26th International Congress on Modelling and Simulation, Adelaide, SA, Australia, 30 Nov–4 Dec 2025
mssanz.org.au/modsim2025

Employing image processing and anomaly detection for faulty solar cells

Mohan Sai Kiran Bonu[a], Ali Ghahremani[b*], Asef Nazari[a], Mohammad Ghasemi[a], Dhananjay Thiruvady[a], Rosemary Van Der Meer[a]

[a] School of Information Technology, [b] School of Engineering, Deakin University, Geelong, 3216, VIC, Australia, Email: s223047002 @deakin.edu.au*

Abstract: With the global transition to renewable energy, solar power has become a key contributor to the green energy movement. However, the efficiency of solar panels can decline over time due to physical faults such as cracks, hotspots, and electrical disconnections, many of which often go unnoticed in large-scale installations. Traditional inspection methods are manual, labour-intensive, and prone to human error, making them unsuitable for long-term, wide-scale use. Therefore, this study investigates the potential of modern computer vision techniques, specifically the YOLOv12 object detection model, to automate fault detection in solar photovoltaic (PV) panels. The research begins with an evaluation of convolutional neural network-based image classification on electroluminescence (EL) images, then shifts to real-time object-level detection using RGB images with bounding box annotations. The PV-Multi-Defect dataset was used, and its XML annotations were converted to YOLO-compatible TXT format. A balanced distribution of samples across five defect classes was ensured, and training was performed using the YOLOv12n model, selected for its lightweight design suited to low-resource environments. After 500 training epochs, the results confirm the effectiveness of YOLOv12 in identifying multiple fault types and support its potential for real-time deployment. The obtained results demonstrate that advanced object detection models can be used effectively for automated PV inspection.

Keywords: *Renewable energy, Solar photovoltaic panels, Image processing, Anomaly detection, Object detection model*

1. INTRODUCTION

The global transition toward net-zero emissions positions solar photovoltaic systems as the backbone of renewable generation. Photovoltaic (PV) panels remain the most prevalent technology due to their modularity, declining cost, and scalability (IEA, 2024). However, the efficiency of the panels does degrade over time. Scratches, broken cells, hotspots, and electrical disconnections are common faults in such systems, significantly compromising energy output, safety, and operational efficiency (Deitsch et al., 2019).

Manual visual checks or electroluminescence (EL) imaging are time-consuming and costly methods. These methods require experts to identify anomalies frame by frame, making the process slow, subjective, and difficult to scale (Velasco-Sánchez et al., 2025). This is especially problematic in utility-scale solar farms where thousands of panels need regular monitoring. The limitations of conventional methods point to the need for fast, accurate, and automated inspection systems to support grid reliability and economic performance. This need has driven the use of computer vision and deep learning for scalable fault detection (Masita et al., 2025).

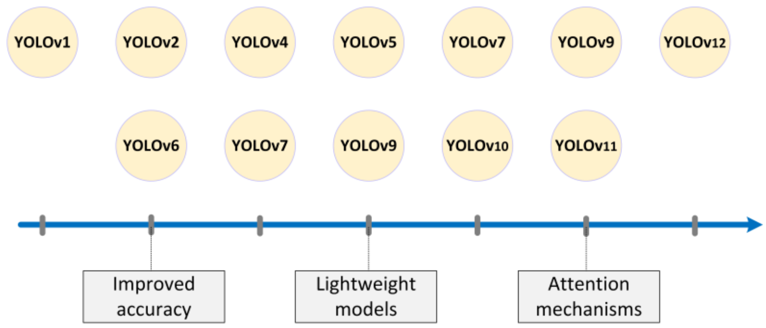

Figure 1 Evolution of YOLO from v1 to v12

Among modern object detectors, the YOLO (You Only Look Once) family is known for combining speed and accuracy. Figure 1 illustrates the evolution of YOLO object detection algorithms over the past decade. It has evolved through several versions designed for real-time performance. While earlier versions, such as YOLOv5 have been applied in PV diagnostics, the more recent YOLOv12 has not yet been explored in the literature for solar fault detection, even though it offers improvements in model size and inference speed (Tian et al., 2025). As shown in the figure, the YOLO models are compared based on three key factors: improved accuracy, lightweight architecture, and attention mechanisms. In the second row, YOLOv7 and YOLOv9 are intentionally repeated to enable direct comparison with YOLOv6, YOLOv10, and YOLOv11 under these criteria.

This study aims to implement and assess YOLOv12n, the lightweight variant of YOLOv12, on a public dataset of annotated solar panel images. The purpose of the study is to focus on the potential of YOLOv12 for deployment in low-resource settings and evaluate its ability to detect multiple types of defects using both EL and RGB images. Employing the proposed model would shift PV monitoring from manual checks to scalable real-time inspection.

2. LITERATURE REVIEW

Traditional solar panel inspection relies on manual processes, which are prone to human error, especially in large-scale solar installations. Deep learning approaches represent a significant advancement in this field (Maddileti et al., 2025; Xue et al., 2025). A foundational study by Deitsch et al. (2019) employed convolutional neural networks (CNNs) to automatically classify photovoltaic cell defects from electroluminescence images. Their work demonstrated that CNNs outperformed rule-based image analysis, particularly in binary classification tasks distinguishing defective from non-defective cells. However, the approach lacked object-level localization, which is essential for pinpointing fault regions in PV modules. This limitation motivated further research into integrating real-time object detection methods for solar inspection. According to the study conducted by Masita et al. (2025), CNN-based classification provided the groundwork. However, real-time localized PV fault detection only became practical with the development of one-stage detectors such as YOLO.

The YOLO framework has progressively established new standards in real-time object detection performance. Initial versions, including YOLOv1 and YOLOv3 pioneered grid-based detection methodologies. YOLOv5 subsequently emerged as a mainstream solution by achieving an optimal equilibrium between detection speed and accuracy. The latest iteration, YOLOv12, advances this legacy significantly through an attention-focused architecture enabling versatile vision tasks such as detection, classification and instance segmentation (Tian et al., 2025). The architecture integrates advanced backbone designs such as RepNCSPELAN4 and is optimized to achieve high inference speed, suitable for real-time deployment. Among the available variants, YOLOv12n offers an ultra-lightweight configuration. This model's computational efficiency makes it well-suited for edge computing applications, including drone-assisted photovoltaic inspections and embedded diagnostic platforms.

Recent benchmarks show the YOLO architecture consistently achieving state-of-the-art performance in photovoltaic fault detection. GBH YOLOv5, presented by Li et al. (2023), exemplifies this trend by combining Ghost Convolution with a refined prediction head. When tested on the PV Multi Defect dataset, this model accurately detected five types of solar panel failures using bounding box annotations. In contrast, the most recent study by Tian et al. (2025) evaluated YOLO versions up to YOLOv12, comparing their latency and accuracy across a different set of datasets. As shown in Figure 2, YOLOv12 achieved the best overall performance among the tested versions. However, YOLOv12 has not yet been evaluated on the PV Multi Defect dataset, which Li et al.'s study demonstrated to be a strong benchmark for YOLO-based photovoltaic fault detection. This presents a clear opportunity to investigate whether the latest architecture can further improve inference speed and generalization while maintaining or surpassing detection accuracy on this specific dataset.

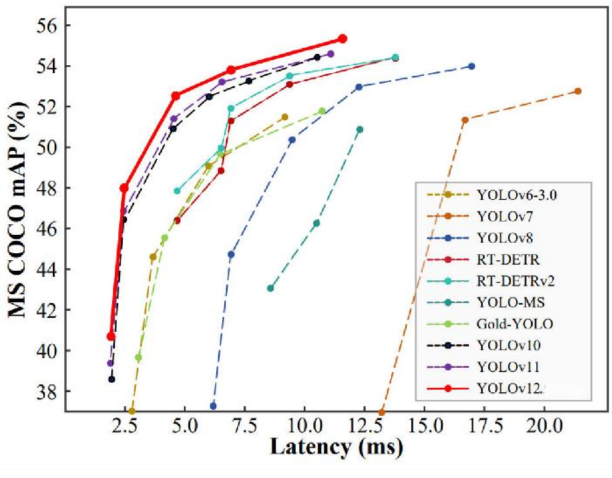

Figure 2 YOLO comparison in terms of latency-accuracy (Tian et al., 2025)

Earlier photovoltaic datasets, such as ELPV (Deitsch et al., 2019), lacked bounding box annotations, which limited their suitability for object detection tasks. In contrast, newer datasets like PV-Multi-Defect provide high-quality annotations, enabling more robust benchmarking (Li et al., 2023). YOLO models, particularly v5 and v12, are optimized for fast inference, with YOLOv12n offering enhanced deployability in real-time scenarios due to its minimal resource footprint (Tian et al., 2025). Complex solar panel images frequently contain multiple, overlapping faults, and state-of-the-art object detectors are now being evaluated for their ability to address such challenges (Li et al., 2023). Recent studies also advocate deeper visual evaluations using PR curves, ROC–AUC plots, and class-wise confusion matrices to assess per-class model robustness (Masita et al., 2025; Li et al., 2023). To address these gaps, the main contributions of this study are as follows. We implement and evaluate YOLOv12n, the lightweight variant of YOLOv12, on a public dataset of annotated solar panel images. In addition, we assess the capabilities of YOLOv12n in detecting five key fault types: scratch, broken cell, hotspot, black border, and no electricity. Finally, we compare its performance against YOLOv5, the baseline model used in the GBH-YOLOv5 base paper.

3. METHODOLOGY

This study employs an experimental and comparative research design. Its primary objective is to assess the newly released YOLOv12 model's effectiveness for object-level detection of photovoltaic panel defects utilizing a real-world annotated dataset. The design facilitates not only model training and validation but also enables a rigorous performance comparison against the established YOLOv5 baseline. This comparison specifically evaluates potential enhancements in inference speed, precision, recall and mean Average Precision (mAP) metrics (Maddileti et al., 2025).

The study utilizes the PV Multi Defect dataset on the basis of Li et al. (2023). It comprises 1,105 annotated solar cell RGB images across five defect categories: Scratch, Broken, No electricity, Black border, and Hot spot. Each image features XML files containing bounding box coordinates, facilitating object-level detection rather than basic classification. The dataset offers a robust benchmark for multi-defect detection research while accommodating real-world variability essential for scalable experimentation. Figure 3 details the annotation distribution across these defect classes.

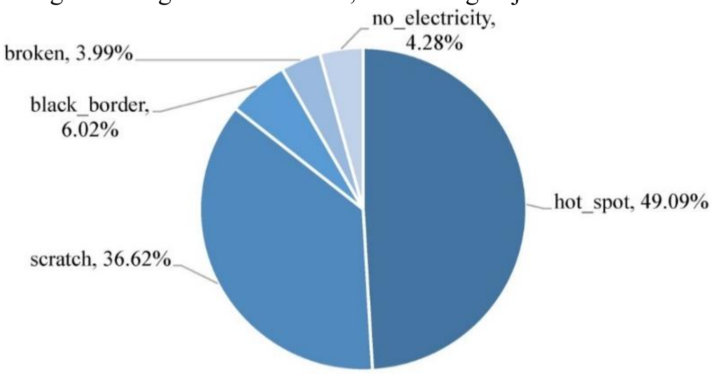

Figure 3 PV-Multi-Defect Dataset annotation distribution (Li et al., 2023)

Substantial preprocessing adapted the PV Multi Defect dataset for YOLO compatibility. A custom Python parser first transformed XML annotations into YOLO TXT format, adhering strictly to specification requirements: class identifier, normalized center coordinates, width and height dimensions all scaled 0 to 1. Second, balanced data partitioning allocated 70% of images for training (773 samples), 20% for validation (222 samples), and 10% for testing (110 samples), preserving proportional class representation. Third, the directory organization systematically structured images and labels into dedicated images/train, images/val, images/test folders with corresponding labels directories. These preprocessing stages ensured annotation integrity and seamless compatibility with Ultralytics' YOLOv12 implementation.

Model selection prioritized the YOLOv12n architecture, the lightweight variant within the YOLOv12 family (Tian et al., 2025). This model proved ideal given its reduced parameter count, minimized memory requirements, and suitability for real-time deployment. It enabled accelerated inference speeds on resource-constrained hardware, specifically the Tesla T4 GPU available via Google Colab. The adoption of YOLOv12 was further justified by its attention-focused architectural enhancements over predecessors and its demonstrated efficacy in established object detection benchmarks (Tian et al., 2025).

Training proceeded through multiple phases to evaluate the impact of extended training durations. Initial training covered 50, 100, and 200 epochs, subsequently extended to 500 epochs. Early stopping was configured with a patience threshold of 100 epochs, permitting automatic termination upon validation metric stabilization. The training configuration employed a batch size of 16, an image resolution of 640×640 pixels, and default learning rate parameters. Data augmentation leveraged intrinsic YOLOv12 framework capabilities, incorporating mosaic augmentation, horizontal flipping, scaling, and copy-paste blending techniques. These intrinsic augmentations enhanced model generalization without requiring supplementary manual scripting.

Performance assessment utilized established object detection metrics: Precision, Recall, mAP@0.5, and mAP@0.5–0.95. The mAP metric provides a consolidated measure of detection accuracy and localization precision. Complementing numerical metrics, diverse visual evaluation tools were employed: confusion matrices, precision-recall curves, F1 score plots, and class-specific ROC AUC curves. These tools facilitated granular performance analysis across defect classes, enabling detailed examination of both prevalent and underrepresented fault types.

4. RESULTS AND ANALYSIS

A rigorous evaluation protocol trained the YOLOv12n model across 500 maximum epochs. Early stopping activated at epoch 412 after 100 consecutive epochs without performance gains. This mechanism prevented overfitting while securing optimal weight preservation.

Training execution occurred via Google Colab's Tesla T4 GPU backend, leveraging YOLOv12n's performance efficiency balance. The lightweight architecture incorporated native Ultralytics framework optimizations, including mosaic augmentation and automatic anchor learning. One of the most informative visual components was the training and validation loss curves (Figure 4), which demonstrated steady convergence patterns across all monitored metrics. These curves clearly depict the progressive reduction of box loss, classification loss, and distribution-focused loss (DFL) over 412 epochs prior to early stopping. While the training losses showed consistent improvement, validation losses exhibited slightly more fluctuation, which is a common occurrence attributed to the scale limitations of the PV-Multi-Defect dataset. Despite these minor oscillations, both curves maintained a robust downward trend, confirming effective model learning without significant overfitting. This consistent reduction in validation losses underscores the model’s stable convergence despite dataset constraints.

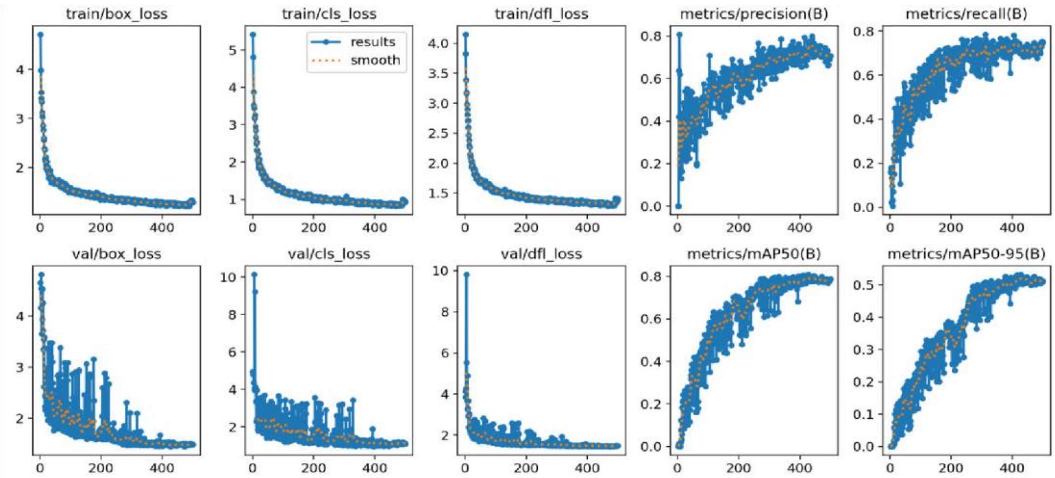

Figure 4 Training and Validation losses for 500 epochs

The final evaluation utilized the 110 sample test split encompassing all five defect classes. Model performance was quantified through precision, recall, and mAP metrics. The obtained results are presented in Table 1.

Table 1 Performance metrics for the YOLOv12n model

Metric	Best value (%)
Precision	85.61
Recall	79.87
mAP@0.5	85.68
mAP@0.5-0.95	56.2

F1 score and precision recall curves (Figures 5 and 6) demonstrated superior accuracy and threshold stability for hot spot, black border, and no electricity defect classes. The F1-confidence curve (Figure 5) offers valuable insight into the model's optimal confidence threshold for all defect classes. This curve visualizes the relationship between the F1 score, defined as the harmonic mean of precision and recall, and varying confidence thresholds. The thick blue curve, which aggregates performance across all classes, reached its maximum at a threshold of approximately 0.42, achieving an F1 score of 0.74. This peak suggests a practical decision boundary suitable for deployment scenarios. Class-specific curves further revealed that hot spot and no electricity defects consistently maintained high F1 scores over a wide range of thresholds, highlighting the strong learnability and reliable detection of these defect types. Additionally, the Precision-Recall (PR) curve (Figure 6) illustrated the trade-off between recall and precision across the different defect categories. The high area under the curve (AUC) values for most classes, particularly no electricity (0.948) and black border (0.854), demonstrated that the model could maintain both high recall and high precision simultaneously. This capability is especially critical for applications such as solar fault diagnosis, where reliable detection directly affects system safety and performance (Lu et al., 2024).

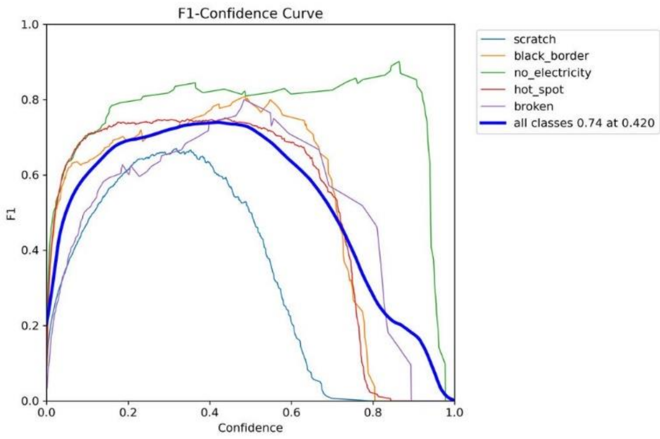

Figure 6 F1-Confidence curves for all classes

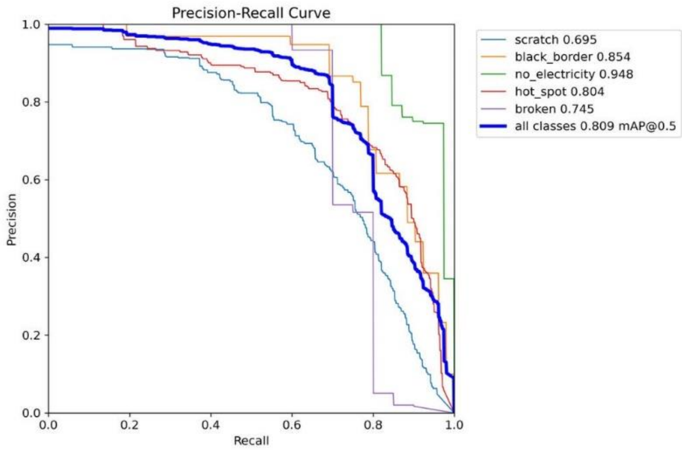

Figure 5 Class wise Precision-Recall Curve with mAP values

On the classification side, the confusion matrix (Figure 7) provided a clear summary of the model's prediction performance for each class. It showed strong true positive rates for hot spot and no electricity, while scratch exhibited more misclassifications. Similar to Li et al.'s study using YOLOv5, scratch defects remain challenging to detect due to their small size (about 4 × 32 pixels, less than 0.08% of the original image). While Li et al.'s GBH-YOLOv5 improved detection of these tiny targets, our YOLOv12 model still shows more misclassifications for scratches compared to other classes, likely because of their subtle appearance and resemblance to natural panel textures. This highlights that detecting small, faint defects remains difficult even with improved models. The matrix was normalized to account for class imbalance, allowing clearer interpretation of class-level trends. These images displayed the model's outputs overlaid on test samples, showing bounding boxes with class labels and confidence scores.

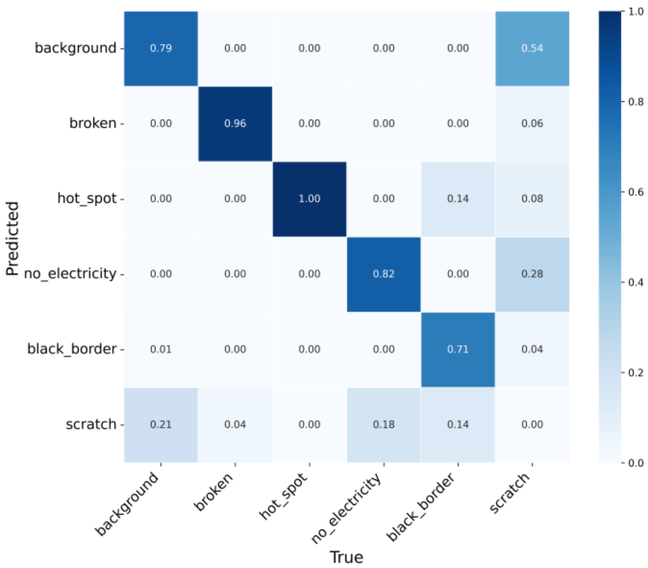

Figure 7 Normalized confusion matrix on test set

The predictions were spatially precise, with boxes closely matching the annotated defect regions. Even in cases involving multiple defects or challenging lighting and surface textures, YOLOv12n maintained consistent and coherent predictions. This qualitative validation reinforces the reliability of the quantitative evaluation.

Benchmarking against the GBH YOLOv5 baseline revealed significant improvements, which are shown in Table 2. YOLOv12n surpassed YOLOv5s across all metrics despite its reduced computational footprint. This performance gain highlights the efficacy of YOLOv12's attention-focused architecture for edge deployment scenarios, particularly drone-assisted inspection and embedded monitoring systems. ROC AUC analysis further validated model robustness, showing notable competence in detecting infrequent classes (broken, black border) – indicating strong handling of long tail data distributions.

Table 2 Comparison between YOLOv5s from the base paper and YOLOv12n results

Model	**mAP@0.5 (%)**	**Precision (%)**	**Recall (%)**
YOLOv5s	78.1 ± 0.06	83.2 ± 0.05	73.4 ± 0.06
YOLOv12n	85.7	85.6	79.9

The study highlights the strong performance of YOLOv12n on the PV-Multi-Defect dataset. The model achieved high precision, recall, and F1 scores, particularly for hot spot and no electricity defects, which had distinct visual patterns. An overall mAP@0.5 of 85.68% confirms reliable localization and classification. YOLOv12n demonstrated robustness under real-world challenges such as texture variation, lighting changes, and partial occlusion. Its lightweight design allowed fast inference, making it well-suited for edge applications like drone-based inspections. Scratches showed lower performance, with a peak mAP@0.5 of 69.5% and precision below the top classes, likely due to their fine, low-contrast features that blend with the background or mimic normal panel textures. Even so, the model performed better than random, indicating room for improvement with enhanced data quality or preprocessing. Overall, YOLOv12n proves effective and scalable for real-time photovoltaic fault detection.

From a practical perspective, the automated detection and localization of multiple fault types in solar panels can significantly improve maintenance efficiency in renewable energy systems. Manual inspection is often slow and labor-intensive, and prone to human error. These challenges are amplified in large solar farms. Using computer vision models like YOLOv12 allows for faster and more consistent inspections, which improves system reliability, reduces downtime, and lowers operational costs. While the findings are promising, it is important to recognize the limitations of this study. A primary constraint is the dataset. The PV-Multi-Defect dataset includes just over 1100 annotated images. Although it captures a range of defect types, it may not fully reflect the variety of real-world conditions found across different solar installations. Class imbalances also posed a challenge. Categories such as broken had relatively few samples, which may have limited the model's ability to generalize for those faults.

5. CONCLUSION AND FUTURE STUDIES

This study evaluated the effectiveness of YOLOv12 for detecting faults in solar photovoltaic (PV) cells using the PV-Multi-Defect dataset. By implementing YOLOv12n on Google Colab, a complete pipeline was developed from preprocessing to evaluation. The model achieved strong performance with a peak precision of 85.6%, recall of 79.9%, and mAP@0.5 of 85.7%, confirming its suitability for real-time, lightweight PV fault detection. Compared to the GBH YOLOv5 baseline, YOLOv12 provided similar accuracy with faster inference and a smaller model size. The model showed robustness across diverse lighting, defect types, and panel structures. While it handled defects like hot spot and no electricity confidently, subtle classes such as scratches remained challenging. Still, the multi-class detection pipeline proved effective and reliable. Future work will focus on enhancing both the technical performance and practical utility of the detection system. A key direction involves moving beyond uniform class treatment to incorporate a consequence-based ranking for the detected faults. This approach will differentiate safety-critical classes, such as hot spots that pose a fire hazard, from defects that primarily impact system performance, such as scratches or electrical disconnections. To operationalize this, future research will explore techniques like weighted loss functions that prioritize detection accuracy for critical faults during model training, and the development of a tiered alert system within the pipeline to trigger immediate responses for high-risk defects. Technically, improvements will also include employing larger YOLOv12 variants to enhance accuracy for subtle defects and expanding the dataset with more real-world images, with a specific focus on addressing the imbalance of critical fault classes. Further investigations will involve integrating hierarchical image classification and attention mechanisms, similar to Boone-Sifuentes (2022), to improve the discrimination of complex and faint anomalies.

1. REFERENCES

Al-Otum, H. M. (2023). Deep learning-based automated defect classification in electroluminescence images of solar panels. *Advanced Engineering Informatics*, *58*, 102147.

Boone-Sifuentes, T., Bouadjenek, M. R., Razzak, I., Hacid, H., & Nazari, A. (2022). A mask-based output layer for multi-level hierarchical classification. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (pp. 3833-3837).

Deitsch, S., Christlein, V., Berger, S., Buerhop-Lutz, C., Maier, A., Gallwitz, F., & Riess, C. (2019). Automatic classification of defective photovoltaic module cells in electroluminescence images. *Solar Energy*, *185*, 455-468.

Ding, S., Jing, W., Chen, H., & Chen, C. (2024). Yolo Based Defects Detection Algorithm for EL in PV Modules with Focal and Efficient IoU Loss. Applied Sciences (2076-3417), 14(17).

Ghahremani, A., Adams, S. D., Norton, M., Khoo, S. Y., & Kouzani, A. Z. (2025). Advancements in AI-Driven detection and localisation of solar panel defects. *Advanced Engineering Informatics*, *64*, 103104.

Ghahremani, A., Adams, S. D., Norton, M., Khoo, S. Y., & Kouzani, A. Z. (2025). Detecting defects in solar panels using the YOLO v10 and v11 algorithms. *Electronics*, *14*(2), 344.

International Energy Agency, "Renewables 2023 – Analysis," IEA, Jan. 2024. https://www.iea.org/reports/renewables-2023

Jahn, U., Herz, M., Köntges, M., Parlevliet, D., Paggi, M., Tsanakas, I., ... & Tanahashi, T. (2018). *Review on infrared and electroluminescence imaging for PV field applications: International Energy Agency Photovoltaic Power Systems Programme*. International Energy Agency.

K. Masita, A. Hasan, T. Shongwe, and H. A. Hilal, "Deep learning in defects detection of PV modules: A review," Solar Energy Advances, vol. 5, p. 100090, 2025, doi: https://doi.org/10.1016/j.seja.2025.100090.

Karakan, A. (2025). Detection of Defective Solar Panel Cells in Electroluminescence Images with Deep Learning. *Sustainability (2071-1050)*, *17*(3).

Li, L., Wang, Z., & Zhang, T. (2023). Gbh-yolov5: Ghost convolution with bottleneckcsp and tiny target prediction head incorporating yolov5 for pv panel defect detection. *Electronics*, *12*(3), 561.

Lu, S. D., Liu, H. D., Wang, M. H., & Wu, C. C. (2024). A novel strategy for multitype fault diagnosis in photovoltaic systems using multiple regression analysis and support vector machines. *Energy Reports*, *12*, 2824-2844.

Maddileti, N. S., Namburi, R., Raj, R. D. A., Yanamala, R. M. R., & Pallakonda, A. (2025). DCGAN-Driven Minority Class Augmentation for Lightweight YOLO-Based Photovoltaic Defect Localization Suitable for Edge Deployment. *IEEE Transactions on Device and Materials Reliability*.

Tian, Y., Ye, Q., & Doermann, D. (2025). Yolov12: Attention-centric real-time object detectors. *arXiv preprint arXiv:2502.12524*.

Velasco-Sánchez, E., Recalde, L. F., Guevara, B. S., Varela-Aldás, J., Candelas, F. A., Puente, S. T., & Gandolfo, D. C. (2024). Visual servoing nmpc applied to uavs for photovoltaic array inspection. *IEEE Robotics and Automation Letters*, *9*(3), 2766-2773.

Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 7464-7475).

Xue, H., Liu, L., Wu, Q., He, J., & Fan, Y. (2025). Defect Detection Algorithm for Photovoltaic Cells Based on SEC-YOLOv8. *Processes*, *13*(8), 2425.

