



Comparative Analysis of Convolutional Neural Network Architectures for Automated Waste Classification

Sakshi Anil Parkhe¹

Department of Computer Science,

MIT College of Railway Engineering and Research, Barshi, Solapur, India¹

Abstract: Effective waste classification is a critical component of modern smart waste management systems, enabling automated sorting and recycling at scale. This paper presents a comprehensive comparative study of six convolutional neural network (CNN) architectures (VGG16, ResNet50, MobileNetV3, EfficientNet-B0, AlexNet, and a custom-designed CNN) for the task of automated waste classification into six categories: cardboard, glass, metal, paper, plastic, and trash. Each model is evaluated on a standardized dataset of 2,527 images under identical experimental conditions, including data augmentation, learning rate scheduling, and transfer learning where applicable. Performance is assessed using accuracy, precision, recall, F1-score, receiver operating characteristic (ROC) curves, and precision-recall (PR) curves. Experimental results demonstrate that the custom CNN achieves the highest classification accuracy of 95.05% with a macro-average ROC-AUC of 0.9974, outperforming established architectures. ResNet50 and MobileNetV3 follow closely at 94.07% and 93.87%, respectively, offering favorable trade-offs between accuracy and computational complexity. The analysis further includes per-class performance evaluation, confusion matrix analysis, and computational efficiency comparison, providing actionable insights for deploying waste classification systems in resource-constrained environments.

Keywords: waste classification, convolutional neural networks, transfer learning, deep learning, image classification, recycling, smart waste management

I. INTRODUCTION

The global waste crisis has intensified in recent decades, with the World Bank estimating that annual waste generation will increase by 70% from 2016 levels to reach 3.40 billion tonnes by 2050 [1]. Rapid urbanization, industrial expansion, and changing consumption patterns have exacerbated this challenge, making effective waste management a pressing environmental and public health priority. A critical bottleneck in the waste management pipeline is the sorting and classification process, which has traditionally relied on manual labor. Manual sorting is not only labor-intensive and costly but also prone to human error, inconsistent quality, and exposure to hazardous materials [2].

Automated waste classification using computer vision and deep learning offers a promising solution to these challenges. By leveraging convolutional neural networks (CNNs), waste items can be identified and categorized from images in real-time, enabling robotic sorting systems to operate with high throughput and accuracy [3]. The deployment of such systems in recycling facilities can significantly reduce operational costs, improve sorting purity, and increase the overall recycling rate. Furthermore, smart waste bins equipped with embedded classification capabilities can help end-users correctly separate recyclables from non-recyclables at the point of disposal, thereby reducing contamination in recycling streams [4].

Over the past decade, numerous CNN architectures have been proposed in the computer vision literature, each offering distinct trade-offs between representational capacity, computational efficiency, and memory footprint. Early architectures such as AlexNet [7] and VGG16 [8] demonstrated the power of deep feature hierarchies but introduced substantial computational overhead. More recent designs including ResNet50 [9], MobileNetV3 [10], and EfficientNet-B0 [11] have achieved state-of-the-art performance while incorporating architectural innovations (residual connections, depthwise separable convolutions, and compound scaling) that reduce the computational cost of inference.

Despite the growing body of work on waste classification, there remains a need for systematic comparative studies that evaluate multiple CNN architectures on identical datasets and under controlled experimental conditions. Many existing studies report results on different datasets, use different preprocessing pipelines, or adopt different evaluation metrics, making direct comparison of model performance difficult [5]. Additionally, few studies incorporate both ROC and



precision-recall analysis alongside standard accuracy metrics, despite the importance of these curves for understanding model behavior across different operating points and class distributions [6].

This paper addresses these gaps by conducting a rigorous comparative evaluation of six CNN architectures for six-class waste classification. The contributions of this work are three-fold: (1) a standardized experimental protocol ensuring fair comparison across all models, (2) a comprehensive evaluation framework incorporating accuracy, precision, recall, F1-score, ROC curves, and precision-recall curves, and (3) a detailed analysis of per-class performance, computational efficiency, and practical deployment considerations. The remainder of this paper is organized as follows. Section II reviews related work. Section III describes the methodology. Section IV presents experimental results. Section V discusses findings, and Section VI concludes the paper.

II. RELATED WORK

A. CNN Architectures for Image Classification

The application of deep learning to image classification has evolved rapidly since the breakthrough performance of AlexNet in the 2012 ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [7]. AlexNet introduced a deep architecture with five convolutional layers and three fully connected layers, establishing the effectiveness of deep feature learning combined with GPU acceleration. VGG16 [8] demonstrated that depth, achieved through the systematic use of 3×3 convolutional filters, is a critical factor in representational power. However, the pure sequential architecture of VGG16 results in a large parameter count of approximately 138 million, making it computationally expensive for both training and inference.

ResNet50 [9] introduced residual connections that enable the training of much deeper networks by addressing the vanishing gradient problem. With 50 layers and approximately 25.6 million parameters, ResNet50 achieves significantly higher accuracy than VGG16 while remaining computationally tractable. MobileNetV3 [10] was designed specifically for mobile and edge deployment, employing depthwise separable convolutions and hardware-aware neural architecture search to achieve a favorable accuracy-efficiency trade-off. EfficientNet-B0 [11] introduced compound scaling, which uniformly scales network depth, width, and resolution using a principled coefficient, achieving state-of-the-art accuracy with fewer parameters than comparable architectures.

B. Waste Classification Using Deep Learning

The application of CNNs to waste classification has been explored in several studies. Chu et al. [3] proposed a waste classification system using a custom deep CNN architecture trained on the TrashNet dataset, achieving classification accuracy exceeding 85%. Aditya et al. [4] developed a smart waste classification system for urban environments using transfer learning with VGG16 and ResNet, reporting improved performance over training from scratch. Yang et al. [5] conducted a survey of deep learning methods for waste classification and highlighted the need for standardized benchmarks and evaluation protocols.

More recently, researchers have explored lightweight architectures suitable for deployment on edge devices. Birla et al. [12] compared MobileNet variants with deeper architectures on waste datasets, finding that MobileNetV2 achieved competitive performance with substantially lower latency. Ozkaya et al. [13] evaluated EfficientNet-B0 for waste classification and reported strong performance relative to its parameter count. However, these studies typically evaluate a limited number of architectures, often without systematic comparison across all major CNN families.

III. METHODOLOGY

A. Dataset

The experiments in this study use a waste image dataset comprising 2,527 images distributed across six waste categories: cardboard, glass, metal, paper, plastic, and trash. The dataset is the publicly available TrashNet dataset [15], a standard benchmark for waste classification research, and contains images captured under diverse real-world conditions, including varying lighting, background complexity, and object orientations. Table I summarizes the class distribution of the dataset.



TABLE I DATASET DISTRIBUTION ACROSS WASTE CATEGORIES

Category	Train	Test
Cardboard	341	85
Glass	438	109
Metal	372	93
Paper	504	126
Plastic	447	112
Trash	110	27
Total	2,212	552

For training and evaluation, the dataset was divided into training and test sets using an 80/20 stratified split. This ensures that the class distribution is preserved in both subsets and that the test set provides a representative estimate of generalization performance.

B. Data Preprocessing and Augmentation

All input images were resized to 224×224 pixels to match the input requirements of the pretrained models. Pixel values were normalized using channel-wise mean and standard deviation statistics derived from the ImageNet dataset ($\mu = [0.485, 0.456, 0.406]$, $\sigma = [0.229, 0.224, 0.225]$). To improve generalization and mitigate overfitting, the following data augmentation strategies were applied during training:

- Random horizontal flip with probability 0.5
- Random rotation in the range $[-15^\circ, 15^\circ]$
- Random resized crop with scale factor in $[0.8, 1.0]$
- Color jitter with brightness, contrast, saturation, and hue perturbations

C. Model Architectures

Six CNN architectures were evaluated in this study, spanning a range of design philosophies from early deep learning to modern efficient architectures:

AlexNet [7]: One of the earliest deep CNN architectures, consisting of five convolutional layers with ReLU activations and three fully connected layers. While pioneering, AlexNet has approximately 61 million parameters and employs a relatively simple architecture with limited feature extraction capacity.

VGG16 [8]: A 16-layer network using only 3×3 convolutional filters stacked sequentially, followed by three fully connected layers. VGG16 has approximately 138 million parameters and is known for its strong feature extraction capabilities but high computational cost.

ResNet50 [9]: A 50-layer residual network employing skip connections that enable effective gradient flow during training of very deep networks. With approximately 25.6 million parameters, ResNet50 achieves a favorable balance between depth and computational efficiency.

MobileNetV3 [10]: An efficient architecture designed for mobile deployment using depthwise separable convolutions and hardware-aware architecture search. MobileNetV3-Small contains approximately 2.9 million parameters, making it the lightest architecture in this comparison.

EfficientNet-B0 [11]: The baseline model of the EfficientNet family, using compound scaling to balance depth, width, and resolution. EfficientNet-B0 has approximately 5.3 million parameters and achieves competitive accuracy with efficient resource utilization.

Custom CNN: A purpose-built five-layer convolutional architecture designed for the waste classification task, incorporating batch normalization, dropout regularization, and global average pooling. The architecture comprises five convolutional blocks (each with Conv2D, BatchNorm, ReLU, and MaxPool), followed by a fully connected classifier head with dropout ($p = 0.5$).

D. Transfer Learning and Training Configuration

For all pre-trained architectures (VGG16, ResNet50, MobileNetV3, and EfficientNet-B0), transfer learning was employed by initializing the convolutional base with ImageNet pre-trained weights and replacing the final classification



head with a six-neuron fully connected layer matching the number of waste categories. The convolutional base was fine-tuned end-to-end during training. The Custom CNN and AlexNet were trained from scratch.

All models were trained using the Adam optimizer [14] with an initial learning rate of $\eta = 1 \times 10^{-4}$ and a cosine annealing learning rate schedule. The batch size was set to 32, and the categorical cross-entropy loss function was used. Training was conducted for 30 epochs with early stopping (patience of 5 epochs) based on validation loss.

E. Evaluation Metrics

Model performance was evaluated using the following metrics:

- Accuracy: The proportion of correctly classified images.
- Precision: $P = TP/(TP + FP)$, measuring the fraction of positive predictions that are correct.
- Recall: $R = TP/(TP + FN)$, measuring the fraction of actual positives correctly identified.
- F1-Score: The harmonic mean of precision and recall: $F1 = 2PR/(P + R)$.
- ROC-AUC: Area under the receiver operating characteristic curve.
- PR Curves: Precision-recall trade-off analysis for each class and macro-average.

IV. EXPERIMENTAL RESULTS

A. Overall Classification Performance

Table II presents the overall classification performance of all six models. The custom CNN achieves the highest test accuracy of 95.05%, followed by ResNet50 at 94.07% and MobileNetV3 at 93.87%. EfficientNet-B0 achieves 91.90%, while VGG16 and AlexNet trail at 86.58% and 83.16%, respectively.

TABLE II COMPREHENSIVE PERFORMANCE COMPARISON OF CNN ARCHITECTURES FOR WASTE CLASSIFICATION

Model	Accuracy (%)	Precision	Recall	F1-Score	ROC-AUC	Params (M)	Training
Custom CNN	95.05	0.95	0.95	0.95	0.9974	1.8	Scratch
ResNet50	94.07	0.94	0.94	0.94	0.9950	25.6	Transfer
MobileNetV3	93.87	0.94	0.94	0.94	0.9935	2.9	Transfer
EfficientNet-B0	91.90	0.92	0.92	0.92	0.9880	5.3	Transfer
VGG16	86.58	0.87	0.87	0.87	0.9700	138.0	Transfer
AlexNet	83.16	0.84	0.83	0.83	0.9550	61.0	Scratch

A notable finding is that the custom CNN, despite having only 1.8 million parameters and being trained from scratch, outperforms all transfer learning-based models. This suggests that the architecture was well-suited for the specific feature distribution of the waste dataset. Among the transfer learning models, ResNet50 and MobileNetV3 demonstrate that modern architectures with efficient design achieve near-optimal performance while offering significantly reduced computational requirements compared to VGG16.

The macro-average ROC-AUC scores reveal that all models achieve strong discriminative ability, with AUC values exceeding 0.95. The custom CNN attains an AUC of 0.9974, indicating near-perfect class separation. Even the lowest-performing model, AlexNet, achieves an AUC of 0.9550, confirming that CNN-based approaches are fundamentally well-suited for the waste classification task.

B. Per-Class Performance Analysis

Table III presents the per-class F1-scores for each model across the six waste categories. Several notable patterns emerge from this analysis.



TABLE III PER-CLASS F1-SCORES FOR EACH MODEL ACROSS WASTE CATEGORIES

Model	Cardboard	Glass	Metal	Paper	Plastic	Trash
Custom CNN	0.99	0.93	0.98	0.99	0.96	0.96
ResNet50	0.98	0.92	0.96	0.97	0.95	0.91
MobileNetV3	0.97	0.93	0.95	0.97	0.94	0.90
EfficientNet-B0	0.95	0.89	0.93	0.96	0.91	0.85
VGG16	0.91	0.83	0.88	0.92	0.85	0.75
AlexNet	0.87	0.78	0.86	0.90	0.81	0.69

Across all models, the “paper” category consistently achieves the highest F1-scores, likely due to the distinctive visual characteristics of paper waste, including its flat profile, consistent color distribution, and textural patterns. Conversely, the “trash” category proves most challenging, with F1-scores ranging from 0.69 (AlexNet) to 0.96 (Custom CNN). The high intra-class variability of trash items—which may include diverse objects such as food waste, ceramics, and mixed materials—contributes to this difficulty.

The glass category also presents notable challenges, particularly for VGG16 ($F1 = 0.83$) and AlexNet ($F1 = 0.78$). This can be attributed to the transparency and reflectivity of glass, which introduces significant appearance variation depending on lighting conditions and background context. The residual connections in ResNet50 and the efficient feature extraction in MobileNetV3 help mitigate these challenges, achieving F1-scores of 0.92 and 0.93, respectively.

C. ROC Curve Analysis

Figs. 1 and 2 present the receiver operating characteristic (ROC) curves for all evaluated models. Fig. 1 shows the per-class ROC curves for the best-performing Custom CNN model. All six classes achieve AUC values of at least 0.995, with cardboard, paper, and metal reaching 0.999. The trash class, while slightly lower at 0.995, still demonstrates excellent discriminative performance.

Fig. 2 compares the macro-average ROC curves across all six models. The curves clearly demonstrate the performance hierarchy identified in the accuracy analysis. The Custom CNN, ResNet50, and MobileNetV3 curves cluster near the top-left corner, indicating high sensitivity at low false positive rates. The separation between VGG16/AlexNet and the top three models is more pronounced in the low-FPR region, which is critical for practical applications where false positives carry operational costs.




Fig. 1. Per-class ROC curves for the Custom CNN model. All classes achieve AUC values of at least 0.995.




Fig. 2. Macro-average ROC curves comparing all six models. The Custom CNN, ResNet50, and MobileNetV3 achieve near-perfect discrimination.

D. Precision-Recall Curve Analysis

Figs. 3 and 4 present the precision-recall (PR) curves, which provide additional insight into model behavior, particularly for imbalanced or difficult classes. PR curves are especially informative when the positive class is rare, as they are more sensitive to performance differences than ROC curves in such scenarios [6].

Fig. 3 shows the per-class PR curves for the Custom CNN. Cardboard and paper maintain high precision across all recall levels, reflecting their distinctive visual features. The trash class exhibits the steepest precision decline at higher recall values, consistent with the higher intra-class variability observed in this category.

Fig. 4 compares the macro-average PR curves across all models. The performance gap between models is more visible in the PR space than in the ROC space, particularly at high recall values. AlexNet shows a pronounced precision drop beyond recall of 0.7, indicating that a significant fraction of its high-confidence predictions are incorrect for harder samples.




Fig. 3. Per-class precision-recall curves for the Custom CNN model.




Fig. 4. Macro-average precision-recall curves comparing all six models.

E. Confusion Matrix Analysis

Fig. 5 presents the confusion matrices for the best-performing (Custom CNN) and worst-performing (AlexNet) models, providing a detailed view of per-class classification behavior.

For the Custom CNN (Fig. 5(a)), the confusion matrix reveals near-diagonal predictions with minimal off-diagonal elements. The most common misclassification occurs between glass and trash (2 instances), and between glass and plastic (2 instances), which is expected given the visual overlap between transparent or semi-transparent glass and plastic items, and the heterogeneous nature of trash.

For AlexNet (Fig. 5(b)), the confusion matrix shows significantly more off-diagonal activity. Notable confusion patterns include: trash being misclassified as glass (4 instances) and plastic (3 instances), glass being misclassified as plastic (3 instances), and plastic being misclassified as glass (3 instances). These patterns highlight the shared visual characteristics among certain waste categories and the limited feature extraction capacity of the AlexNet architecture.




Fig. 5. Confusion matrices for (a) Custom CNN and (b) AlexNet. The Custom CNN exhibits near-diagonal predictions with minimal misclassification, while AlexNet shows significantly more off-diagonal confusion, particularly between glass, plastic, and trash categories.

F. Training Dynamics

Fig. 6 illustrates the training loss and accuracy curves for all six models over 30 epochs. Several key observations emerge from the training dynamics analysis.

The Custom CNN exhibits the fastest convergence, reaching 90% training accuracy by epoch 10 and plateauing near 96% by epoch 20. ResNet50 and MobileNetV3 converge at similar rates, benefiting from their ImageNet pre-trained



feature extractors. AlexNet shows the slowest convergence, requiring approximately 20 epochs to reach 80% accuracy and exhibiting higher loss throughout training. The transfer learning models (VGG16, ResNet50, MobileNetV3, EfficientNet-B0) consistently show lower initial loss compared to the models trained from scratch (Custom CNN, AlexNet), owing to their pre-trained feature representations. However, the Custom CNN quickly overcomes this initial disadvantage and surpasses all transfer learning models by epoch 15, demonstrating the effectiveness of its task-specific architecture design.




Fig. 6. Training dynamics over 30 epochs showing (a) loss curves and (b) accuracy curves for all six models. The Custom CNN, despite training from scratch, achieves the fastest convergence and highest final accuracy.

V. DISCUSSION

A. Architecture Design Implications

The experimental results reveal several important implications for the design of CNN-based waste classification systems. First, the superior performance of the Custom CNN over all transfer learning models suggests that task-specific architecture design, when informed by the characteristics of the target domain, can yield architectures that outperform general-purpose feature extractors even without the benefit of large-scale pre-training. The Custom CNN's use of batch normalization, dropout, and global average pooling provides effective regularization that prevents overfitting despite the relatively small training set.

Second, the strong performance of MobileNetV3 (93.87% accuracy with only 2.9 million parameters) demonstrates that efficient architectures designed for mobile deployment are viable for waste classification. This finding has significant practical implications, as waste classification systems deployed in smart bins or mobile sorting robots require models that can run efficiently on edge devices with limited computational resources. The minimal accuracy gap between MobileNetV3 and the top-performing models (1.18% below the Custom CNN) suggests that the efficiency gains far outweigh the marginal accuracy loss in most deployment scenarios.

Third, the relatively poor performance of VGG16 (86.58%) and AlexNet (83.16%), despite VGG16 having the largest parameter count (138M), underscores the importance of architectural innovation beyond depth and parameter count. The sequential nature of VGG16 and the limited receptive field of AlexNet restrict their ability to capture the diverse visual features present in waste images.

B. Computational Efficiency and Deployment Considerations

Table II lists each model's parameter count, which directly affects memory requirements and inference latency. For deployment scenarios with strict computational constraints, such as embedded systems in smart waste bins, MobileNetV3 (2.9M parameters) offers the best accuracy-per-parameter ratio among the transfer learning models. The Custom CNN (1.8M parameters) provides the highest absolute accuracy while maintaining a compact footprint.

For high-throughput industrial sorting facilities where computational resources are less constrained, ResNet50 or the Custom CNN may be preferred due to their higher accuracy. In cloud-based classification services where images are processed on powerful servers, the marginal computational cost of larger models is negligible, making accuracy the primary consideration.

C. Limitations and Future Work

Several limitations should be noted. First, the dataset used in this study, while representative, is relatively small (2,527 images), which may limit the statistical power of the performance comparisons. Future work should validate these



findings on larger and more diverse waste datasets. Second, the ROC and PR curves presented in Sections IV-C and IV-D were generated from model predictions on the test set, and the exact curve shapes may vary with different train-test splits. Cross-validation-based curve analysis would provide more robust estimates. Third, this study evaluates static image classification and does not consider temporal information that could be leveraged in video-based sorting systems.

VI. CONCLUSION

This paper presented a comprehensive comparative analysis of six CNN architectures for automated waste classification across six categories. Through standardized experimental evaluation incorporating accuracy metrics, per-class F1-score analysis, ROC curves, precision-recall curves, and confusion matrix examination, the following key findings were established:

The custom CNN achieved the highest classification accuracy of 95.05% with a macro-average ROC-AUC of 0.9974, demonstrating that task-specific architectural design can outperform general-purpose transfer learning approaches. ResNet50 (94.07%) and MobileNetV3 (93.87%) emerged as the best-performing transfer learning models, with MobileNetV3 offering a particularly attractive trade-off between accuracy and computational efficiency for edge deployment. VGG16 and AlexNet, despite being pioneering architectures, showed significant performance gaps compared to modern designs, achieving 86.58% and 83.16% accuracy, respectively. The per-class analysis revealed that paper and cardboard are the easiest categories to classify, while trash and glass present the greatest challenges due to high intra-class variability and visual overlap with other categories. The ROC and PR curve analysis provided deeper insight into model discriminative ability, confirming that all models achieve strong class separation but with notable performance differences at high recall operating points.

These findings provide practical guidance for practitioners deploying waste classification systems: MobileNetV3 is recommended for resource constrained environments, ResNet50 for balanced performance, and custom architectures for maximum accuracy when training data and computational budget are available. Future work should explore ensemble methods, attention mechanisms, and multi-modal classification incorporating depth or spectral information to further improve performance on challenging waste categories.

REFERENCES

- [1] S. Kaza, L. Yao, P. Bhada-Tata, and F. Van Woerden, "What a waste 2.0: A global snapshot of solid waste management to 2050," World Bank, Washington, DC, Tech. Rep., 2018.
- [2] W. Raghupathi and V. Raghupathi, "Big data analytics in healthcare: Promise and potential," *Health Information Science and Systems*, vol. 2, no. 1, p. 3, 2014.
- [3] Y. Chu, C. Huang, X. Xie, B. Tan, S. Kamel, and A. Kharrat, "Deep learning based automatic waste classification system," in *Proc. IEEE Int. Conf. Big Data*, 2018, pp. 5098–5103.
- [4] L. Aditya, A. R. Rao, G. S. Bhat, and S. N. Omka, "Smart bin for smart city-An IoT approach," in *Proc. Int. Conf. IoT Social, Mobile, Analytics and Cloud (ISMAC)*, 2020, pp. 418–423.
- [5] Q. Yang, Y. Mao, Z. Li, and G. Li, "A comprehensive survey of waste classification using deep learning," *Waste Management*, vol. 143, pp. 310–325, 2022.
- [6] J. Davis and M. Goadrich, "The relationship between precision-recall and ROC curves," in *Proc. ICML*, 2006, pp. 233–240.
- [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NeurIPS*, 2012, pp. 1097–1105.
- [8] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, 2016, pp. 770–778.
- [10] A. Howard, M. Sandler, G. Chen, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhuang, R. Andri, V. Roth, M. Adam, and F. Xiao, "Searching for MobileNetV3," in *Proc. IEEE/CVF ICCV*, 2019, pp. 1314–1324.
- [11] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019, pp. 6105–6114.
- [12] S. Birla, R. Kumar, and R. Kaur, "Waste classification using deep learning: A comparative study," *Journal of Environmental Management*, vol. 314, p. 115046, 2022.
- [13] U. Ozkaya and S. Yigit, "Deep learning-based waste classification system using EfficientNet-B0," *Journal of Air Waste Management Association*, vol. 71, no. 8, pp. 1004–1014, 2021.
- [14] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [15] G. Thung and M. Yang, "Classification of trash for recyclability status," *CS229 Project Report*, Stanford University, 2016.