

Impact of contamination variability on convolutional neural network accuracy in recycling classification

Aarav Rao¹, Sarita Mohan²

¹ T.C. Jasper High School, Plano, Texas

² Technical Program Manager, Outcomes, Plano, Texas

SUMMARY

Recycling plays an important role in reducing waste and saving resources. But if recyclables aren't sorted correctly, a lot of material that could have been reused ends up in landfills. Even small bits of food or grease can contaminate recyclable items and make them hard to process. The efficacy of recycling programs relies upon effective sorting with machine learning solutions, increasingly providing solutions. Although convolutional neural networks (CNN) offer a baseline for automated waste classification, recycling bins are often contaminated. Contamination, whether food scraps, oils, or a persistent failure of the consumers to understand non-recyclability, alters the critical visual features needed for CNN classification, leading to higher misclassification rates. Therefore, this study aimed to determine whether increased levels of contamination may decrease classification accuracy by increasing the percentage of contaminant-recognizable recyclables classified as non-recyclables. Accordingly, we trained a CNN on images of recyclables, non-recyclables, and contaminated recyclables of varying contamination levels, then tested a held-out dataset to measure classification accuracy based on contamination severity. Our results indicated a significant inverse relationship between contamination levels and classification accuracy, suggesting that higher contamination is associated with reduced CNN performance. These findings suggest that contamination substantially interferes with feature extraction in CNN-based classification, deepening the need for improved preprocessing strategies in machine-learning-assisted waste management systems.

INTRODUCTION

Waste is increasing at an alarming rate globally. People around the world generate approximately two billion tons of trash annually (1). A significant portion of municipal waste, especially paper and cardboard, glass, metals, and certain plastics, is recyclable; however, contamination within recycling streams continues to limit effective processing (2). Each year, millions of tons of contamination, from food and grease to other non-recyclables, decrease the quality of recyclables and add to the costs of material recovery facilities (MRFs) (2). According to the National Solid Wastes Management Association, contamination within recycling streams occurs anywhere from 10% to 25% of recycling streams across the United States depending on local recycling habits and regulations (3). These high contamination rates contribute to increased landfill waste as MRFs reject contaminated material,

with over half of U.S. municipal solid waste ultimately ending up in landfills (8).

With the advancement of artificial intelligence (AI), researchers have developed a variety of deep learning models, including recurrent neural networks (RNNs), generative adversarial networks (GANs), and transformers - for computer vision tasks. However, convolutional neural networks (CNNs) remain the preferred architecture for image classification because they can learn the representation hierarchy of features inherently from original image inputs (4). CNNs have achieved over 90% accuracy in classifying items in recyclable bins under ideal circumstances (5). However, because CNNs are highly feature-dependent, their accuracy may decline if the training dataset lacks sufficient examples of contaminated waste (5). Without exposure to a diverse range of contamination patterns during training, CNNs fail to generalize, leading to misclassification, especially when contamination blurs into the visual boundaries between recyclable and non-recyclable items (5).

A previous study on feature occlusion, which is removing or hiding a piece of an image or object during the training to assess how well a model works or to evaluate what features are relevant, have found that CNNs are highly sensitive to distortions that obscure key object characteristics (6). Therefore, one can argue that contamination is an inappropriate feature of occlusion because it alters the color, texture, and shape of an image, causing misclassification if not adequately represented in training. Despite these insights, little research has specifically investigated how varying levels of contamination impact CNN-based classification accuracy in waste management.

Our study analyzes the impact of contamination on CNN classification accuracy of a trained model for recycling. We hypothesized that increased levels of contamination decreases classification performance as features relevant to classification are obscured. To test this, we trained convolutional neural networks on images of recyclable, non-recyclable, and contaminated materials at 25%, 50%, and 75% contamination level and then measured how accurately the models could classify a separate set of test images. Our results showed that as contamination increased, classification accuracy dropped significantly, with EfficientNetB0 performing about 5-10% better than MobileNetV2 despite slightly longer processing times. These findings suggest that even advanced AI systems struggle when recycling streams are contaminated. Overall, this study shows that improving AI-based sorting and reducing contamination could make recycling more reliable and sustainable on a larger scale. By presenting substantial

data of classification accuracy of the model at various rates of contamination, our project exposed the failings of trained, machine learning based solutions for effective recycling, and pushes for the need for a more robust model and contamination reduction strategies.

RESULTS

We hypothesized that increasing levels of contamination would reduce the classification accuracy of CNNs for recyclable materials. To test this, we created a dataset of more than 100,000 images from typical recyclables—such as plastic bottles, soda cans, and shoe boxes—with four levels of contamination (**Figure 1**). Each image was either clean (the object only), or had low (smudges, food residue), moderate (grease half covering the object), or high (covered with left-over food and sandwiches on top) levels of contamination. For the purpose of this study, contamination levels were assessed based on how much contamination obscured the dominant visual characteristics of the item. The dataset was split as 70% for training and 30% for testing, with training performed on clean images only. A CNN with two convolutional layers, max pooling, and a dense output layer was trained for 25 epochs with categorical cross-entropy as the loss function and Adam as the optimizer. The assessment was how classification accuracy relied on progressive levels of contamination.

The CNN classified clean recyclables with 94.5% accuracy; however, as contamination increased, accuracy declined (**Figure 2**). Specifically, the model achieved 85.7% for low contamination, 72.3% for moderate contamination, and 58.9% for high contamination levels. Contamination had a significant effect on classification accuracy (one-way ANOVA, $F(3, 196) = 42.78, p = 1.11 \times 10^{-16}$). A Tukey's Honestly Significant Difference (HSD) post-hoc test showed that all levels of contamination were significantly different from each other ($p < 0.05$), except for moderate versus high contamination, (**Figure 2**; $p = 0.072$).

In addition to overall classification accuracy, the model was evaluated using common performance metrics including precision, recall, and F1 score. Confusion matrices were also constructed to visualize classification performance across contamination levels and waste categories (**Figure 3**). These metrics highlighted that false positives increased at higher contamination levels, particularly for contaminated recyclables misclassified as clean recyclables.

DISCUSSION

We demonstrated that contamination significantly reduces CNN classification accuracy in waste sorting. Our findings show that contamination impairs the model's ability to correctly identify recyclable materials, supporting the statistical results obtained in the ANOVA. The most notable drops occurred between clean and low contamination, and between low and moderate contamination levels, showing that even small amounts of contamination progressively obscure key visual features needed for correct classification. Once a certain threshold of contamination is passed, the model struggles to distinguish recyclables from non-recyclables, reinforcing the idea that contamination, such as color, texture, and shape acts as an occlusion that increases misclassification likelihood.

Our data shows how contamination may impact real-life recycling efforts. When a contaminated recyclable is sorted into the wrong category, visual distortions, such as color (i.e., grease stains), texture (i.e., food remnants), or shape (i.e., crinkled packaging), can distract the CNN from recognizing key identifying features. Similar sorting systems using CNN exist in the world. For instance, automated material recovery facilities have started to use AI-based sorting systems in international recycling centers, including facilities in Japan. These systems face significant challenges and obstacles however when they process contaminated items. They face this hurdle due to reduced image recognition accuracy (10). Ultimately, these distortions lead to misclassification, where non-recyclables enter clean streams. Our findings suggest that implementing targeted preprocessing, such as contaminant detection, removal, or human oversight, could directly improve CNN sorting accuracy, consistent with a prior study on occlusion effects in deep learning models (6).

A potential solution to contamination-related misclassification is to implement a pre-sorting step that removes contaminants before machine sorting. More advanced preprocessing approaches, such as image segmentation and contamination detection algorithms, could boost CNN accuracy by allowing the model to focus only on uncontaminated parts of a product. In addition, a hybrid AI-human sorting system could serve as a safeguard against erroneous classifications, preventing major sorting errors.

A limitation of our study is that it was conducted under

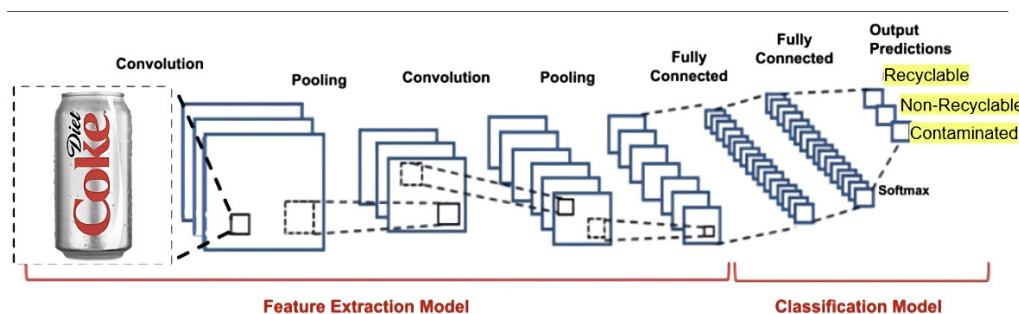


Figure 1. CNN architecture and preprocessing steps. Input images were typical waste items, including plastic bottles, soda cans, cardboard boxes, and food wrappers, and were resized and standardized before entering the network. The dotted-line boxes represent preprocessing and feature extraction, while the solid-line boxes represent classification stages. The convolutional neural network included three convolutional layers, each followed by a rectified linear unit (ReLU) activation and max pooling, prior to fully connected layers for final classification. The final model was based on EfficientNetB0 with additional Squeeze-and-Excitation (SE) blocks to enhance sensitivity to contamination features.

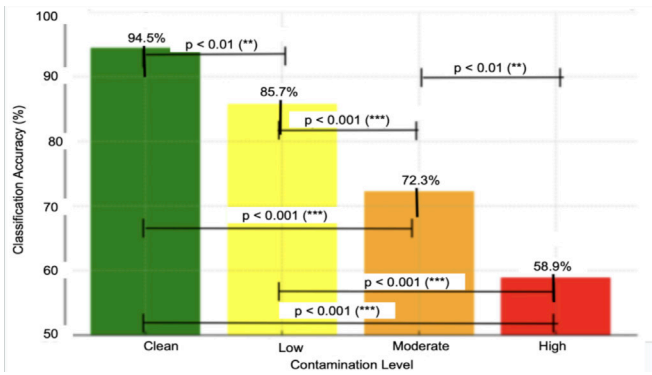


Figure 2. Impact of contamination on CNN classification accuracy. The y-axis represents mean classification accuracy, defined as the proportion of correctly predicted waste categories across 10 experimental runs. The x-axis represents contamination levels, categorized as clean (no visible residue), low (<25% coverage with contaminants such as grease or food), moderate (25-50% coverage), and high (>50% coverage). Error bars represent ± 1 standard deviation. A one-way ANOVA ($\alpha = 0.05$) followed by Tukey's Honestly Significant Difference (HSD) post-hoc tests was conducted to compare groups. Asterisks (*) denote statistically significant pairwise differences.

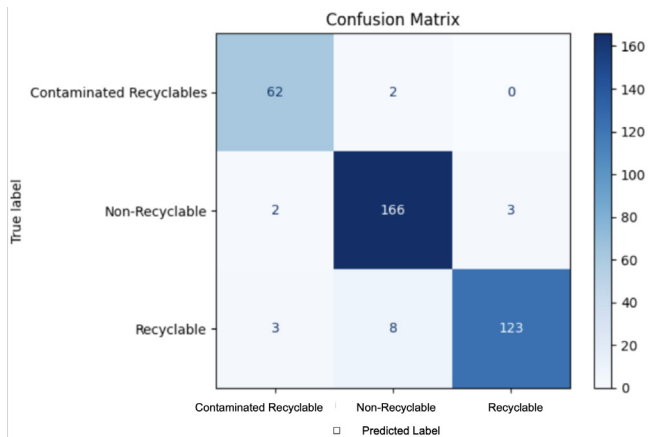


Figure 3. Classification accuracy across three waste categories. The confusion matrix was used to visualize the model's performance across each waste category (Recyclable, Non-Recyclable, and Contaminated Recyclable) and identify misclassifications. The color scale (0-160) indicates the number of images predicted for each class, with darker colors representing higher counts. The x-axis represents the predicted labels, while the y-axis represents the true labels.

controlled laboratory conditions. In real-world settings, variations in lighting, object positioning, and contamination types may further impact classification accuracy. Future research should explore adaptive CNN architectures capable of learning from real-world waste conditions and implement data augmentation techniques to improve generalizability. While contamination level was the primary variable studied, future studies should explore whether the material in question (such as paper versus plastic versus metal) affects classification as well. If trained images are given additional category labels based on materials, the study could find whether the model has more difficulty with certain materials—such as reflective aluminum surface versus a textured

cardboard one. This would allow researchers to target augmentation or training specifically for those materials but still allow the model to perform properly across all forms of waste.

Our research suggests that contamination reduces CNN classification accuracy for waste recovery. Future work should involve more adaptive approaches, such as hyperspectral imaging to detect contaminants at different wavelengths, and fluid heuristics, which take the contamination into account automatically after detection. In addition, more extensive pre-processing (i.e. segmentation or aggressive filtering) and ensemble approaches could reduce the impact of contamination even more and create classifiers that adapt better to the real-world recycling environment. Overall, this study shows how AI can be used to make recycling smarter and more efficient. If machines can better recognize waste, it could make recycling safer for workers, reduce pollution, and help create a more sustainable future. While this study focused exclusively on contamination, future research should also investigate other influencing factors, such as material type or texture, to provide a more comprehensive understanding of CNN model behavior in waste classification.

MATERIALS AND METHODS

Dataset preparation

A dataset of more than 100,000 labeled images was generated by the author for this study using photographs of common household waste collected from local recycling and trash bins. The dataset contained three categories: clean recyclables (e.g., bottles, cans, paper), non-recyclables (e.g., plastic bags, food wrappers), and contaminated recyclables. Contamination levels were simulated by adding substances such as grease stains, food remnants, and dirt to otherwise clean items, with low, moderate, and high levels defined by the visible extent of coverage on the object surface.

The contamination levels present in the images reflect reality: food (sauce, oil), dirt, and grease, the primary reasons for rejecting materials in the recycling stream. Accordingly, these contaminants were artificially created from clean recyclables to generate low, medium, and high contaminated scenarios. Contamination was determined by percentage of the recyclable covered; low contamination was smudges or speckles (less than 25% coverage), medium was obvious (25-50% coverage), and high contamination was excessive (greater than 50% coverage). All images were taken inside a white box under an LED lamp to reduce external noise unrelated to contamination. Once compiled, the dataset was pre-processed for image resizing (224×224 pixels) and pixel value standardization.

CNN model construction and data analysis

A CNN model with three convolutional layers, each followed by a rectified linear unit (ReLU) activation and max pooling, was implemented using Python 3.10 with TensorFlow 2.12 and Keras 2.12. The model was trained using categorical cross-entropy loss and an Adam optimizer for 50 epochs, with a training-validation-test split of 70-15-15%. Model performance was evaluated using classification accuracy, precision, recall, F1-score, confusion matrices, and statistical analyses, including a one-way ANOVA with post-hoc Tukey's tests.

The classification model was created using EfficientNetB0 as a base, a CNN architecture that was also pre-trained on Im-

geNet (7). Squeeze-and-Excitation (SE) blocks were overlaid to the architecture to ensure the final classifier would be more sensitive to features of contamination. The subsequent architecture was called “ContaminetLite” and the overlay of the SE blocks permitted the recognition of final spatial features more relevant to the changes induced by contaminants. All development and implementation were carried out solely by the author, and the complete codebase has been publicly shared via CodeOcean (9). This all ran on TensorFlow and Keras.

Experimental procedure

All experiments were carried out in triplicate: each model configuration trained and tested on three different instances. Mean classification accuracy and standard deviations from the experiments were calculated. The mean evaluation metrics were also calculated with standard deviations from the experiments.

Statistical analysis

To further examine differences between specific contamination levels, classification accuracy was compared using a one-way ANOVA ($\alpha = 0.05$), followed by post-hoc Tukey’s Honest Significant Difference (HSD) tests to determine pairwise significance.

ACKNOWLEDGEMENTS

We want to thank Sachin Satishkumar, a student at the University of Maryland, for helping us get started with machine learning and for assisting with the technical aspects of the project. We would also like to acknowledge Abhijit Choudhary for offering helpful suggestions and advice from a business perspective that helped guide the project.

Received: February 15, 2025

Accepted: September 6, 2025

Published: April 30, 2026

REFERENCES

1. Grand View Research, Inc. “Waste Management Industry Growth to Soar with Rising Demand for Recycling Services - Grand View Research, Inc..” *PR Newswire: Press Release Distribution, Targeting, Monitoring and Marketing*, Cision PR Newswire, 11 Apr. 2023, www.prnewswire.com/news-releases/waste-management-industry-growth-to-soar-with-rising-demand-for-recycling-services---grand-view-research-inc-301794243.html. Accessed 15 Oct. 2025.
2. Brunn, Michael. “Nearly 17% of Recycling Is Wasted Due to Contamination.” *RECYCLING Magazine*, 26 Oct. 2022, www.recycling-magazine.com/2022/10/26/nearly-17-of-recycling-is-wasted-due-to-contamination/. Accessed 15 Oct. 2025.
3. “14 Recycling Contamination Facts That Will Blow Your Mind.” *Rubicon*, 6 Feb. 2023, www.rubicon.com/blog/recycling-contamination-facts/. Accessed 15 Oct. 2025.
4. Lecun, Y., et al. “Gradient-Based Learning Applied to Document Recognition.” *Proceedings of the IEEE*, vol. 86, no. 11, 1998, pp. 2278–324. <https://doi.org/10.1109/5.726791>.
5. Sultana, Ramana. “Trash and recyclable material

- identification using convolutional neural networks (CNN).” UNCG, <https://libres.uncg.edu/ir/wcu/f/Sultana2020.pdf>. Accessed 15 Oct. 2025.
6. Garcia-Garcia, Alberto, et al. “A Study of the Effect of Noise and Occlusion on the Accuracy of Convolutional Neural Networks Applied to 3D Object Recognition.” *Computer Vision and Image Understanding*, vol. 164, Nov. 2017, pp. 124–34. <https://doi.org/10.1016/j.cviu.2017.06.006>.
7. Tan, Mingxing, and Quoc V. Le. “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks.” *Proceedings of the 36th International Conference on Machine Learning*, 2019, pp. 6105-6114. <https://doi.org/10.48550/arXiv.1905.11946>.
8. Advancing Sustainable Materials Management: 2017 Fact ..., www.epa.gov/sites/default/files/2019-11/documents/2017_facts_and_figures_fact_sheet_final.pdf. Accessed 15 Oct. 2025.
9. Aarav Rao and Nishit Anand. “Smart Waste Classification Using CNNs with SE Blocks on Edge Devices.” *Code Ocean*, version 2.0, 2025, <https://doi.org/10.24433/CO.8327446.V2>.
10. Cheng, Tianhao, et al. “Optimizing Waste Sorting for Sustainability: An AI-Powered Robotic Solution for Beverage Container Recycling.” *Sustainability*, vol. 16, no. 23, 2024, p. 10155. <https://doi.org/10.3390/su162310155>.

Copyright: © 2026 Rao and Mohan. All JEI articles are distributed under the Creative Commons Attribution Noncommercial No Derivatives 4.0 International License. This means that you are free to share, redistribute, remix, transform, or build upon the material for any purpose, provided that you credit the original author and source, include a link to the license, indicate any changes that were made, and make no representation that JEI or the original author(s) endorse you or your use of the work. The full details of the license are available at <https://creativecommons.org/licenses/by-nc-nd/4.0/deed.en>.