ELECTRICAL ENGINEERING AND ENERGY

OPEN ACCESS

Defect Classification of Medium-Voltage Overhead Line Conductors Using YOLO11

Oya Kilci[1], Mucahid Mustafa Saritas[2], and Murat Koklu[3]

[1] Department of Mechatronics, Vocational School of Trade and Industry, KTO Karatay University, Konya, Türkiye
[2] Department of Computer Engineering, Graduate School of Natural and Applied Sciences, Selcuk University, Konya, Türkiye
[3] Department of Computer Engineering, Technology Faculty, Selcuk University, Konya, Türkiye

Submitted: 28.09.2026
Accepted : 12.11.2026
Online First: 19.112026

Corresponding author,
Murat Koklu
mkoklu@selcuk.edu.tr

DOI:10.64470/elene.2026.50

RESEARCH ARTICLE

Copyright, Authors, Distributed under Creative Commons CC-BY 4.0

Abstract: Defects in medium-voltage overhead line conductors may adversely affect the reliability and continuity of electrical distribution networks, highlighting the importance of effective automated inspection methods. This study investigates whether region-of-interest (ROI) cropping improves the classification of overhead conductor conditions compared with conventional full-image input strategies. A two-branch classification framework based on the YOLO11 architecture was constructed using a publicly available three-class image dataset representing medium-voltage overhead line conductors under normal, surface-defect, and strand-fraying conditions. The dataset was partitioned into 1,440 training, 180 validation, and 180 test images. While the full-image branch classified the original images directly, the ROI branch employed ground-truth bounding-box annotations to extract individual conductor regions, yielding 323 object-level samples in the held-out test split. YOLO11n-cls, YOLO11s-cls, and YOLO11m-cls were evaluated using accuracy together with macro-averaged precision, recall, and F1-score. Full-image classification yielded accuracies of 94.44%, 93.89%, and 93.89%, with corresponding macro F1-scores of 93.91%, 93.92%, and 93.48%, respectively. ROI-based classification yielded corresponding accuracies of 98.76%, 97.21%, and 97.52% and macro F1-scores of 98.56%, 96.87%, and 97.40%, with descriptive accuracy differences of 4.32, 3.32, and 3.63 percentage points relative to the full-image results, respectively. YOLO11n-cls achieved the highest observed ROI-based accuracy and F1-score under the evaluated experimental setting. Because the two branches operated on different evaluation units (whole images and object-level crops), the comparison is interpreted as a strategy-level analysis rather than a direct paired evaluation.

Keywords Deep learning, Medium-voltage overhead line, Overhead conductor, Defect classification, Region of interest, YOLO11.

1. Introduction

The reliable operation of overhead electrical distribution networks depends strongly on the structural condition of their conductors. Surface deterioration, strand damage, and localized structural abnormalities can affect conductor integrity and may increase the risk of deterioration or service disruption. Visual

inspection is therefore an important component of overhead-line condition assessment. However, conventional inspection procedures may be time-consuming and dependent on human observation, particularly when defects are small or difficult to distinguish from the surrounding scene. Advances in deep learning have created new opportunities for automated image-based inspection of overhead electrical infrastructure.

The increasing need for reliable visual condition assessment has emphasized the importance of robust and automated defect detection and classification systems (Hussain, 2023). Traditional visual inspection methods, although widely used, are increasingly supported by automated solutions because of their limitations in consistency, speed, and accuracy, particularly when dealing with subtle or low-contrast defects (Meng et al., 2025). This transition is particularly relevant to overhead line inspection and conductor condition assessment, where localized surface and strand-level abnormalities may be difficult to identify reliably using conventional visual inspection alone (Meng et al., 2025).

Early automated defect detection approaches primarily relied on conventional computer vision techniques, including thresholding, edge detection, morphological operations, and handcrafted feature extraction. However, these methods often struggle in complex visual inspection environments characterized by varying illumination conditions, diverse defect appearances, and complex background information. The emergence of deep learning, particularly convolutional neural networks (CNNs), has substantially improved defect detection performance by enabling automatic hierarchical feature learning directly from raw image data (Khanam et al., 2024). Models such as ResNet, EfficientNet, and MobileNet have demonstrated strong performance in classification tasks due to their ability to extract complex and discriminative features. However, when defects occupy only a small portion of the image, surrounding scene information may complicate the extraction of discriminative defect features.

To address these limitations, object detection frameworks such as the YOLO (You Only Look Once) family have gained widespread attention due to their ability to perform localization and classification simultaneously within a unified architecture (Mao & Hong, 2025). These models simultaneously predict bounding boxes and class labels, enabling efficient real-time visual inspection (Nahar et al., 2026). The YOLO framework has evolved significantly from its early versions to more advanced architectures such as YOLO5, YOLO8, and YOLO11, incorporating various architectural improvements to enhance both inference speed and detection accuracy. This evolution has enabled a shift from traditional full-image classification toward more refined object-level analysis, where defects are localized and analyzed individually. Compared to two-stage detectors such as Faster R-CNN, one-stage detectors like YOLO provide superior inference speed with competitive accuracy, making them suitable for applications requiring rapid and automated visual analysis (Lu et al., 2024; Tian et al., 2025; Wang et al., 2026).

Recent advancements in YOLO-based models have focused on improving feature extraction, multi-scale feature fusion, and attention mechanisms to enhance defect detection performance. For instance, YOLOv5 and subsequent models incorporate advanced backbone structures, improved loss functions, and efficient feature aggregation strategies to achieve higher robustness across diverse defect types (Wan et al., 2024). Further improvements in models such as YOLOv7 and YOLOv8 include optimized feature fusion strategies and lightweight architectures, enabling efficient real-time detection without compromising accuracy (Lema et al., 2025). In addition, modern variants such as YOLO11 introduce enhanced feature extraction modules, attention mechanisms, and multi-scale processing capabilities, significantly improving performance in challenging visual defect-detection scenarios (Li et al., 2025; Song et al., 2025).

Region-of-interest (ROI)-based preprocessing enables models to focus on local features by reducing the surrounding background information presented to the classifier (Pan et al., 2024). In visually complex inspection scenes, ROI-based methods are particularly useful for detecting small or low-contrast defects. Studies such as Chen et al. (2025) demonstrate that integrating region-based processing with attention

mechanisms significantly improves localization accuracy and detection performance in printed circuit board (PCB) inspection tasks.

Moreover, advanced architectures such as YOLOv11-WBD incorporate wavelet decomposition, cross-attention mechanisms, and dilated convolution structures to enhance feature representation by separating low-frequency and high-frequency components (Guan et al., 2025). Similarly, attention-enhanced feature fusion mechanisms and adaptive feature focusing strategies have been shown to significantly improve small object detection performance in complex environments (Yan et al., 2025; W. Zhang et al., 2026; Y. Zhang et al., 2021). Additionally, multi-scale feature fusion approaches have been successfully applied in industrial inspection tasks such as PCB defect detection, where defects often exhibit complex patterns and small spatial scales (Bai & Xu, 2025).

ROI-based preprocessing restricts the classifier input to selected image regions; however, its computational cost depends on the number of crops processed per source image (Su et al., 2022). Furthermore, recent studies combining ROI-based strategies with advanced YOLO architectures, including YOLOv8 and YOLO11, have demonstrated improved sensitivity to fine-grained defect patterns and enhanced classification robustness in visual defect-detection applications (Bai & Xu, 2025; Chen et al., 2025; Wang et al., 2026).

Despite these advancements, a clear research gap remains in the direct comparison of full-image and ROI-based classification strategies for the visual assessment of overhead line conductor defects using modern YOLO11 classification models. Most existing studies focus on either full-image classification or object detection, while fewer studies quantify how object-focused cropping affects classification performance when the model family is held constant. This distinction is particularly relevant for overhead line imagery because localized conductor defects may occupy only a limited portion of the captured scene, while surrounding visual information may influence classification performance.

Accordingly, this study develops a comparative framework for classifying medium-voltage overhead line conductor conditions using the same three YOLO11 classification scales under two input strategies: original full images and annotation-derived ROI crops. The study establishes a full-image baseline, converts YOLO bounding boxes into object-level classification samples, and compares accuracy and macro-averaged precision, recall, and F1-score for YOLO11n, YOLO11s, and YOLO11m. In addition to reporting the observed model rankings under the evaluated data split, the analysis explicitly describes the different evaluation units used by the two strategies and the limitations of the comparison.

2. Materials and Methods

A systematic two-branch workflow was designed to compare full-image and ROI-based classification of medium-voltage overhead line conductor conditions. First, the publicly available dataset was prepared using its original train, validation, and test partition. In the full-image branch, each original source image was used as a single classification sample without annotation-derived ROI extraction. In the ROI branch, the YOLO-format bounding-box annotations were used to extract individual object regions, and each crop inherited the class label of its corresponding annotation. The resulting datasets were evaluated with YOLO11n-cls, YOLO11s-cls, and YOLO11m-cls. Final performance was assessed using accuracy and macro-averaged precision, recall, and F1-score. The comparative workflow is illustrated in Figure 1.

2.1. Dataset

The dataset used in this study was created by Masters Workspace and is publicly available as "project Dataset" on Roboflow Universe (Masters Workspace, 2026). It is distributed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Version 3 of the dataset, identified as "original auto-orient rebalanced", was used in this study and contains annotated conductor regions belonging to three classes: normal, surface_defect, and wire_fray. The source annotations use numeric class identifiers, which were mapped to descriptive labels in this study as follows: 0 = normal, 1 = surface_defect, and 2 = wire_fray. In

the context of the present study, these classes represent normal conductor condition, conductor surface defects, and conductor strand fraying, respectively. Representative full-image examples from the three classes, together with their annotated bounding boxes, are shown in Figure 2.

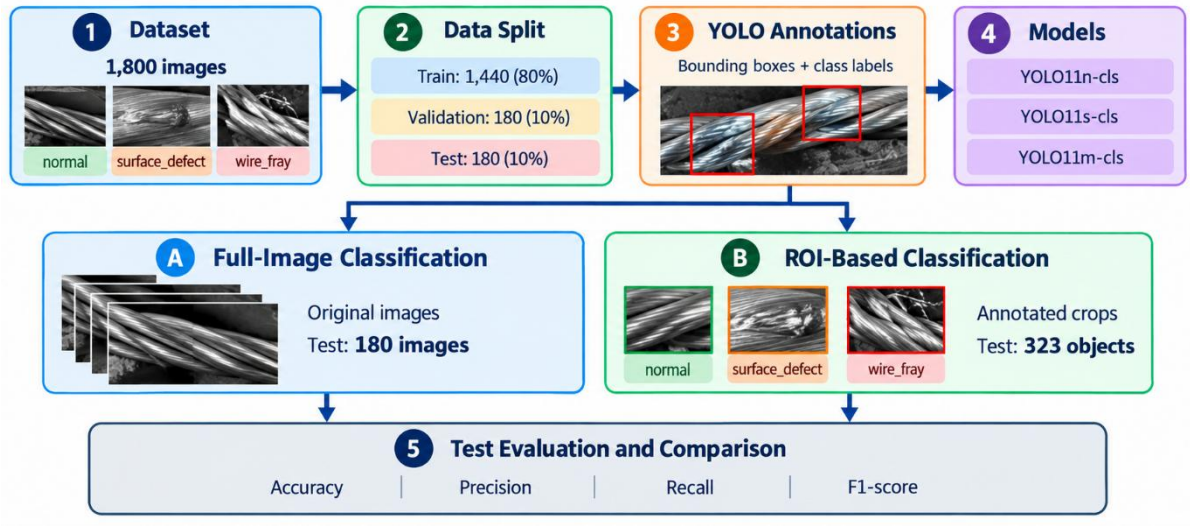

Figure 1. Flow diagram of the comparative full-image and ROI-based YOLO11 overhead conductor defect classification framework

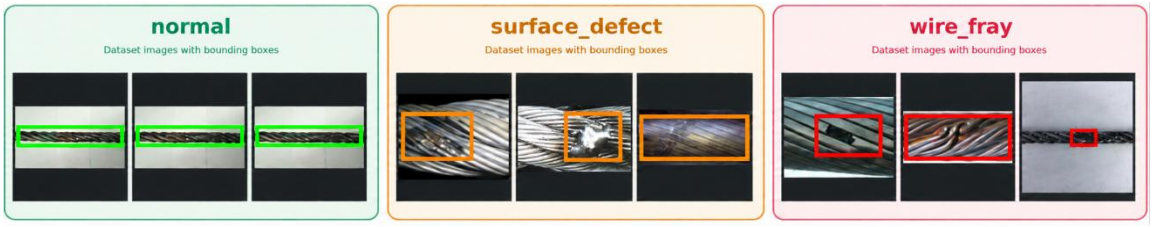

Figure 2. Representative full-image examples from the dataset with annotated bounding boxes

The dataset contains 1,800 source images, divided into training, validation, and test subsets comprising 1,440, 180, and 180 images, respectively. This 80/10/10 hold-out partition was preserved when constructing both the full-image and ROI datasets. Each ROI crop remained in the same subset as its source image. The classification datasets used three class names: normal, surface_defect, and wire_fray. The source images were annotated in YOLO format, with each annotation containing a numeric class identifier and normalized bounding-box coordinates.

Two input strategies were generated from these annotations. In the full-image strategy, each original image was used directly for classification and assigned the class of its first annotation. This rule also applied to images containing annotations from multiple classes. In the ROI-based strategy, each annotated bounding box was extracted as an independent crop and assigned to the corresponding annotation class. Because individual source images could contain multiple annotated regions, this procedure yielded 2,392 training, 327 validation, and 323 test crops, totaling 3,042 ROI samples.

Full-image evaluation was conducted on 180 test images, whereas ROI-based evaluation was conducted on 323 test crops. The strategies therefore used different evaluation units derived from the same held-out source-image subset. Their results were compared as alternative input-processing approaches rather than as paired predictions on identical samples. Table 1 presents the class counts for the training, validation, and test splits at both the image and ROI levels. Representative ROI crops extracted from the annotated bounding boxes shown in Figure 2 are presented in Figure 3.

Table 1. Class counts in the training, validation, and test splits at the image and ROI levels

Level	Split	normal	surface_defect	wire_fray	Total
Image	Training	794	503	143	1,440
Image	Validation	101	63	16	180
Image	Test	96	62	22	180
ROI	Training	971	606	815	2,392
ROI	Validation	147	69	111	327
ROI	Test	118	76	129	323



Figure 3. Corresponding ROI crops extracted from the annotated bounding boxes shown in Figure 2

2.2. Full-image and ROI-based input strategies

For the full-image baseline, the original source images were used directly as classification inputs without annotation-derived ROI extraction. Because the source dataset was originally prepared for object detection, individual images could contain more than one annotation record, whereas the classification task required a single image-level label. During conversion to the full-image classification dataset, the class label associated with the first annotation record was therefore assigned as the image-level target when multiple annotations were present. This deterministic conversion rule was applied consistently to all source images. The resulting full-image test set contained 180 image-level samples.

For the ROI branch, each YOLO bounding box was used to crop the corresponding object region from the source image, and the crop inherited the bounding-box class label. Multiple ROI samples could therefore originate from one source image. The ROI test set contained 323 crops. Importantly, these ROIs were generated from the ground-truth annotation boxes rather than from detector-predicted boxes; consequently, the experiment evaluates the effect of object-focused input representation on classification rather than end-to-end detection performance. Bounding-box coordinates were converted to integer pixel coordinates and clipped to the source-image boundaries. No additional margins or padding were applied during ROI extraction. During training, random resized cropping produced 224 × 224-pixel inputs using bilinear interpolation. During validation and testing, the shorter edge was resized to 224 pixels while preserving the aspect ratio, followed by a central crop of 224 × 224 pixels. No letterbox padding was used.

The two branches were compared at the strategy level using the same three YOLO11 model scales. Since their evaluation units differ (180 images versus 323 object crops), numerical differences are reported descriptively and are not treated as a paired sample-level significance test.

2.3. YOLO11 classification models

YOLO (You Only Look Once) has evolved from a real-time object detection framework into a broader family of deep learning models supporting multiple computer vision tasks, including image classification. Its efficient feature extraction capability and computationally lightweight design have contributed to its widespread use in automated image analysis applications (Ali & Zhang, 2024; Kilci & Koklu, 2025a).

In addition to object detection, recent YOLO architectures can be adapted for image classification, where the objective is to assign an input image to one of the predefined classes. This capability makes YOLO-based classification models suitable for visual defect recognition tasks in which both computational efficiency and reliable feature extraction are important (Kilci & Koklu, 2025b; Sirisha et al., 2023).

YOLO11 incorporates architectural components designed to improve feature extraction and representation efficiency. Among these components, the C3k2 module supports efficient hierarchical feature learning, while the C2PSA mechanism enhances feature representation by emphasizing informative spatial regions. These mechanisms contribute to the extraction of discriminative visual features required for image classification tasks (Song et al., 2025; Zhong et al., 2025).

For overhead conductor defect classification, effective feature representation is particularly important because defect-related visual patterns may be small, localized, or visually similar to normal conductor surface structures. Deep learning architectures capable of extracting hierarchical visual information have therefore become increasingly important for automated surface inspection and defect recognition applications (Khanam et al., 2024). YOLO11 is available in different scales, including nano (YOLO11n), small (YOLO11s), medium (YOLO11m), and larger variants. These variants provide different trade-offs between computational complexity and representation capacity.

In the present study, YOLO11n-cls, YOLO11s-cls, and YOLO11m-cls were selected to investigate whether increasing model scale provides additional classification benefits for overhead conductor defect recognition. The same three YOLO11 classification variants were applied to both the full-image and ROI-based input strategies. This experimental design maintains the model family across the two branches and enables the effect of input representation to be examined independently of changes in the selected classification architecture.

2.4. Experimental protocol

For both branches, the models were trained for 50 epochs with a batch size of 16 and an input resolution of 224 × 224 pixels. The augmentation settings included RandAugment, horizontal flipping with a probability of 0.5, and random erasing with a probability of 0.4. Weight decay was set to 0.0005, with a three-epoch warm-up period. Automatic optimizer selection was enabled. Because the total number of training iterations was below 10,000 in both branches (approximately 4,500 for the full-image branch and 7,500 for the ROI branch), this procedure selected the AdamW optimizer with an initial learning rate of 0.001429 for both branches. All models were trained on an Intel Core i7-13700H CPU.

3. Performance Evaluations

Performance was evaluated separately for the full-image and ROI-based branches. The full-image branch was assessed on 180 image-level test samples, while the ROI branch was assessed on 323 object-level crops derived from the same held-out source split. Accuracy, macro-averaged precision, macro-averaged recall, and macro-averaged F1-score were used to summarize predictive performance. Confusion-matrix analysis was performed for both branches to examine class-specific correct predictions and error patterns.

3.1. Confusion matrix and performance metrics

A confusion matrix summarizes the relationship between true and predicted class labels and provides a detailed view of classification performance for each class. The general structure of the three-class confusion matrix is illustrated in Figure 4. In this study, the three-class confusion matrix includes the classes normal, surface_defect, and wire_fray. Correct predictions are represented by the diagonal elements of the matrix, whereas off-diagonal elements indicate misclassifications between classes (Hayta et al., 2023; Koklu et al., 2026). The confusion matrices were used to examine class-specific prediction patterns and identify the most common classification errors (Dogansah et al., 2026; Kursun & Koklu, 2025; Saritas et al., 2025). For each class, true positives (TP) represent samples correctly assigned to that class, false positives (FP) represent samples from other classes incorrectly assigned to that class, and false negatives (FN) represent samples belonging to that class but incorrectly assigned to another class. True negatives (TN) represent samples from the other classes correctly predicted as not belonging to the class under consideration (Citak & Koklu, 2026). These quantities were determined separately for each class using a one-vs-rest approach. In overhead

conductor condition assessment, misclassification of defective samples is particularly important because missed abnormalities may reduce the reliability of automated inspection and condition-monitoring systems (Koklu et al., 2012).

The classification models were evaluated using accuracy, precision, recall, and F1-score. Accuracy represents the proportion of correctly classified samples among all test samples (overall accuracy), i.e., the sum of the diagonal elements of the confusion matrix divided by the total number of test samples. Precision indicates the reliability of the predictions assigned to each class, whereas recall measures the ability of the model to correctly identify the samples belonging to each class. For each class, the F1-score is the harmonic mean of its precision and recall. Precision, recall, and F1-score were first calculated separately for each class and then averaged across the three classes (macro-averaging), so that each class contributed equally to the reported performance (Haque et al., 2025). Accordingly, the reported macro F1-score is the arithmetic mean of the class-wise F1-scores rather than the harmonic mean of macro precision and macro recall. The mathematical definitions of the evaluation metrics are presented in Table 2 (Akbas et al., 2025; Isgor & Koklu, 2025).

Table 2. Mathematical definitions of the performance evaluation metrics

Metrics	Formula		
Accuracy	$\frac{\sum_{i=1}^{K} tp_i}{N}$	(1)	
Macro Precision	$\frac{\sum_{i=1}^{K} \frac{tp_i}{tp_i + fp_i}}{K}$	(2)	tp_i = True Positives for class i fp_i = False Positives for class i
Macro Recall	$\frac{\sum_{i=1}^{K} \frac{tp_i}{tp_i + fn_i}}{K}$	(3)	fn_i = False Negatives for class i K = Total number of classes
Macro F1-score	$\frac{1}{K}\sum_{i=1}^{K} \frac{2tp_i}{2tp_i + fp_i + fn_i}$	(4)	N = Total number of test samples

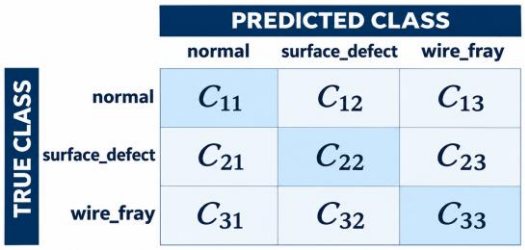

Figure 4. General structure of the three-class confusion matrix

4. Results and Discussion

The comparative evaluation considered three YOLO11 classification architectures (YOLO11n, YOLO11s, and YOLO11m) under two input strategies: full images and annotation-derived ROI crops. The full-image baseline was evaluated on 180 original test images, whereas the ROI branch was evaluated on 323 object-level crops from the same held-out source split. Table 3 reports accuracy and macro-averaged precision, recall, and F1-score for all six model-input combinations. Figures 5 and 6 provide the corresponding class-wise confusion patterns for the full-image and ROI-based branches.

Figure 5 presents the full-image confusion matrices for YOLO11n, YOLO11s, and YOLO11m on the 180-image test set. YOLO11n correctly classified 95 of 96 normal images, 55 of 62 surface_defect images, and 20 of 22 wire_fray images, yielding 170 correct predictions overall. YOLO11s correctly classified 95 normal, 54 surface_defect, and 20 wire_fray images, whereas YOLO11m correctly classified 95 normal, 55 surface_defect, and 19 wire_fray images. The dominant full-image errors occurred when surface_defect

samples were predicted as normal, suggesting that visual information outside the annotated conductor region may reduce the separability of subtle surface-level abnormalities.

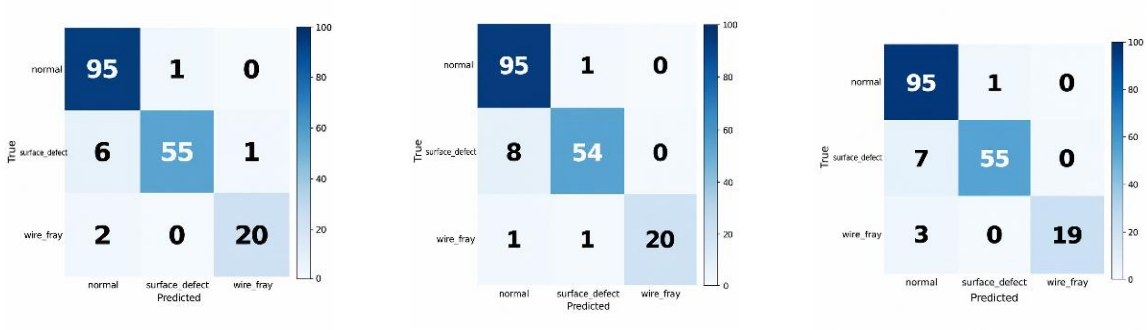

(a) (b) (c)

Figure 5. Full-image confusion matrices of (a) YOLO11n, (b) YOLO11s, and (c) YOLO11m

Figure 6 presents the ROI-based confusion matrices for the three YOLO11 models. YOLO11n achieved the highest observed ROI accuracy in this experiment, correctly classifying 116 of 118 normal samples, 74 of 76 surface_defect samples, and all 129 wire_fray samples. Four ROI samples were misclassified. These observations are limited to the evaluated test split and do not establish consistent superiority over the other models.



(a) (b) (c)

Figure 6. ROI-based confusion matrices of (a) YOLO11n, (b) YOLO11s, and (c) YOLO11m

YOLO11s correctly classified 114 normal, 72 surface_defect, and 128 wire_fray ROI samples, while YOLO11m correctly classified 115 normal, 74 surface_defect, and 126 wire_fray ROI samples. The remaining errors were concentrated mainly between normal and surface_defect or involved a small number of wire_fray samples. These patterns suggest that some surface-level irregularities share visual characteristics with both intact conductor surface patterns and more pronounced strand-level structural damage, whereas the ROI-based representation generally preserves highly discriminative local details.

The full-image baselines achieved similar overall performance across the three model scales. YOLO11n obtained the highest full-image accuracy at 94.44%, with macro precision of 95.23%, macro recall of 92.86%, and macro F1-score of 93.91%. YOLO11s and YOLO11m both achieved 93.89% accuracy. Their macro F1-scores were 93.92% and 93.48%, respectively. These small differences do not establish a consistent relationship between model scale and classification performance.

ROI-based classification produced higher reported test metrics for every YOLO11 variant. YOLO11n reached 98.76% accuracy, 98.56% macro precision, 98.56% macro recall, and 98.56% macro F1-score. YOLO11s achieved 97.21% accuracy and 96.87% macro F1-score, while YOLO11m achieved 97.52% accuracy and 97.40% macro F1-score. Relative to the corresponding full-image baselines, the reported accuracy differences were +4.32 percentage points for YOLO11n, +3.32 points for YOLO11s, and +3.63 points for YOLO11m. The corresponding macro F1 differences were +4.65, +2.95, and +3.92 percentage points, respectively.

Table 3. Comparative test performance of full-image and ROI-based YOLO11 classification models

Input	Model	Acc. (%)	Prec. (%)	Recall (%)	F1 (%)
Full image	YOLO11n	94.44	95.23	92.86	93.91
Full image	YOLO11s	93.89	95.92	92.32	93.92
Full image	YOLO11m	93.89	96.23	91.34	93.48
ROI	YOLO11n	98.76	98.56	98.56	98.56
ROI	YOLO11s	97.21	96.88	96.86	96.87
ROI	YOLO11m	97.52	97.31	97.50	97.40

The largest changes were observed in macro recall. The ROI branch exceeded the full-image branch by 5.70 percentage points for YOLO11n, 4.54 points for YOLO11s, and 6.16 points for YOLO11m. This pattern is consistent with the intended role of ROI cropping, which reduces unrelated image regions and concentrates the classifier on the conductor region associated with the annotation. YOLO11n obtained the highest observed ROI accuracy and macro F1-score under the evaluated experimental setting; however, the small differences between the models do not establish a reliable model ranking. From the perspective of overhead line inspection, the advantage of ROI-based input representation is particularly relevant when the conductor or localized defect occupies only a limited portion of the captured image. Under such conditions, full-image classification requires the model to process both the target conductor region and additional scene information. Annotation-derived cropping provides a controlled means of isolating the visually relevant region and therefore helps clarify how strongly the classification models benefit from conductor-focused input representation.

The comparison should nevertheless be interpreted with two methodological qualifications. First, the full-image branch uses 180 image-level samples, whereas the ROI branch uses 323 object-level crops generated from those source images, so the reported improvements do not represent a one-to-one paired test on identical sample units. Second, the ROI crops were generated from ground-truth bounding-box annotations rather than from predicted detector boxes. Accordingly, the present experiment isolates the classification benefit of object-focused input preparation, but it does not yet quantify localization errors or end-to-end detector-plus-classifier performance. Furthermore, the results are based on a single random seed and a fixed data split, and variability across repeated runs was not assessed. One test sample corresponds to approximately 0.56 percentage points in full-image accuracy and 0.31 percentage points in ROI accuracy. YOLO11n correctly classified only one more full image than either larger model, and five and four more ROI samples than YOLO11s and YOLO11m, respectively. Therefore, these differences should be interpreted cautiously. Evaluation across multiple random seeds and data splits is needed to assess the robustness of the model rankings.

5. Conclusion

This study presented a comparative evaluation of full-image and annotation-derived ROI-based classification of medium-voltage overhead line conductor conditions using YOLO11n, YOLO11s, and YOLO11m. The original 1,800-image dataset was preserved as an image-level baseline, while YOLO-format bounding boxes were additionally used to generate object-level crops for the ROI branch. This design enabled the effect of input representation to be examined while maintaining the same model family across both strategies.

On the 180-image full-image test set, YOLO11n achieved the highest accuracy of 94.44%, while YOLO11s and YOLO11m each achieved 93.89%. On the 323-crop ROI test set, the corresponding accuracies were 98.76%, 97.21%, and 97.52%, respectively. YOLO11n also achieved the highest ROI-based macro precision, recall, and F1-score, all at 98.56%. The observed accuracy differences between the ROI and full-image strategies were 4.32 percentage points for YOLO11n, 3.32 percentage points for YOLO11s, and 3.63 percentage points for YOLO11m.

These findings show that object-focused cropping was associated with higher classification performance for localized overhead conductor abnormalities by reducing surrounding scene information and

concentrating feature extraction on the annotated conductor region. YOLO11n achieved competitive classification performance relative to the larger variants in this experiment. However, the single-seed, fixed-split evaluation does not establish that a smaller model will consistently outperform larger models. Repeated experiments are required before drawing broader conclusions about model selection.

This study also has important limitations. The full-image and ROI branches use different evaluation units, and the ROI test set contains multiple crops that may originate from the same source image. Therefore, the observed performance differences should be interpreted as a strategy-level comparison rather than as a strictly paired gain. In addition, ROI extraction relies on ground-truth bounding boxes; consequently, the reported results represent classification performance under accurate localization rather than complete end-to-end detection performance. Future work should evaluate performance across multiple random seeds and data splits, assess detector-predicted ROIs, employ grouped or source-image-aware validation to account for dependence among crops, test larger and more diverse conductor datasets, and investigate edge deployment and explainable artificial intelligence methods.

Declaration of Ethical Standards

As the authors of this study, we declare that we comply with all ethical standards.

This study did not involve human participants, animals, or personal identifiable data. The dataset used in this study is publicly available and was used in accordance with its Creative Commons Attribution 4.0 (CC BY 4.0) license. Therefore, ethics committee approval was not required. The authors declare that the study was conducted in accordance with generally accepted ethical standards for scientific research.

Credit Authorship Contribution Statement

O. Kilci: Conceptualization, Methodology, Data Curation, Software, Validation, Formal analysis, Visualization, Writing – Original Draft, Writing – Review & Editing.

M. M. Saritas: Methodology, Software, Validation, Formal analysis, Visualization, Writing – Review & Editing.

M. Koklu: Conceptualization, Methodology, Supervision, Validation, Writing – Review & Editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this study.

Acknowledgements

The authors would like to thank the creators and contributors of the publicly available dataset used in this study.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Data Availability

The source dataset is publicly available on Roboflow Universe at https://universe.roboflow.com/masters-workspace-yqrrt/project-p314o-2stwo/dataset/3 under the CC BY 4.0 license. The ROI subsets were derived from the source images using the provided YOLO-format bounding-box annotations. The numeric source class identifiers were mapped to the descriptive labels normal, surface_defect, and wire_fray in the classification datasets.

References

Akbas, I., Kilci, O., & Koklu, M. (2025). Deep learning based classification of real and synthetic animal images. *Intelligent Methods in Engineering Sciences*, *4*(3), 114–128. https://doi.org/10.58190/imiens.2025.158

Ali, M. L., & Zhang, Z. (2024). The YOLO framework: A comprehensive review of evolution, applications, and benchmarks in object detection. *Computers*, *13*(12), 336. https://doi.org/10.3390/computers13120336

Bai, L., & Xu, W. H. (2025). Improved printed circuit board defect detection scheme. *Scientific Reports*, *15*(1), 2389. https://doi.org/10.1038/s41598-025-85245-2

Chen, W., Ke, X., & Meng, S. (2025). Small defect detection in printed circuit boards based on the multiscale edge strengthening and an improved YOLOv10. *Scientific Reports*, *15*(1), 36445. https://doi.org/10.1038/s41598-025-20387-x

Citak, A. C., & Koklu, M. (2026). Filtering out defects in open-source meat freshness datasets: A new leak-free benchmark and a comparative analysis of four different mobile deep learning models. *Journal of Food Science*, *91*(10), e71554. https://doi.org/10.1111/1750-3841.71554

Dogansah, Y., Cizmeci, İ. H., & Koklu, M. (2026). Deep learning-based lung cancer classification and Grad-CAM++ with LIME-supported explainability analysis. *International Journal of Applied Methods in Electronics and Computers*, *14*(1), 34–48. https://doi.org/10.58190/ijamec.2026.164

Guan, L., Zhang, H., Zhou, Y., Du, X., & Li, M. (2025). YOLOv11-WBD: A wavelet-bidirectional network with dilated perception for robust metal surface defect detection. *PLOS ONE*, *20*(9), e0331025. https://doi.org/10.1371/journal.pone.0331025

Haque, Z., Koklu, M., Mirza, M. A., Omar, M., & Mukhayyokhon, S. (2025). Mitigating spatial scale loss in CNN-based fine-grained image classification: Application to date fruit grading. *Impact in Agriculture*, *1*, 1–15. https://doi.org/10.65500/agriculture-2025-005

Hayta, E., Gencturk, B., Ergen, C., & Koklu, M. (2023). Predicting future demand analysis in the logistics sector using machine learning methods. *Intelligent Methods in Engineering Sciences*, *2*(4), 102–114. https://doi.org/10.58190/imiens.2023.70

Hussain, M. (2023). YOLO-v1 to YOLO-v8, the rise of YOLO and its complementary nature toward digital manufacturing and industrial defect detection. *Machines*, *11*(7), 677. https://doi.org/10.3390/machines11070677

Isgor, B., & Koklu, M. (2025). Lightweight hybrid model for bone fracture detection using MobileNetV2 feature extraction and ensemble learning. *Journal of Future Artificial Intelligence and Technologies*, *2*(3), 521–533. https://doi.org/10.62411/faith.3048-3719-284

Khanam, R., Hussain, M., Hill, R., & Allen, P. (2024). A comprehensive review of convolutional neural networks for defect detection in industrial applications. *IEEE Access*, *12*, 94250–94295. https://doi.org/10.1109/ACCESS.2024.3425166

Kilci, O., & Koklu, M. (2025a). Automated classification of biscuit quality using YOLOv8 models in food industry. *Food Analytical Methods*, *18*(5), 815–829. https://doi.org/10.1007/s12161-025-02755-5

Kilci, O., & Koklu, M. (2025b). Machine learning-based detection of solar panel surface defects using deep features from InceptionV3. In *4th International Conference on Trends in Advanced Research*, (pp. 59–67), July 4–5, 2025, Konya, Türkiye.

Koklu, M., Isgor, B., & Tumer, C. (2026). Hybrid feature-based two-stage framework for audio deepfake detection and generative model attribution. *Advances in Engineering and Intelligence Systems*, *5*(1), 241–259. https://doi.org/10.22034/aeis.2026.566863.1400

Koklu, M., Kahramanli, H., & Allahverdi, N. (2012). A new approach to classification rule extraction problem by the real value coding. *International Journal of Innovative Computing, Information and Control*, *8*(9), 6303–6315.

Kursun, R., & Koklu, M. (2025). Classification of eggplant diseases using feature extraction with AlexNet and random forest. In *4th International Conference on Contemporary Academic Research (pp. 243–254), February 22–23, 2025* , Konya, Türkiye.

Lema, D. G., Sánchez-González, L., Usamentiaga, R., & delaCalle, F. J. (2025). Benchmarking deep learning models for surface defect detection: A reproducible and statistically-rigorous approach. *Journal of Intelligent Manufacturing*, 37, 3001–3018.. https://doi.org/10.1007/s10845-025-02672-8

Li, X., Fan, Z., Liu, Q., & Wan, X. (2025). DSP-YOLO: An improved YOLO11-based method for steel surface defect detection. *Measurement Science and Technology*, *36*(9), 096201. https://doi.org/10.1088/1361-6501/adfc88

Lu, J., Yu, M., & Liu, J. (2024). Lightweight strip steel defect detection algorithm based on improved YOLOv7. *Scientific Reports*, *14*(1), 13267. https://doi.org/10.1038/s41598-024-64080-x

Mao, M., & Hong, M. (2025). YOLO object detection for real-time fabric defect inspection in the textile industry: A review of YOLOv1 to YOLOv11. *Sensors*, *25*(7), 2270. https://doi.org/10.3390/s25072270

Masters Workspace. (2026). *Project dataset* (Version 3) [Data set]. Roboflow Universe. https://universe.roboflow.com/masters-workspace-yqrrt/project-p314o-2stwo/dataset/3

Meng, J., Guo, L., Hao, W., & Jain, D. K. (2025). A surface defect detection method for electronic products based on improved YOLOv11. *PLOS ONE*, *20*(10), e0334333. https://doi.org/10.1371/journal.pone.0334333

Nahar, L., Awrangjeb, M., & Islam, M. S. (2026). AI-enabled defect detection in industrial products: A comprehensive survey, key insights and future research challenges. *Advanced Engineering Informatics*, *69*, 104067. https://doi.org/10.1016/j.aei.2025.104067

Pan, W., Sun, X., Wang, Y., Cao, Y., Lang, Y., & Qian, Y. (2024). Enhanced photovoltaic panel defect detection via adaptive complementary fusion in YOLO-ACF. *Scientific Reports*, *14*(1), 26425. https://doi.org/10.1038/s41598-024-75772-9

Saritas, M. M., Kilci, O., & Koklu, M. (2025). Evaluation of CNN models for multi-class gear fault detection using waveform images. In *Proceedings of International Conference on Intelligent Systems and New Applications* 3, 31–40. https://doi.org/10.58190/icisna.2025.137

Sirisha, U., Praveen, S. P., Srinivasu, P. N., Barsocchi, P., & Bhoi, A. K. (2023). Statistical analysis of design aspects of various YOLO-based deep learning models for object detection. *International Journal of Computational Intelligence Systems*, *16*(1), 126. https://doi.org/10.1007/s44196-023-00302-w

Song, X., Zuo, G., Wang, X., & Xie, J. (2025). Automatic recognition of environmental hazards in river and lake ecosystems using deep learning. *Frontiers in Environmental Science*, *13*, 1657930. https://doi.org/10.3389/fenvs.2025.1657930

Su, B., Chen, H., & Zhou, Z. (2022). BAF-Detector: An efficient CNN-based detector for photovoltaic cell defect detection. *IEEE Transactions on Industrial Electronics*, *69*(3), 3161–3171. https://doi.org/10.1109/TIE.2021.3070507

Tian, J.-H., Feng, X.-F., Li, F., Xian, Q.-L., Jia, Z.-H., & Liu, J.-L. (2025). An improved YOLOv5n algorithm for detecting surface defects in industrial components. *Scientific Reports*, *15*(1), 9756. https://doi.org/10.1038/s41598-025-94109-8

Wan, W., Wang, L., Wang, B., Yu, H., Shi, K., & Liu, G. (2024). Space to depth convolution bundled with coordinate attention for detecting surface defects. *Signal, Image and Video Processing*, *18*(5), 4861–4874. https://doi.org/10.1007/s11760-024-03122-3

Wang, J., Gong, Y., & Wang, J. (2026). Research on improved model for PCB image defect detection based on YOLOv11. *IEEE Access*, *14*, 40788–40804. https://doi.org/10.1109/ACCESS.2026.3672639

Yan, W., Yan, H., Xu, N., Su, X., & Jiang, A. (2025). Real-time detection algorithm for industrial PCB defects based on HS-LSKA and global-EMA. *Journal of Electronic Imaging*, *34*(3), 033005. https://doi.org/10.1117/1.JEI.34.3.033005

Zhang, W., Lin, B., Wei, S., & Wu, J. (2026). Small target detection of surface defects on PCB boards: An improved YOLO method integrating attention mechanism and multi-scale feature focusing. *AI EDAM*, *40*, e2. https://doi.org/10.1017/S0890060426100213

Zhang, Y., Xie, F., Huang, L., Shi, J., Yang, J., & Li, Z. (2021). A lightweight one-stage defect detection network for small object based on dual attention mechanism and PAFPN. *Frontiers in Physics*, *9*, 708097. https://doi.org/10.3389/fphy.2021.708097

Zhong, M., Wei, L., & Mo, H. (2025). Cotton pest and disease diagnosis via YOLOv11-based deep learning and knowledge graphs: A real-time voice-enabled edge solution. *Frontiers in Plant Science*, *16*, 1671755. https://doi.org/10.3389/fpls.2025.1671755