Tiako P.F. (ed) *Intelligent and Sustainable Solutions*. Chronicle of Computing. OkIP. CAIS25#06

https://doi.org/ 10.55432/978-1-6692-0011-6_11

YOLO algorithms for Solar Panel and Boiler Detection from Aerial Imagery

Douglas Koomson[a*], James Abugre[a], Stephen Moore[b]

[a]*Ashesi University, 1 University Avenue, Berekusu, Ghana*
[b]*University of Cape Coast, University Avenue, Cape Coast, Ghana*

ABSTRACT

The integration of solar panels and boilers, considered key renewable energy technologies, remains crucial for addressing global sustainable energy challenges. Aerial imagery from drones and satellites plays an important role in monitoring and assessing these infrastructures. This study presents a comparative evaluation of three You Only Look Once (YOLO) model architectures, specifically YOLOv8-Large, YOLOv10-Large and YOLOv11-Large for the detection of solar panels and boilers. Through Exploratory Data Analysis (EDA), the patterns and characteristics of the dataset were identified, providing a strong basis for developing a robust object detection model. YOLOv10-Large demonstrated the best performance, achieving a precision of 89.2%, a recall of 81.3%, mAP (50) of 83.1%, and an F1 score of 85.0%. This advancement prioritizes efficiency and reliability as the backbone, contributing to more sustainable and resource-efficient practices in industries that rely on aerial imagery analysis.

Keywords: YOLO, Solar Panel Detection, Boiler Detection, Deep Learning, Drone Imagery, Satellite Imagery

I. INTRODUCTION

The global shift toward renewable energy has led to the widespread adoption of solar technologies. With a current worldwide installed solar photovoltaic (PV) capacity surpassing 2 terawatts (TW) and averaging annual installations of around 500 gigawatts (GW) (Apeh, Meyer, & Overen, 2022; Urbina, 2022). Solar panels and boilers have become crucial components in this transition, forming a vital part of the vast pool of renewable technologies, playing a significant role in both residential and commercial energy systems.

In regions where centralized tracking is limited, monitoring and mapping the deployment of solar infrastructure is crucial for both investment planning and sustainability assessments. Accurate detection of installed solar panels and boilers provides key insights into renewable energy adoption trends, informing rural electrification and energy expansion projects. Traditional survey-based methods are time-consuming and often lack scalability. As a result, aerial imagery (i.e. satellite and drone) analysis has become a viable yet resource-intensive alternative for assessing this infrastructure at a larger scale.

Significant advancements in Artificial Intelligence (AI) for object detection algorithms, such as the You Only Look Once (YOLO) family (Redmon, Divvala, Girshick & Farhadi, 2016), have opened possibilities for automated infrastructure monitoring. These models can identify and localize objects in real time, making them suitable for detecting solar energy installations from overhead imagery. This paper presents a comparative analysis of three YOLO architectures — YOLOv8-large (v8-l), YOLOv10-large (v10-l), and YOLOv11-large (v11-l) — for detecting solar panels and boilers. While this work assesses the presence and distribution of solar panels and boilers, previous studies focused on detecting solar panels, faults, or damages (Khanam, Asghar, & Hussain, 2025; Malik, Saxena, Raj, Singh, & Kumar, 2024; Pu et al., 2025; Yin, Lingxin, Maohuan, Qianlai, & Xiaosong, 2023).

The remainder of this paper is organized as follows: the literature review is presented in Section II, the dataset and model design are discussed in Section III, the results and discussion are presented in Section IV, and the conclusions are drawn in Section V.

II. LITERATURE REVIEW

A. Overview of YOLO

YOLO is a widely used deep learning (DL) algorithm that employs a classification and regression-based object detection method, giving the algorithm its core strengths. It is modelled on a straightforward structure, small size, and fast computational speed. Newer versions of the YOLO object detection algorithm have been released since its initial release in 2015, aiming to address challenges in the previous version (Diwan, Anirudh, & Tembhurne, 2023). The mode of operation of this algorithm is that it transforms input images to equal sizes and then divides them into a total of S x S network cells of the same size, where objects can be detected in each cell. A network cell predicts a target object if its centre is located within the network. Each network cell can contain N detection boxes, where each box determines its position and a corresponding prediction score. This score indicates the likelihood of detecting a target within the predicted network cell. Multiple detection boxes can exist in a network cell; the algorithm will automatically select the highest-scoring target category for prediction. The prediction scores of each box within the cell are evaluated, and the one with the highest confidence level is selected, which ensures that the most likely object is identified and classified accurately. The current released model at the time of this writing in the YOLO series is the YOLOv11. This literature review will focus on the three YOLO architectures used in this study. Figure 1 below outlines the YOLO versions and their corresponding release dates.

* Corresponding author E-mail: douglas.koomson@ashesi.edu.gh

Presented at the 2025 OkIP International Conference
on Automated and Intelligent Systems (CAIS)
in Oklahoma City, OK, USA, and Online, on October 1, 2025

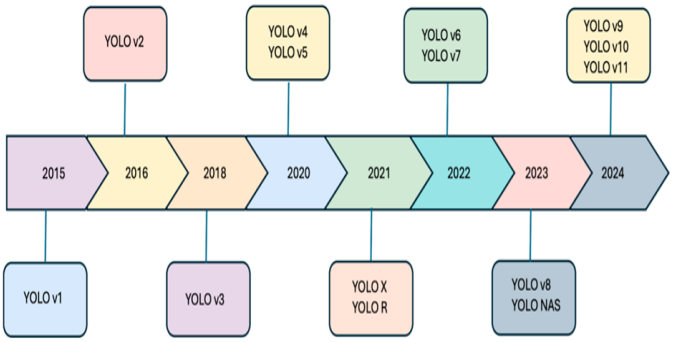

Fig. 1. Timeline of YOLO Models Release

1) YOLOv8: YOLOv8 represents a significant evolution in the YOLO object detection family, introducing several architectural and functional enhancements over its predecessors. It features a modified CSPDarknet backbone integrated with the C2f module, which enables more profound feature fusion by concatenating outputs from all bottleneck layers, thereby enhancing contextual understanding without inflating computational load (Sohan, Sai Ram, & Rami Reddy, 2024). The algorithm employs an anchor-free detection head, simplifying bounding box predictions and accelerating inference by reducing post-processing overhead, such as Non-Maximum Suppression (NMS). It supports computer vision tasks, including object detection, segmentation, classification, and pose estimation, with pre-trained models available for each task. The model is available in scalable variants, ranging from YOLOv8-nano to YOLOv8-extra-large. They cater to diverse deployment needs, ranging from edge devices to high-performance servers. Evaluated on the COCO and Roboflow 100 benchmarks, YOLOv8 consistently outperformed previous models, YOLOv5 and YOLOv7, in terms of mAP and speed. This model at that time demonstrated superior generalization across diverse domains.

2) YOLOv10: The YOLOv10 is an end-to-end real-time object detector that introduces architectural and training innovations to enhance performance and efficiency (Wang et al., 2024). Addressing the limitations of traditional YOLO models, which rely on NMS for post-processing, YOLOv10 implements a novel, consistent dual assignment strategy. This approach integrates one-to-many supervision for training and one-to-one prediction for inference, eliminating the need for NMS and enabling accurate end-to-end deployment. The architecture's redesign employed a holistic, efficiency-accuracy-driven model strategy, incorporating lightweight classification heads, spatial-channel decoupled downsampling, and rank-guided block allocation. These modifications improved accuracy, reduced latency, and decreased computational overhead. Experimental results on the COCO benchmark demonstrate that YOLOv10 consistently outperforms previous YOLO versions and other state-of-the-art detectors across multiple model scales, offering superior trade-offs between speed, accuracy, and parameter efficiency. Figure 2 shows the architecture of the consistent dual assignments for NMS-free training.

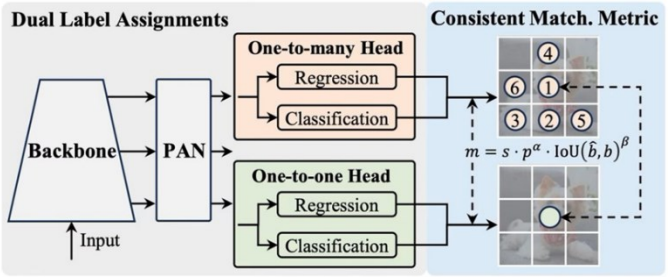

Fig. 2. Consistent dual assignments for NMS-free training

3) YOLOv11: This highly versatile model can be utilized for various computer vision tasks, including object detection, image classification, and instance segmentation, making it a concrete solution for diverse vision applications (Jocher, Qiu, & Chaurasia, 2023). The algorithm introduces new features that enhance its performance, flexibility, and efficiency. This model can handle various vision tasks, mainly detection, segmentation, pose estimation, tracking, and classification (Khanam & Hussain, 2024). The multiple variants (nano to x-large) encompass edge and high-performance deployments. The YOLOv11-medium variant achieved a higher mAP than its counterpart, YOLOv8-medium, using 22% fewer parameters. Benchmarking results show that YOLOv11 exhibits superior scaling efficiency across all model sizes, with faster inference times and improved mAP(50–95) scores compared to other models. YOLOv11 also outperforms earlier versions, particularly in low-latency regimes (2–6 ms), achieving real-time speeds and maintaining accuracy. These enhancements position YOLOv11 as a state-of-the-art model capable of delivering precise, efficient, and scalable object detection solutions across a wide range of real-world applications, including autonomous systems, healthcare, surveillance, and industrial automation (Yang, Lan, & Wang, 2025). Figure 3 shows a comparative performance test for all YOLO-scaled models (Ultralytics, 2025a).

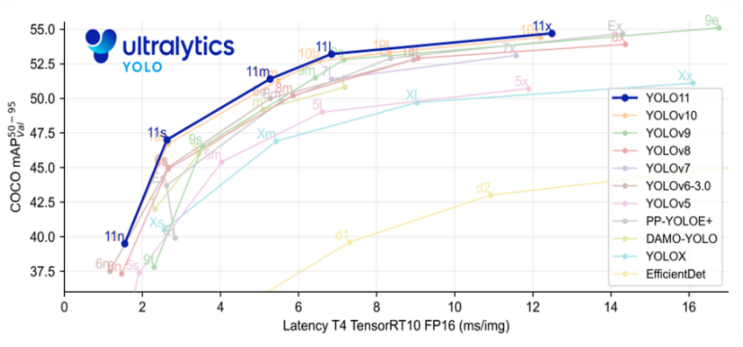

Fig. 3. Comparative Performance Curve of YOLO Versions

B. *Review of Related Works*

Utilizing YOLO for studies related to solar energy is crucial. In large-scale solar farms for energy production, Unmanned Aerial Vehicles (UAVs) are used for conducting field inspections. Pu et al. (2025) proposed a refined UAV-based wind and solar power station inspection system using the YOLOv8 object detection model. Their method focused on the sparse spatial distribution of PV modules and wind turbine blades, capturing aerial images using drones with zoom lenses. The YOLOv8 model achieved an mAP(50) of 85.4% and demonstrated a 6.2% improvement in mAP(50:95) compared to the base YOLOv8. The proposed system demonstrated strong adaptability and accuracy across varying terrain and lighting conditions. The paper by Khanam, Asghar & Hussain (2025) conducted a comparative evaluation of YOLOv5, YOLOv8, and YOLOv11 for automated defect detection in PV modules. Their study utilized a dataset of 6,493 annotated solar panel images, covering four classes: bird droppings, cracks, dust accumulation, and non-defective panels. Model evaluation conducted used mAP, R, and inference speed as metrics. YOLOv5 and v8 achieved the fastest inference speed (7.1 ms/image) with high precision (94.1%) for crack detection and a recall rate of 79.2% for bird-dropping defects respectively. Overall, the highest accuracy was achieved in experimentation with an mAP(50) of 93.4% by the v11 model. The study emphasized the importance of balancing accuracy and speed for defect detection systems in large-scale solar farms. In another research by Malik, Saxena, Raj, Singh, & Kumar (2024), a YOLO-based fault detection system was proposed to identify standard PV module defects, such as cracks, hotspots, and bypass diode failures, using thermal imagery. Their study compared the performance of YOLOv5 through YOLOv9 on a custom dataset created via infrared thermography. Among these, a GELAN (Generalised Efficient Layer Aggregation Network)-based variant achieved the highest mAP of 70.4%, outperforming traditional YOLO architectures. The method automates visual inspection by leveraging DL, thereby reducing manual effort and enhancing the accuracy and consistency of defect diagnosis across large-scale solar installations. Yin, Lingxin, Maohuan, Qianlai & Xiaosong (2023) developed a PV-YOLO, a lightweight model for PV panel fault detection, by modifying the original YOLOX architecture. They replaced the CSPDarknet53 backbone with a transformer-based PVTv2 network. They added a CBAM attention module to improve the identification of fine-grained faults in IR images, a revised label assignment mechanism, and the SIoU loss function to boost convergence and accuracy. Experimental results showed that PV-YOLO achieved a peak mAP of 92.56%, greater than YOLOX by 10.46%. Clark & Pacifici (2023) introduced a high-resolution satellite imagery dataset tailored for solar panel detection. The dataset includes 2,542 annotated solar panel objects captured at native 31 cm and enhanced 15.5 cm resolutions from WorldView-3 satellite imagery over Southern Germany. Annotated using QGIS, with confidence levels assigned via cross-verification in Google Earth and designed to aid in small object detection for the dataset, it was optimized for the YOLOv4 architecture and enables the evaluation of resolution-performance trade-offs.

Through this comprehensive literature review, we identified that limited work has been conducted on detecting solar panels, with the focus remaining mainly on fault detection. Some traditional image processing and AI techniques have been used for real-time detection of solar panels (Arnaudo et al., 2023; Liao and Lu, 2020; Malof, Hou, Collins, Bradbury, & Newell, 2015; Parhar et al., 2022; Umar, Qureshi, & Nawaz, 2024; Vega Díaz, Vlaminck, Lefkaditis, Orjuela Vargas, & Luong, 2020), but there is limited work involving solar boilers. Hence, a YOLO model that incorporates both technologies to address these gaps in current research.

III. METHODOLOGY

This section gives a detailed overview of the model development process from data preprocessing to model evaluation.

A. *Dataset*

1) Description: The dataset used for experimentation was licensed under CC-BY-4.0 from Lacuna Fund and listed as open source on Openstat (MAIDI, 2025). The images were presented in multiple resolutions, initially comprising 2,125 satellite and 9,202 drone images taken across several regions in Madagascar. A total of 3,308 images were selected at random. These images were annotated with two classes: solar panel and solar boiler, totalling 6,746 class instances. Some images contained multiple classes. With wide variation in size, shape, color, orientation, sharpness, zoom, location and lighting, the dataset demonstrates the model's robustness to these varying conditions. Figure 4 shows aerial images with (a) solar panels, (b) boilers, and (c) both.

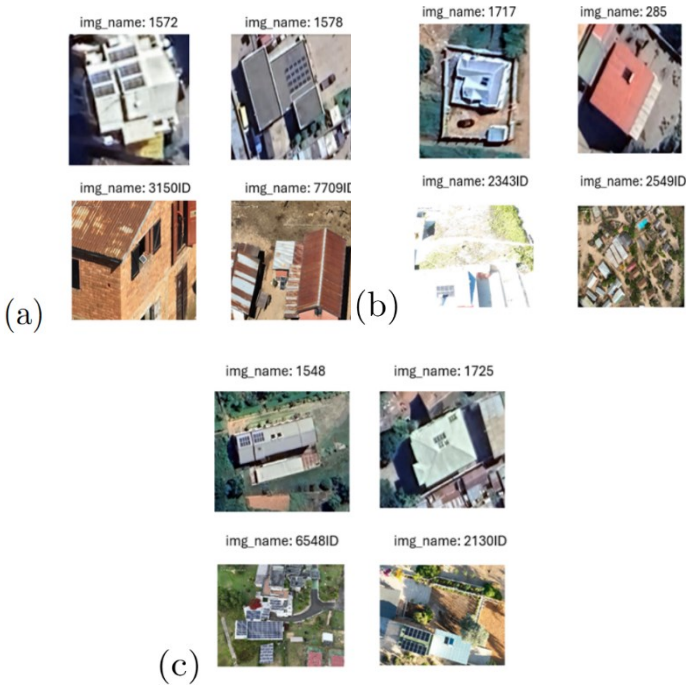

Fig. 4. (a) Solar Panels (b) Boilers (c) Both

2) Data Preprocessing: The adopted Exploratory Data Analysis (EDA) of the aerial imagery data influences the model development for accurate class detection. Cross-Balance Analysis (CBA) was conducted to assess the bias of the dataset. We found the panels comprised 80% of the classes, representing an unbalanced dataset. Data augmentation techniques, such as flipping, rotation, and scaling, were employed to mitigate a skewed distribution and enhance generalizability. Second, by examining the size distribution, we identified images that exceeded reasonable memory limits for processing and discarded them to ensure efficient memory usage, thereby optimizing computational resources and speeding up the training process. The class distribution before and after augmentation is shown in Figure 5.

B. Model

The three models were initialized using pre-trained weights, leveraging transfer learning to benefit from the features discovered on large datasets, such as COCO. The models were trained with a batch size of 16, 100 epochs and an image size of 640x640. The training was conducted using Google Colab servers on an A100 GPU and 15 GB of RAM to accelerate the training process. Efficient hyperparameter tuning using Ray Tune was performed to obtain the optimal values during training (Ultralytics, 2025b). The model's performance was evaluated on the training dataset using Precision (P), Recall (R), F1 score, and Mean Average Precision (mAP) as the primary performance metrics. The v10-l model emerged as the best-performing architecture. In the next section, a detailed model comparison is done.

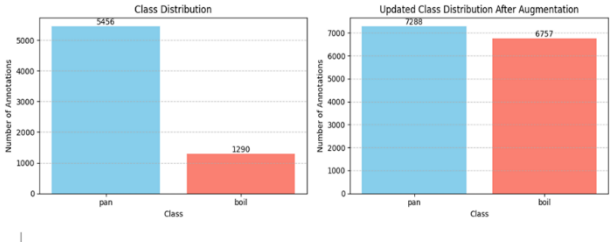

Fig. 5. Class Distribution Before and after Augmentation

IV. RESULTS AND DISCUSSION

We performed experiments to (a) train and validate, and (b) test on unseen data. The training and validation stage provided and updated the weights to achieve optimal performance, while the testing evaluated the real-time performance of the built model.

A. Training and Validation

The experiments demonstrated that the YOLOv10-l model delivered the best results in terms of both detection performance and computational efficiency, as shown in Table 1. YOLOv10-l achieved the highest scores in all metrics except for recall (81.2%). An overall precision of 89.2% on the test dataset, outperforming v11-l (88.8%) and v8-l (86.9%), which were recorded. YOLOv10-l demonstrated the highest mAP(50) and mAP(50:95) at 83.1% and 77.6% respectively. These results emphasize the YOLOv10 model's capability to balance computational complexity and detection accuracy, which is essential for real-time deployment in solar monitoring applications. The v10-l's high precision advantages are preferred for real-time scenarios. The relatively smaller model size of YOLOv10 reduces the need for expensive hardware, allowing for real-time monitoring in remote locations and resulting in lower operational costs. Figure 6 shows the training graphs for the losses and metrics of YOLOv11-l.

TABLE I. SUMMARY OF TRAINING AND VALIDATION METRICS

Model	P	R	mAP(50)	mAP(50:95)	F1 Score
YOLOv8-l	0.869	**0.813**	0.827	0.768	0.840
YOLOv10-l	**0.892**	0.812	**0.831**	**0.776**	**0.850**
YOLOv11-l	0.888	0.810	0.825	0.767	0.847

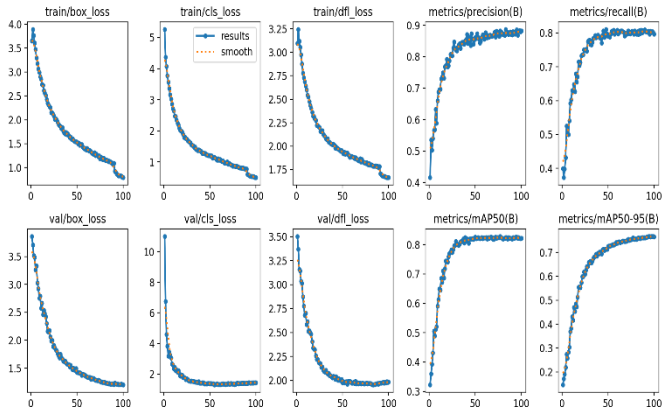

Fig. 6. Results of YOLOv11-l Training

Confusion matrices provide valuable insights into model performance, specifically regarding correct predictions (true positives (TP) and true negatives (TN)), and errors (false positives (FP) and false negatives (FN)). Figure 7 shows the confusion matrix of YOLOv10-l. The top left value represents the panel True Positives (TP), indicating that 67% of them were correctly identified as such by the model. The bottom left value signifies False Positives (FP), where 33% of the solar panels were misclassified. The middle value shows that the model predicts an almost perfect score of 99% for the boiler class. The model also incorrectly identifies more non-existent panels (93%) than boilers (7%). This suggests the model tends to over-predict panels. Finally, the bottom right, which shows white, indicates False Negatives (FN) at 0%, indicating that there were no instances where either panel or boiler was unidentified. The matrix suggests that while the model excels in identifying boilers, it requires improvement to reduce overpredictions of panels, thereby enhancing overall accuracy.

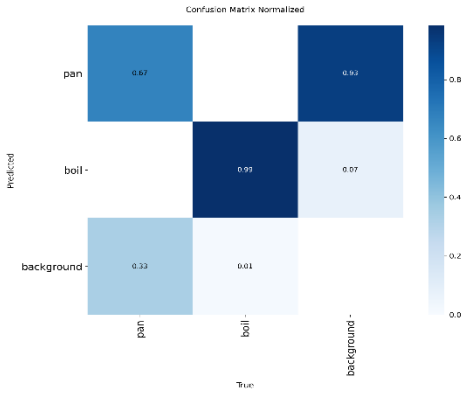

Fig. 7. Confusion Matrix for YOLOv10-l

B. *Test Results*

The test on the image was done at an acceptable confidence interval of 0.25. This test confirmed the training results as the v10 showed superior detection capabilities compared to the other models. To ensure its application in tracking real-time large solar projects, the test image used was a solar farm. It detected more than half of the solar panels in the solar farm as shown in Figure 9. This demonstrates that this model can be utilized for the detection and monitoring of large-scale solar projects. The confidence in the detection of the boiler was high, suggesting a balanced performance of the boiler class across all models. Figures 8 and 10 show the detection results for YOLOv8 and YOLOv11 respectively.

Fig. 8. YOLOv8 Detection test at a Confidence threshold = 0.25

Fig. 9. YOLOv10 Detection test at a Confidence threshold = 0.25

Fig. 10. YOLOv11 Detection test at a Confidence threshold = 0.25

V. CONCLUSION

In conclusion, developing an object detection model for solar panel and boiler detection and classification presents a transformative opportunity for the energy sector by providing the tools needed for key energy players to map out the adoption of these solar technologies. We successfully identified YOLOv10-l as the best-performing model, achieving a precision of **0.892**, a recall of **0.813**, mAP(50) of **0.831**, and an F1 score of **0.850**. This paper successfully showed the use of this model for wide-scale monitoring. This study benchmarks modern object detection models within the context of solar energy aiming to monitor renewable energy expansion through remote sensing. Utilizing a balanced dataset, experimenting with alternative data augmentation techniques to optimize predictive capabilities, and investigating what caused the superior performance of the v10-l over the v11-l are areas to explore for future research.

REFERENCES

Apeh, O. O., Meyer, E. L., & Overen, O. K. (2022). Contributions of solar photovoltaic systems to environmental and socioeconomic aspects of national development—A review. *Energies*, *15*(16), 5963.

Arnaudo, E., Blanco, G., Monti, A., Bianco, G., Monaco, C., Pasquali, P., & Dominici, F. (2023). A comparative evaluation of deep learning techniques for photovoltaic panel detection from aerial images. *IEEE Access*, *11*, 47579-47594.

Clark, C. N., & Pacifici, F. (2023). A solar panel dataset of very high resolution satellite imagery to support the Sustainable Development Goals. *Scientific Data*, *10*(1), 636.

Diwan, T., Anirudh, G., & Tembhurne, J. V. (2023). Object detection using YOLO: challenges, architectural successors, datasets and applications. *multimedia Tools and Applications*, *82*(6), 9243-9275.

Jocher, G., Qiu, J., & Chaurasia, A. Ultralytics YOLO, 2023. Original-Date: 2022-09-11T16: 39: 45Z. *URL https://github.com/ultralytics/Ultralytics*.

Khanam, R., & Hussain, M. (2024). Yolov11: An overview of the key architectural enhancements. *arXiv preprint arXiv:2410.17725*.

Khanam, R., Asghar, T., & Hussain, M. (2025, February). Comparative performance evaluation of yolov5, yolov8, and yolov11 for solar panel defect detection. In *Solar* (Vol. 5, No. 1, p. 6). MDPI.

Liao, K. C., & Lu, J. H. (2020, April). Using Matlab real-time image analysis for solar panel fault detection with UAV. In *Journal of Physics: Conference Series* (Vol. 1509, No. 1, p. 012010). IOP Publishing.

MAIDI. (2025, April 16). *Data on solar energy and labeling of images of photovoltaic panels in Madagascar*. Openstat Madagascar. Retrieved April 29, 2025, from https://openstat-madagascar.com/bdd/energie-et-environnement/131-donnees-sur-l-energie-solaire-et-labellisation-d-images-de-panneaux-photovoltaiques-a-madagascar

Malik, P., Saxena, V., Raj, S., Singh, S., & Kumar, S. (2024, September). Fault detection of the solar photovoltaic modules using yolo models. In *2024 IEEE Region 10 Symposium (TENSYMP)* (pp. 1-6). IEEE.

Malof, J. M., Hou, R., Collins, L. M., Bradbury, K., & Newell, R. (2015, November). Automatic solar photovoltaic panel detection in satellite imagery. In *2015 International Conference on Renewable Energy Research and Applications (ICRERA)* (pp. 1428-1431). IEEE.

Parhar, P., Sawasaki, R., Todeschini, A., Vahabi, H., Nusaputra, N., & Vergara, F. (2022). HyperionSolarNet: Solar panel detection from aerial images. *arXiv preprint arXiv:2201.02107*.

Pu, J., Zhang, Q., Zhao, W., Zhang, W., Qin, Z., & Zhang, Y. (2025). Research on refined UAV inspection method of wind/solar power stations based on YOLOv8. *Journal of Engineering and Applied Science*, *72*(1), 10.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 779-788).

Sohan, M., Sai Ram, T., & Rami Reddy, C. V. (2024). A review on yolov8 and its advancements. In *International Conference on Data Intelligence and Cognitive Informatics* (pp. 529-545). Springer, Singapore.

Ultralytics. (2025a, May 28). *Models supported by ultralytics*. Ultralytics YOLO Docs. https://docs.ultralytics.com/models/

Ultralytics. (2025b, May 24). *Ray Tune*. https://docs.ultralytics.com/integrations/ray-tune/

Umar, S., Qureshi, M. S., & Nawaz, M. U. (2024). Thermal imaging and AI in solar panel defect identification. *International Journal of Advanced Engineering Technologies and Innovations*, *1*(3), 73-95.

Urbina, A. (2022). Scenarios for solar electricity at the TeraWatt scale. In *Sustainable Solar Electricity* (pp. 3-17). Cham: Springer International Publishing.

Vega Díaz, J. J., Vlaminck, M., Lefkaditis, D., Orjuela Vargas, S. A., & Luong, H. (2020). Solar panel detection within complex backgrounds using thermal images acquired by UAVs. *Sensors*, *20*(21), 6219.

Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., & Han, J. (2024). Yolov10: Real-time end-to-end object detection. *Advances in Neural Information Processing Systems*, *37*, 107984-108011.

Yang, Z., Lan, X., & Wang, H. (2025). Comparative analysis of YOLO series algorithms for UAV-based highway distress inspection: Performance and application insights. *Sensors*, *25*(5), 1475.

Yin, W., Lingxin, S., Maohuan, L., Qianlai, S., & Xiaosong, L. (2023). PV-YOLO: lightweight YOLO for photovoltaic panel fault detection. *Ieee Access*, *11*, 10966-10976.