

Anatoliy Sachenko, Pavlo Radiuk, Mykola Lysyi, Oleksandr Melnychenko, Oleg Zastavnyy

ARCHITECTURE OF CYBERPHYSICAL SYSTEMS FOR UAV-BASED LATE-FUSION DEFECT DETECTION IN PHOTOVOLTAIC MODULES

The **subject** of research is the architectural advancement of inspection systems for large-scale solar power plants. As global solar infrastructure expands, reliance on manual or offline analytical methods creates significant operational bottlenecks. The **goal** of research is to improve the operational utility of unmanned aerial vehicle (UAV)-based photovoltaic module inspection by developing a cyber-physical system (CPS) architecture. It integrates onboard deep learning, edge nodes, cloud analytics, and Supervisory Control and Data Acquisition (SCADA)-aware decision-making into a single coordinated workflow. The **tasks** of research: 1) formalise a multi-tiered CPS architecture (UAV-edge-cloud) and define interfaces for data, geo-tags, and alarms; 2) develop and validate an onboard thermographic detection pipeline with palette-aware fusion; and 3) integrate detection results with a SCADA-aware logic layer for hazard inference and fire risk mitigation. The **methods** of research: Computer vision and deep learning (YOLOv11) are used for onboard defect segmentation. Model ensembling via a late-fusion strategy for M2 and M3 thermal palettes mitigates domain shift. RTK-supported spatial clustering algorithms ensure precise geo-indexing and deduplication, and deterministic Boolean logic assesses fire risks based on bypass diode states. The **results** obtained with five-fold cross-validation shows the proposed architecture significantly outperforms single-modality baselines. The onboard YOLOv11 model achieved a macro mAP@0.5 of 0.91 and 0.90 for M2 and M3 palettes, respectively. The late-fusion ensemble elevated mAP@0.5 for cracks to 0.96 and delamination to 0.95. It reduced end-to-end per-frame processing latency from 4.235 s to 2.858 s. Field validation demonstrated an error of 0.71 defects per inspected string compared to manual counts. Sensitivity analysis highlighted that a 10 m flight altitude provides an optimal balance, yielding 93% precision and 90% recall. **Conclusions:** Treating UAV inspection as an integrated cyber-physical service improves defect detection. This offers a scalable, real-time solution for preventive maintenance and automated fire-risk mitigation in renewable energy.

Keywords: cyber-physical systems; photovoltaic modules; unmanned aerial vehicles; infrared thermography; edge computing; deep learning.

Introduction

Photovoltaic (PV) energy has evolved from a niche technology into a vital component of modern power systems. At the same time, photovoltaic power plants remain vulnerable to harsh environmental conditions and a long list of faults that are commonplace but collectively costly [1]. Microcracks can form during manufacturing, transportation, installation, or hail; they can lead to inactive cell areas, hot spots, and energy losses [2]. Aging and delamination of the encapsulant can lead to moisture accumulation and accelerate corrosion. Contamination reduces the intensity of solar radiation on the cells and alters the thermal regime, while damaged connections or bypass diodes can cause abnormal heating, posing a safety risk. Therefore, reliability-focused maintenance is mandatory [3]: predictive and preventive inspections directly impact energy output and the plant's service life.

Infrared (IR) thermography is one of the most informative methods for non-contact diagnostics of photovoltaic systems. It provides spatial temperature maps that reveal electrical mismatches, increased series

resistance, shunting, and localized degradation. Reviews show that thermography can detect a wide range of photovoltaic system anomalies, but also emphasize that its diagnostic value depends on data acquisition conditions, assumptions regarding emissivity, inspection geometry, and the ability to distinguish defect signatures from reflections and environmental gradients [4, 5]. Comparative studies also show that defects can manifest differently under lighting conditions and in darkness, meaning that the inspection protocol is just as important as the camera itself [6].

Unmanned aerial vehicles (UAVs) have recently become the standard platform for large-scale photothermal imaging. UAVs can cover industrial-scale fields in a matter of minutes and inspect hard-to-reach rooftop installations [7, 8]. Many studies report the operational advantages of UAV thermography, including improved coverage and reduced inspection time [9], but also note practical limitations: flight altitude and speed control spatial resolution and motion blur; time of day and weather affect thermal contrast; and specular reflections from glass surfaces can create misleading hot spots [10, 11]. These limitations are not mere nuisances;

they are the reason why “laboratory” image analysis pipelines often fail when used in the field.

Automation is the second major trend. Deep learning has achieved high performance in classifying and localizing defects in photovoltaic systems on infrared and RGB images, and surveys describe a rapidly growing set of architectures ranging from classical convolution-based classifiers to modern single-stage detectors and segmentation networks [12, 13]. Recent UAV-oriented approaches use variants of the YOLO family to detect modules and anomalies in thermal and visible data [14, 15], while methods based on thermal video demonstrate that automation can be extended from individual frames to temporal streams [16]. However, a persistent limitation is that many published solutions implicitly assume an offline workflow: data is collected, downloaded, and processed in batches [12, 17]. In operational photovoltaic power plants, the inspection process is a closed-loop system that must meet requirements for latency, throughput, and reliability, as well as provide results that can be traced back to physical modules and utilized by operators [18].

This observation calls for the creation of a system-level architecture. The photovoltaic plant inspection pipeline is not merely a computer vision model; it is a cyber-physical system (CPS) [19] that integrates sensing, computing, communication, and control functions. The UAV captures physical measurements, but the station requires cyber actions: alerts, maintenance tickets, and safety procedures. Edge nodes offer a natural architectural solution: they can reduce bandwidth by processing data close to the source and provide fast responses without relying on a cloud connection [20]. However, despite the maturity of edge nodes in industrial monitoring, UAV solutions for inspecting photovoltaic stations rarely provide an explicit KFS architecture that integrates on-board inference, edge aggregation, long-term cloud analysis, and SCADA-aware decision-making logic [21]. Similarly, the literature on photovoltaic thermography typically emphasizes detection accuracy while underestimating the georeferencing and deduplication steps necessary to avoid false positives and alarm fatigue.

Therefore, the problem addressed in this article lies in the need to develop and validate an accurate and operationally implementable UAV-based PV plant defect monitoring architecture. The technical gap lies not simply in “detecting more defects,” but in creating a reliable, comprehensive service with sequential georeferencing, real-time calculations on onboard

hardware, and direct integration with security surveillance and monitoring systems. This is particularly relevant for fire prevention, as abnormal heating and faults in shunt diodes can quickly turn local defects into critical incidents. The integration of UAV-based detection with SCADA logic aligns with recent trends that view unmanned technologies and artificial intelligence as integral components of power plant safety monitoring systems [22, 23].

The aim of this study is to improve the operational utility of UAV-based photovoltaic module inspections by developing a cyber-physical architecture that integrates embedded deep learning, edge nodes, cloud analytics, and SCADA-aware decision-making logic into a single coordinated workflow.

To achieve this goal, the paper proposes three main scientific contributions, namely:

- 1) a multi-level CFA (UAV-periphery-cloud) architecture for inspecting photovoltaic modules, which clearly defines data flow, responsibilities, and decision-making points, including RTK-based geotagging and deduplication to create traceable defect events at the module level;
- 2) a thermography-oriented on-board inference pipeline based on the YOLOv11 segmentation model, with post-processing that accounts for the color palette and an original late-merging ensemble between two representations of thermograms, ensuring high detection quality and significantly reducing end-to-end processing time;
- 3) integration of detection results with an interpretable logic layer that accounts for SCADA states (in particular, shunt diodes) for automated hazard inference and rapid fire risk mitigation.

Related Works

Research on UAV-based monitoring of photovoltaic module defects spans sensor physics, computer vision, and systems engineering. This section briefly discusses the place of this work in this context and explains how the proposed approach to cyber-physical operation differs from approaches based solely on detectors.

Thermography is widely used to detect hot spots, anomalies, and signs of degradation in photovoltaic modules. Comprehensive reviews emphasize that thermal signals are influenced by solar irradiance, ambient temperature, wind, and viewing angle, and discuss recommended inspection procedures and limitations [4, 5].

Gallardo-Saavedra et al. [6] compared lighting and darkness conditions and demonstrated that defect detection depends on operating conditions, highlighting the need for consistent flight protocols and contextual information, such as solar irradiance and environmental conditions. The practical application of UAVs has been evaluated in field studies that highlight how flight altitude and motion control affect spatial resolution and, consequently, the ability to detect anomalies [7, 11]. These findings prompt a focus on the sensitivity of flight parameters and reflection-aware planning.

Deep learning for the diagnosis of photovoltaic systems has progressed from the classification of cropped module images to detection and segmentation at the power plant level. Reviews summarize the expanding use of neural networks, object detectors, and hybrid models in the monitoring of photovoltaic systems [17, 24]. Early work by Dotenko et al. [25] demonstrated the automatic detection of photovoltaic modules and defect analysis in aerial infrared images using statistical tests and classical vision, reporting high defect classification accuracy while emphasizing the importance of robust preprocessing. Later works increasingly rely on single-stage detectors. Di Tommaso et al. [14] proposed a multi-stage pipeline using YOLOv3 for detecting modules and anomalies in both infrared and visible images, which improved automation but still required careful separation of tasks and datasets. Xi et al. [18] introduced ST-YOLO for detecting faults in photovoltaic systems on infrared images and reported a very high mAP@0.5 under their experimental conditions. Other studies focus on investigating thermographic features and reliability, including deep learning strategies for fault diagnosis from thermograms [26] and segmentation-based methods for identifying modules and defects [27]. These works prompt the use of a detector with segmentation capabilities (YOLOv11-seg) and attention to palette-induced domain shift.

Although thermography is central, many implementations use RGB images to provide context and reduce false positives. RGB-based deep learning is used to classify signs of faults and operational issues in photovoltaic stations [28]. Multimodal pipelines can also support human interpretation, as operators can check the RGB context when thermal features are ambiguous [29]. Therefore, the proposed architecture involves synchronized RGB and thermal imaging, although the defect detector in this paper focuses on thermal masks.

From a systems perspective, UAV inspection is an example of distributed sensing with severe constraints [30]. Edge nodes are often recommended for the Industrial Internet of Things because they reduce dependence on the cloud and provide low-latency solutions [20]. Tang et al. [22] demonstrated a deep learning architecture at the cloud edge for detecting linear defects in large photovoltaic power plants, showing that distributed computing can support plant-level monitoring and reduce data transmission volume. However, as noted by Ferlito et al. [23], many edge-of-the-cloud studies focus solely on “edge detection” and do not explicitly model feedback to plant operations. This study extends this line of research by adding RTK-based georeferencing, duplicate suppression, and SCADA-aware fire risk logic, motivated by the observation that photovoltaic system defects have implications for both productivity and safety [21]. Thus, the technical gap lies not simply in detecting a greater number of defects, but in creating a comprehensive cyber-physical architecture capable of functioning in real time and integrating with the production process.

Thus, the literature review confirms that, to overcome existing limitations, it is necessary to develop and validate a UAV-based defect monitoring architecture for photovoltaic systems that is accurate, georeferenced, and fully functional under real-world operating conditions. Based on the analysis, the overall research objective was broken down into 5 interrelated tasks.

1) Formalize a multi-level UAS architecture (UAS-periphery-cloud) and define protocols and interfaces for the efficient exchange of data, geotags, and alarms.

2) Develop and adapt a lightweight deep learning model (YOLOv11) to perform accurate on-board segmentation of photovoltaic module defects directly in real time.

3) Propose and validate an embedded thermal detection pipeline that uses an ensemble strategy (late merging of M2 and M3 palettes) to mitigate domain shift and improve the accuracy of recognizing different types of damage.

4) Develop spatial clustering algorithms with RTK navigation support to ensure reliable georeferencing and eliminate duplication (deduplication) of detected defects at the edge computing level.

5) Integrate the obtained detection results with a deterministic logic layer based on SCADA data for automated fire risk assessment and instant generation of critical alerts to operators.

Methods and Materials

This section describes the proposed approach in detail. It covers the KFS architecture, sensor and communication design, flight planning with mapping, dataset construction, on-board deep learning, edge/cloud analytics, and SCADA-informed decision-making. The core principle is to treat defect detection as a cyber-physical workflow: sensors, computations, communications, and actions must be designed in concert.

System Architecture and Operational Cycle

In this study, the inspection of photovoltaic systems is implemented as a closed-loop cycle, starting with measurement and ending with action at the power plant level. The CPS is organized into three computational levels (Fig. 1): the UAV level, the peripheral level deployed at the photovoltaic power plant, and the cloud level for long-term analysis. The operator level interacts with the CPS via a control panel and SCADA/HMI interfaces.

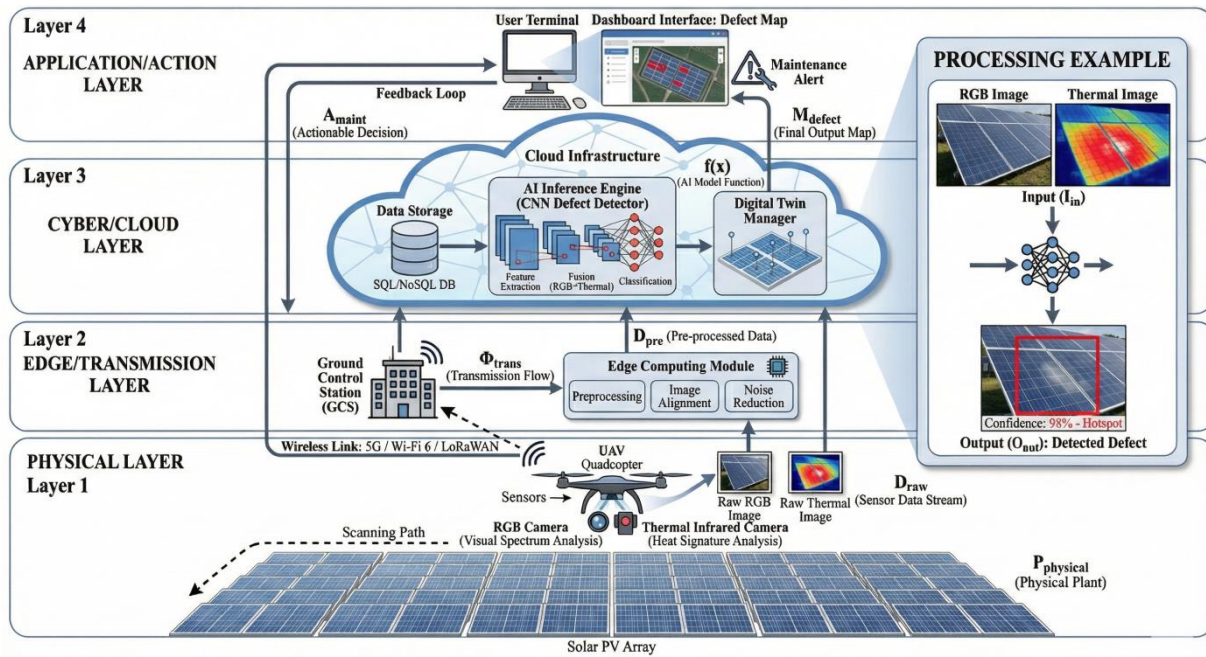


Fig. 1. Architecture of a cyber-physical system for monitoring defects in photovoltaic systems based on UAVs

The UAV level is responsible for tasks requiring rapid response. It captures synchronized RGB and thermal images and performs on-board computations to detect potential defects. The edge layer is responsible for tasks related to the plant context: it performs geotagging using RTK, removes duplicates of recurring detections, aggregates them into events at the module level, and interacts with SCADA to trigger alarms and initiate maintenance. The cloud stores mission artifacts and provides dashboards for trend analysis, module history, and inspection reports.

Let the task generate a continuous stream of frames $\{I_t\}_{t=1}^T$, where each frame contains an RGB image, I_t^{rgb} a thermal image I_t^{th} , and a navigation record n_t (RTK position, altitude, heading, and timestamp). The built-in inference system performs a set of detections as follows:

$$\mathcal{D}_t = \left\{ (c_k, M_k, s_k, g_k) \right\}_{k=1}^{K_t}, \quad (1)$$

where c_k is the defect class (crack, contamination, delamination), M_k is the segmentation mask, s_k is the confidence score, and g_k is the geotag (calculated location of the module in the object's coordinates).

The edge server receives \mathcal{D}_t and performs (i) spatial clustering and temporal deduplication, (ii) association of detections with module IDs and strings, and (iii) generation of a status report that maps detections and SCADA signals to operational actions. The cloud stores mission artifacts and provides dashboards for trend analysis, module history, and audit reports.

UAV Platform, Sensors, and Communication

KFS provides a UAV platform capable of stable flight along predefined patterns (grid or corridor scanning) and equipped with a combined RGB/thermal payload. The thermal channel provides the primary signal

for detecting abnormal heating patterns; the RGB channel provides contextual information and supports operator verification. Although the architecture is platform-independent, the stated implementation is comparable to a Matrice-class UAV with a dual-sensor payload gimbal (RGB + thermal imager) and RTK positioning.

From an ORM perspective, georeferencing is just as important as the imagery. RTK positioning reduces drift and ensures consistent mapping of detected objects onto physical modules throughout missions. In addition, precise timestamps allow for the synchronization of the thermal camera, RGB camera, and flight controller. The UAV records navigation data for each frame, including RTK position and altitude. This data is subsequently used for deduplication and associating detections with the plant layout.

Communication is treated as a limited resource. Streaming thermal video at full resolution is often impossible via standard UAV channels, especially in large factories. Instead, the UAV transmits compact

detection messages according to formula (1) to a peripheral server in near real-time. Full-resolution images can be cached on board and downloaded after the flight or selectively downloaded when the edge server requests evidence. This design aligns with the general logic of edge nodes in industrial monitoring: local processing reduces bandwidth, and decisions at the plant level can be made independently of cloud connectivity [20].

Choosing a viewpoint to account for reflections

Thermography of photovoltaic modules using UAVs is susceptible to specular reflections from glass surfaces. If the camera is pointed toward a specular reflection of the sun, the thermogram may be obscured by glare, leading to the appearance of false hot spots and distortion of the thermal contrast required for defect detection. Therefore, to mitigate this effect, an observation point selection algorithm that accounts for reflection has been added to the RTK geometry and the sun's position (Fig. 2).




Fig. 2. Geometric model for selecting a viewing point that accounts for reflection. Given the plane of the module n and the direction of the sunbeam s , the direction of specular reflection r is calculated. The desired camera viewing direction v^* is obtained by deviating from r (Rodrigues rotation) to reduce glare while maintaining a sufficient angle of incidence for thermal contrast

Let the three points on the module plane that have an RTK relationship be $p_1, p_2, p_3 \in \mathbb{R}^3$ (e.g., the corners of a representative module). The normal vector of the module is calculated as follows:

$$n = \frac{(p_2 - p_1) \times (p_3 - p_1)}{\|(p_2 - p_1) \times (p_3 - p_1)\|}. \quad (2)$$

Let s be the single vector pointing from the module to the Sun (obtained from the solar ephemeris model at time t). The direction of specular reflection is as follows:

$$r = s - 2(n^\top s)n. \quad (3)$$

If the camera's viewing direction coincides with r , specular reflections dominate. The target's viewing direction v^* is determined by rotating r around the normal n by an angle θ , chosen such that the reflection direction lies outside the camera's field of view. Using the Rodriguez rotation formula, the inverse vector takes the following form:

$$v^* = r \cos \theta + (n \times r) \sin \theta + n(n^\top r)(1 - \cos \theta). \quad (4)$$

Algorithm 1 summarizes the procedure for selecting the viewpoint. In practice, Θ may be a small set of possible

angles (e.g., $\pm 5^\circ$, $\pm 10^\circ$), and the glare assessment may be approximated by the angular distance between v_θ and r relative to the camera's field of view. The key

point is that the algorithm is lightweight and compatible with real-time decision-making at the edge.

Algorithm 1. Selection of a viewing point accounting for reflection for UAV thermography

Requirements: Points on the RTK plane p_1, p_2, p_3 , sun direction s , possible rotation angles Θ

Implementation: Select the viewing angle v^* that minimizes glare

- 1: Calculate the surface normal n using formula (2)
- 2: Calculate the reflection direction r using formula (3)
- 3: **For all** θ in Θ , **perform**
- 4: Calculate the possible v_θ using formula (4)
- 5: Estimate the glare level (expected reflection in the camera's field of view)
- 6: end for
- 7: Select $v^* = \arg \min \{ \text{GlareScore}(v_\theta) \}$
- 8: **return** v^*

Building a dataset and presenting thermograms

Thermal cameras often support multiple color palettes that represent temperature ranges as colors for visualization. Although these palettes are designed for human interpretation, they also alter the statistical distribution of pixel values fed into the training model. The experiments consider two representations of thermograms: a two-color palette (labeled M2) and a three-color palette (labeled M3). The key difference lies in the effective temperature range highlighted by the visualization. In the described configuration, M3 uses a narrower dynamic range of approximately 15°C , which increases sensitivity to small temperature contrasts (on the order of $0.1\text{--}0.5^\circ\text{C}$) but may saturate extreme anomalies. M2 uses a wider range of approximately 35°C , which reduces sensitivity to small hot spots but covers the entire spectrum from minor heating inconsistencies to serious damage. M2 and M3 are considered as two reference representations of thermograms obtained from the same physical scenes to assess robustness to palette-induced domain differences and to motivate selective merging.

To develop and validate the proposed approach, two publicly available UAV thermography datasets were additionally used to provide additional defect coverage:

1. STHS-277 [31]: This dataset contains 277 full-frame thermal images capturing defects in the form of “snail trails” and hotspots, and includes environmental metadata. We extend the provided annotations by adding bounding boxes for all instances of photovoltaic panels to ensure comprehensive detector training.

2. PVF-10 [32]: This is a large-scale, high-resolution UAV thermal imaging dataset containing 5,579 annotated patches of individual solar panels from eight power plants and a detailed taxonomy of ten defect classes.

For the on-board segmentation experiments discussed in this article, we prepared a compact dataset of masks with three classes (cracks, contamination, delamination) by combining field thermography with selected samples from the aforementioned publicly available datasets. The images were annotated using the Computer Vision Annotation Tool (CVAT) [33] and exported as polygonal masks suitable for segmentation training.

Since the minority classes are small, the evaluation is based on stratified five-fold cross-validation to avoid inaccurate point estimates from a small test set. Table 1 shows the composition of the dataset and the approximate partitioning used in the evaluation.

Table 1. Dataset structure and approximate partitioning used for stratified 5-fold cross-validation (number per partition)

Class	Defect Type	Total images	Training/Validation/Testing per component (count)
1	Crack	20	14/2/4
2	Contamination	10	7/1/2
3	Delamination	170	122/14/34
Total		200	143/17/40

On-board deep learning: YOLOv11 segmentation

The UAV layer uses a lightweight deep learning model for defect detection and segmentation. A segmentation model from the YOLO family was chosen because it provides a good balance between accuracy and speed and has well-developed deployment tools. Specifically, YOLOv11-seg [34] is used as the base network. The model outputs bounding boxes and pixel-level masks for each detected defect. Masks are valuable in the inspection of photovoltaic systems because they allow for the estimation of the defect's area and shape, which improves prioritization and supports temporal tracking.

The detector is trained on the annotated dataset described in Table 1. To mitigate overfitting under conditions of severe class imbalance, training employs transfer learning and class-aware sampling: the network is initialized from a pre-trained YOLO checkpoint and fine-tuned using strong geometric and photometric augmentations (random affine transformations, flipping, blurring, and contrast/brightness fluctuations), while instances of minority classes (cracks and contaminants) are downsampled during mini-batch construction. Public thermal imaging datasets (STHS-277 and PVF-10) [31, 32] are used to expand the thermal distribution during fine-tuning.

During inference, the model generates a set of potential objects $\{(b_k, M_k, s_k, c_k)\}$, where b_k is the bounding box and M_k is the mask. Non-maximum suppression (NMS) removes duplicate bounding boxes, and a similar suppression at the mask level is applied to reduce overlap between instance masks. The post-processing stage is particularly important because UAV video streams contain many nearly identical frames, and thermal palettes can create color artifacts that cause false detections.

Mask similarity, size-based routing, and post-processing

The built-in output provides instance masks that must be filtered and merged to obtain stable defect events. In this study, mask overlap is based on the intersection over union (IoU), which is calculated as follows:

$$\text{IoU}(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (5)$$

Algorithm 2. Embedded processing of a three-color thermogram (M3)

Input: Thermogram I in a three-color palette, YOLOv11-seg model, thresholds $(\tau_c, \tau_{\text{IoU}})$

Execution: Filtered set of defects $\mathcal{D}_{\text{small}}$

Mask-level noise suppression (Mask-NMS) removes duplicates when the IoU exceeds a threshold, retaining predictions with a higher confidence level. Mask-NMS is important in UAV inspection because adjacent frames often contain the same defect with slightly offset masks.

In addition to removing duplicates, the FCS performs defect routing by size between two thermograms. Let S_{def} denote the area (in pixels) of the detected defect mask, and S_{cell} denote the area of a single photodiode at the same image scale. The relative defect area is determined as follows:

$$R_d = S_{\text{def}} / S_{\text{cell}}. \quad (6)$$

Intuitively, one can assume that defects smaller than a single cell behave as localized hot spots and benefit from the higher contrast of M3, whereas larger anomalies are more reliably captured in the M2 thermogram with a wider range.

Using R_d , two mutually exclusive sets of defects are formed as follows:

$$\mathcal{D}_{\text{small}} = \{d \in \mathcal{D}^{\text{M3}} \mid R_d(d) < 1.0\}, \quad (7)$$

$$\mathcal{D}_{\text{large}} = \{d \in \mathcal{D}^{\text{M2}} \mid R_d(d) \geq 1.0\}, \quad (8)$$

and the final set of detections is determined by the formula:

$$\mathcal{D}_{\text{final}} = \mathcal{D}_{\text{small}} \cup \mathcal{D}_{\text{large}}. \quad (9)$$

Since the routing criteria are complementary and Mask-NMS is applied within each branch, the resulting sets of $\mathcal{D}^{\text{M3}} \cap \mathcal{D}^{\text{M2}} = \emptyset$ meet practical requirements, and the number of duplicate alarms is reduced. The application of Mask-NMS and selective merging allowed for a reduction in the number of false alarms by approximately 8% in sequences resembling real-world conditions.

This study identifies two palette-specific post-processing pipelines.

M3 processing (three-color palette). The three-color palette increases sensitivity to small anomalies but can cause banding. Therefore, the M3 branch targets small defects according to formula (7) and uses a conservative IoU threshold (0.4) in Mask-NMS. Algorithm 2 summarizes the steps.

- 1: Convert the thermogram into a multi-channel image (palette separation)
- 2: Run YOLOv11-seg to obtain candidate masks and confidence scores.
- 3: Apply confidence filtering: keep masks with $s_k \geq \tau_c$ (e.g., $\tau_c = 0.25$)
- 4: Compute R_d for each mask using formula (6); keep only instances with $R_d < 1$
- 5: Apply Mask-NMS using an IoU threshold of τ_{IoU} (e.g., 0.4)
- 6: Detect geotags using the RTK navigation log n_t
- 7: **Return** \mathcal{D}_{small}

M2 processing (two-color palette). The two-color palette covers a wider temperature range and supports the detection of larger anomalies. The M2 branch targets

large defects according to formula (8) and uses a slightly higher Mask-NMS IoU threshold (0.5). Algorithm 3 outlines the steps.

Algorithm 3. Built-in processing of a two-color thermogram (M2)

Input: Thermogram I in a two-color palette, YOLOv11-seg model

Execution: Filtered set of defects \mathcal{D}_{large}

- 1: Normalize the thermal image and run YOLOv11-seg
- 2: Apply confidence filtering and Mask-NMS (IoU threshold 0.5)
- 3: Calculate the R_d for each mask using formula (6); keep only instances with $R_d \geq 1$
- 4: Detect geotags using RTK navigation logs n_t
- 5: Return \mathcal{D}_{large}

Implementation of late merging in thermograms

Domain differences caused by the palette prompt the merging of models. Instead of training a large multi-domain model, a lightweight late-merging ensemble is implemented based on size-routed branches defined in formulas (7)–(9). The detector is applied to both representations, generating \mathcal{D}_t^{M2} and \mathcal{D}_t^{M3} , and the edge level combines them according to the set union rule as follows:

$$\mathcal{D}_t^{ens} \equiv \mathcal{D}_{final} = \mathcal{D}_{small} \cup \mathcal{D}_{large}. \quad (10)$$

This selective integration preserves the branch that best suits the expected thermal contrast regime: M3 provides a set of small defects, while M2 provides a set of large defects. The edge level further suppresses duplicates between frames using RTK georeferencing (Section 3.8). The design is compatible with real-time operation, as it does not require joint training or a more powerful foundation; additional costs are limited to running the detector on two representations.

Georeferencing and deduplication at the edge

The UAV mission enables repeated observations of the same modules in consecutive frames and along overlapping flight paths. Without deduplication, a single defect could be counted multiple times, leading to an excessive number of alerts and reduced operator confidence. Therefore, georeferencing and deduplication using RTK are performed at the edge.

Each detection in \mathcal{D}_t is mapped to the approximate location of the module using the RTK UAV's position, altitude, and camera geometry. The exact projection depends on camera calibration and gimbal orientation; the current implementation uses a simplified mapping sufficient for clustering, with module-level granularity. Additionally, the association is refined using a site layout map.

To calculate distances between geotagged detections, the Haversine distance between latitude/longitude coordinates is used as follows:

$$d = 2R \arcsin \left(\sqrt{\sin^2 \left(\frac{\Delta \varphi}{2} \right) + \cos \varphi_1 \cos \varphi_2 \sin^2 \left(\frac{\Delta \lambda}{2} \right)} \right), \quad (11)$$

where r is the Earth's radius, φ is the latitude, and λ is the longitude.

Detections are grouped using a distance threshold ϵ , derived from the ground sampling distance and the expected RTK error.

Within each cluster, the edge server aggregates detections over a specific time interval, saves representative frame evidence, and associates the cluster with a module identifier or a string based on the plant layout. The result is a set of defect events at the module level, rather than detections at the frame level.

Condition Assessment Incorporating SCADA and Fire Risk Logic

Defect detection becomes operationally useful when it translates into actions at the plant level. KFS includes a decision-making layer that combines UAV detections with electrical and controlled signals to determine operational states and hazardous conditions. The core idea is to express safety-related conditions as

understandable logic that engineers can verify and integrate into SCADA rule engines.

Binary variables are defined as X_1 , X_2 , and X_3 . These variables indicate the presence of a fault in circuits 1–3 of the module group (or the corresponding section of the electrical diagram). The variables X_{11} , X_{12} , and X_{13} indicate whether the bypass diodes in the respective circuits are functioning correctly (1 indicates correct operation). The hazard output Y is defined as:

$$Y = X_1 \overline{X_{11}} + X_2 \overline{X_{12}} + X_3 \overline{X_{13}}, \quad (12)$$

which reflects the engineering intuition that defects combined with faulty bypass diodes increase the risk of overheating. A complete report on the research results is presented in Table 2.

Table 2. $X_1 - X_3$ indicate defects in three circuits; $X_{11} - X_{13}$ indicate that the bypass diode is functioning correctly; $Y = 1$ denotes a hazardous state

#	X_1	X_2	X_3	X_{11}	X_{12}	X_{13}	Y	Status description
1	0	0	0	0	0	0	0	normal operation
2	0	0	0	0	0	1	0	normal operation
3	0	0	0	0	1	0	0	normal operation
4	0	0	0	0	1	1	0	normal operation
5	0	0	0	1	0	0	0	normal operation
6	0	0	0	1	0	1	0	normal operation
7	0	0	0	1	1	0	0	normal operation
8	0	0	0	1	1	1	0	normal operation
9	1	0	0	0	0	0	1	Danger; bypass diode failure
10	0	1	0	0	0	0	1	Danger; bypass diode failure
11	0	0	1	0	0	0	1	Danger; bypass diode failure
12	1	1	0	0	0	0	1	Danger; bypass diode failure
13	1	0	1	0	0	0	1	Danger; failure of the bypass diode
14	0	1	1	0	0	0	1	Danger; bypass diode failure
15	1	1	1	0	0	0	1	Danger; bypass diode failure
16	1	0	0	1	0	0	0	The bypass diode blocks the circuit
17	0	1	0	0	1	0	0	The bypass diode blocks the circuit
18	0	0	1	0	0	1	0	The bypass diode blocks the circuit
19	1	1	1	1	1	1	0	Bypass diodes block the circuits

To enable additional safety sensors (such as smoke or flame detectors), the binary variables X_{21} and X_{22} are added, representing the sensor triggers. The extended hazard logic is formalized as follows:

$$Y = X_1 \overline{X_{11}} + X_2 \overline{X_{12}} + X_3 \overline{X_{13}} + X_1 X_{21} + X_2 X_{22}. \quad (13)$$

The edge server evaluates these expressions in real time and forwards alarm signals to the SCADA/HMI level, where standard procedures (e.g., dispatch, shutdown recommendation) can be executed.

Temperature Context Using the NOCT Approximation

To interpret thermal anomalies, the KFS can optionally estimate the expected module temperature

under nominal conditions using the nominal operating cell temperature (NOCT) approximation. Given an ambient temperature of T_{amb} and solar irradiance G (W/m^2), the expected module temperature is approximated by the formula:

$$T_m = T_{amb} + \frac{NOCT - 20}{800} G. \quad (14)$$

This estimate does not replace thermography; it provides context for whether the observed temperatures are plausible under current insolation and weather conditions, which can help identify instances where reflections or temporary shading are prevalent.

Research and Results

This section briefly describes a practical example of the presented study and provides quantitative results for the proposed defect detection process using KFS. This study focuses on experiments with real data and excludes synthetic experiments. The results are organized into the following categories: (i) detection quality on two thermograms (M2 and M3), (ii) the ensemble averaging effect, (iii) hardware characteristics and a comparison of YOLO models, (iv) practical validation compared to manual inspection, and (v) sensitivity to flight and weather parameters.

Experiment setup

The experiments were conducted in a closed-loop KFS cycle: (i) UAV surveying, (ii) on-board inference modeling, (iii) aggregation at the edge, and (iv) SCADA-related events, as described in Section 3.1. An industrial-grade corporate-class UAV (DJI Matrice 300 RTK) was used for surveying, providing a flight range of up to 15 km. The aircraft is equipped with a gimbal-mounted dual sensor system comprising a thermal infrared camera and a 4K RGB camera. The sensor captures synchronized thermal images and high-resolution contextual RGB frames, each of which is geotagged using the UAV's RTK GNSS with centimeter-level accuracy. Similar UAV-based thermal imaging configurations have been documented for the inspection and quality control of photovoltaic power plants [35]. The gimbal imaging system is controlled by a specialized low-power video processor (Ambarella H2 SoC) [36], which supports 4K 60 FPS video encoding. All image frames were acquired and preprocessed using OpenCV v4.9.0 [37] prior to analysis.

The flight mission was configured using specialized software (e.g., DJI Pilot 2) as a grid mission, ensuring 80% frontal and 70% lateral frame overlap, taking into account the tilt angle of the panels. The UAV control mode was semi-automatic: the operator performed the launch and general visual monitoring of flight safety, while the UAV carried out object search, positioning, and task execution (imaging along a specified route) completely autonomously. Return to base (RTH) was initiated automatically upon completion of the scanning route or when the battery charge reached a critical level (below 20%).

A real-time object detection model ran continuously on board the UAV to identify anomalies in photovoltaic panels from the thermal video stream. Ultralytics v8.3.193 YOLO models [34] were deployed and

implemented using the PyTorch v2.7.0 deep learning library [38]. Prior to deployment, the segmentation model was finely tuned on the selected thermal dataset in Table 1 using stratified five-fold cross-validation; all detection metrics listed correspond to the mean \pm standard deviation across all folds, unless otherwise specified. Training was performed on a workstation equipped with a single NVIDIA RTX 3060 (12 GB) graphics processing unit, using 300 epochs, an input resolution of 640, and a batch size of 16. Inference was performed on an embedded artificial intelligence computing platform integrated into the UAV's payload. In our configuration, the onboard inference engine uses the Qualcomm QCS605 system-on-chip [39], which provides hardware acceleration for neural processing of computer vision tasks.

During flight, the UAV transmitted compact detection messages, as defined in Equation (1), via a secure wireless communication channel (e.g., OcuSync Enterprise at 2.4/5.8 GHz) to a ground server. The mission's effectiveness was demonstrated under real-world conditions involving adverse weather factors and radio interference. In particular, even in strong winds (up to 15 m/s) or in the presence of significant electromagnetic interference from high-voltage power lines – which typically cause delays or temporary loss of video stream transmission to the operator's console – on-board processing of YOLOv11 on the Qualcomm QCS605 processor ensured stability. The system reliably cached segmentation results and full frames locally for subsequent auditing. This ensured uninterrupted task execution and the preservation of critical defect data regardless of the quality and stability of the radio link.

A ground server running on the NVIDIA Jetson Orin platform [40] collected incoming detection data and performed spatial clustering to eliminate duplicates (Section 3.8). The clustering radius was set to account for the positioning uncertainty of the UAV's RTK navigation; specifically, it was chosen to correspond to an accuracy of ~ 2 cm, which can be achieved using inexpensive dual-frequency RTK receivers such as the u-blox ZED-F9P [41]. To maximize throughput and minimize latency, the YOLO inference model was optimized on a Jetson edge server using NVIDIA TensorRT v10.9.0 [42]. Finally, the consolidated and deduplicated detection events were transmitted to the SCADA system in real time, enabling automatic monitoring and the generation of alerts for the operator.

Defect detection performance is evaluated using standard detection metrics [43]. For a given class, precision

and recall are used. Detection quality is summarized using average precision (AP) and mean average precision (mAP). This study reports the average precision (mAP@0.5), i.e., AP, calculated at an IoU threshold of 0.5 and averaged across classes. For field validation against manual monitoring, the root mean square error (RMSE) between automatic and manual defect counts is reported as defects per verified photovoltaic circuit.

Performance characteristics are determined by inference throughput (frames per second), end-to-end processing time, and sensitivity to flight and

environmental parameters (altitude, speed, time of day, weather). In Section 4.5, the accuracy metric is used to represent the proportion of correct defect detections among all detected defects based on selected IoU/confidence thresholds.

Quality as detected on two reference thermal images

Table 3 summarizes the results by class for YOLOv11-seg on M2 and M3, presented as the mean \pm standard deviation based on stratified five-fold cross-validation.

Table 3. Comprehensive comparison of YOLOv11-seg productivity on the M2 (two-color palette) and M3 (three-color palette) thermograms, presented as mean \pm standard deviation based on stratified 5-fold cross-validation

Class	M2 thermogram (2 colors)			M3 thermal image (3 colors)		
	mAP@0.5	Accuracy	Reproduction	mAP@0.5	Accuracy	Reproduction
Crack	0.93 \pm 0.04	0.91 \pm 0.05	0.89 \pm 0.06	0.87 \pm 0.07	0.86 \pm 0.07	0.84 \pm 0.08
Pollution	0.90 \pm 0.06	0.88 \pm 0.07	0.86 \pm 0.08	0.90 \pm 0.06	0.88 \pm 0.07	0.86 \pm 0.08
Stratification	0.89 \pm 0.02	0.87 \pm 0.02	0.85 \pm 0.03	0.93 \pm 0.02	0.90 \pm 0.02	0.88 \pm 0.03
Macro-average value	0.91 \pm 0.03	0.89 \pm 0.03	0.87 \pm 0.04	0.90 \pm 0.04	0.88 \pm 0.04	0.86 \pm 0.05

On the two-color palette (M2), the model achieves a macro-mAP@0.5 of 0.91 \pm 0.03, a precision of 0.89 \pm 0.03, and a reproducibility of 0.87 \pm 0.04. On the three-color palette (M3), the macro-mAP@0.5 is 0.90 \pm 0.04 with accuracy of 0.88 \pm 0.04 and reproducibility of 0.86 \pm 0.05. These average values are close, but the behavior across classes differs. Crack

detection is stronger on M2 (mAP@0.5 0.93 \pm 0.04) than on M3 (0.87 \pm 0.07), whereas delamination detection is strongest on M3 (0.93 \pm 0.02). This pattern confirms the practical observation that thermogram palettes act as separate domains and that one representation may emphasize certain defect features at the expense of others (Fig. 3).





Fig. 3. Detection quality assessment (average mAP@0.5 across five runs) for YOLOv11-seg on different thermograms: a) the two-color palette (M2) is excellent for crack detection; b) the three-color palette (M3) demonstrates excellent results in delamination detection

Impact of late merging set and runtime

The late merging mode improves reliability without requiring a more powerful network. Table 4 shows the mAP@0.5 metrics before and after merging (mean \pm standard deviation across runs).

mAP@0.5 benefits the most, improving from 0.93 \pm 0.04 to 0.96 \pm 0.03. This is of practical importance because crack-like patterns can be barely visible and can be confused with thermal noise or reflections. Delamination improves from 0.93 \pm 0.02 to 0.95 \pm 0.02, while contamination remains stable. Figure 4 shows these changes.

Table 4. Effect of merging on mAP@0.5 (before and after), presented as mean ± standard deviation across stratified 5-fold cross-validation

Class	Before	After Merger
Crack	0.93 ± 0.04	0.96 ± 0.03
Soiling	0.90 ± 0.06	0.90 ± 0.06
Delamination	0.93 ± 0.02	0.95 ± 0.02

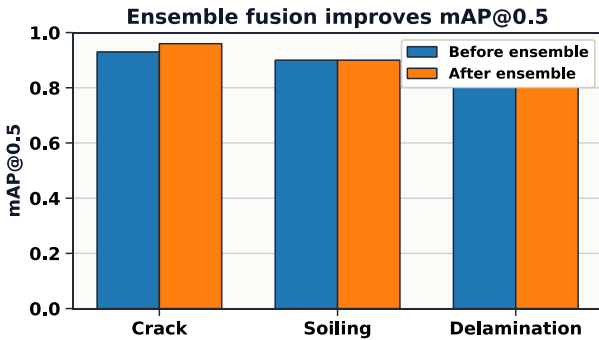


Fig. 4. Improvement in the mAP@0.5 after late merging of M2 and M3 detections

The pipeline approach also improves execution time. Thanks to early detection merging and the avoidance of repeated post-processing, the processing

Table 5. Comparison of typical edge hardware for embedded inference. The mAP@0.5 listed correspond to the deployment configuration (processor and numerical precision), not just the hardware

Platform	Output Accuracy	Throughput (FPS)	mAP@0.5	Accuracy
NVIDIA Jetson Orin Nano	FP16 (TensorRT)	100	0.95	0.93
Ambarella H2	INT8 (compressed)	60	0.70	0.75
Qualcomm QCS605	INT8 (SNPE)	80	0.85	0.82

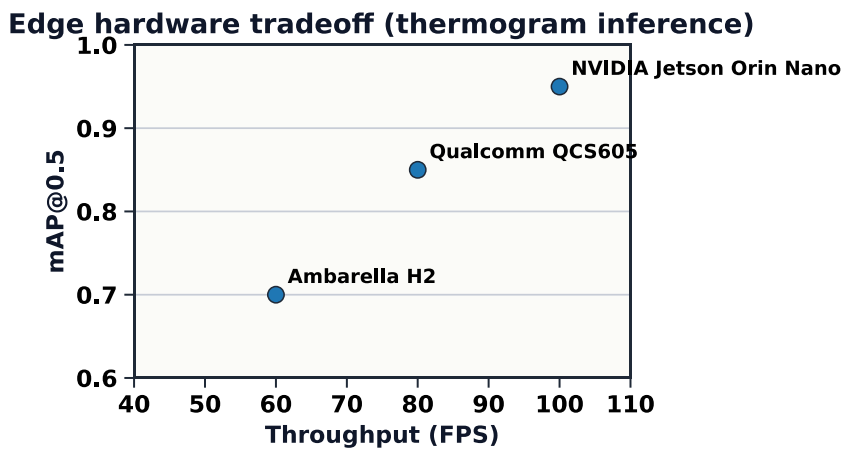


Fig. 5. Trade-off between throughput and detection quality for typical embedded platforms

To further justify the choice of YOLOv11, several generations of YOLO were trained and evaluated

time for each frame was reduced from 4.235 s to 2.858 s. Although this latency depends on the hardware and implementation, its reduction demonstrates that careful pipeline design can improve both accuracy and performance, which is critical for the operation of the KFS.

Comparison of Edge Hardware and the YOLO Model

Implementing deep learning on UAVs imposes strict constraints on computational power, weight, and power consumption. Table 5 provides an overview of typical embedded platforms and highlights a practical trade-off: throughput determines the maximum possible flight speed and frame rate (see Fig. 5).

It is important to note that detection quality (mAP/accuracy) is a property of the deployed model and numerical precision, not just the hardware. Therefore, the values of mAP@0.5 in Table 5 correspond to a deployment configuration specific to the platform (e.g., FP16 versus INT8 quantization and, if necessary, additional compression to meet memory/latency constraints).

under a unified protocol. Table 6 shows accuracy, reproducibility, and mAP@0.5, as well as training time.

YOLOv11 demonstrates the highest overall productivity (accuracy 0.96, reproducibility 0.95, mAP@0.5 0.93), although training time is longer than for some smaller

models. The results indicate that for embedded inference, YOLOv11 provides a favorable balance between accuracy and manageable deployment complexity.

Table 6. Quantitative comparison of YOLO model generations (segmentation). Training time is measured over 300 epochs on a single NVIDIA RTX 3060 (12 GB) GPU with an input resolution of 640 and a batch size of 16

Model	Accuracy	Recall	mAP@0.5	Training time (min)
YOLOv5	0.92	0.90	0.86	55
YOLOv8	0.91	0.88	0.83	45
YOLOv9	0.94	0.80	0.90	59
YOLOv10	0.94	0.86	0.90	132
YOLOv11	0.96	0.95	0.93	106

On-site data verification versus manual monitoring

The KFS, designed for plant operation, must match the results of a human-conducted field inspection. The performance of the AIC is verified through manual monitoring at two photovoltaic installations: a rooftop and a ground-mounted site. The evaluation compares the automatic defect count after edge-level deduplication with the manual count obtained from expert inspection. The test set consisted of approximately 2,500 photovoltaic modules across two sites and included 145 expert-verified defect cases; the agreement was quantified using an RMSE of 0.71 defects per inspected photovoltaic string.

Beyond the overall error, field tests highlight the role of KFS components that are not visible in offline metrics. The built-in inference alone yields repeated detections of the same defect in adjacent frames; RTK-based clustering and edge-level deduplication reduce this redundancy and generate module-level events suitable for reporting. Similarly, when abnormal heating is detected, the logic supporting SCADA can elevate the event's priority, especially if the bypass diode status indicates an elevated risk. These elements are critical for operational implementation, as operators require stable counts, clear module identifiers, and actionable alerts, rather than a stream of raw detections.

Sensitivity to Flight and Environmental Parameters

Thermal imaging is limited by physical factors. Therefore, flight altitude, speed, time of day, and weather conditions are analyzed in terms of their impact

on detection quality. The results quantitatively evaluate the operating parameters of the proposed thermal imaging system.


Altitude provides an optimal balance between resolution and coverage (Table 7). At an altitude of 5 m, detection accuracy reaches 98% with a reproducibility of 96% (Fig. 6, a), but coverage is limited, and mission duration increases. At a height of 10 m, accuracy remains high (93%, reproducibility – 90%), and coverage increases approximately 4–5 times, making it acceptable for operation. At a height of 15 m, accuracy decreases to 84% (reproducibility 79%), indicating insufficient spatial resolution for detecting minor defects.

Flight speed affects motion blur and temporal redundancy (Table 7, Fig. 6, b). At speeds of 3–7 m/s, accuracy remains above 91% and reproducibility above 88%, indicating that the system can handle moderate speeds when the camera shutter and stabilization are sufficient. At a speed of 10 m/s, accuracy drops to 85%, and reproducibility to 82%, which may be unacceptable for safety-critical monitoring.

Environmental conditions are also important (Table 8). Surveys conducted during midday (12:00–14:00) yield the best results (96% accuracy, 94% reproducibility), corresponding to stronger radiation and thermal contrast (Fig. 6, c). Cloudy conditions improve performance (96% accuracy, 94% reproducibility) compared to clear skies (92% accuracy, 89% reproducibility), likely because diffuse lighting reduces specular reflections (Fig. 6, d). These results are consistent with previous observations that reflections may dominate thermograms under clear conditions and that careful flight planning is necessary [7, 11].

Table 7. Sensitivity to height and speed (accuracy/reproducibility)

Height			Speed		
m	Accuracy (%)	Reproducibility (%)	m/s	Accuracy (%)	Repeatability (%)
5	98	96	3	94	92
10	93	90	5	93	90
15	84	79	7	91	88
			10	85	82

**Fig. 6.** Sensitivity of detection to changes in operating conditions and the environment. The results show that an increase in a) flight altitude and b) speed degrades detection quality, requiring strict adherence to flight profiles. Environmental factors also affect performance: c) the time from 12:00 to 14:00 and d) cloudy weather provide optimal accuracy and reproducibility**Table 8.** Sensitivity to time of day and weather (accuracy/reproducibility)

Time of day			Weather		
Window	Accuracy (%)	Reproducibility (%)	Conditions	Accuracy (%)	Reproducibility (%)
8:00–10:00	93	90	Clear	92	89
12:00–14:00	96	94			
5:00 PM–7:00 PM	92	89	Cloudy	96	94

Quantitative comparison with the current state of the art

The obtained results are compared with typical approaches described in the literature. A direct comparison

is imperfect, as the datasets, defect definitions, and evaluation protocols differ. However, Table 9 provides a quantitative basis and explains how the proposed KFS differs in scope, which is further illustrated in Fig. 7.

Table 9. Quantitative comparison with published approaches (according to the respective authors)

Work	Modality	Model / approach	Task	Reported metric
Dotenko et al. [25]	IR	Statistical tests + vision	Defect classification	F1 = 93.88%
Di Tommaso et al. [14]	IR	Multi-stage YOLOv3	Hotspot detection	AP@0.5: 66.9%
Xi et al. [18]	IR	ST-YOLO	Detection of faults in photovoltaic systems	mAP@0.5 = 96.6%
This work	IR + RGB	YOLOv11-seg ensemble + KFS	Defect segmentation + decision making	mAP@0.5s up to 96% (cracks)

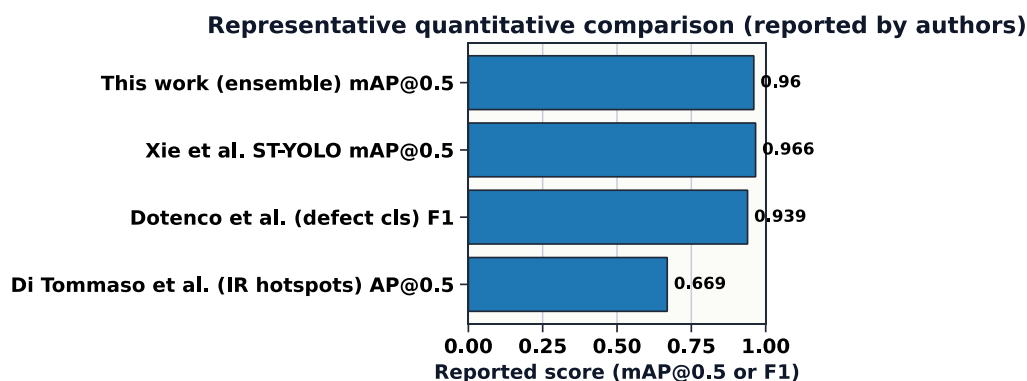


Fig. 7. Illustrative quantitative comparison (indicators shown). Values are not strictly comparable across datasets; the figure contextualizes the magnitude of the results

The literature notes that a high mAP can be achieved under favorable conditions, for example, using ST-YOLO on infrared images [18]. The proposed ensemble achieves a comparable mAP@0.5 for the crack class (0.96), incorporating system-level components necessary for the facility's operation: georeferencing, deduplication, and SCADA-aware decision-making logic. Multi-stage approaches, such as those by Di Tommaso et al. [14], report a significantly lower AP@0.5 anomaly detection rate on hotspots (66.9%), illustrating that anomaly detection can be more challenging than module detection and that performance is strongly influenced by the dataset and defect type. Dotenko et al. [25] reported high efficiency in defect classification using classical vision and statistical tests, which remains relevant as a baseline metric and demonstrates that robust preprocessing is crucial in infrared aerial photography.

Discussion

The presented results confirm two interrelated theses. First, a modern detector with a segmentation function can achieve high accuracy in defect localization on UAV thermal images even with a small training dataset, provided that post-processing requirements and operational constraints are met. Second, and more importantly, the operational utility of UAV inspection depends on a cyber-physical architecture that links detection results with georeferenced, deduplicated events and plant monitoring.

From a computer vision perspective, palette-specific results illustrate a subtle but important phenomenon: the representation of the thermogram is part of the model's input domain. The two-color palette (M2) highlights

certain gradients and enables more effective crack detection, while the three-color palette (M3) enhances delamination detection. This is consistent with broader observations in photothermal thermography that the choice of acquisition and visualization methods affects defect contrast and, consequently, their detection [4, 6]. The late-merging ensemble improves reliability by leveraging this complementarity. The improvement for cracks (from 0.93 to 0.96 mAP@0.5) is particularly important, as cracks are often early-stage defects whose thermal signals can be weak. The reduction in runtime (from 4.235 s to 2.858 s) further underscores that pipeline design can yield tangible benefits.

From a systems perspective, the peripheral layer is not a luxury. Without RTK-based georeferencing and deduplication, raw UAV detections would yield inflated metrics and unstable alarms. Therefore, the architecture treats georeferencing as a first-class signal, consistent with field studies that highlight the practical difficulty of tracing anomalies back to modules and chains [7, 11]. The edge layer also provides a natural integration point for SCADA signals, such as bypass diode status. The hazard Boolean logic, formalized by equations (12)–(13), is intentionally interpretable: plant engineers can verify, adjust, and integrate it into existing rule mechanisms. This contrasts with “black box” alarm models, which can be difficult to certify for safety monitoring purposes.

Sensitivity analysis offers another argument in favor of KFS framing. The sharp drop in productivity at a height of 15 m and a speed of 10 m/s indicates that the detector's productivity is limited by physics and optics, not just the model's power. Similarly, better productivity in cloudy conditions suggests that reflections may be the primary source of false alarms.

These observations prompt the inclusion of reflection-aware planning (Section 3.3) and the use of contextual temperature models (Section 3.10) as part of the overall system. In other words, the “best model” alone is insufficient; the system must actively generate its own data through flight planning and quality control.

A key limitation of this evaluation is the small and unbalanced labeled dataset for minority classes. Even with stratified five-fold cross-validation, the crack and contamination classes contain only 20 and 10 labeled images, respectively, leading to greater variance (Tables 3–4). Therefore, the reported mean values should be interpreted as feasibility indicators rather than definitive benchmarks. Expanding labeled datasets and adopting standardized public benchmarks for UAV thermography with consistent metadata would significantly improve statistical reliability and reproducibility.

The proposed approach also has drawbacks and trade-offs. The architecture complicates deployment: a peripheral server must be maintained on-site, and reliable communication between the UAV and the periphery must be ensured. The approach relies on RTK positioning and continuous camera calibration; failure of these subsystems can degrade georeferencing and reduce confidence in module-level mapping. Although transfer learning and augmentation mitigate the need for retraining, generalization to new defect types, new thermal cameras, or different palettes remains a challenge. Future work should explore palette-independent representations or temperature-calibrated models that operate directly on radiometric data rather than on reconstructed palettes.

Several limitations point to open research directions. First, thermography remains sensitive to environmental conditions; a key challenge is to incorporate radiation intensity, wind, and ambient temperature into training models in a principled manner, potentially using physically grounded characteristics or uncertainty-aware inference. Second, current SCADA logic is based on explicit Boolean rules; hybrid approaches combining learned risk assessment with interpretable rules can increase flexibility while maintaining auditability. Third, a simplified projection is currently used during the georeferencing stage; improvements in geometric modeling (e.g., through camera calibration and 3D facility models) could reduce association errors in dense installations. Finally, there is a broader issue

of comparative analysis: the photovoltaic installation community would benefit from standardized, publicly available UAV thermography datasets with consistent labels and metadata (flight altitude, solar irradiance, camera parameters), which would enable reproducible comparisons between systems.

Conclusions

This paper proposes an architecture for cyber-physical systems designed to monitor photovoltaic power plants, which successfully bridges the gap between raw computer vision data and practical operational maintenance. By integrating real-time UAV-based YOLOv11 segmentation with an intelligent edge computing layer for RTK-supported geotagging and duplicate suppression, the system effectively transforms fragmented output results into trackable defect events at the module level. Experimental results confirm the superior performance of the proposed late fusion ensemble, which significantly improved crack detection to 0,96mAP@0.5 and delamination detection to 0,95mAP@0.5, while reducing the processing latency per frame to 2.858 s. Field validation further demonstrated consistency with manual inspection records, achieving an RMSE of 0.71 defects per inspected PV row and identifying a flight altitude of 10 m as the most balanced operating point with 93% accuracy and 90% reproducibility. The unique integration of SCADA-aware interpretive logic enables real-time fire risk assessment based on the status of the bypass diode. Despite these successes, notable limitations of this study include sensitivity to specular reflections under clear sky conditions, reduced productivity at flight speeds exceeding 7 m/s, and statistically unstable estimates for a minority of defect classes due to the limited amount of labeled data. The dependence on precise RTK calibration for module association also poses a challenge for deployment in dense array configurations. Overall, the proposed architecture offers a scalable, safety-oriented solution for automating the lifecycle management of large-scale solar assets.

Future research will focus on palette-independent radiometric models to reduce the impact of ambient noise, larger-scale benchmark-based assessments to reduce uncertainty, and predictive maintenance algorithms utilizing historical SCADA trends.

Conflict of interest

The authors declare that they have no conflicts of interest, including financial, personal, authorship, or any other conflicts that could influence the research or the results published in this article.

Funding

This research was supported by the Ministry of Education and Science of Ukraine and funded by the European Union's external aid instrument as part of Ukraine's commitments under the EU Framework Program for Research and Innovation "Horizon 2020". The work was carried out as part of the research project "Intelligent Defect Detection System for Green Energy Facilities Using UAVs", state registration number 0124U004665 (2024–2026).

Data availability

Data will be provided upon reasonable request. The base datasets (STHS-277 and PVF-10) are publicly

available. Analysis scripts may be provided by the authors via email upon reasonable request. The proprietary dataset obtained during field validation is not publicly available due to commercial confidentiality.

Use of artificial intelligence

The authors confirm that they did not use artificial intelligence technology in the creation of this paper.

Acknowledgments

The authors express their sincere gratitude to the Ministry of Education and Science of Ukraine for its support. This research was funded by the European Union's external aid instrument under the "Horizon 2020" research and innovation program. The work was carried out as part of the project "Intelligent Defect Detection System for Green Energy Facilities Using UAVs" (state registration number 0124U004665).

All authors have read and approved the published version of this manuscript.

References

1. Sun, L., Sun, Y. (2022), "Photovoltaic power forecasting based on artificial neural network and ultraviolet index", *International Journal of Computing*, Vol. 21, No. 2, 1531158 p. DOI: <https://doi.org/10.47839/ijc.21.2.2583>
2. Sinha, A., Sulas-Kern, D. B., Owen-Bellini, M., Spinella, L., Uličná, Š., Ayala, R., Murphy, R., Walker, C., Hacke, P. (2024), "A review of degradation and reliability analysis of a solar PV module", *IEEE Access*, Vol. 12, pp. 185036-185056. DOI: <https://doi.org/10.1109/ACCESS.2024.3432394>
3. Lynnyk, R., Vysotska, V., Matseliukh, Y., Burov, Y., Demkiv, L., Sachenko, A., Shylynska, I., Yevseyeva, I., Bihun, O. (2020), "DDoS attacks analysis based on machine learning in challenges of global changes", *Proceedings of the 2nd International Workshop on Modern Machine Learning Technologies and Data Science (MoMLT+DS 2020)*, Volume 2631, CEUR.org, Aachen, pp. 159-171, available at: <https://ceur-ws.org/Vol-2631/paper12.pdf> (last accessed 16.01.2025).
4. Aghaei, M., Kolahi, M., Nedaei, A., Venkatesh, N. S., Esmailifar, S. M. (2025), "Autonomous intelligent monitoring of photovoltaic systems: An in-depth multidisciplinary review", *Progress in Photovoltaics: Research and Applications*, Vol. 33, No. 3, pp. 381-409. DOI: <https://doi.org/10.1002/pip.3859>
5. Khatri, A., Khadka, S., Lamichhane, N., Shrestha, R. (2025), "A comprehensive review of infrared thermography and deep learning applications for solar photovoltaic systems", *Infrared Physics & Technology*, Vol. 148, 105878 p. DOI: <https://doi.org/10.1016/j.infrared.2025.105878>
6. Gallardo-Saavedra, S., Hernández-Callejo, L., Alonso-García, M. d. C., Muñoz-Cruzado-Alba, J., Ballestín-Fuertes, J. (2020), "Infrared thermography for the detection and characterization of photovoltaic defects: Comparison between illumination and dark conditions", *Sensors*, Vol. 20, No. 16, 4395 p. DOI: <https://doi.org/10.3390/s20164395>
7. Prasshanth, C. V., Deepak, B. B. V. L., Vijayakumar, N., Prabadevi, B., Alhussen, F. A. (2025), "Fault detection in photovoltaic systems using unmanned aerial vehicle-captured images and rough set theory", *Solar Energy*, Vol. 290, 113348 p. DOI: <https://doi.org/10.1016/j.solener.2025.113348>
8. Moctezuma, D., Valentín-Coronado, L. M. (2025), "Deep-learning-based electrical fault detection in photovoltaic modules through aerial infrared imaging: Addressing data complexity", *Computación y Sistemas*, Vol. 29, No. 1, pp. 65-75. DOI: <https://doi.org/10.13053/cys-29-1-5531>

9. Svystun, S., Scislo, L., Pawlik, M., Melnychenko, O., Radiuk, P., Savenko, O., Sachenko, A. (2025), "DyTAM: Accelerating wind turbine inspections with dynamic UAV trajectory adaptation", *Energies*, Vol. 18, No. 7, 1823 p. DOI: <https://doi.org/10.3390/en18071823>
10. Rouibah, N., El Hammoumi, A., Bouttout, A., Haddad, S., Oukaci, S., Limam, A., Benghanem, M. (2025), "Smart monitoring of photovoltaic energy systems: An IoT-based prototype approach", *Scientific African*, Vol. 30, e02973 p. DOI: <https://doi.org/10.1016/j.sciaf.2025.e02973>
11. Rodriguez-Vazquez, J., Prieto-Centeno, I., Fernandez-Cortizas, M., Pérez-Saura, D., Molina, M., Campoy, P. (2024), "Real-time object detection for autonomous solar farm inspection via UAVs", *Sensors*, Vol. 24, No. 3, 777 p. DOI: <https://doi.org/10.3390/s24030777>
12. Noura, H., Sun, Y., Taghezouit, B., Dairi, A. (2024), "Explainable artificial intelligence of tree-based algorithms for fault detection and diagnosis in grid-connected photovoltaic systems". *Engineering Applications of Artificial Intelligence*, Vol. 139, 109503 p. DOI: <https://doi.org/10.1016/j.engappai.2024.109503>
13. Bodyanskiy, Y., Deineko, A., Skorik, V., Brodetskiy, F. (2022), "Deep neural network with adaptive parametric rectified linear units and its fast learning", *International Journal of Computing*, Vol. 21, No. 1, pp. 11-18. DOI: <https://doi.org/10.47839/ijc.21.1.2512>
14. Di Tommaso, A., Betti, A., Fontanelli, G., Michelozzi, B. (2022), "A multi-stage model based on YOLOv3 for defect detection in PV panels based on IR and visible imaging by unmanned aerial vehicle", *Renewable Energy*, Vol. 193, pp. 941-962. DOI: <https://doi.org/10.1016/j.renene.2022.04.046>
15. Melnychenko, O., Scislo, L., Savenko, O., Sachenko, A., Radiuk, P. (2024), "Intelligent integrated system for fruit detection using multi-UAV imaging and deep learning", *Sensors*, Vol. 24, No. 6, 1913 p. DOI: <https://doi.org/10.3390/s24061913>
16. Aljafari, B., Satpathy, P. R., Thanikanti, S. B., Nwulu, N. (2024), "Supervised classification and fault detection in grid-connected PV systems using 1D-CNN: Simulation and real-time validation", *Energy Reports*, Vol. 12, pp. 2156-2178. DOI: <https://doi.org/10.1016/j.egy.2024.08.008>
17. Setiawan, E. A., Fathurrahman, M. (2025), "Exploring insights on deep learning-based photovoltaic fault detection for monofacial and bifacial modules using thermography", *International Journal of Cognitive Computing in Engineering*, Vol. 6, pp. 495-507. DOI: <https://doi.org/10.1016/j.ijcce.2025.04.001>
18. Xie, H., Yuan, B., Hu, C., Gao, Y., Wang, F., Wang, C., Wang, Y., Chu, P. (2024), "ST-YOLO: A defect detection method for photovoltaic modules based on infrared thermal imaging and machine vision technology", *PLOS ONE*, Vol. 19, No. 12, pp. 1-19. DOI: <https://doi.org/10.1371/journal.pone.0310742>
19. Carni, D. L., Grimaldi, D., Lamonaca, F., Nigro, L., Sciammarella, P. F. (2017), "From distributed measurement systems to cyber-physical systems: A design approach", *International Journal of Computing*, Vol. 16, No. 2, pp. 66-73. DOI: <https://doi.org/10.47839/ijc.16.2.882>
20. Boucif, O. H., Toral-Cruz, H. (2025), "Artificial intelligence of things for solar energy monitoring and control", *Applied Sciences*, Vol. 15, No. 11, 6019 p. DOI: <https://doi.org/10.3390/app15116019>
21. Lysyi, A., Sachenko, A., Radiuk, P., Lysyi, M., Melnychenko, O., Ishchuk, O., Savenko, O. (2025), "Enhanced fire hazard detection in solar power plants: An integrated UAV, AI, and SCADA-based approach", *Radioelectronic and Computer Systems*, Vol. 2025, No. 2, pp. 99-117. DOI: <https://doi.org/10.32620/reks.2025.2.06>
22. Tang, W., Xu, J., Wang, X., Li, H., Chen, J., Zhang, Y. (2022), "Deep learning-based detection of linear defects in large-scale photovoltaic plants using an edge-cloud architecture", *Solar Energy*, Vol. 237, pp. 144-161. DOI: <https://doi.org/10.1016/j.solener.2021.11.016>
23. Ferlito, S., Ippolito, S., Santagata, C., Schiattarella, P., Di Francia, G. (2024), "A study on an IoT-based SCADA system for photovoltaic utility plants", *Electronics*, Vol. 13, No. 11, 2065 p. DOI: <https://doi.org/10.3390/electronics13112065>
24. Masita, K., Hasan, A., Shongwe, T., Abu Hilal, H. (2025), "Deep learning in defects detection of PV modules: A review", *Solar Energy Advances*, Vol. 5, 100090 p. DOI: <https://doi.org/10.1016/j.seja.2025.100090>
25. Dotenco, S., Dalsass, M., Winkler, L., Würzner, T., Brabec, C., Maier, A., Gallwitz, F. (2016), "Automatic detection and analysis of photovoltaic modules in aerial infrared imagery", *Proceedings of the 2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, IEEE, New York, NY, USA, pp. 1-9. DOI: <https://doi.org/10.1109/WACV.2016.7477658>
26. Ebied, M. A., Munshi, A., Alhuzali, S. A., El-sotouhy, M. M., Shehta, A. I., Elborlsy, M. S. (2025), "Advanced deep learning modeling to enhance detection of defective photovoltaic cells in electroluminescence images", *Scientific Reports*, Vol. 15, 31640 p. DOI: <https://doi.org/10.1038/s41598-025-14478-y>
27. Barraz, Z., Sebari, I., Lamrini, N., Ait El Kadi, K., Ait Abdelmoula, I. (2025), "Fast and automatic solar module geo-labeling for optimized large-scale photovoltaic systems inspection from UAV thermal imagery using deep learning segmentation", *Cleaner Engineering and Technology*, Vol. 28, 101048 p. DOI: <https://doi.org/10.1016/j.clet.2025.101048>
28. Abdelsattar, M., AbdelMoety, A., Emad-Eldeen, A. (2025), "ResNet-based image processing approach for precise detection of cracks in photovoltaic panels", *Scientific Reports*, Vol. 15, 24356 p. DOI: <https://doi.org/10.1038/s41598-025-09101-z>

29. Svystun, S., Melnychenko, O., Radiuk, P., Savenko, O., Sachenko, A., Lysyi, A. (2024), "Thermal and RGB images work better together in wind turbine damage detection", *International Journal of Computing*, Vol. 23, No. 4, pp. 526-535. DOI: <https://doi.org/10.47839/ijc.23.4.3752>
30. Zheng, Y., Shcherbakova, G., Rusyn, B., Sachenko, A., Volkova, N., Kliushnikov, I., Antoshchuk, S. (2025), "Wavelet transform cluster analysis of UAV images for sustainable development of smart regions due to inspecting transport infrastructure", *Sustainability*, Vol. 17, No. 3, 927 p. DOI: <https://doi.org/10.3390/su17030927>
31. Alfaro-Mejía, E., Loaiza-Correa, H., Franco-Mejía, E., Restrepo-Girón, A. D., Nope-Rodríguez, S. E. (2019), "Dataset for recognition of snail trails and hot spot failures in monocrystalline Si solar panels", *Data in Brief*, Vol. 26, 104441 p. DOI: <https://doi.org/10.1016/j.dib.2019.104441>
32. Wang, B., Chen, Q., Wang, M., Chen, Y., Zhang, Z., Liu, X., Gao, W., Zhang, Y., Zhang, H. (2024), "PVF-10: A high-resolution unmanned aerial vehicle thermal infrared image dataset for fine-grained photovoltaic fault classification", *Applied Energy*, Vol. 376, 124187 p. DOI: <https://doi.org/10.1016/j.apenergy.2024.124187>
33. CVAT.ai Corporation, (2024), "CVAT: Computer vision annotation tool (Version 2.16.1)", DOI: <https://doi.org/10.5281/zenodo.12771595>
34. Jocher, G., Qiu, J., Chaurasia, A. (2025), "Ultralytics YOLO", available at: <https://github.com/ultralytics/ultralytics> (last accessed 16.01.2026).
35. European Commission, (2018), "Standards for the assessment of the environmental performance of photovoltaic modules, power conversion equipment and photovoltaic systems", Publications Office of the European Union, Luxembourg, available at: 10.2760/89830
36. Ambarella, Inc. (2016), "Ambarella introduces low power 4K ultra HD SoCs for sports and flying cameras", available at: <https://www.ambarella.com/news/ambarella-introduces-low-power-4k-ultra-hd-socs-for-sports-and-flying-cameras/> (last accessed 16.01.2026).
37. Bradski, G. (2000), "The OpenCV library", *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, Vol. 25, No. 11, pp. 120-123.
38. Paszke, A. et al. (2019), "PyTorch: An imperative style, high-performance deep learning library", *Advances in Neural Information Processing Systems, Volume 32*, Curran Associates, Inc., pp. 8024-8035, available at: https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf (last accessed 16.01.2026).
39. Qualcomm Technologies, Inc. (2020), "Qualcomm QCS603/QCS605 SoCs for IoT: Product brief", available at: https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/qcs603.605-socs-product-brief_87-pg764-1-c.pdf (last accessed 16.01.2026).
40. NVIDIA Corporation, (2024), "Jetson AGX Orin: Technical brief", available at: <https://www.nvidia.com/en-us/lp/embedded-computing/robotics-edge-ai-tech-brief/> (last accessed 16.01.2026).
41. Robustelli, U., Cutugno, M., Pugliano, G. (2023), "Low-cost GNSS and PPP-RTK: Investigating the capabilities of the u-blox ZED-F9P module", *Sensors*, Vol. 23, No. 13, 6074 p. DOI: <https://doi.org/10.3390/s23136074>
42. NVIDIA Corporation, (2024), "NVIDIA TensorRT documentation", available at: <https://docs.nvidia.com/deeplearning/tensorrt/10.9.0/getting-started/release-notes.html> (last accessed 16.01.2026).
43. Rainio, O., Teuvo, J., Klén, R. (2024), "Evaluation metrics and statistical tests for machine learning", *Scientific Reports*, Vol. 14, No. 1, 6086 p. DOI: <https://doi.org/10.1038/s41598-024-56706-x>

Received (Надійшла) 18.01.2026

Accepted for publication (Прийнята до друку) 18.02.2026

Publication date (Дата публікації) 30.03.2026

Відомості про авторів / About the Authors

Саченко Анатолій Олексійович – доктор технічних наук, професор, Західноукраїнський національний університет, директор Науково-дослідного інституту інтелектуальних комп'ютерних систем; Тернопіль, Україна; Радомський університет імені Казімежа Пуласького; Радом, Польща;

Anatoliy Sachenko – Doctor of Technical Sciences, Professor, West Ukrainian National University, Director of the Research Institute for Intelligent Computer Systems; Ternopil, Ukraine; Kazimierz Pulaski University of Radom; Radom, Poland;

e-mail: as@wunu.ua

ORCID ID: <https://orcid.org/0000-0002-0907-3682>

Scopus ID: <https://www.scopus.com/authid/detail.uri?authorId=35518445600>

Радюк Павло Михайлович – доктор філософії, Хмельницький національний університет, доцент кафедри комп'ютерних наук; Хмельницький, Україна;

Pavlo Radiuk – PhD in Computer Science, Khmelnytskyi National University, Associate Professor at the Department of Computer Science; Khmelnytskyi, Ukraine;

e-mail: radiukp@khnmu.edu.ua

ORCID ID: <https://orcid.org/0000-0003-3609-112X>

Scopus ID: <https://www.scopus.com/authid/detail.uri?authorId=57216894492>

Лисий Микола Іванович – доктор технічних наук, професор, Національна академія Державної прикордонної служби України імені Богдана Хмельницького; Хмельницький, Україна;

Mykola Lysyi – Doctor of Technical Sciences, Professor, Bohdan Khmelnytskyi National Academy of the State Border Guard Service of Ukraine; Khmelnytskyi, Ukraine;

e-mail: lisiy3152@ukr.net

ORCID ID: <https://orcid.org/0000-0002-9858-706X>

Scopus ID: <https://www.scopus.com/authid/detail.uri?authorId=59482244100>

Мельниченко Олександр Вікторович – доктор філософії, Хмельницький національний університет, старший викладач кафедри комп'ютерної інженерії та інформаційних систем; Хмельницький, Україна;

Oleksandr Melnychenko – PhD in Computer Science, Khmelnytskyi National University, Senior Lecturer at the Department of Computer Engineering and Information Systems; Khmelnytskyi, Ukraine;

e-mail: melnychenko@khnmu.edu.ua

ORCID ID: <https://orcid.org/0000-0001-8565-7092>

Scopus ID: <https://www.scopus.com/authid/detail.uri?authorId=45961382900>

Заставний Олег Михайлович – кандидат технічних наук, Західноукраїнський національний університет, старший викладач кафедри спеціалізованих комп'ютерних систем; Тернопіль, Україна;

Oleg Zastavnyi – Candidate of Technical Sciences, West Ukrainian National University, Senior Lecturer at the Department of Specialized Computer Systems; Ternopil, Ukraine;

e-mail: o.zastavnyi@wunu.edu.ua

ORCID ID: <https://orcid.org/0000-0001-8630-8791>

Scopus ID: <https://www.scopus.com/authid/detail.uri?authorId=8366871500>

АРХІТЕКТУРА КІБЕРФІЗИЧНИХ СИСТЕМ ДЛЯ ВИЯВЛЕННЯ ДЕФЕКТІВ У ФОТОЕЛЕКТРИЧНИХ МОДУЛЯХ НА ОСНОВІ ПІЗЬНОГО ЗЛИТТЯ ДАНИХ ІЗ БПЛА

Предмет дослідження – архітектурне вдосконалення систем інспектування великих сонячних електростанцій на основі БПЛА. **Метою дослідження** є поліпшення операційної корисності інспектування фотоелектричних модулів на основі БПЛА через розроблення кіберфізичної архітектури, яка інтегрує вбудоване глибоке навчання, периферійні вузли, хмарну аналітику та логіку прийняття рішень з урахуванням SCADA в єдиний скоординований робочий процес. **Завдання дослідження:** 1) формалізувати багаторівневу архітектуру КФС (БПЛА-периферія-хмара) та визначити інтерфейси для даних, гео-тегів та тривоги; 2) розробити та валідувати вбудований конвеєр термографічного виявлення з обробкою та злиттям з урахуванням палітри; 3) інтегрувати результати виявлення з логічним шаром, що враховує SCADA, для інференції небезпеки та зниження ризику пожеж. **Методи дослідження:** методи комп'ютерного зору та глибокого навчання (YOLOv11) для бортової сегментації дефектів; ансамблювання моделей (стратегія пізнього злиття M2 та M3 палітр) для нівелювання зсуву домену; алгоритми просторової кластеризації з підтримкою RTK для геоіндексації та дедуплікації; детермінована булева логіка для визначення пожежних ризиків на основі станів шунтуючих діодів. **Результати** п'ятикратної перехресної валідації свідчать, що запропонована архітектура значно перевершує одноmodalні підходи. Бортова модель YOLOv11 досягла показників mAP@0.5 на рівні 0,91 та 0,90 для палітр M2 та M3 відповідно. Ансамбль пізнього злиття підвищив mAP@0.5 для тріщин до 0,96, а розшарування – до 0,95. Запропонована архітектура зменшила наскрізну затримку оброблення кадру із 4,235 с до 2,858 с. Польова перевірка продемонструвала RMSE 0,71 дефекту на інспектований рядок порівняно з ручним підрахуванням. Аналіз чутливості

виявив, що висота польоту 10 м забезпечує оптимальне узгодження між роздільною здатністю та охопленням (точність 93 %, відтворюваність 90 %). **Висновки:** дослідження демонструє, що розгляд інспектування БПЛА як інтегрованого кіберфізичного сервісу суттєво покращує точність виявлення дефектів. Запропонована архітектура кіберфізичних систем є масштабованим рішенням для профілактичного обслуговування та автоматизованого зниження ризику пожеж у відновлюваній енергетиці.

Ключові слова: кіберфізичні системи; фотоелектричні модулі; безпілотні літальні апарати; інфрачервона термографія; периферійні обчислення; глибоке навчання.

Бібліографічні описи / Bibliographic descriptions

Саченко А. О., Радюк П. М., Лисий М. І., Мельниченко О. В., Архітектура кіберфізичних систем для виявлення дефектів у фотоелектричних модулях на основі пізнього злиття даних із БПЛА. *Сучасний стан наукових досліджень та технологій в промисловості*. 2026. № 1 (35). С. 79–99. DOI: <https://doi.org/10.30837/2522-9818.2026.1.079>

Sachenko, A., Radiuk, P., Lysyi, M., Melnychenko, O., Zastavnyy, O. (2026), "Architecture of cyberphysical systems for UAV-based late-fusion defect detection in photovoltaic modules", *Innovative Technologies and Scientific Solutions for Industries*, No. 1 (35), P. 79–99. DOI: <https://doi.org/10.30837/2522-9818.2026.1.079>
