
DESIGN AND PERFORMANCE EVALUATION OF AN INTELLIGENT UAV-BASED VISION SYSTEM FOR AUTOMATED POWER TRANSMISSION LINE COMPONENT INSPECTION USING YOLOV8S

Asmau Sanusi Gumbi¹, S. Salisu¹, A.s Adamu^{1,*}, Muhammad Dikko Almustaha²,
Sulaiman Garba³

¹Ahmadu Bello University Zaria

²Department of Electronics and Telecommunications Engineering

³Nigerian Communication Communications, Abuja

Email: najagumbi@gmail.com, sanialisu@abu.edu.ng, abubakaras@abu.edu.ng*, md.almustapha@gmail.com,
Sgarba@abu.edu.ng

*Corresponding author: A.s Adamu

DOI: <https://doi.org/10.67203/sp.ICESID.2026>

ABSTRACT

Power transmission line inspection requires the identification and localisation of infrastructure components from aerial imagery. Conventional inspection approaches can be labour-intensive, costly and constrained by difficult terrain, while UAVs provide a practical means of acquiring high-resolution imagery for subsequent analysis (Alcayde-García et al., 2022; Qi et al., 2022; Mendu & Mbuli, 2025). This study develops and evaluates a YOLOv8s-based vision system for automated detection of power transmission line components using the Inspection of Power Line Assets Dataset (InSPLAD). The predefined InSPLAD training and validation partitions were retained, with COCO annotations converted to YOLO format and images prepared at 640×640 pixels. A pretrained YOLOv8s model was fine-tuned for 50 epochs using transfer learning in Google Colaboratory with an NVIDIA Tesla T4 GPU. The best checkpoint was obtained at Epoch 48. On the predefined validation subset, the selected model achieved 89.5% Precision, 86.2% Recall, 88.4% mAP@0.5, 72.1% mAP@0.5:0.95, and an F1-score of 87.68%. The results establish an empirical performance baseline for the stated YOLOv8s-InSPLAD configuration. The findings are limited to offline evaluation on the predefined validation subset and should not be interpreted as evidence of onboard UAV performance, real-time field operation, or generalisation to external datasets.

Keywords: UAV inspection; power transmission lines; object detection; YOLOv8s; InSPLAD; deep learning; computer vision

1. INTRODUCTION

Power transmission networks depend on the condition of overhead lines and associated components for reliable electricity delivery. Conductors, insulators, spacers, dampers, clamps and suspension fittings are exposed to mechanical loading, environmental conditions and ageing, making periodic inspection an important maintenance activity (Alcayde-García et al., 2022; Qi et al., 2022). Traditional inspection methods include ground patrols, manual inspection and helicopter surveys. These approaches can be labour-intensive, costly and constrained by terrain and operational safety considerations (Alhassan et al., 2020; Xu et al., 2023).

UAVs have expanded the options for transmission-line inspection by enabling aerial acquisition of detailed RGB imagery from viewpoints that may be difficult to obtain through ground inspection. Reviews of UAV applications in power-line inspection describe their use for infrastructure observation, condition assessment and automated image analysis (Shakhatreh et al., 2019; Qi et al., 2022; Mohsan et al., 2023; Mendu & Mbuli, 2025). The resulting image volume also motivates automated computer-vision methods capable of detecting and localising multiple components within inspection imagery.

Deep learning object detectors, particularly the YOLO family, have been widely applied to visual detection because they formulate localisation and classification within a unified detection pipeline (Redmon et al., 2016; Terven et al., 2023). YOLOv8 introduced an anchor-free detection

formulation and a revised architecture implemented by Ultralytics, with model variants including YOLOv8s for smaller computational footprints (Jocher et al., 2023; Terven et al., 2023). These characteristics make YOLOv8s a suitable baseline for investigating automated component detection in UAV imagery without introducing architectural modifications.

Recent transmission-line inspection studies have increasingly modified YOLO architectures to address small objects, complex backgrounds, foreign objects, low-light conditions and other inspection challenges (Qiang et al., 2023; Shao et al., 2024; Li et al., 2025; Ait El Haj et al., 2026; Zheng et al., 2026; Sheng et al., 2026; Farooq et al., 2026; Zhu & Pei, 2026). While these studies demonstrate continued progress in specialised detectors, their numerical results are not directly rankable because datasets, target classes, model configurations, training procedures and evaluation protocols differ. The present study therefore focuses on establishing a reproducible detection baseline for a standard YOLOv8s configuration on InSPLAD rather than claiming universal superiority.

The objective of this paper is to design and evaluate an automated UAV-image component detection system based on pretrained YOLOv8s and the InSPLAD dataset. The evaluation is deliberately restricted to detection performance: Precision, Recall, mAP@0.5, mAP@0.5:0.95, confusion-matrix analysis, Precision–Recall analysis and F1-score analysis. The study does not evaluate onboard UAV deployment, real-time field performance, defect severity or predictive-maintenance decisions.

2. RELATED WORK

UAV-based transmission-line inspection has developed from image acquisition toward automated interpretation. Public datasets have supported this transition by providing annotated imagery for training and benchmarking object detectors. InSPLAD was developed as a large-scale dataset for inspection of power-line assets using UAV imagery, while other datasets such as TTPLA and PLAD provide complementary coverage of transmission towers and power-line assets (André et al., 2023; Vieira e Silva et al., 2023; Jiang et al., 2020; Zhang et al., 2021).

The YOLO family has been repeatedly adapted to transmission-line inspection. Qiang et al. (2023) developed an improved YOLOv8s model for transmission-line fault detection using attention, feature-fusion and loss-function modifications. Shao et al. (2024) proposed TL-YOLO for foreign-object detection on power transmission lines using an improved YOLOv8 framework. Li et al. (2025) proposed DI-YOLOv5 for dense small-object detection using wavelet-based processing and coordinate attention. These studies illustrate a common research direction: modifying a baseline detector to address task-specific visual challenges.

More recent work has investigated lightweight and specialised YOLO-based systems. Ait El Haj et al. (2026) proposed YOLOv8-ECC α for power-line asset inspection using CoordConv, Efficient Channel Attention and Alpha-IOU loss. Zheng et al. (2026) developed YOLO-MLSL for insulator-defect detection, while Sheng et al. (2026) proposed SDS-YOLO for foreign-object detection from UAV images. Farooq et al. (2026) investigated a Transformer-Enhanced YOLOv8 framework for small and complex insulator defects, and Zhu and Pei (2026) developed LCDD-YOLO for low-light transmission-line defect detection. Du et al. (2026) further investigated multimodal RGB and infrared imagery for insulator-defect detection. Collectively, these studies show the diversity of model modifications and datasets used in the field.

The present work differs in scope. Instead of introducing a modified detection architecture, it evaluates the standard YOLOv8s detector after transfer learning on the predefined InSPLAD partitions. The purpose is to quantify the detection behaviour of this specific configuration and establish a defensible empirical baseline. Comparisons with related studies are treated as contextual rather than as a direct ranking, because reported metrics are generated under different experimental conditions.

3. MATERIALS AND METHODS

3.1 Dataset

The study used the Inspection of Power Line Assets Dataset (InSPLAD), a UAV-image dataset containing annotated transmission-line assets (André et al., 2023; Vieira e Silva et al., 2023).

The dataset used in the experiment contains 18 component classes. The predefined training and validation partitions were retained rather than creating a new random split. The training subset contained 7,946 images and the validation subset contained 2,753 images. No independent external test dataset was introduced.

3.2 Dataset Verification and Preparation

Dataset preparation followed the procedure documented in the thesis. The available imagery and annotations were checked for correspondence, annotation format and class alignment before model training. The original COCO annotation files (`instances_train.json` and `instances_val.json`) were converted to YOLO text annotations, and a dataset configuration file was prepared for the Ultralytics framework. The predefined train/validation directory structure was retained. Images were prepared at 640×640 pixels, and the built-in YOLOv8 augmentation pipeline was used during training. Randomly selected annotated images were visually inspected to confirm that converted bounding boxes and class labels aligned with the corresponding components.

3.3 YOLOv8s Detection Model

YOLOv8s was selected as the detection model and used without architectural modification. YOLO detectors jointly perform object localisation and classification within a single detection framework (Redmon et al., 2016). YOLOv8 uses an anchor-free formulation, a C2f-based backbone, feature fusion through the neck, and a decoupled detection head (Jocher et al., 2023; Terven et al., 2023). In this study, the model received the prepared UAV images and produced bounding-box coordinates, class labels and confidence scores for detected transmission-line components.

Spatial correspondence between a predicted bounding box and its ground-truth counterpart was assessed using Intersection over Union (IoU), defined as the area of intersection divided by the area of union of the two boxes. Standard detection metrics were then calculated from the model predictions on the validation subset (Everingham et al., 2010; Padilla et al., 2021).

3.4 Transfer Learning and Training

Transfer learning was employed by fine-tuning pretrained YOLOv8s weights on the verified InSPLAD training subset. Training was conducted for 50 epochs with a batch size of 16 in Google Colaboratory using an NVIDIA Tesla T4 GPU. The implementation used Python 3.12 and Ultralytics YOLO 8.4.96. The optimizer was automatically selected as AdamW, and the random seed was set to 42. Validation was performed after each epoch and the best-performing checkpoint was retained as `best.pt`. The best checkpoint occurred at Epoch 48, based primarily on the highest recorded `mAP@0.5:0.95` in the validation results.

Table 1: Experimental configuration.

Parameter	Value
Dataset	InSPLAD
Number of classes	18
Training images	7,946
Validation images	2,753
Detection model	YOLOv8s
Learning strategy	Transfer learning
Framework	Ultralytics YOLO 8.4.96
Programming language	Python 3.12
Computing platform	Google Colaboratory
GPU	NVIDIA Tesla T4
Training epochs	50
Batch size	16
Optimizer	AdamW (automatic selection)
Random seed	42

3.5 Performance Evaluation

The trained model was evaluated on the predefined InSPLAD validation subset. Precision was calculated as $TP/(TP+FP)$, and Recall as $TP/(TP+FN)$, where TP denotes true positives, FP false positives and FN false negatives. Average Precision was obtained from the Precision–Recall relationship, and mAP was calculated as the mean of the class-wise AP values. The study reports mAP@0.5 and mAP@0.5:0.95, the latter averaging performance across IoU thresholds from 0.50 to 0.95 in increments of 0.05 (Padilla et al., 2021). The evaluation also examined the Precision–Recall curves, confusion matrix, normalized confusion matrix and F1-score curve to provide complementary views of class-wise and threshold-dependent detection behaviour.

4. RESULTS AND DISCUSSION

4.1 Dataset Verification

Before training, the dataset was checked for image–annotation correspondence, class alignment, bounding-box characteristics and spatial distribution. The resulting verification plots provide an empirical view of the data used in the experiment. Figure 1 shows the distribution of object classes and bounding-box characteristics in the InSPLAD dataset. The plots indicate variation in class frequency, object size and object position, confirming that the validation and training material contains components appearing at different locations and scales.




Figure 1: InSPLAD dataset distribution: object classes and bounding-box characteristics.

4.2 Training Performance

The model was trained for 50 epochs, with validation monitoring after each epoch. Figure 2 shows the training and validation loss and metric curves. The recorded losses declined over training while the evaluation metrics increased and stabilised toward the later epochs. The best checkpoint was obtained at Epoch 48, rather than assuming that the final epoch necessarily represented the best validation state. The thesis records the progression from Epoch 1 to Epoch 50 as a reduction in training and validation losses together with increases in Precision, Recall, mAP@0.5 and mAP@0.5:0.95.




Figure 2: Training and validation performance over 50 training epochs.

4.3 Detection Results

Figure 3 presents representative detections from the InSPLAD validation subset. The detector produces bounding boxes, class labels and confidence scores for multiple transmission-line components within the UAV imagery. These examples provide qualitative evidence that the trained model can localise several component types under the visual conditions represented in the validation data. They are illustrative only; quantitative performance is assessed in the following sections.




Figure 3: Representative detection results on the InSPLAD validation dataset.

4.4 Quantitative Detection Performance

The selected YOLOv8s checkpoint achieved 89.5% Precision, 86.2% Recall, 88.4% mAP@0.5 and 72.1% mAP@0.5:0.95 on the predefined InSPLAD validation subset. The F1-score was 87.68%. Together, these values indicate that the evaluated configuration achieved a substantial level of detection performance across the 18 component classes represented in the validation subset. The difference between mAP@0.5 and mAP@0.5:0.95 also provides a stricter view of localisation performance when increasingly demanding IoU thresholds are applied.

Table 2: Overall detection performance of YOLOv8s on the InSPLAD validation subset.

Metric	Best checkpoint
Precision	89.5%
Recall	86.2%
mAP@0.5	88.4%
mAP@0.5:0.95	72.1%

4.5 Precision–Recall Analysis

Figure 4 presents the Precision–Recall curves for the evaluated component classes. The curves show how precision changes as recall varies with the detection threshold. The overall behaviour is consistent with the aggregate Precision, Recall and mAP values reported in Table 2. Because the classes are represented with different frequencies and visual characteristics, the curves should be interpreted as class-specific evidence rather than as a single universal operating characteristic.




Figure 4: Precision–Recall curves for the transmission-line component classes.

4.6 Confusion Matrix Analysis

The confusion matrix in Figure 5 provides a class-wise view of predicted labels against ground-truth labels. The concentration of detections along the main diagonal indicates that many validation instances were assigned to their corresponding classes, while off-diagonal entries identify cases of class confusion. Such errors are particularly relevant where component categories share similar visual characteristics. The matrix therefore complements aggregate metrics by showing where classification errors occur across the evaluated classes.




Figure 5: Confusion matrix of the proposed YOLOv8s model.

4.7 Normalized Confusion Matrix and F1-Score

The normalized confusion matrix in Figure 6 expresses class-wise outcomes on a normalized scale, making differences between classes easier to inspect when their numbers of validation instances are unequal. The diagonal concentration across the classes indicates that correct classification accounts for a substantial share of the evaluated predictions, while the remaining off-diagonal proportions identify the relative pattern of class confusion.




Figure 6: Normalized confusion matrix of the proposed YOLOv8s model.

Figure 7 shows the F1-score as a function of confidence threshold. The F1-score combines Precision and Recall and is useful for examining the balance between false positives and false negatives. The reported final F1-score of 87.68% indicates a strong balance between the two measures for the evaluated configuration. This result is consistent with the 89.5% Precision and 86.2% Recall reported in Table 2, while remaining bounded by the validation dataset and evaluation protocol used in the experiment.




Figure 7: F1-score versus confidence threshold for the proposed YOLOv8s model.

4.8 Comparison with Related Studies

The detection results can be placed in the context of recent transmission-line inspection research, but numerical comparison should not be interpreted as a direct ranking. Qiang et al. (2023), for example, modified YOLOv8s for transmission-line fault detection, while Shao et al. (2024) and Li et al. (2025) addressed foreign-object and dense-small-object detection through architectural changes. Ait El Haj et al. (2026), Zheng et al. (2026), Sheng et al. (2026), Farooq et al. (2026) and Zhu and Pei (2026) similarly targeted specific inspection challenges through architectural, lightweight, transformer-based or low-light adaptations.

The present experiment instead evaluates a standard YOLOv8s configuration with transfer learning and dataset preparation on the predefined InSPLAD partitions. Its contribution is therefore an empirical baseline for this particular model–dataset–training configuration. The 89.5% Precision, 86.2% Recall, 88.4% mAP@0.5 and 72.1% mAP@0.5:0.95 should be understood within that experimental boundary. Differences from published results may arise from target definitions, dataset variants, class distributions, preprocessing, model variants, optimisation settings and evaluation protocols; consequently, the values do not establish superiority or inferiority across studies.

5. LIMITATIONS

The study has several limitations. First, evaluation was conducted using a single public dataset and its predefined validation subset; no independent external test dataset was used. Second, the experiments were conducted offline using stored UAV imagery rather than on an onboard UAV platform, and real-time field performance was not measured. Third, the study focused on component detection and classification rather than defect severity assessment or predictive-maintenance decision making. Finally, generalisation to transmission-line imagery outside the evaluated dataset was not established. These boundaries should be considered when interpreting the reported results and when comparing them with studies using different datasets or deployment settings.

6. CONCLUSION

This paper presented the design and performance evaluation of a YOLOv8s-based vision system for automated power transmission line component detection using the InSPLAD dataset. The predefined training and validation partitions were retained, COCO annotations were converted to

YOLO format, and a pretrained YOLOv8s model was fine-tuned for 50 epochs using transfer learning. Training was performed in Google Colaboratory with an NVIDIA Tesla T4 GPU, and the best validation checkpoint occurred at Epoch 48.

On the predefined InSPLAD validation subset, the selected checkpoint achieved 89.5% Precision, 86.2% Recall, 88.4% mAP@0.5, 72.1% mAP@0.5:0.95 and an F1-score of 87.68%. These results establish an empirical baseline for the stated YOLOv8s–InSPLAD configuration and show that the standard detector can provide measurable component-detection performance without architectural modification. The results do not, however, establish real-time UAV operation, onboard deployment, external-dataset generalisation or production readiness. Future work should therefore evaluate the detector on additional datasets and embedded UAV platforms and extend the task toward defect-specific assessment where appropriate.

REFERENCES

- [1] Alcayde-García, F., Salmerón-Manzano, E., Montero, M. A., Alcayde, A., & Manzano-Agugliaro, F. (2022). Power transmission lines: Worldwide research trends. *Energies*, 15(16), 5777. <https://doi.org/10.3390/en15165777>
- [2] Alhassan, A., Zhang, X., Shen, G., & Irfan, M. (2020). Power transmission line inspection robots: A review, trends and challenges for future research. *International Journal of Electrical Power & Energy Systems*, 118, 105862. <https://doi.org/10.1016/j.ijepes.2020.105862>
- [3] André, G., Kirchkopf, S., Osmani, V., & Schulz, D. (2023). InSPLAD: A large-scale dataset for inspection of power line assets using UAV imagery. *Data in Brief*, 49, 109395. <https://doi.org/10.1016/j.dib.2023.109395>
- [4] Ait El Haj, R., Benelmostafa, B.-E., & Medromi, H. (2026). YOLOv8-ECCα: Enhancing object detection for power line asset inspection under real-world visual constraints. *Algorithms*, 19(1), 66. <https://doi.org/10.3390/a19010066>
- [5] Du, Z., Geng, Y., Wang, H., Zhang, H., & Hu, Y. (2026). Aerial insulator defect detection method based on CWSP-YOLO. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-40782-2>
- [6] Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2010). The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88, 303–338. <https://doi.org/10.1007/s11263-009-0275-4>
- [7] Farooq, U., Yang, F., & Shaikh, J. A. (2026). Transformer enhanced based YOLOv8 integration: A hybrid deep learning framework for intelligent insulator defect detection in high-voltage transmission systems. *Frontiers in Artificial Intelligence*, 8, 1732616. <https://doi.org/10.3389/frai.2025.1732616>
- [8] Jiang, H., Zhang, J., Li, Y., & Liu, X. (2020). TTPLA: A dataset for transmission tower and power line inspection. *IEEE Access*, 8, 164043–164055. <https://doi.org/10.1109/ACCESS.2020.3021820>
- [9] Jocher, G., Chaurasia, A., & Qiu, J. (2023). Ultralytics YOLOv8 (Version 8.0.0) [Computer software]. GitHub. <https://github.com/ultralytics/ultralytics>
- [10] Li, Z.-X., Wang, Y.-L., & Wang, F. (2025). DI-YOLOv5: An improved dual-wavelet-based YOLOv5 for dense small object detection. *IEEE/CAA Journal of Automatica Sinica*, 12(2), 457–459. <https://doi.org/10.1109/JAS.2024.124368>
- [11] Mendu, B., & Mbuli, N. (2025). State-of-the-art review on the application of unmanned aerial vehicles (UAVs) in power line inspections: Current innovations, trends, and future prospects. *Drones*, 9(4), 265. <https://doi.org/10.3390/drones9040265>
- [12] Mohsan, S. A. H., Othman, N. Q. H., Li, Y., Alsharif, M. H., & Khan, M. A. (2023). Unmanned aerial vehicles (UAVs): Practical aspects, applications, open challenges, security issues, and future trends. *Intelligent Service Robotics*, 16, 109–137. <https://doi.org/10.1007/s11370-022-00452-4>

- [13] Padilla, R., Passos, W. L., Dias, T. L. B., Netto, S. L., & da Silva, E. A. B. (2021). A comparative analysis of object detection metrics with a companion open-source toolkit. *Electronics*, 10(3), 279. <https://doi.org/10.3390/electronics10030279>
- [14] Qi, J., Song, D., Shang, H., Wang, N., & Hua, C. (2022). UAV-based inspection of overhead power lines: A review of technologies and applications. *Energy Reports*, 8, 1040–1055. <https://doi.org/10.1016/j.egy.2021.12.053>
- [15] Qiang, H., Tao, Z., Ye, B., Yang, R., & Xu, W. (2023). Transmission line fault detection and classification based on improved YOLOv8s. *Electronics*, 12(21), 4537. <https://doi.org/10.3390/electronics12214537>
- [16] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (pp. 779–788). IEEE. <https://doi.org/10.1109/CVPR.2016.91>
- [17] Shao, Y., Zhang, R., Lv, C., Luo, Z., & Che, M. (2024). TL-YOLO: Foreign-object detection on power transmission line based on improved YOLOv8. *Electronics*, 13(8), 1543. <https://doi.org/10.3390/electronics13081543>
- [18] Sheng, J., Wu, S., Liu, G., Ma, X., Deng, S., & Xu, H. (2026). SDS-YOLO: Drone-based foreign object detection model for power lines using an enhanced YOLOv8n approach. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-59200-8>
- [19] Shakhatreh, H., Sawalmeh, A. H., Al-Fuqaha, A., Dou, Z., Almaita, E., Khalil, I., Othman, N. S., Khreishah, A., & Guizani, M. (2019). Unmanned Aerial Vehicles (UAVs): A survey on civil applications and key research challenges. *IEEE Access*, 7, 48572–48634. <https://doi.org/10.1109/ACCESS.2019.2909530>
- [20] Terven, J., Córdova-Esparza, D.-M., & Romero-González, J.-A. (2023). A comprehensive review of YOLO architectures in computer vision: From YOLOv1 to YOLOv8 and YOLO-NAS. *Machine Learning and Knowledge Extraction*, 5(4), 1680–1716. <https://doi.org/10.3390/make5040083>
- [21] Vieira e Silva, A. L. B., Felix, H. C., Simões, F. P. M., Teichrieb, V., Santos, M. M., Santiago, H., Sgotti, V., & Lott Neto, H. (2023). InsPLAD: A dataset and benchmark for power line asset inspection in UAV images. *arXiv:2311.01619*.
- [22] Xu, B., Zhao, Y., Wang, T., & Chen, Q. (2023). Development of power transmission line detection technology based on unmanned aerial vehicle image vision. *SN Applied Sciences*, 5, 72. <https://doi.org/10.1007/s42452-023-05299-7>
- [23] Zhang, Y., Li, X., Wang, J., & Sun, H. (2021). PLAD: A power line assets dataset for UAV-based inspection. *Remote Sensing*, 13(18), 3672. <https://doi.org/10.3390/rs13183672>
- [24] Zheng, R., Duan, G., Cao, X., & Wang, H. (2026). Transmission line insulator defect detection method based on YOLO-MLSL model. *Energies*, 19(10), 2305. <https://doi.org/10.3390/en19102305>
- [25] Zhu, Y., & Pei, S. (2026). Transmission line defect detection in low-light environments using image adversarial synthesis and the LCDD-YOLO algorithm. *International Journal of Intelligent Systems*, 2026, Article 2674074. <https://doi.org/10.1155/int/2674074>