

Olha Zalutska

Assistant of Computer Science Department

Oleksandr Mazurets

Ph.D. in Engineering Science, Associate Professor of Computer Science Department

Maryna Molchanova

Ph.D. Senior Lecturer of Computer Science Department

Khmelnytskyi National University, Ukraine

EFFICIENCY ANALYSIS OF WRECKING WASTE CLASSIFICATION USING NEURAL NETWORK

Abstract. This study examines the efficiency of a neural approach to identifying and classifying wrecking waste in photographic data. Building on a two-factor design that couples scene-level detection with per-class binary refinement, we quantify where accuracy is earned, where it is lost, and which components: architecture, data composition, augmentation, and arbitration – most strongly govern the outcome. Using a composite dataset of ten material classes, we report that the hybrid method attains up to 97.8% classification accuracy, with substantial gains over a detector-only baseline on heterogenous or texture-confusable categories. Detector metrics such as $mAP50 \approx 0.746$ and $mAP50-95 \approx 0.669$ confirm reliable localization, while binary residual classifiers close the gap in label assignment, lifting macro F_1 from ≈ 0.751 to ≈ 0.974 . These findings indicate that efficiency, measured as correct material routing per unit inference cost, is maximized when detection and classification are decoupled yet reconciled through a calibrated arbitration rule.

C&D waste is visually complex due to overlapping fragments, contamination, and shared texture–color statistics [1]. End-to-end detectors with multi-class heads often show cross-class interference and unstable false positives, whereas a split pipeline – with detection followed by one-vs-rest specialist classifiers – improves stability and adaptability [2].

This study compares the efficiency and robustness of a two-stage approach against a detector-only system. The pipeline starts with YOLOv11 trained on a curated dataset of ten C&D classes (brick, concrete, foam, general waste, gypsum board, pipes, plastic, stone, tile, wood). Detector outputs (640×640 inputs) provide bounding boxes and coarse class hypotheses. Cropped fragments are then processed by binary ResNet50/50V2 classifiers trained per class with restrained augmentation – rotations, affine jitter, and mild photometric variation simulating realistic conditions. During inference, a calibrated consensus validates the binary decision if its posterior exceeds a threshold and aligns with the detector; otherwise, the detector’s label prevails. Efficiency is evaluated at detection (localization, coverage) and classification (Accuracy, Precision, Recall, F_1 , macro-averaged) levels.

The image corpus integrates an object-detection dataset collected at an industrial sorting facility together with stacked-waste and RGB-D supplements to improve coverage of concrete, brick, wood, and related classes. This composite was split into detector training/validation sets and, separately, into per-class binary subdatasets.

YOLOv11 achieves robust localization with overall $mAP50 \approx 0.746$ and $mAP50-95 \approx 0.669$, as well as mean IoU near 0.897, indicating that box proposals are

generally well placed even under heterogeneous backgrounds. Classwise performance varies: categories with distinctive macro-texture and geometry (brick, concrete, gypsum board, tile) maintain high Precision–Recall pairs, whereas visually heterogeneous or small-scale classes (foam, pipes, plastic) degrade, a pattern typical of detectors operating under partial occlusion and surface contamination. These detector characteristics are desirable for throughput few missed regions and reliable boxes yet they leave residual label noise that impacts downstream routing costs.

The binary ResNet stage yields decisive improvements in attribution. On a stratified evaluation, the detector’s macro metrics centered near Accuracy ≈ 0.601 and F1 ≈ 0.751 . After refinement, macro Accuracy rises to ≈ 0.975 and macro F1 to ≈ 0.974 , with the largest deltas on historically weak classes. For foam, F1 improves from ≈ 0.544 to ≈ 0.988 ; for pipes, from ≈ 0.450 to ≈ 0.959 ; for plastic, from ≈ 0.455 to ≈ 0.937 . Even for strong classes such as brick and concrete, residual classifiers consolidate Recall without materially harming Precision. Aggregate best-case accuracy over the ten classes reaches $\approx 97.8\%$, evidencing that the detector’s spatial prior plus specialist refinement is more sample-efficient than increasing detector capacity alone.

Training-time sweeps over epochs reveal a typical U-shaped risk curve: gains accrue through 10-15 epochs for most classes, while further epochs occasionally erode minority-class generalization, especially for categories with higher intra-class variability such as plastic or pipes. Comparisons with MobileNetV2 show that compact backbones can approach ResNet performance on simpler classes, but deeper residual models better capture micro-texture cues that separate confusable materials. From an efficiency standpoint, this suggests a mixed strategy: retain MobileNet-scale specialists for easy, high-frequency categories and reserve ResNet-scale specialists for the hard tail to maximize accuracy per compute budget.

References:

1. M. Molchanova, V. Didur, O. Mazurets, O. Sobko, O. Zakharkevich, Method for Construction and Demolition Waste Classification Using Two-Factor Neural Network Image Analysis, in: CEUR Workshop Proceedings, 3970, 2025, pp. 168–182.
2. O. Hladun, O. Zalutska, V. Klimenko, O. Mazurets, Research on the Effectiveness of Classifying the Remains of Destroyed Buildings Using MobileNetV3 Neural Network Architecture, in: Innovations in Science: From Theoretical Foundations to Practical Impact, Proceedings of the 1st International Scientific and Practical Conference, Antwerp, Belgium, 2025, pp. 158–162.