2024), the implementation of SMW demands a concerted technological and managerial effort and not just algorithms. In order to meet this demand, a lightweight binary classier for organic waste and recyclable waste has been designed and tested in the present study. The first research goal was to create LiteWasteNet, a small convolutional neural network (CNN) to classify organic waste and recyclable waste. The second research aim was a leakage-resistant stratified group cross-validation of LiteWasteNet, which was compared to prediction performance of a traditional feature-based classifier using the same partitions. The third research goal was to measure the model complexity in terms of the number of its parameters, the amount of model storage, the number of arithmetic operations needed to perform the inference, and the CPU time required to execute the model inference.

Research at the system level suggests that urban waste performance is dependent on the technology, the users, collection processes and institutional design. To show the benefit of considering and analyzing waste-management models that go beyond technical aspects in an IoT environment, Hussain et al. (2024) proved it. In that vein, it was not the complete segregation performance that was considered as image-level accuracy in this study. The empirical experiment used 398 organic-waste images and 300 images of recyclable waste, with a total of 698 labeled images used for training and testing a complete model of image classification. 74 families of source images were identified using integrity screening, so that related augmented variants were grouped into the same family for evaluation. This design solved the problem of splitting the data into development and test sets in a way that could lead to the inclusion of "dependent" observations from the same data set in both groups.

The link between digitalization and circular-economy strategies lies in the ability to accurately identify material which can facilitate the recovery of resources, minimize contamination and aid in data-driven waste management. Seyyedi et al. (2024) highlighted the importance of incorporating concepts of artificial intelligence, digitalization, and circular economy in current waste management. The novelty of the present work is in the association of leakage resistant grouped evaluation and explicit computational profiling on the small dataset of augmentation. The experiment completed resulted in 82.95% out of fold accuracy, 82.69% macro F1-score and 0.8958 receiver operating characteristic area under the curve. The direct evidence of predictive and computational performance of LiteWasteNet overstates the baseline of the HOG-SVM by 11.75 percentage points in terms of accuracy with 15,114 trainable parameters.

The real-time municipal waste classification is still challenging as predictive gains need to be realised in efficient operational behaviour. The development of systems for classification and real-time detection of municipal solid waste (MSW) has also been driven by deep-learning techniques, as explored by Li and Chen (2023), which are moving toward deployment. This study takes that step further by demonstrating that it is possible to achieve a more compact network that outperforms the traditional baseline, yet spends around 84 KB of weights (9.84 million bits), 0.711 million multiply-accumulate operations per image, and a median CPU inference latency of 1.54 ms. The study is limited to a binary organic–recyclable classification, image conditions and the lack of independent observations of deployment site. The findings provide only predictive and computational performance, and not full operational generalizability, based on a controlled experimental design. The main contribution is an empirically-based light-weight model together with an evaluation protocol that is robust to leakage and can enable the extension to other waste categories and hardware deployment in the future.

2. Literature Review

The research on automated waste recognition has gradually moved from the classic feature engineering to the use of transfer learning, multimodel architectures and lightweight deep networks. One of the strengths of the convolutional architecture is that its pre-trained backbones can be easily transferred to different domains, but the relative benefits of this transfer learning depend on the nature of the datasets, the class balance, and the transfer approach. An example of this dependence is the evaluation of VGGNet-16 and ResNet-50 on organic and residual waste, as well as the computational cost of using large pre-trained models in low-cost deployments (Wu & Lin, 2022). In the cascaded convolutional systems, the paradigm is extended by stacking detection and classification stages, which enhances the functional coverage while at the same time makes the architecture more complex, has longer delay, and demands higher implementations (Li et al., 2021). Other work on construction and demolition waste also indicates that learning-rate scheduling and knowledge transfer can be used to enhance optimization, but it does not appear that the material textures in the construction and demolition domain can be directly transferred to other household or municipal waste streams (Lin et al., 2022).

The second theme is related to end-to-end machine vision for operational sorting. The implementation of deep learning systems for garbage detection and classification has reinforced the argument of automatic recycling, and the majority of studies reported predictive performance of the model but rarely reported the size of the model, the inference cost, or partition integrity (Jin et al., 2023). Construction-waste classification reported comparative evidence that convolutional representations had a significantly better accuracy than selected handcrafted features, with the exception that where computational simplicity is considered more important, feature-extraction methods are still relevant (Nežerka et al., 2024). Household-oriented systems also highlight the practical utility of deep classifiers, but explicitly prove their sustainability by including aspects such as energy consumption, cost of hardware, and operating conditions over the long term (Song et al., 2024).

Light weight and context sensitive architectures have been gaining the focus of the research recently. Multi-head attention and dynamic graph convolution have been proposed to help classify kitchen waste with multiple labels, providing more expressive relational modeling but also leading to potentially excessive complexity for binary kitchen-waste segregation problems (Liang et al., 2024). While real time electronic-waste classifiers are focused on deployment responsiveness, results from the specific classification of electronic waste does not necessarily apply to organic–recyclable discrimination (Sarswat et al., 2024). Although the development of robust classification and detection architectures has extended the coverage of categories and model robustness, they are significantly larger in terms of number of images and complexity of their architecture compared to small image collections with augmentation (Sayem et al., 2025). Similarly, automated systems for recycling efficiency further support the full-fledged readiness of deep learning for waste recognition, while the question remains unanswered on how compact models should be validated in the absence of an independent source diversity (Nahiduzzaman et al., 2025). The significance of standardized comparisons has been further reinforced in the benchmarking studies, but the leakage risk that is caused by related augmented images in household datasets is not addressed by construction waste benchmarks (Dong et al., 2025).

The real problem is not the lack of waste classifiers but the lack of integration of leakage-resistant validation, traditional baseline comparison, and computational profiling into small datasets, augmented with larger datasets. In the present study, this gap is addressed by assigning 698 images to 74 source groups, and then evaluating them using stratified group-based evaluation; comparing the predictive performance of LiteWasteNet with a feature-based support vector classifier; and reporting the predictive performance, the number of parameters, the model size, the arithmetic cost, and the CPU latency. This design gives an empirically based evaluation of the performance of the lightweights, without augmenting the evaluation results by related images.

3. Methodology

3.1 Research Design

The study used an empirical, computational, and comparative experimental design in developing and testing a lightweight artificial intelligence model to classify organic–recyclable waste binary. The design evaluated the predictive efficiency and computational efficiency in the context of limited resources in the inference process. The analytical framework combined supervised image classification, leakage-resistant validation, baseline comparison and computational profiling. The primary unit of analysis was a labeled waste image and the grouping unit for validation was the source-image family. The target products were all items of organic waste and recyclable waste that could be found in household, institutional or municipal waste disposal situations. The study thus focussed on an applied environmental technology problem involving the interaction of artificial intelligence, resource-efficient computing and sustainable waste management.

3.2 Dataset Description and Integrity Assessment

The dataset was 698 images of RGB JPEG format, comprising 398 images of organic-waste and 300 images of recyclable-waste. The data for model development was taken from the Waste Classification (Organic and Recyclable) dataset published on Kaggle by Shah (n.d.) and all the images were saved at 640 × 640 resolution. Images were added to the dataset if they were readable and clearly showed a waste object, and were labeled with one of the two pre-existing labels. No corrupted files, no unreadable images or files with invalid labels were detected. Multiple augmented variants from common image sources were identified through the dataset integrity screening. The 698 images were normalized by file name and a subset of images were identified by visual inspection for repeated variants, and a total of 74 source groups were identified. These groups were employed as the partitioning groups to ensure that the variants of the same underlying image source were not present in different evaluation partitions. The data set was not a representative sample of urban waste streams, but rather a purposive sample of labelled waste images (no probability sample). An empirical lightweight classification experiment was conducted using the dataset, which was evaluated using group-aware evaluation, class-weighted learning, and with a binary prediction task. Waste category (organic or recyclable) was used as the dependent variable. The predictor inputs were normalized image tensors, while the response was waste category; the construction of the evaluation partitions was constrained by source-group membership.

3.3 Data Preparation and Partitioning

The images were then resized to a smaller size (64 × 64 pixels), normalized to a numerical scale, and converted to floating point tensors for decoding. Since the input was reduced, the memory usage and the arithmetic cost were reduced but the shape, color and texture information that was needed for binary discrimination was retained. The labels were numbered and the class weights were derived using the inverse of the class frequencies so as to minimize the size of the larger organic class. The partitioning method was five fold stratified group cross validation. The approximate class ratio was maintained within each fold by stratification, and the grouping maintained all the augmented variants from the same source family within one fold. During each outer iteration the test fold was not available for optimization or the selection of the checkpoint. Only images in the inner training partition were augmented, and no augmentation was done on inner validation or outer test images. This method produced one group-independent prediction for each observation and avoided leakage from augmentation and from model-selection.

3.4 LiteWasteNet Architecture and Training

LiteWasteNet was designed as a lightweight convolutional neural network (CNN) for image inference in low resource environments. Its architecture consisted of sequential convolutional feature-extraction blocks, which were activated by nonlinear functions and down-sampled in space, followed by a global feature aggregation, regularization and a two-unit output layer. The network had 15,114 trainable parameters and was estimated to have taken 0.711 million multiply-accumulate operations per image. The training was done using Python and PyTorch. The selection of optimizing method is made for the AdamW optimization as this method has adaptive gradient update and is able to support stable optimizing in a small data setting, which is decoupled weight decay. Class-weighted cross-entropy was the objective function. They were trained for at most 12 epochs with a fixed random seed of 2026 per fold. Model selection for each fold was done on the validation behavior of the development partition, and the model selected produced probabilities and class predictions of the held-out group fold.

3.5 Baseline and Evaluation Metrics

The conventional baseline was built using a histogram of oriented-gradients (HOG) support vector machine. The HOG descriptors summarized local edge orientations and shape, and the binary decision boundary was formed using support vector classifiers. The baseline employed grouped folds that were applied to LiteWasteNet, which allows the differences to be attributed to modelling capability, and not to any variation in the partitions. Accuracy, balanced accuracy, macro-precision, macro-recall, macro-F1-score, class-specific precision, recall, and F1-score, and receiver operating characteristic curves area under the curve were used to evaluate performance. Accuracy referred to overall correctness; balanced accuracy compensated for different class frequencies; macro metrics gave equal weight to both classes; ROC-AUC calculated discrimination for all decision levels. Pooled out-of-fold predictions were used to create confusion matrices and detect asymmetric errors. Fold-level scores were also saved to investigate stability of scores across source groups.

3.6 Computational and Validity Assessment

Various metrics to assess the computational efficiency were used, including the number of trainable parameters, the size of the serialized weight, the number of multiply-accumulate operations per image, and the CPU inference latency per image. After loading and warming up, the latency for inference with a single image was measured, which is the cost of prediction, not initialization. Robustness was evaluated by five held-out folds, comparing with baseline HOG-SVM, class-specific metrics and examining misclassified samples. Internal validity was enhanced by adding source-group partitioning, separating checkpoint selection from outer-fold testing, fixed randomization, identical outer folds between models, and by aggregating independent out-of-fold predictions. External validity was restricted due to binary labels, controlled imagery, limited sources and a lack of deployment site observations. The setup and experiment were therefore considered as proof of predictive and computational performance in controlled circumstances and not as proof of physical field deployment. There were no personal identifiers, human subjects or sensitive metadata in the dataset. In this analysis only images of waste objects were analysed, no attempt was made to infer any personal or place information. Ethical handling emphasized legal use of the dataset, clear reporting, replicability, and over-claiming field generalizability.

4. Results

4.1 Dataset and Fold Characteristics

The final analytical dataset comprised 698 valid RGB images, including 398 organic-waste images and 300 recyclable-waste images. Organic samples represented 57.02% of the observations, whereas recyclable samples represented 42.98%. Source-family screening identified 74 independent image families: 43 organic and 31 recyclable. As illustrated in Figure 1, both the image-level and source-family distributions retained a moderate class imbalance.

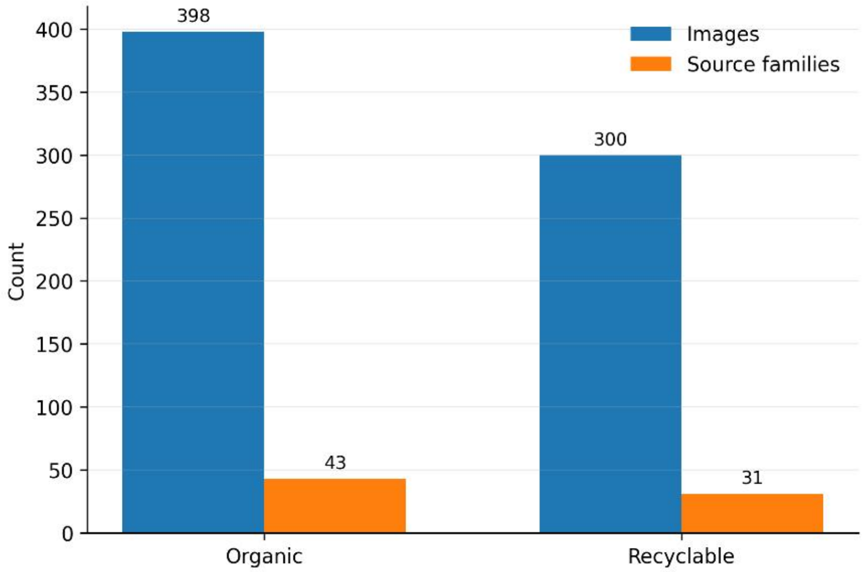

Figure 1. Distribution of images and source families by waste class

Figure 1 shows that the organic class contained 98 more images and 12 more source families than the recyclable class. The image-to-family ratios were 9.26 for organic waste and 9.68 for recyclable waste.

The five-fold stratified group partition produced held-out folds containing either 139 or 140 images. Each fold included between 14 and 16 source families, with no source-family overlap between development and held-out partitions. Table 1 presents the complete partition characteristics.

Table 1. Stratified group cross-validation partition characteristics

Fold	Development images	Held-out images	Development families	Held-out families	Held-out organic	Held-out recyclable	Family overlap
1	558	140	60	14	79	61	0
2	559	139	60	14	80	59	0
3	558	140	58	16	80	60	0
4	558	140	60	14	79	61	0
5	559	139	58	16	80	59	0

Table 1 shows that the held-out class composition remained consistent across folds, with 79–80 organic images and 59–61 recyclable images per partition. The zero-overlap result confirms that related augmented variants were not distributed across development and evaluation data.

4.2 LiteWasteNet Predictive Performance

Across the pooled out-of-fold predictions, LiteWasteNet correctly classified 579 of 698 images, corresponding to an accuracy of 82.95%. Balanced accuracy was 82.88%, while macro precision, macro recall, and macro F1-score were 82.57%, 82.88%, and 82.69%, respectively. The model produced a receiver operating characteristic area under the curve of 0.8958.

Fold-level performance varied across the five independent held-out partitions. Accuracy ranged from 72.86% in Fold 3 to 89.93% in Fold 5, while macro F1-score ranged from 72.72% to 89.81%. ROC-AUC remained above 0.82 in every fold. As illustrated in Figure 2, Fold 5 produced the highest values for all three displayed measures.

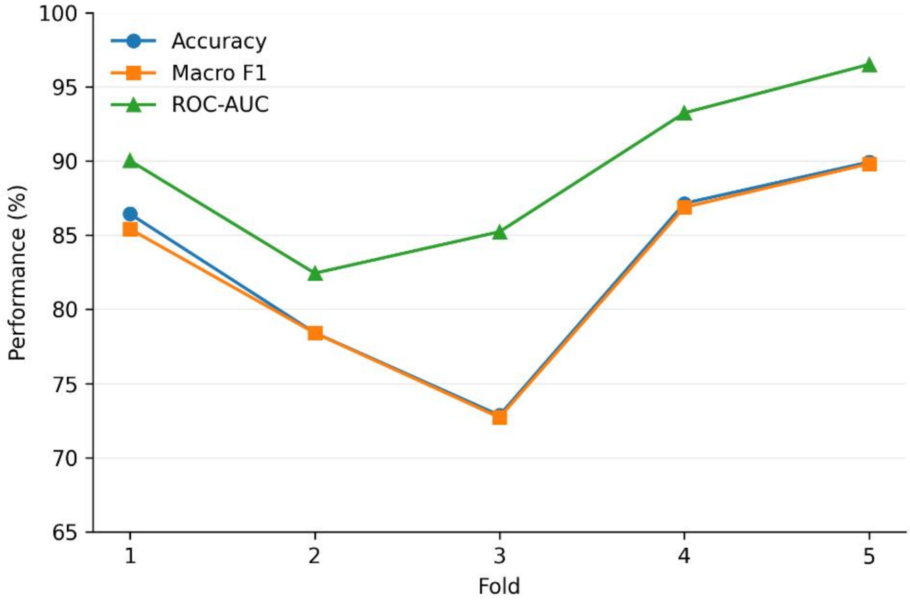

Figure 2. Fold-level predictive performance of LiteWasteNet

Figure 2 shows that accuracy and macro F1-score followed closely aligned fold-level patterns. ROC-AUC ranged from 82.44% to 96.50% and exceeded the corresponding threshold-dependent metrics in four of the five folds.

The unweighted fold means were 82.95% for accuracy, 82.95% for balanced accuracy, 82.64% for macro F1-score, and 89.48% for ROC-AUC. Their fold-level standard deviations were 7.08, 6.58, 6.95, and 5.73 percentage points, respectively. The small difference between the pooled and fold-mean estimates resulted from held-out partitions containing either 139 or 140 observations.

4.3 Class-Specific Outcomes

Organic-waste classification produced precision of 86.23%, recall of 83.42%, and an F1-score of 84.80%. Recyclable-waste classification produced precision of 78.91%, recall of 82.33%, and an F1-score of 80.59%. The organic class therefore exceeded the recyclable class by 7.32 percentage points in precision and 4.21 percentage points in F1-score, whereas recyclable recall was 1.09 percentage points lower.

Table 2 presents the pooled overall and class-specific evaluation statistics generated from all 698 out-of-fold predictions.

Table 2. Aggregate and class-specific LiteWasteNet performance

Evaluation level	Metric	Value (%)	Support
Overall	Accuracy	82.95	698
Overall	Balanced accuracy	82.88	698
Overall	Macro precision	82.57	698
Overall	Macro recall	82.88	698
Overall	Macro F1-score	82.69	698
Overall	ROC-AUC	89.58	698
Organic	Precision	86.23	398
Organic	Recall	83.42	398
Organic	F1-score	84.80	398

Recyclable	Precision	78.91	300
Recyclable	Recall	82.33	300
Recyclable	F1-score	80.59	300

Table 2 shows that class-specific recall differed by only 1.09 percentage points, consistent with the balanced accuracy of 82.88%. The largest class-level difference occurred in precision, where organic predictions exceeded recyclable predictions by 7.32 percentage points.

4.4 Confusion-Matrix Analysis

The pooled confusion matrix contained 332 correctly identified organic images and 247 correctly identified recyclable images. Among the 398 organic observations, 66 were predicted as recyclable. Among the 300 recyclable observations, 53 were predicted as organic. Figure 3 presents the distribution of correct and incorrect predictions.

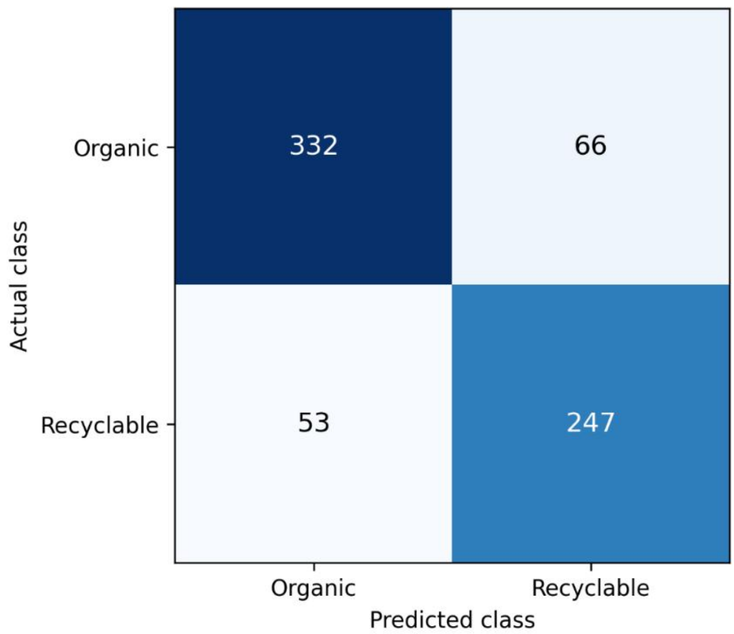

Figure 3. Pooled out-of-fold confusion matrix for LiteWasteNet

Figure 3 shows that 579 observations appeared on the main diagonal and 119 appeared in the off-diagonal cells. The class-specific error rates were 16.58% for organic waste and 17.67% for recyclable waste.

The model generated 385 organic predictions and 313 recyclable predictions. Of the organic predictions, 332 were correct and 53 were incorrect, yielding an organic precision of 86.23%. Of the recyclable predictions, 247 were correct and 66 were incorrect, yielding a recyclable precision of 78.91%. False recyclable assignments accounted for 55.46% of all errors, while false organic assignments accounted for 44.54%.

4.5 Baseline Comparison

The HOG–SVM baseline correctly classified 497 of 698 observations and produced an aggregate accuracy of 71.20%. Its balanced accuracy was 69.50%, macro precision was 70.97%, macro recall was 69.50%, macro F1-score was 69.75%, and ROC-AUC was 0.7804. Under the identical grouped folds, LiteWasteNet exceeded the baseline across every aggregate evaluation measure.

LiteWasteNet improved accuracy by 11.75 percentage points, balanced accuracy by 13.38 percentage points, macro precision by 11.60 percentage points, macro recall by 13.38 percentage points, macro F1-score by 12.94 percentage points, and ROC-AUC by 11.54 percentage points. The comparative values are illustrated in Figure 4.

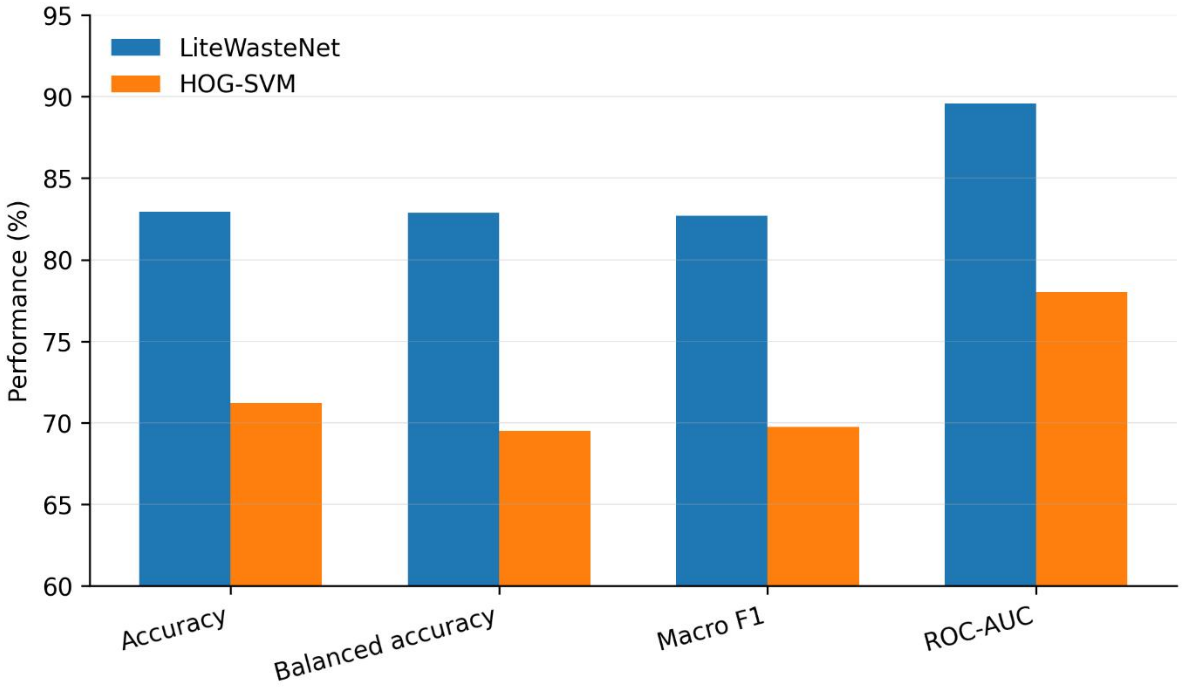

Figure 4. Predictive performance comparison between LiteWasteNet and HOG–SVM

Figure 4 shows that LiteWasteNet achieved higher accuracy, balanced accuracy, macro F1-score, and ROC-AUC than HOG–SVM. The largest absolute difference occurred for balanced accuracy, followed by macro F1-score.

At the fold level, LiteWasteNet accuracy exceeded HOG–SVM accuracy in Folds 1, 2, 4, and 5. Fold 3 was the only partition in which HOG–SVM accuracy was higher, recording 75.00% compared with 72.86% for LiteWasteNet. LiteWasteNet nevertheless produced a higher Fold 3 ROC-AUC of 85.23%, compared with 79.19% for the baseline. The baseline confusion matrix contained 325 correctly classified organic images, 172 correctly classified recyclable images, 73 organic-to-recyclable errors, and 128 recyclable-to-organic errors.

4.6 Computational Efficiency and Robustness

LiteWasteNet contained 15,114 trainable parameters and required 710,784 multiply-accumulate operations for one 64 × 64 RGB image. The serialized model weights occupied 85,239 bytes, equivalent to approximately 83.24 KiB. Median single-image CPU inference latency was 1.54 ms, while the 95th-percentile latency was 2.47 ms. Mean latency was 3.88 ms, corresponding to an estimated throughput of 257.53 images per second under the measurement environment.

The complete five-fold experimental procedure required 27.25 seconds of elapsed runtime, including training, validation, checkpoint handling, and evaluation. Individual fold-training durations ranged from 4.19 to 6.72 seconds, with a mean of 5.12 seconds. The retained checkpoints occurred between epochs 6 and 12, with a mean selected epoch of 9.4. Table 3 presents the measured computational and robustness indicators.

Table 3. Computational efficiency and cross-fold robustness of LiteWasteNet

Indicator	Measured value
Trainable parameters	15,114
Serialized weight size	85,239 bytes
Weight size	83.24 KiB
MACs per image	710,784
Median CPU latency	1.54 ms
Mean CPU latency	3.88 ms

95th-percentile CPU latency	2.47 ms
Estimated throughput	257.53 images/s
Total five-fold training time	27.25 s
Mean training time per fold	5.12 s
Selected epoch range	6–12
Fold accuracy range	72.86%–89.93%
Fold macro F1-score range	72.72%–89.81%
Fold ROC-AUC range	82.44%–96.50%

Table 3 shows that the trained model combined a sub-100 KiB weight file with fewer than one million MACs per prediction. All held-out folds produced ROC-AUC values above 82%, while fold-level accuracy remained above 72%.

5. Discussion

The results show that a simple convolutional structure can attain binary waste discrimination, and it is possible to achieve a comparable computational complexity to resources-constrained monitoring. The agreement between overall and balanced accuracy shows that the overall accuracy was not driven by the organic class, and similar recall between classes suggests that the class weightings and grouped validation contributed to the sensitivity. This is important because waste-classification studies do not reveal the accuracy of the different minority classifications. The outcome reaffirms the argument that the effectiveness of prediction should be evaluated based on class-sensitive measures and leakage-free evaluation as opposed to a summary score, as mentioned in the reviews of computer vision for solid-waste sorting (Lu & Chen, 2022).

Compared to the HOG–SVM baseline, LiteWasteNet was able to learn combinations of color, texture, and shape more effectively. This is in line with the ability of deep networks to build hierarchical visual features, though for a small architecture and a small binary dataset. This is important for resource-constrained waste-classification systems, since the real-world application is a tradeoff between maximum recognition accuracy and computation efficiency, as opposed to the maximum accuracy achieved by larger architectures (Fang et al., 2023). With compact parameter count, storage requirement, and low inference latency, the study demonstrates a backbone that is not computationally intensive is sufficient to achieve predictive performance.

The source-family-aware cross-validation is a methodological contribution. Since there were some augmented variants here, created from the same original images, random splitting would have potentially caused similar images to appear in both the development set and evaluation set. This dependence was mitigated by grouped partitioning, which yielded a conservative estimate of generalization. This is to overcome a disadvantage of small waste-image studies in which transfer learning can be very effective but with very small independent scene diversity (Wu & Lin, 2022). The variability of the folds reveals that performance is not only related to the families of sources excluded, but also to the partition in which the analysis is done; hence, robustness needs to be reported for every partition and not just for one where performance is good.

Class-specific difference indicates limitations of the binary taxonomy. Organic waste had higher precision values probably due to the texture and shape of biological waste. Recyclable waste was more heterogeneous, including packaging with different colors, surfaces and material compositions which often overlapped with organic-looking containers or contaminated materials. This is consistent with previous work showing that convolutional techniques are more effective than some feature extraction techniques, with the need to consider material variability and context of the image (Nežerka et al., 2024). Therefore, errors are not simply an algorithm problem and also are a sign of ambiguity in general category definitions.

The final product is a visual classification module validated in the field, which can be used in future waste-monitoring systems for household, institutional and municipal waste management systems, all of which are resource constrained. Its small parameter and storage size make the technical feasibility of local inference feasible, but it needs to be tested by hardware to determine the performance of the device. However, this needs to be matched to the local recycling regulations, particularly for contaminated or multi layered packaging. Household oriented systems like DEEPBIN prove the value of automated classification in practice, but also emphasize the importance of linking the performance of algorithms to the conditions of user interaction and sustained deployment (Song et al., 2024).

There are some potential limitations to the external validity of the study. The dataset was very small, consisted of controlled imagery and stock-style imagery, and was very limited in number of independent source families. It does not include mixed waste, occlusion, poor illumination, motion blur, sensor noise or repeated disposal at a true disposal site. Furthermore, a low-cost device's energy efficiency or performance is not determined by CPU latency. Future work should involve capturing and analyzing independent urban field images, evaluating the quantized models on the edge devices, measuring the power consumption and the cost of the entire system, and expanding the taxonomy to cover residual waste, hazardous waste, glass waste, metal waste, paper waste, electronic waste, and sanitary waste. Benchmarking on different datasets and real sorting environments is required to validate the practical efficiency advantage under demanding sorting environments (Dong et al., 2025).

6. Conclusion

In this study, LiteWasteNet has been designed and tested as a lightweight AI model for the binary classification of organic and recyclable waste in a leakage-resistant approach plus a source-family-aware validation strategy. The results prove that with a carefully designed lightweight CNN it is possible to achieve reliable class discrimination with a small computational cost, and that it can significantly outperform a baseline based on HOG–SVM with the same evaluation condition. From a methodological perspective, the study makes further contributions to waste classification research, highlighting the crucial role of grouped cross-validation, class-sensitive metrics and computational profiling in the context of small, augmentation-heavy datasets. In practical application, the model is a computational basis for organic–recyclable image classification based on empirical evaluation that is suitable for limited storage size and low arithmetic demand requirements and has fast CPU inference speed. The external validation should be carried out in 3 phases: independent validation of the images from the waste, match the binary labels with the local recycling laws and deployment of a quantized model on deployment devices (low cost edge hardware) with direct measurement of latency, memory consumption, power consumption, reliability and total system cost. For future work, the taxonomy should be extended to include residual, hazardous, sanitary, electronic, paper, metal, glass, and mixed waste, should add object detection to cluttered scenes, and should evaluate the system performance when it is being used in an occluded, contaminated, poorly-lit, and continuous mode. Before a system is considered operationally validated, it is important that it be subjected to additional testing in a longitudinal field trial to investigate the system's disposal compliance, contamination reduction, user acceptance, and maintenance requirements as well. This assessment would help to decide if the achieved computational performance would carry over to the operational waste-monitoring environment.

References

1. Lu, W., & Chen, J. (2022). Computer vision for solid waste sorting: A critical review of academic research. Waste Management, 142, 29-43.
2. Lin, K., Zhao, Y., Kuo, J. H., Deng, H., Cui, F., Zhang, Z., ... & Wang, T. (2022). Toward smarter management and recovery of municipal solid waste: A critical review on deep learning approaches. Journal of Cleaner Production, 346, 130943.
3. Mao, W. L., Chen, W. C., Fathurrahman, H. I. K., & Lin, Y. H. (2022). Deep learning networks for real-time regional domestic waste detection. Journal of Cleaner Production, 344, 131096.
4. Fang, B., Yu, J., Chen, Z., Osman, A. I., Farghali, M., Ihara, I., ... & Yap, P. S. (2023). Artificial intelligence for waste management in smart cities: a review. Environmental Chemistry Letters, 21(4), 1959-1989.
5. Ahmed, M. M., Hassanien, E., & Hassanien, A. E. (2023). IoT-based intelligent waste management system. Neural Computing and Applications, 35(32), 23551-23579.
6. Kannan, D., Khademolqorani, S., Janatyan, N., & Alavi, S. (2024). Smart waste management 4.0: The transition from a systematic review to an integrated framework. Waste Management, 174, 1-14.
7. Hussain, I., Elomri, A., Kerbache, L., & El Omri, A. (2024). Smart city solutions: Comparative analysis of waste management models in IoT-enabled environments using multiagent simulation. Sustainable Cities and Society, 103, 105247.

8. Seyyedi, S. R., Kowsari, E., Gheibi, M., Chinnappan, A., & Ramakrishna, S. (2024). A comprehensive review integration of digitalization and circular economy in waste management by adopting artificial intelligence approaches: Towards a simulation model. Journal of Cleaner Production, 460, 142584.
9. Li, N., & Chen, Y. (2023). Municipal solid waste classification and real-time detection using deep learning methods. urban climate, 49, 101462.
10. Wu, F., & Lin, H. (2022). Effect of transfer learning on the performance of VGGNet-16 and ResNet-50 for the classification of organic and residual waste. Frontiers in Environmental Science, 10, 1043843.
11. Li, J., Chen, J., Sheng, B., Li, P., Yang, P., Feng, D. D., & Qi, J. (2021). Automatic detection and classification system of domestic waste via multimodel cascaded convolutional neural network. IEEE transactions on industrial informatics, 18(1), 163-173.
12. Lin, K., Zhou, T., Gao, X., Li, Z., Duan, H., Wu, H., ... & Zhao, Y. (2022). Deep convolutional neural networks for construction and demolition waste classification: VGGNet structures, cyclical learning rate, and knowledge transfer. Journal of Environmental Management, 318, 115501.
13. Jin, S., Yang, Z., Królczykg, G., Liu, X., Gardoni, P., & Li, Z. (2023). Garbage detection and classification using a new deep learning-based machine vision system as a tool for sustainable waste recycling. Waste Management, 162, 123-130.
14. Nežerka, V., Zbíral, T., & Trejbal, J. (2024). Machine-learning-assisted classification of construction and demolition waste fragments using computer vision: Convolution versus extraction of selected features. Expert Systems with Applications, 238, 121568.
15. Song, Y., He, X., Tang, X., Yin, B., Du, J., Liu, J., ... & Geng, S. (2024). DEEPBIN: Deep learning based garbage classification for households using sustainable natural technologies. Journal of Grid Computing, 22(1), 2.
16. Liang, Q., Li, J., Qin, H., Liu, M., Xiao, X., Zhang, D., ... & Zhang, D. (2024). MHA-DGCLN: multi-head attention-driven dynamic graph convolutional lightweight network for multi-label image classification of kitchen waste: Q. Liang et al. Applied Intelligence, 54(24), 13057-13074.
17. Sarswat, P. K., Singh, R. S., & Pathapati, S. V. S. H. (2024). Real time electronic-waste classification algorithms using the computer vision based on convolutional neural network (cnn): Enhanced environmental incentives. Resources, Conservation and Recycling, 207, 107651.
18. Sayem, F. R., Islam, M. S. B., Naznine, M., Nashbat, M., Hasan-Zia, M., Kunju, A. K. A., ... & Chowdhury, M. E. (2025). Enhancing waste sorting and recycling efficiency: robust deep learning-based approach for classification and detection. Neural Computing and Applications, 37(6), 4567-4583.
19. Nahiduzzaman, M., Ahamed, M. F., Naznine, M., Karim, M. J., Kibria, H. B., Ayari, M. A., ... & Haider, J. (2025). An automated waste classification system using deep learning techniques: Toward efficient waste recycling and environmental sustainability. Knowledge-Based Systems, 310, 113028.
20. Dong, Z., Yuan, L., Yang, B., Xue, F., & Lu, W. (2025). Benchmarking computer vision models for automated construction waste sorting. Resources, Conservation and Recycling, 213, 108026.
21. Shah, P. (n.d.). Waste classification (organic and recyclable) [Data set]. Kaggle. https://www.kaggle.com/datasets/preetishah/waste-classificationorganic-and-recyclable