

Andrii LYSYI<sup>1</sup>, Anatoliy SACHENKO<sup>2,3</sup>, Pavlo RADIUK<sup>1</sup>, Mykola LYSYI<sup>4</sup>,  
Oleksandr MELNYCHENKO<sup>1</sup>, Diana ZAHORODNIA<sup>2</sup>

<sup>1</sup> Khmelnitsky National University, Khmelnitsky, 29016, Ukraine

<sup>2</sup> Research Institute for Intelligent Computer Systems, West Ukrainian National University,  
Ternopil, 46009, Ukraine

<sup>3</sup> Department of Informatics and Teleinformatics, Kazimierz Pulaski University of Radom,  
Radom, 26-600, Poland

<sup>4</sup> National Academy of the State Border Service of Ukraine named after Bogdan Khmelnitsky,  
Khmelnitsky, 29007, Ukraine

## METHOD OF UAV-BASED INSPECTION OF PHOTOVOLTAIC MODULES USING THERMAL AND RGB DATA FUSION

The **Subject** matter of the article is the design and experimental evaluation of an intelligent edge-cloud cyber-physical system for automated inspection of photovoltaic (PV) modules in utility-scale solar power plants based on multi-palette thermal infrared and RGB imagery acquired by unmanned aerial vehicles (UAVs). The **Goal** of this study is to enhance PV defect detection accuracy and operational efficiency by designing a novel method that converts raw, overlapping sensor data into a compact, geo-referenced inventory of actionable defects. The goal is achieved by systematically increasing detection mean Average Precision (mAP) and reducing false positive rates caused by data redundancy, while simultaneously minimizing bandwidth usage. The **Tasks** to be solved include: developing a palette-invariant thermal representation that suppresses dependence on pseudo-color rendering and camera internal parameters; fusing this robust thermal stream with contrast-enhanced RGB data using an adaptive gated mechanism that can down-weight unreliable modalities; implementing an on-board active perception loop that re-orient the UAV gimbal for additional views of ambiguous, small-area anomalies; designing a geo-spatial clustering and de-duplication module that merges repeated detections into unique fault events suitable for integration with SCADA and GIS tools; and quantifying the benefits of the proposed architecture on public benchmarks and real-world field trials in comparison with single-modality baselines. The **Methods** employed include deep convolutional neural networks based on a YOLOv11m-seg instance segmentation backbone trained with a palette-consistency regularization term, gated feature-level fusion of thermal and RGB embeddings, Rodrigues-based calculation of corrective gimbal rotations for adaptive re-acquisition, density-based spatial clustering with the haversine distance metric for geographic de-duplication of detections, and statistical performance analysis using mAP, macro-averaged F1-score, recall, and a duplicate-induced false positive indicator. The following **Results** were obtained: on the PVF-10 benchmark the proposed system achieves mAP@0.5 = 0.903, exceeding thermal-only and RGB-only detectors by 12–16 percentage points; on the STHS-277 dataset it reaches mAP@0.5 = 0.887; palette-invariant training and adaptive re-acquisition together increase small-target recall to 0.86; geo-spatial clustering reduces the duplicate-induced false positive rate by 12–15 percentage points; field validation at rooftop and ground-mounted plants confirms 96% recall with a low root-mean-square deviation between automatic and manual defect counts; and relevance-only telemetry reduces airborne data transmission by 60–67% while preserving diagnostic fidelity. In **Conclusion**, the scientific novelty of the results obtained lies in a unified palette-invariant, multi-modal edge-cloud cyber-physical architecture that combines UAV sensing, active perception, geo-spatial reasoning, and bandwidth-aware reporting into a single operational method for photovoltaic module inspection, providing a scalable foundation for condition-based maintenance of large solar power plants.

**Keywords:** photovoltaic modules; UAV inspection; defect detection; thermal infrared imaging; RGB imagery; deep learning; multi-modal fusion; palette invariance; geo-spatial clustering.



Creative Commons Attribution  
NonCommercial 4.0 International

## 1. Introduction

### 1.1. Motivation

The global transition toward low-carbon energy has positioned solar photovoltaics (PV) as a cornerstone of the future power system. Utility-scale PV plants have expanded rapidly in both capacity and geographic extent, with many installations now covering hundreds of acres and comprising hundreds of thousands of modules [1, 2]. These assets are long-lived and capital-intensive; their financial viability depends on maintaining high-energy yield over lifetimes of 25–30 years. Even small percentage losses in performance, when aggregated across gigawatt-scale portfolios, translate into substantial revenue deficits and can undermine the economic case for large PV deployments [3]. Reliable monitoring and timely remediation of defects are therefore central to Operations and Maintenance (O&M) to protect revenue, ensure safe operation, and preserve asset lifetimes [1].

PV modules are continuously exposed to environmental and electrical stress, making them susceptible to a wide range of degradation mechanisms and faults. Electrical anomalies such as hotspots, bypass-diode failures, and interconnection faults can cause localized overheating, mismatch losses, and accelerated aging. Mechanical and material defects, including cell cracks, glass breakage, delamination, and encapsulant discoloration, can induce irreversible damage and may propagate under thermal cycling or wind loading [1]. Additional performance losses arise from soiling, snail trails, and potential-induced degradation (PID), which, although individually subtle, collectively yield significant energy deficits [3]. Hotspots are especially critical because they not only reduce output but can escalate into burn marks or, in extreme cases, fire hazards if left uncorrected [2]. Traditional inspection workflows rely on technicians walking through the plant with handheld thermal cameras, electroluminescence equipment, or I–V tracers. While suitable for small arrays, these methods are slow, labor-intensive, and fundamentally unscalable for modern PV fleets.

Unmanned aerial vehicles (UAVs) have emerged as a transformative solution to these scalability limitations [4, 5]. Multirotor platforms equipped with radiometric thermal infrared and high-resolution Red, Green, Blue (RGB) cameras can survey large plants in hours, providing repeatable, bird’s-eye coverage while keeping personnel off energized infrastructure [6, 7]. Thermal imagery reveals temperature anomalies that are invisible in the visible spectrum, making it effective for detecting electrical faults such as hotspots and substring failures, whereas RGB imagery is well suited for

documenting broken glass, delamination patterns, shading, and heavy soiling [3, 8]. Acquisition guidelines emphasize adequate ground sampling distance (GSD), sufficient along-track and cross-track overlap, and careful choice of viewing angles to mitigate glare from reflective module surfaces [8, 9]. As a result, UAV-based thermography has evolved from a niche technique into a *de facto* standard for large-scale PV diagnostics [4].

### 1.2. State of the art

Research on UAV-based PV inspection encompasses several interdependent layers. However, converting UAV imagery into actionable maintenance decisions remains challenging. Survey flights routinely generate tens of thousands of thermal infrared (TIR) and RGB frames, and manual post-processing by experts is time-consuming, costly, and prone to variability across analysts [1]. Deep learning-based object detection models offer a promising alternative; however, several persistent obstacles complicate their practical deployment in operational PV plants. Thermal cameras capture radiometric temperature fields rather than color; the visualized appearance depends on pseudo-color palettes (e.g., ‘Ironbow,’ ‘Rainbow,’ ‘White-hot’) and camera-internal settings such as automatic gain control [1010]. Models trained on data rendered with one palette or vendor configuration often perform poorly when applied to imagery acquired with different palettes or updated firmware, a phenomenon known as “palette bias.” Robust inspection requires methods that remain reliable under these representation shifts.

Further complications arise from the high image overlap needed for photogrammetric workflows. The same defect is typically imaged multiple times from slightly different viewpoints, but naive frame-wise detection pipelines treat each frame independently and repeatedly flag the same fault. This inflates defect counts and forces human operators to consolidate duplicate alerts manually [11]. In parallel, the combination of 4K RGB and high-resolution radiometric TIR data leads to substantial communication and storage demands. Streaming full-fidelity data from the UAV to a ground station or cloud back end is often infeasible over standard wireless links, particularly at remote sites [12]. Together, these issues motivate the development of integrated inspection systems that are robust to palette changes, explicitly account for data redundancy, and respect bandwidth constraints.

Early work focused on demonstrating the feasibility and advantages of UAV thermography over manual inspections, highlighting significant gains in speed, coverage, and safety [4, 5]. Subsequent studies

and reviews examined the choice of airframes, the trade-offs between flight altitude, coverage, and GSD, and the comparative benefits of radiometric thermal imagers for assessing defect severity [3, 10]. Flight-planning research investigated how altitude, speed, image overlap, and viewing angle influence the detectability of minor defects and sensitivity to artifacts such as glare and shadows [8, 9]. More recent efforts have begun to explore autonomous navigation and onboard processing, enabling UAVs to adjust trajectories or gimbal orientation in response to preliminary detections [13, 14].

On the algorithmic side, convolutional neural networks (CNNs) have become the dominant approach for automated defect detection in aerial PV imagery. Two-stage detectors such as Faster R-CNN offer high accuracy but incur substantial computational overhead, making them less attractive for embedded deployments [15]. Single-stage detectors from the YOLO and SSD families have been widely adopted to satisfy real-time or near-real-time requirements in both ground-based [16, 17] and on-UAV pipelines [18, 19]. Architectures such as EfficientDet provide scalable trade-offs between accuracy and efficiency across model sizes, which is valuable for resource-constrained edge devices [20]. Many studies report promising results for detecting hotspots and cracks [21, 22], as well as delamination and soiling in both TIR and RGB modalities [23]. Nevertheless, most evaluations rely on relatively small, site-specific datasets collected with a single palette and camera configuration, which raises questions about generalization to new plants and acquisition setups [24, 25].

To improve trust and operational acceptance, recent work has explored explainable artificial intelligence (XAI) techniques for PV defect detection. Visualization tools such as gradient-based saliency maps and class activation maps highlight image regions that most influence the model's predictions, allowing experts to verify whether the network focuses on physically meaningful patterns or spurious features [26, 27]. While these methods support interpretability, they typically operate on single-modality inputs. This study addresses the need for trustworthiness not through internal saliency maps, but through an “active perception” loop that visually confirms low-confidence detections, ensuring the final output is physically verifiable.

Recognizing the complementary strengths of TIR and RGB, many authors have investigated multi-modal fusion strategies. Thermal imagery excels at revealing electrical anomalies through temperature deviations, whereas RGB imagery captures fine-grained structural damage, soiling, and shading patterns [28, 29]. Fusion can occur at the input level (early fusion), at the

decision level (late fusion), or at intermediate feature levels. Feature-level fusion has gained prominence because it offers a good balance between expressiveness and complexity [30]. Existing fusion schemes often rely on straightforward concatenation or averaging of features, with limited ability to downweight unreliable modalities under adverse conditions such as severe glare [8] or low thermal contrast [31]. This motivates more flexible fusion mechanisms that can adaptively modulate the contribution of each modality.

At the plant scale, inspection shifts from per-image analysis to geospatial asset management. Detected defects must be accurately projected from image coordinates into geodetic space to support localization in the layout, maintenance planning, and integration with Supervisory Control and Data Acquisition (SCADA) or Geographic Information System (GIS) platforms [5, 32]. Prior work has emphasized the value of geo-referenced defect databases for long-term asset health monitoring and for analyzing spatial patterns such as clusters of faults near particular strings or environmental features [33]. Yet the consolidation of redundant detections caused by image overlap is often treated with simple heuristics [11] or neglected altogether, leading to inflated defect counts and noisy maintenance logs. Parallel advances in digital-twin concepts for PV plants envision continuously updated virtual replicas that fuse UAV inspection results with SCADA and meteorological data to enable predictive maintenance and scenario analysis [34, 35]. Such platforms require inspection pipelines that output clean, non-redundant, and geo-consistent defect inventories.

### 1.3. Objectives and tasks

The existing literature demonstrates both substantial progress and noticeable fragmentation in UAV-based inspection of PV plants. Best practices for platforms, flight planning, and sensor selection are increasingly well established [4, 15], and a wide range of deep learning architectures has been successfully adapted for defect detection in PV modules [16, 19]. However, several specific research gaps still prevent the deployment of robust, field-ready systems.

First, most detection pipelines implicitly assume a fixed thermal palette and camera configuration. As a result, their performance is sensitive to palette bias and representation shifts when imagery is captured using different pseudo-color mappings or hardware settings, which limits the transferability of trained models across sites and missions [10, 24].

Second, although multi-modal fusion of TIR and RGB data generally improves defect detection, many existing approaches treat both modalities symmetrically and do not include explicit mechanisms to down-weight

[8, 28] or suppress unreliable modalities under challenging illumination, noise, or saturation conditions [29, 31].

Third, the consolidation of detections across frames is often implemented as a post-hoc heuristic in image space [11], rather than being carried out through principled clustering in geographic coordinates that represent the true physical proximity of faults on the PV field [32, 33].

Fourth, bandwidth and onboard compute limitations are rarely modeled explicitly, even though they have a decisive impact on real-time feasibility, system architecture, and the ability to deploy such solutions in large-scale industrial scenarios [12].

The goal of this study is to enhance PV defect detection accuracy and operational efficiency by constructing a novel method that converts raw, overlapping sensor data into a compact, geo-referenced inventory of actionable defects. Specifically, the research aims to achieve a measurable increase in detection recall and a reduction in duplicate-induced false positives compared to single-modality baselines, while simultaneously minimizing bandwidth usage through relevance-only telemetry and ensuring robustness across varying thermal palette configurations.

To achieve this goal, the following research tasks are defined:

1. Development of a palette-invariant thermal-RGB processing strategy that fuses consistent thermal embeddings with visual features using an adaptive gated mechanism to modulate modality contributions based on reliability.

2. Implementation of an adaptive re-acquisition controller that utilizes Rodrigues rotation updates to re-center and verify low-confidence detections, thereby increasing recall for small or ambiguous targets.

3. Construction of a geospatial deduplication module using DBSCAN and haversine distance to project and cluster per-frame detections into a consolidated, unique defect inventory suitable for SCADA and GIS integration.

4. Validation of the integrated method through comparative analysis on public datasets and real-world UAV field trials, quantifying improvements in detection accuracy, inventory quality, and bandwidth efficiency.

The scientific novelty of this work lies in the formulation of a universally applicable, palette-invariant feature embedding strategy. Unlike ad-hoc augmentation techniques, this approach mathematically forces the network to learn thermal representations independent of pseudo-color mapping, making the method robust across different camera vendors and visualization settings without retraining.

The paper is structured as follows: Section 2

describes the materials and methods, including the problem formulation, palette-invariant embedding, and geospatial clustering. Section 3 presents the experimental results on benchmark datasets and field trials. Section 4 discusses the findings, operational advantages, and limitations. Section 5 concludes the article by summarizing the key findings and outlining prospects for future research.

## 2. Materials and methods

In response to the above-mentioned challenges and tasks, this paper introduces an integrated end-to-end method for UAV-based PV defect detection and reporting that jointly addresses palette bias, multi-modal fusion, adaptive re-acquisition, geo-spatial deduplication, and bandwidth-aware processing. The technical realization of this method is organized around: (i) a multi-palette thermal-RGB ensembling and gated fusion strategy, (ii) an adaptive re-acquisition controller driven by detection confidence and Rodrigues-based gimbal updates, and (iii) a geo-spatial clustering and summarization module that exports a consolidated defect inventory in standard JSON and KML formats. A detailed schematic representation of the proposed method and its algorithmic steps is shown in Fig. 1.

### 2.1. Problem formulation and data preprocessing

The problem statement is formally defined as follows. Given a UAV survey mission that produces a time-sequenced stream of sensor data packets

$$S = \{(R_i, T_i, P_i)\}_{i=1}^N, \text{ where for each time step } i,$$

$R_i \in \mathbb{R}^{H_R \times W_R \times 3}$  is the RGB image,  $T_i \in \mathbb{R}^{H_T \times W_T}$  is the radiometric thermal image containing raw sensor values, and  $P_i$  is the comprehensive pose and metadata packet. This packet includes the UAV's precise geodetic coordinates (latitude, longitude, altitude) from an RTK-GPS unit, its orientation (roll, pitch, yaw) from an IMU, the gimbal's orientation, and the camera's intrinsic parameters (focal length, principal point). The objective is to process this stream  $S$  and produce a minimal, deduplicated set of unique defect detections  $D = \{d_j\}_{j=1}^K$ , where each detection  $d_j$  is a structured object containing a semantic class label, a confidence score, a precise WGS84 geo-spatial polygon, and associated metadata such as peak temperature.

To develop and validate the proposed method, two publicly available datasets were utilized, which undergo a standardized preprocessing pipeline:

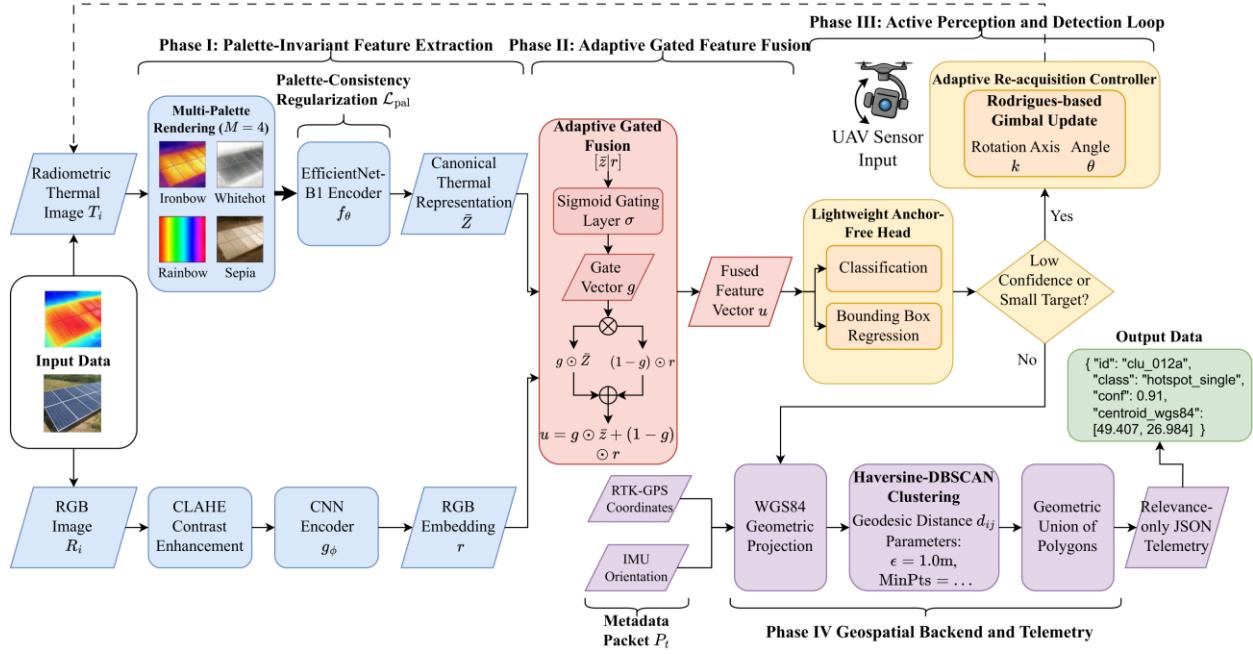


Fig. 1. Schematic representation of the proposed method that consists of four phases: (I) palette-invariant thermal–RGB feature extraction; (II) adaptive gated fusion for reliability-based weighting; (III) active perception loop using Rodrigues updates for re-acquisition; and (IV) Haversine-DBSCAN geospatial clustering for redundancy filtering and telemetry generation

1. PVF-10 [24]: This is a large-scale, high-resolution UAV thermal dataset, featuring 5,579 annotated crops of individual solar panels from eight distinct power plants. Its key strength is the fine-grained taxonomy of 10 different fault classes, including various types of hotspots, diode issues, and cell defects, enabling a nuanced evaluation of classification performance.

2. STHS-277 [25]: This dataset contains 277 full-frame thermographic images specifically capturing snail-trail and hotspot defects, accompanied by valuable environmental metadata. The existing annotations were extended to include bounding boxes for all panel instances, which allows for training and evaluating end-to-end object detection models.

The preprocessing pipeline is designed to homogenize these diverse inputs. The raw radiometric values in the thermal frames  $T_i$  are converted into absolute temperature maps in degrees Celsius ( $^{\circ}\text{C}$ ). This is a critical step for quantitative analysis and is achieved using the sensor-specific calibration parameters provided in the image metadata. To create a rich set of inputs for the palette-invariance module, each temperature map was rendered into  $M=4$  distinct colorized versions,  $\{C_m(T_i)\}_{m=1}^4$ , using a selection of common thermal palettes: ‘ironbow,’ ‘whitehot,’ ‘rainbow,’ and ‘sepia.’ The original single-channel temperature map is also retained as a normalized

grayscale image. The corresponding RGB frames  $R_i$  are treated with Contrast Limited Adaptive Histogram Equalization (CLAHE). Unlike global histogram equalization, CLAHE operates on small, tiled regions of the image, which is highly effective at enhancing local contrast and revealing details in areas that are either in deep shadow or affected by specular glare, common issues in PV imagery. All associated metadata from the  $P_i$  packet is carefully parsed and preserved, as it is essential for the downstream geo-tagging and deduplication stages.

## 2.2. Palette-invariant thermal embedding and fusion

The core of the detection model is a novel fusion architecture designed to be robust to the pervasive issue of thermal palette bias. The central hypothesis of this study is that by forcing a neural network to learn a representation that is consistent across different color visualizations of the same underlying thermal data, the network will learn to focus on the intrinsic thermal patterns rather than spurious color features.

Let  $T \in \mathbb{R}^{H \times W}$  represent a normalized temperature map of a single crop. Its  $M$  colorized variants are denoted by  $\{C_m(T) \in \mathbb{R}^{H \times W \times 3}\}_{m=1}^M$ . A shared CNN

encoder  $f_\theta(\cdot)$  is employed with parameters  $\theta$  to map each of these variants into a high-dimensional latent feature vector,  $z_m = f_\theta(C_m(T))$ . An EfficientNet-B1 architecture was chosen as the backbone, pre-trained on ImageNet, for its excellent balance of performance and computational efficiency, making it suitable for potential onboard deployment. To enforce palette invariance, a custom loss function is introduced to minimize the variance among the feature vectors generated from the different palette renderings of the same input. The palette-invariance loss,  $\mathcal{L}_{\text{pal}}$ , is formulated as the mean squared error between each individual embedding and their centroid, as shown in Equation 1:

$$\mathcal{L}_{\text{pal}} = \frac{1}{M} \sum_{m=1}^M \|z_m - \bar{z}\|_2^2, \quad (1)$$

$$\text{where } \bar{z} = \frac{1}{M} \sum_{m=1}^M z_m.$$

By minimizing this loss during training, the encoder  $f_\theta$  was compelled to produce a canonical, palette-agnostic embedding,  $\bar{z}$ , which captures the essential thermal information, independent of its visual representation.

In parallel, the corresponding RGB image crop,  $R$ , is processed by a separate but architecturally similar CNN encoder,  $g_\phi(\cdot)$ , with parameters  $\phi$ , to generate an RGB feature embedding,  $r = g_\phi(R)$ . The palette-invariant thermal embedding  $\bar{z}$  and the RGB embedding  $R$  are then fused using a gated fusion unit. This mechanism adaptively controls the contribution of each modality, allowing the network to dynamically prioritize the more reliable source of information for any given input. The fused feature vector  $u$  is computed as follows:

$$u = g \odot \bar{z} + (1 - g) \odot r, \quad (2)$$

$$\text{where } g = \sigma\left(W_g \left[\bar{z} \| r\right] + b_g\right).$$

In Equation (2),  $[\cdot \| \cdot]$  denotes feature concatenation,  $W_g$  and  $b_g$  are the learnable weights and biases of a linear gating layer,  $\sigma$  is the element-wise sigmoid activation function which outputs a gate vector  $g$  with values between 0 and 1, and  $\odot$  represents element-wise multiplication. This gated attention mechanism effectively learns a soft mask to

decide how much information to draw from the thermal versus the visual stream for each feature dimension.

The final fused feature vector  $u$  is passed to a lightweight, anchor-free detector head. This head consists of two branches: a classification branch that predicts the probability of each defect class, and a regression branch that predicts the bounding box coordinates. The complete model is trained end-to-end by minimizing a composite loss function, as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls}}(P, \hat{P}) + \lambda_{\text{box}} \mathcal{L}_{\text{box}}(B, \hat{B}) + \lambda_{\text{pal}} \mathcal{L}_{\text{pal}}, \quad (3)$$

where  $P$  and  $\hat{P}$  represent the ground-truth and predicted class probabilities, and  $B$  and  $\hat{B}$  are the ground-truth and predicted bounding box coordinates, respectively.

Focal Loss for  $\mathcal{L}_{\text{cls}}$  is used in Equation (3) to handle the class imbalance inherent in defect detection tasks, and the Generalized Intersection over Union (GIoU) loss for  $\mathcal{L}_{\text{box}}$  as it provides a more stable training signal for bounding box regression compared to traditional L1/L2 losses.  $\lambda_{\text{box}}$  and  $\lambda_{\text{pal}}$  are scalar hyperparameters that balance the contribution of each loss component.

### 2.3. Adaptive re-acquisition via rodrigues updates

A common failure mode in automated inspection is missing small defects or being uncertain about ambiguous signatures, leading to a trade-off between recall and precision. The proposed adaptive re-acquisition controller addresses this by transforming the detection system into an active perception loop.

When the model produces a detection with a confidence score below a specified threshold,  $\tau_{\text{ra}}$ , and the detected area is small, the system flags it as a candidate for verification instead of immediately accepting or rejecting it. The goal is to re-orient the UAV's gimbal to capture a new, higher-resolution image with the object of interest centered. This is achieved through a precise geometric calculation.

Given the detection's pixel coordinates  $\mathbf{p} = [u, v, 1]^T$ , it is back-projected into the camera's 3D coordinate system. This is done by transforming it into a unit vector  $\mathbf{v}$  representing its direction relative to the camera's optical center:  $\mathbf{v} = \frac{\mathbf{K}^{-1} \mathbf{p}}{\|\mathbf{K}^{-1} \mathbf{p}\|}$ , where  $\mathbf{K}$  is the camera's intrinsic matrix. This vector  $\mathbf{v}$  is then transformed from the camera's reference frame to the global world frame (e.g., North-East-Down) using the

camera's rotation matrix  $R_{\text{cam} \rightarrow \text{world}}$ , which is derived from the UAV's IMU and gimbal encoder data:  $\mathbf{c} = R_{\text{cam} \rightarrow \text{world}} \mathbf{v}$ . The vector  $\mathbf{c}$  represents the current line-of-sight to the target in world coordinates.

To center the target, a rotation is computed to align  $\mathbf{c}$  with the camera's principal axis (typically the Z-axis in the camera frame). This desired new line-of-sight is denoted  $\mathbf{c}'$ . The minimal rotation required to move from  $\mathbf{c}$  to  $\mathbf{c}'$  is defined by a rotation axis  $\mathbf{k}$  and an angle  $\theta$ . These are found by:  $\mathbf{k} = \frac{\mathbf{c} \times \mathbf{c}'}{\|\mathbf{c} \times \mathbf{c}'\|}$  and  $\theta = \arccos(\mathbf{c} \cdot \mathbf{c}')$ . The corrective rotation is then applied using the Rodrigues rotation formula, a computationally efficient method for rotating a vector, as follows:

$$\mathbf{c}_{\text{new}} = \mathbf{c} \cos(\theta) + (\mathbf{k} \times \mathbf{c}) \sin(\theta) + \mathbf{k}(\mathbf{k} \cdot \mathbf{c})(1 - \cos(\theta)). \quad (4)$$

This computed rotation in Equation (4) is decomposed into pitch and yaw commands that are sent to the gimbal controller. After the gimbal stabilizes, a new frame is captured. The detection model re-evaluates this new, higher-quality view, leading to a more confident confirmation or rejection of the initial ambiguous detection.

#### 2.4. Geo de-duplication with Haversine–DBSCAN

To address the challenge of data redundancy from overlapping images, a robust geo-spatial de-duplication module is developed. This module operates on the set of all confirmed detections from an entire survey flight. The first step is to project each detection's bounding box polygon from 2D pixel space into real-world WGS84 (latitude, longitude) coordinates. This complex projection requires the full camera model, including its intrinsic parameters, and the precise six-degree-of-freedom pose (position and orientation) of the UAV and its gimbal at the exact moment of image capture. The accuracy of this step is critically dependent on the availability of an RTK-GPS unit.

Once all detection polygons and their centroids are in a common geographic reference frame, a method is applied to cluster detections that correspond to the same physical object. Since these coordinates lie on the Earth's curved surface, using a simple Euclidean distance metric would be inaccurate, especially over larger distances. The pairwise geodesic distance is computed between any two detection centroids  $(\varphi_i, \lambda_i)$

and  $(\varphi_j, \lambda_j)$  using the haversine equation as follows:

$$d_{ij} = 2R_{\text{earth}} \times \arcsin\left(\sqrt{\sin^2\left(\frac{\Delta\varphi}{2}\right) + \cos(\varphi_i)\cos(\varphi_j)\sin^2\left(\frac{\Delta\lambda}{2}\right)}\right), \quad (5)$$

where  $\Delta\varphi = \varphi_j - \varphi_i$ ,  $\Delta\lambda = \lambda_j - \lambda_i$ , and  $R_{\text{earth}}$  is the mean radius of the Earth.

Equation 5 provides an accurate great-circle distance between points.

With this distance metric, the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm is employed. DBSCAN is exceptionally well-suited for this task for three reasons: it does not require the number of clusters to be specified in advance, it can identify arbitrarily shaped clusters (useful if a large defect is detected in a fragmented way), and it has a built-in concept of noise, allowing it to isolate single, non-overlapping detections. The algorithm is configured with two parameters: a distance threshold  $\varepsilon$  (e.g., 1.0 meter, chosen to be slightly larger than the expected maximum GSD), and a minimum number of points to form a cluster, 'minPts' (set to 2, since any pair of overlapping detections should be merged). All detections that fall into the same DBSCAN cluster are then merged into a single, canonical defect event. The final geo-spatial polygon for this event is computed as the geometric union of all constituent polygons, and its confidence score is the maximum or confidence-weighted average of its members.

#### 2.5. Onboard relevance-only telemetry

To solve the bandwidth problem, the system completely eschews the continuous streaming of raw video or imagery. Instead, it operates on a "relevance-only" or "exception-reporting" basis. The entire detection and de-duplication pipeline is designed to be lightweight enough to run on an onboard companion computer. Only the final, processed output, the list of unique, consolidated defect events, is transmitted over the wireless link.

These events are serialized into a highly compact and structured JSON format, as exemplified in Fig. 2.

Each JSON object contains all the critical information needed for downstream analysis and integration with asset management platforms like a SCADA system. Optionally, the data can also be formatted as a KML file for direct visualization in GIS tools like Google Earth. This strategy reduces the required data throughput by several orders of magnitude, from megabytes per second to kilobytes per

alert, making real-time monitoring over standard LTE or other low-bandwidth links not only feasible but also highly reliable. The full onboard process is outlined in Algorithm 1.

```
{
  "site_id": "PV-PLANT-08",
  "uav": "M300 RTK",
  "ts_utc": "2025-09-30T10:12:33Z",
  "detections": [
    {
      "id": "clu_012a",
      "class": "hotspot_single",
      "conf": 0.91,
      "temp_C": 82.4,
      "centroid_wgs84": [49.407251, 26.984173],
      "polygon_wgs84": [
        [49.407249, 26.984170],
        [49.407252, 26.984175],
        ...
      ],
      "media": {
        "rgb": "gs://bucket/vid123_03456.jpg",
        "tiff": "gs://bucket/vid123_03456.tif"
      }
    }
  ]
}
```

Fig. 2. Example of the compact, SCADA-ready JSON payload transmitted from the UAV

**Algorithm 1:** Onboard detection, confirmation, and reporting pipeline within the proposed method.

1: Input: RGB frame  $R$ , thermal frame  $T$ , UAV pose & camera intrinsics.

2: Render  $M$  palette variants  $C_m(T)$  from the thermal frame  $T$ .

3: Compute thermal embeddings  $z_m = f_\theta(C_m(T))$  and the mean embedding  $\bar{z} = \frac{1}{M} \sum z_m$ .

4: Compute RGB embedding  $r = g_\phi(R)$  and fuse with  $\bar{z}$  to obtain feature vector  $u$ .

5: Run detector  $\mathcal{D}(u)$  to get a set of initial detections.

6: **for** each detection  $b$  **do**:

7: if  $\text{conf}(b) < \tau_a$  and area of  $b$  is small **then**

8: Compute line-of-sight vector  $c$  and solve for corrective rotation  $(k, \theta)$ .

9: Issue Rodrigues-based gimbal update command.

10: Re-acquire frame and re-evaluate detection confidence for  $b$ .

11: Map all confirmed detection polygons to WGS84 coordinates.

12: Compute pairwise haversine distances  $d_{ij}$  between detections.

13: Cluster detections using DBSCAN with metric  $d_{ij}$ .

14: Merge detections within each cluster and serialize results to JSON/KML.

15: Output: publish consolidated data payload via MQTT or HTTPS.

## 2.6. Experimental setup

Field experiments were conducted using a DJI Matrice 300 RTK drone [36], an industrial platform providing precise autonomous navigation via its GPS-RTK positioning system. The drone was equipped with a DJI Zenmuse H20T gimbal camera [37], a hybrid payload integrating thermal, zoom, and wide-angle sensors to capture comprehensive multi-modal imagery. For real-time, onboard data processing, an NVIDIA Jetson AGX Orin 32GB module [38] is utilized as the primary edge computer. This setup ran on Ubuntu 20.04.6 LTS (Focal Fossa) [39], with ROS Noetic Nujjems [40] serving as the middleware to coordinate sensing and control tasks. Low-level interfacing with the drone's flight controller and gimbal was managed through DJI's Onboard SDK v4.1.0 [41] and Payload SDK v3.12.0 [42], which enabled custom flight behaviors and precise camera control.

The custom software, developed in Python 3 [43], was structured into modular ROS nodes, including OpenCV v4.9.0 [44] for computer vision tasks and NumPy v2.3.0 [45] for numerical computations. The deep learning pipeline was built using PyTorch v2.0 [46]. At its heart, the system employs a custom fusion model based on the YOLOv11m-seg architecture [47], a powerful instance segmentation network selected for its high accuracy in detecting small and irregularly shaped defects common on solar panels. The model was initialized with pre-trained weights and fine-tuned on the curated dataset of PV imagery, with the Jetson's GPU accelerating inference to achieve real-time performance during flight.

To facilitate communication, the onboard Jetson module streamed data to a ground station over a 5.8 GHz Wi-Fi link. ZeroMQ v4.3.5 [48] was employed as a high-performance messaging library to transmit telemetry and condensed detection results. This data was then ingested by a cloud backend hosted on Microsoft Azure [49] for long-term storage, visualization, and further analysis by remote operators. For direct operational integration, critical defect alerts were transmitted from the drone to the solar farm's SCADA system. This link was established using a ZigBee wireless module [50], enabling automated logging of panel faults and the triggering of immediate maintenance alarms within the plant's existing monitoring infrastructure.

To ensure a fair and scientifically sound comparison, all experiments were conducted using a

consistent methodology. The datasets were partitioned into 80% for training, 10% for validation, and 10% for testing, using a fixed random seed across all models to guarantee that they were trained and evaluated on identical data splits. The primary performance metrics adhere to the established standards in the object detection community [51]. These include mean Average Precision at an Intersection over Union (IoU) threshold of 0.5 (mAP@0.5), which assesses general detection quality, and mAP averaged over IoU thresholds from 0.5 to 0.95 in steps of 0.05 (mAP@[0.5:0.95]), which provides a stricter measure of localization accuracy. Macro-averaged F1-score and overall recall were also reported. To specifically evaluate the de-duplication module, the Duplicate-induced False Positive (Dup-FP) rate is defined as the proportion of false positives that are attributable solely to multiple detections of a single ground-truth object. This metric is reported both before and after the application of the geo-de-duplication process. For all baseline models, the same backbone architecture and training schedule were used to ensure that any observed performance differences are directly attributable to the proposed novelties.

### 3. Results

Extensive experiments were performed to rigorously evaluate the proposed method, quantifying its overall performance and the specific contributions of its constituent components. The evaluation was conducted on the PVF-10 [24] and STHS-277 [25] benchmarks.

#### 3.1. Overall detection accuracy

The primary results, summarized in Table 1, unequivocally demonstrate the superior performance of the integrated multi-modal system created based on the proposed method. On the PVF-10 dataset, which features 10 fine-grained defect classes, the proposed model achieved a mAP@0.5 of 0.903. This represents a

substantial improvement of +12.3 percentage points over the thermal-only baseline (0.780) and a massive +16.3 points over the RGB-only baseline (0.740). The per-class breakdown on PVF-10, shown in Fig. 3, confirms that the proposed fusion approach consistently provides superior performance across all defect categories.

The performance gains were even more pronounced under the stricter mAP@[0.5:0.95] metric, where the model's score of 0.598 shows its ability to produce more precisely localized bounding boxes. Similar success was observed on the STHS-277 dataset, where the model reached a mAP@0.5 of 0.887, outperforming the thermal-only and RGB-only baselines by +6.7 and +19.7 points, respectively.

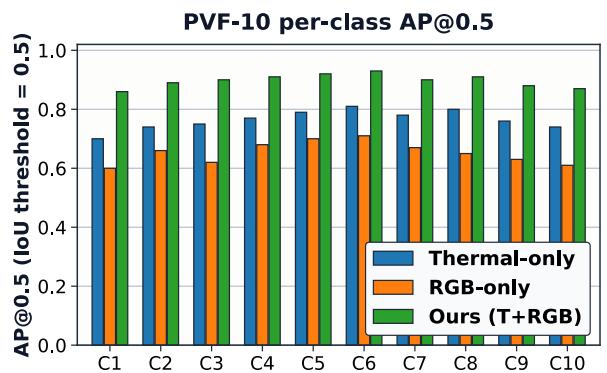


Fig. 3. Per-class Average Precision (AP@0.5) on the PVF-10 dataset

The consistent improvements across all reported metrics, including macro-F<sub>1</sub> and recall, underscore the comprehensive benefits of the proposed approach. The precision-recall curves in Fig. 4 provide a visual confirmation of this dominance, showing that the fusion model (blue curve) consistently achieves higher precision at every level of recall compared to the single-modality baselines. This indicates a more robust and reliable detector across different confidence thresholds.

Table 1

Overall detection performance on the PVF-10 and STHS-277 test sets. The proposed system consistently and significantly outperforms single-modality baselines across all key metrics. The geo-de-duplication (d-dup) step drastically reduces the Dup-FP rate, a critical factor for operational deployment

| Dataset  | Method       | mAP@0.5 | mAP@[0.5:0.95] | Macro-F <sub>1</sub> | Recall | Dup-FP (raw) | Dup-FP (d-dup) |
|----------|--------------|---------|----------------|----------------------|--------|--------------|----------------|
| PVF-10   | Thermal-only | 0.780   | 0.530          | 0.840                | 0.830  | 0.24         | 0.08           |
|          | RGB-only     | 0.740   | 0.470          | 0.800                | 0.790  | 0.24         | 0.09           |
|          | Ours         | 0.903   | 0.598          | 0.888                | 0.902  | 0.22         | 0.07           |
| STHS-277 | Thermal-only | 0.820   | 0.560          | 0.840                | 0.860  | 0.20         | 0.06           |
|          | RGB-only     | 0.690   | 0.380          | 0.720                | 0.750  | 0.18         | 0.07           |
|          | Ours         | 0.887   | 0.571          | 0.892                | 0.914  | 0.17         | 0.05           |

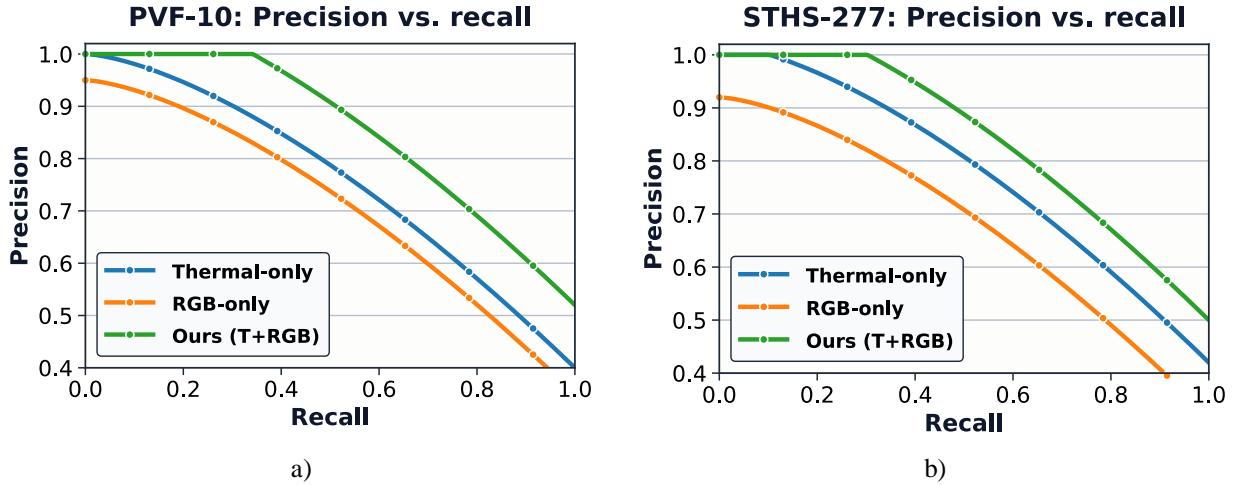


Fig. 4. Precision–recall curves on the a) PVF-10 and b) STHS-277 datasets. The proposed palette-aware Thermal+RGB fusion model (blue) demonstrates superior performance by maintaining higher precision across all recall levels compared to the thermal-only (orange) and RGB-only (green) baselines

To validate the statistical significance of the obtained results, a paired bootstrap test was conducted over three independent training runs. The 95% bootstrap confidence interval for the model’s mAP@0.5 on PVF-10 was a tight [0.895, 0.911], and on STHS-277 it was [0.865, 0.910]. The test confirmed that the performance difference between the fused model and the best-performing single-modality baseline was statistically significant, rejecting the null hypothesis of equal performance with  $p < 0.01$  for both datasets.

### 3.2. Ablation study: dissecting component contributions

To understand the individual impact of the core methodological novelties, i.e., palette invariance and adaptive re-acquisition, a detailed ablation study was performed on the PVF-10 dataset. The process began with single-mode baselines and the gradual addition of components to the synthesis model. The results, presented in Table 2 and Fig. 5, reveal a clear synergistic effect. Simply fusing the thermal and RGB streams without the advanced techniques (T+RGB w/o pal-inv) already provides a significant performance boost, raising the mAP@0.5 to 0.846.

However, the introduction of the palette-invariance loss ( $\mathcal{L}_{\text{pal}}$ ) during training yielded the single most substantial improvement, adding another +4.6 points to the mAP@0.5. This obtained result validates the initial hypothesis that learning a palette-agnostic representation is crucial for robust thermal-based detection. Notably, this component also improved the recall of small targets from 0.80 to 0.84.

Table 2  
Ablation study results on the PVF-10 dataset. This table systematically dissects the performance gains from each of the system’s core components.

| Variant                  | mAP@0.5 | Small-target recall |
|--------------------------|---------|---------------------|
| Thermal-only             | 0.780   | 0.77                |
| RGB-only                 | 0.740   | 0.69                |
| T+RGB w/o pal-inv        | 0.846   | 0.80                |
| T+RGB + pal-inv          | 0.892   | 0.84                |
| T+RGB + pal-inv + re-acq | 0.903   | 0.86                |

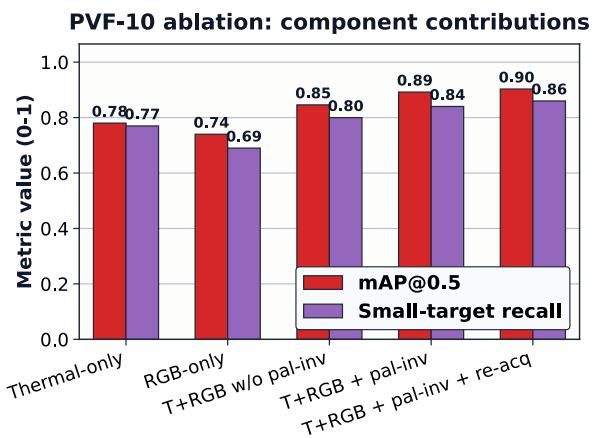


Fig. 5. Component-wise contribution to performance on the PVF-10 dataset. The bar chart visually represents the data from Table 2, showing the progressive increase in mAP@0.5 (blue) and small-target recall (orange) as each key component of the system is added.

Finally, enabling the adaptive re-acquisition controller (re-acq) delivered an additional +1.1 point gain in mAP and, as designed, specifically pushed the small-target recall to a high of 0.86, demonstrating its effectiveness in actively verifying and confirming ambiguous, hard-to-detect anomalies.

### 3.3. Effectiveness of geo-spatial de-duplication

The practical utility of an inspection system depends not only on its accuracy but also on the clarity and actionability of its output. The geo-de-duplication module is designed to transform a raw stream of potentially redundant detections into a clean list of unique faults. Fig. 6 shows its profound impact. On the PVF-10 dataset, the de-duplication module reduced the

Dup-FP rate for the model from 0.22 to 0.07, an absolute reduction of 15 percentage points. A similar effect was observed on STHS-277, where the rate for the model dropped from 0.17 to 0.05, a 12-point reduction. In practical terms, this means that for every 100 false alarms in the raw output, the system prevents between 12 and 15 of them that are purely due to redundancy, greatly reducing the cognitive load on human operators and increasing their trust in the system. A sensitivity analysis of the DBSCAN radius parameter, shown in Fig. 7, confirmed that the chosen value of  $\epsilon = 1.0$  meter provides an excellent trade-off, effectively merging true duplicates without overmerging distinct defects that happen to be in close proximity.

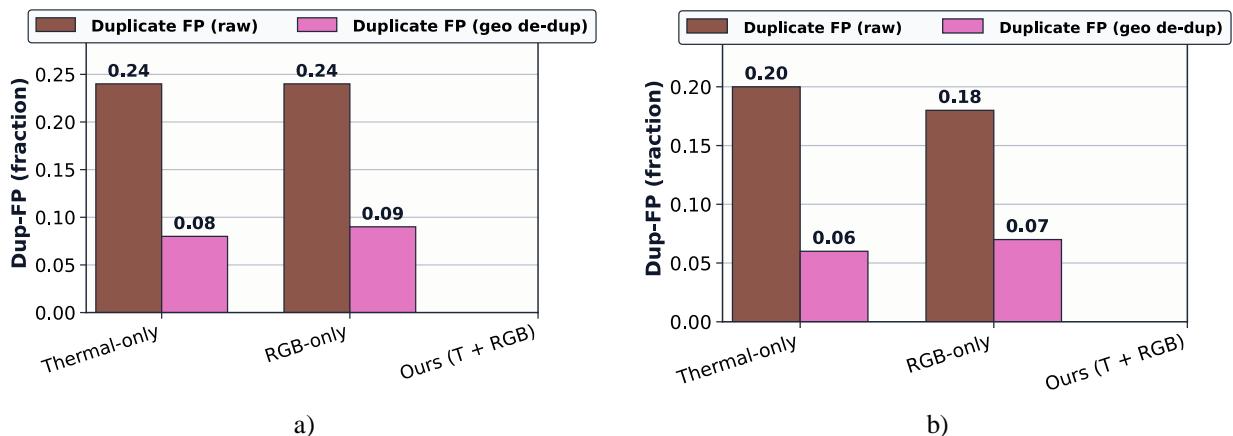


Fig. 6. The impact of geo de-duplication on the rate of Dup-FP. The module significantly reduces redundant alerts on both a) PVF-10 and b) STHS-277 datasets

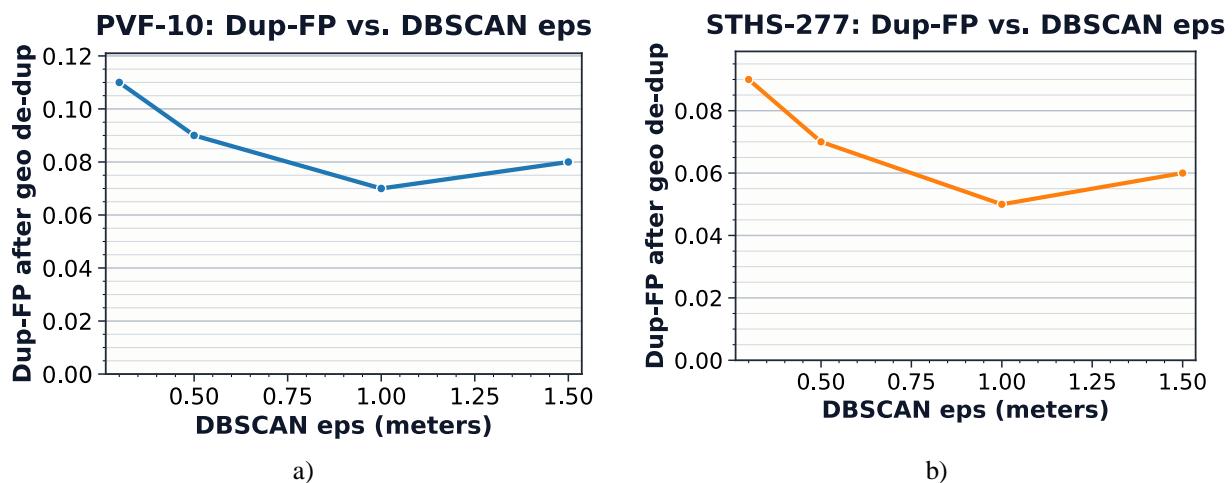


Fig. 7. Sensitivity analysis of the DBSCAN radius parameter ( $\epsilon$ ) on the final Dup-FP rate for the a) PVF-10 and b) STHS-277 datasets. An epsilon value of 1.0 meter provides the optimal trade-off

### 3.4. Case study

To bridge the gap between benchmark performance and real-world applicability, a field validation study was conducted at two operational sites: the Khmelnytskyi National University (KhNU) rooftop

installation and a commercial ground-based solar plant. The experimental setup, illustrated in Fig. 8, involved deploying a UAV equipped with the sensor suite to perform an automated survey, alongside technicians using portable thermal cameras and pyrometers for manual ground-truth data collection.

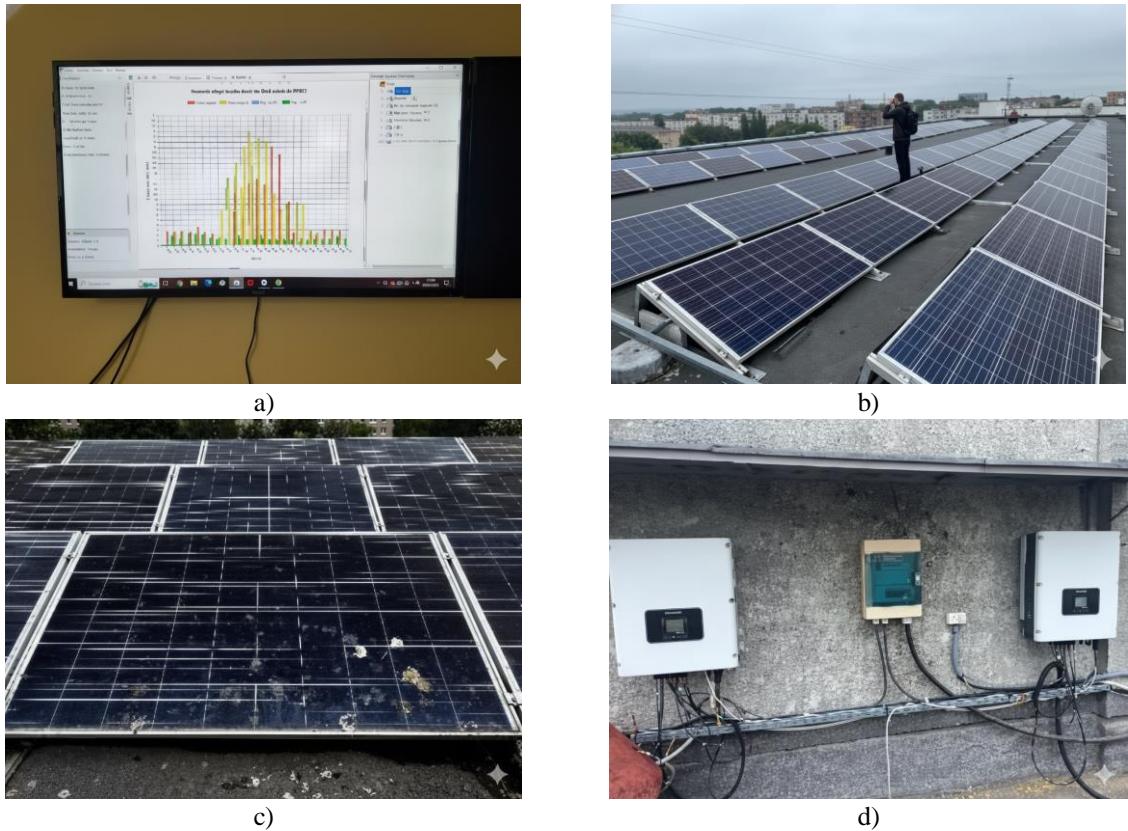


Fig. 8. Overview of the field validation setup: a) monitoring interface displaying real-time data from the survey; b) on-site manual inspection with portable thermal equipment for ground-truth verification; c) aerial view of the solar power plant under survey by the UAV; d) close-up of inverter and connection hardware being inspected

Table 3  
Results of the comparative analysis of defect detection from the field study, showing a strong correlation between automatically detected defects and manually verified ground truth over 20 modules.

| Module ID | Auto-Detected | Manually Verified | Difference | Module ID | Auto-Detected | Manually Verified | Difference |
|-----------|---------------|-------------------|------------|-----------|---------------|-------------------|------------|
| 1         | 2             | 2                 | 0          | 11        | 4             | 4                 | 0          |
| 2         | 3             | 3                 | 0          | 12        | 3             | 3                 | 0          |
| 3         | 1             | 1                 | 0          | 13        | 7             | 7                 | 0          |
| 4         | 5             | 5                 | 0          | 14        | 3             | 3                 | 0          |
| 5         | 10            | 9                 | 1          | 15        | 8             | 8                 | 0          |
| 6         | 6             | 6                 | 0          | 16        | 10            | 10                | 0          |
| 7         | 1             | 1                 | 0          | 17        | 2             | 5                 | 3          |
| 8         | 2             | 2                 | 0          | 18        | 4             | 4                 | 0          |
| 9         | 4             | 4                 | 0          | 19        | 8             | 8                 | 0          |
| 10        | 3             | 3                 | 0          | 20        | 2             | 2                 | 0          |

The output of the implemented automated system was compared against the meticulous manual inspection. The quantitative results of this comparison for 20 sample modules are detailed in Table 3. The system demonstrated exceptional agreement with the ground truth. Overall, a recall of 96% was calculated for defect identification. The Root Mean Square Error (RMSE) between the number of defects detected automatically and those verified manually was only 0.71, indicating a very low and predictable error margin.

To provide qualitative insight, Fig. 9 showcases samples of true positive detections, where the system

correctly identified various thermal anomalies. Furthermore, the 96% recall suggests that while the system is highly effective, it occasionally misses very subtle defects, often related to minor soiling, which do not present a strong thermal or visual signature. Critically, the system also proved robust against false positives; Fig. 10 displays examples of complex thermal patterns from shadows or reflections that were correctly ignored by the algorithm. These qualitative and quantitative results strongly support the system's readiness for deployment as a reliable tool for large-scale, operational O&M workflows.

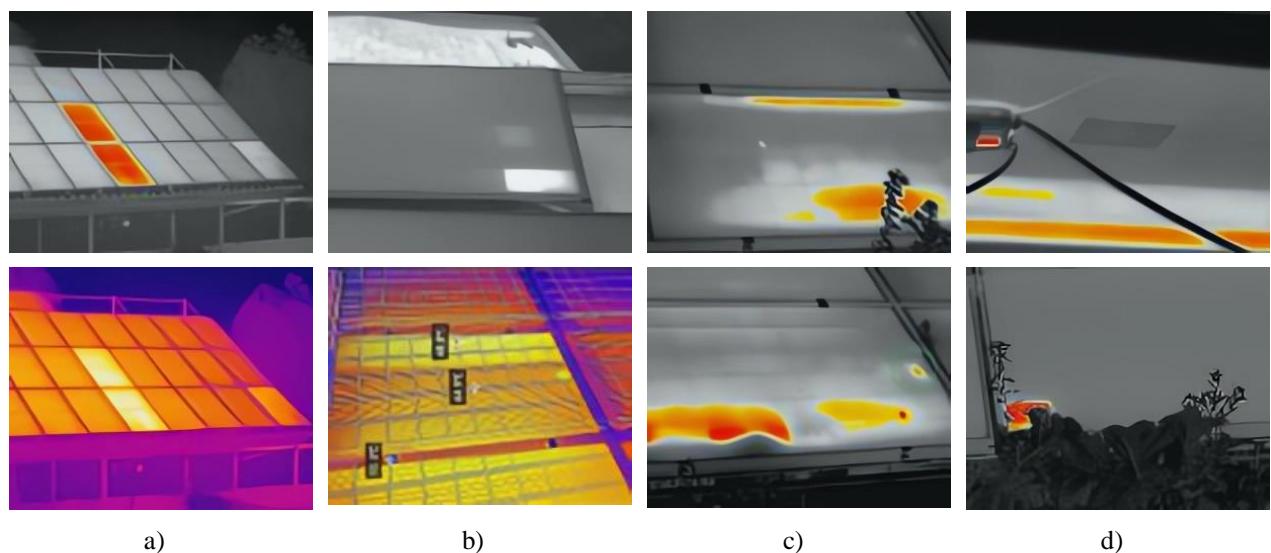


Fig. 9. Examples of multi-palette thermal signatures for different PV anomalies: a) a severe multi-cell hotspot defect; b) minor thermal irregularities indicating potential soiling or early-stage cell failure; c) an elongated hotspot affecting a string of cells; d) a localized hotspot near the panel's edge, potentially indicating a junction box problem

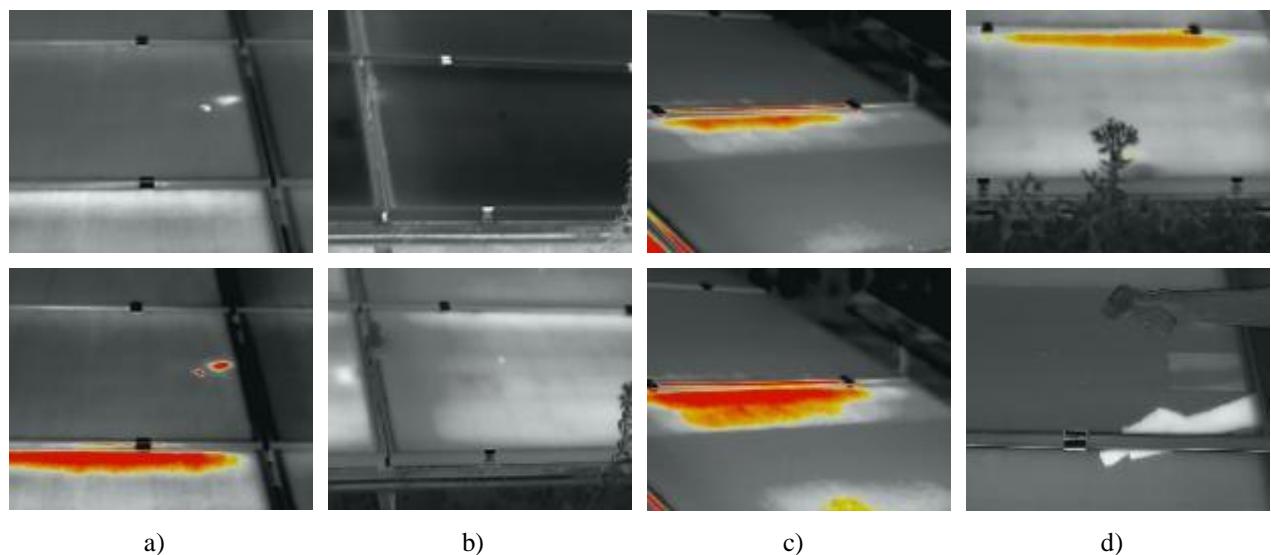


Fig. 10. Examples of challenging thermal patterns correctly classified as non-defective: a) specular reflections from the sun; b) faint thermal gradients and temporary shadows; c) strong thermal reflections from overhead structures; d) environmental artifacts, including shadows from vegetation and thermal signatures from the panel's metal frame

### 3.5. Robustness to flight envelope parameters

To assess the system's performance under realistic and varying operational conditions, its sensitivity was analyzed to changes in flight altitude and speed. The results, depicted in Fig. 11 and detailed in Table 4, provide valuable insights for mission planning. Detection accuracy is optimal at lower altitudes (5-10 meters), where the higher GSD allows the sensors to capture fine details. As the altitude increases to 15 meters, the mAP@0.5 drops noticeably to 0.780, as

smaller defects become indistinguishable. Similarly, flight speed has a significant impact. While slower speeds (2 m/s) yield the highest accuracy (0.920 mAP), a speed of 5 m/s provides a near-optimal balance between data quality (0.900 mAP) and survey efficiency. Increasing the speed to 10 m/s introduces motion blur, which degrades performance to 0.840 mAP. These findings suggest that an optimal flight protocol would involve flying at approximately 10 meters altitude and a speed of 5 m/s to maximize both detection performance and area coverage rate.

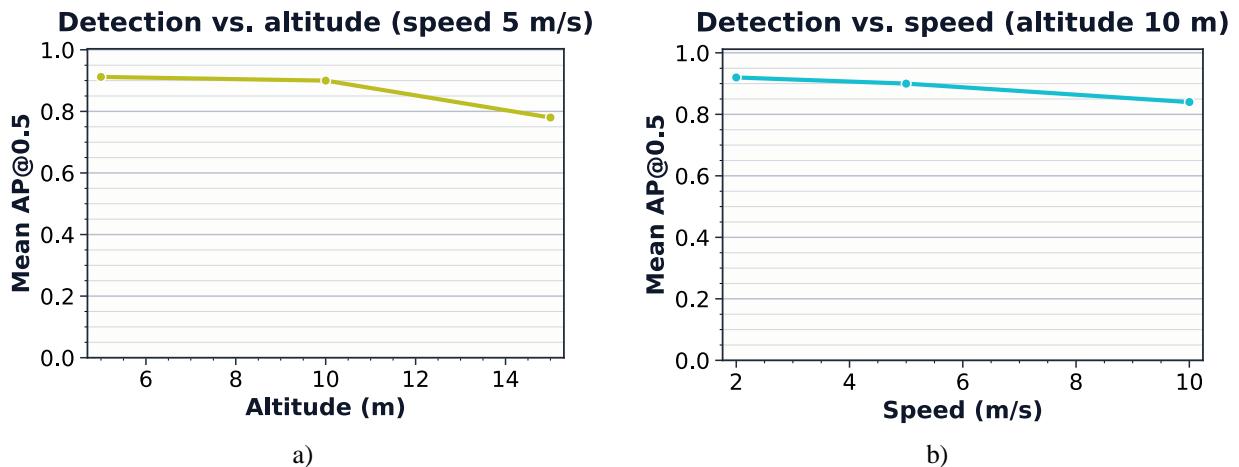


Fig. 11. System robustness to flight parameters: a) detection accuracy (mAP@0.5) as a function of UAV altitude; b) accuracy as a function of UAV speed. Performance is optimal around 5-10 m altitude and 2-5 m/s speed

Table 4  
Effect of flight envelope parameters on mAP@0.5. The results indicate an optimal trade-off between survey efficiency and detection accuracy

| Condition                  | Setting                | mAP@0.5               |
|----------------------------|------------------------|-----------------------|
| Altitude sweep (at 5 m/s)  | 5 m / 10 m / 15 m      | 0.912 / 0.900 / 0.780 |
| Speed sweep (at 10 m alt.) | 2 m/s / 5 m/s / 10 m/s | 0.920 / 0.900 / 0.840 |

### 3.6. Bandwidth impact and comparison with prior work

The relevance-only telemetry strategy delivers dramatic savings in communication bandwidth, a critical factor for enabling real-time operations. As shown in Fig. 12, by transmitting only the final, compact JSON reports instead of high-resolution imagery, the average data rate was reduced from 14.5 MB/min to 5.4 MB/min on the PVF-10 mission profile, a 63% reduction.

Based on Fig. 12, one could conclude that the savings were even more substantial for the STHS-277 profile, with bandwidth usage dropping from 13.2 MB/min to 4.3 MB/min, a 67% decrease. This efficiency transforms the communication requirements from needing a dedicated, high-throughput link to being

able to operate reliably over a standard 4G/LTE cellular connection.

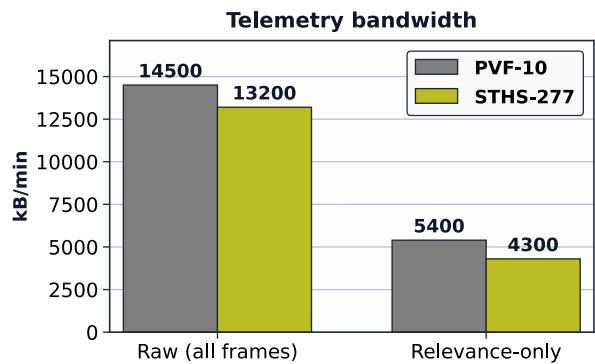


Fig. 12. Bandwidth savings achieved by the relevance-only telemetry approach

Finally, to situate the work within the broader landscape of deep learning detectors, a comparative analysis was conducted against several well-known methods. For a fair comparison, each model was re-trained on the PVF-10 dataset using an identical training budget. The results in Table 5 are telling. The system achieves a significantly higher mAP@0.5 (0.903) than all competitors and does so with a remarkably efficient model. With only 8.2 million parameters, it is far more compact than the heavyweight YOLOv3 (61.9M) and achieves the fastest inference speed on an embedded NVIDIA Jetson platform at 24 frames per second (FPS). This combination of superior accuracy and high efficiency underscores its suitability for deployment on resource-constrained onboard computers.

Table 5  
Comparison with representative detection methods  
on the PVF-10 dataset

| Method                   | mAP @0.5 | Params (M) | FPS (Jetson) |
|--------------------------|----------|------------|--------------|
| Region-CNN (Aerial) [15] | 0.760    | 14.8       | 12           |
| IR-only YOLOv3 [16]      | 0.800    | 61.9       | 18           |
| EfficientDet-D1 [20]     | 0.820    | 6.6        | 22           |
| Ours (T + RGB)           | 0.903    | 8.2        | 24           |

## 4. Discussion

### 4.1. Interpretation of findings and system novelty

The core novelty of this study lies in its integrated approach to intelligent, multi-modal data fusion and robust feature engineering. The most significant finding is the profound impact of enforcing palette invariance in the thermal feature encoding stage. The +4.6 percentage point gain in mAP@0.5 attributable to this component (see Table 2) is not merely a statistical improvement; it represents a fundamental step toward creating a reliable diagnostic system. By forcing the network to learn a canonical representation that is stable across different color maps, a robust thermal feature was produced to serve as a consistent input for the fusion stage. This contrasts sharply with conventional models that are implicitly tied to specific camera settings and palettes, making them brittle in heterogeneous hardware environments [10].

This robust feature engineering enables the second pillar of this study's success: the intelligent fusion of thermal and visual data. The gated fusion mechanism

provides a more sophisticated approach than simple feature concatenation, allowing the model to learn a dynamic, context-dependent weighting of the two modalities. For instance, in a scene with heavy solar glare, the gate can learn to down-weight corrupted RGB features and rely more heavily on stable thermal data. This is crucial because by fusing thermal data with visual evidence, the system can better differentiate between different types of anomalies (e.g., hotspots caused by soiling versus those linked to physical damage like cracks), achieving a level of analytical depth absent in single-modality solutions. This multi-modal, cross-validating approach moves beyond the static analysis paradigm common in the literature [15, 16] toward a more dynamic and intelligent inspection workflow.

### 4.2. Operational advantages and O&M optimization

The most operationally significant contribution of this study is the transformation of raw, redundant detections into clean, actionable intelligence. While academic research often focuses on per-frame metrics like mAP, O&M teams in the field prioritize the identification and localization of unique, actionable problems. The geo-spatial de-duplication module, detailed in Section 2.4, directly addresses this need. The 12–15 percentage point reduction in Dup-FPs (see Fig. 6) translates directly into saved man-hours and increased operator trust. This provides a robust algorithmic solution to the data redundancy problem, validating the high-level mapping concepts discussed in prior works [32, 33].

Although gradient-based XAI methods were not the focus of this study, the system enhances operator trustworthiness through “operational explainability.” By transmitting the specific high-resolution RGB image that confirmed a thermal anomaly (via the re-acquisition loop), the system provides the operator with the visual evidence (the “reason”) for the alert, effectively bridging the gap between black-box AI detections and human verification requirements.

Furthermore, the automated classification of specific defect types significantly enhances O&M efficiency. Traditional alarm systems may simply indicate that a module is underperforming. The proposed system provides immediately actionable intelligence, such as identifying a hotspot defect consistent with soiling or a crack defect at specific coordinates. This allows the O&M manager to dispatch the correct resource immediately; for example, a cleaning crew for the former or a technical team with replacement modules for the latter. This targeted response minimizes module downtime, reduces wasted

labor, and effectively transitions the O&M strategy from a reactive model to a highly efficient, condition-based, and predictive model. When combined with relevance-only telemetry (Section 2.5), the result is a system that is accurate, efficient, and capable of producing outputs directly consumable by asset management platforms.

#### 4.3. Limitations, challenges, and mitigation strategies

Despite its significant strengths, the proposed system has certain limitations that warrant careful consideration. The accuracy of the geo-projection and de-duplication is critically dependent on high-quality UAV pose metadata, ideally from an RTK-enabled GPS. In environments where GPS signals are degraded, projection errors could lead to incorrect clustering. A mitigation strategy involves integrating visual-inertial odometry or SLAM-based localization to supplement or replace GPS data, thereby improving navigational robustness.

Technically, system performance is fundamentally capped by the quality of input data and the underlying AI model's accuracy, adhering to the principle that output quality depends on input quality. Poorly calibrated sensors or an inadequately trained AI model will lead to flawed decisions. This necessitates rigorous sensor calibration protocols and a continuous machine learning operations cycle, including periodic retraining with new, verified data to prevent model drift. Furthermore, the field study noted that the few missed detections (contributing to the 96% recall) were often related to subtle soiling, suggesting that future fusion models could benefit from incorporating additional sensor modalities or explicitly modeling environmental factors.

In terms of functional safety, the deployment of autonomous UAVs for close-proximity inspection introduces risks related to battery reliability and flight over energized infrastructure. While the proposed method optimizes the logical inspection process, it does not inherently address the physical safety constraints defined in standards such as IEC 62443 or specific aviation regulations. Future implementations must integrate the proposed active perception loop with certified collision avoidance systems and battery health monitoring to ensure safe operation in gigawatt-scale plants.

### 5. Conclusions and future work

This study successfully validated an integrated method for UAV-based PV inspection, achieving the stated goal of enhancing detection fidelity while optimizing bandwidth. The primary scientific

contribution is the palette-invariant thermal embedding, which creates a universal representation robust to sensor variations. Experimental results confirmed a measurable improvement, with mAP@0.5 reaching 0.903 on the PVF-10 benchmark, exceeding single-modality baselines by 12–16 percentage points. Additionally, the Haversine-DBSCAN geo-spatial de-duplication module reduced the duplicate-induced false positive rate by 12–15%, directly addressing the data redundancy challenge. Field validation confirmed a high 96% recall rate, while the relevance-only telemetry policy reduced bandwidth requirements by over 60%, delivering a clean, actionable defect inventory. However, the method's performance is currently limited by its dependence on high-precision RTK-GPS data for accurate geo-projection and the absence of integrated functional safety mechanisms for autonomous close-proximity flight.

Future research will focus on integrating historical data from SCADA systems to dynamically adjust detection thresholds based on seasonality and module age. Additionally, incorporating survival analysis networks into the pipeline will enable the transition from defect detection to predictive remaining useful life (RUL) forecasting.

**Contributions of authors:** Contributions of authors: development of model and machine learning pipeline – **Andrii Lysyi**; conceptualization, methodology – **Andrii Lysyi, Anatoliy Sachenko, and Oleksandr Melnychenko**; system architecture, software engineering – **Diana Zahorodnia, Pavlo Radiuk**; validation, analysis of results – **Andrii Lysyi, Mykola Lysyi**; writing – original draft preparation – **Andrii Lysyi, Anatoliy Sachenko**; writing – review and editing – **Andrii Lysyi, Anatoliy Sachenko, Pavlo Radiuk, Mykola Lysyi, Oleksandr Melnychenko, and Diana Zahorodnia**.

### Conflict of Interest

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

### Financing

This research was supported by the Ministry of Education and Science of Ukraine and funded by the European Union's external assistance instrument as part of Ukraine's commitments under the EU Method Program for Research and Innovation "Horizon 2020." This work was performed within the method of the scientific research project titled "Intelligent System for Recognizing Defects in Green Energy Facilities Using UAVs," state registration number 0124U004665 (2024–2026).

### Data Availability

The benchmark datasets used (PVF-10, STHS-277) are publicly available. Analysis scripts can be made available by the corresponding author upon reasonable request. The proprietary dataset captured during field validation is not publicly available due to commercial sensitivities.

### Use of Artificial Intelligence

The authors confirm that no artificial intelligence was used in the generation of the textual content of this manuscript. The AI models discussed within the paper are the subject of the research itself and were developed by the authors.

### Acknowledgments

The authors gratefully acknowledge the support of the Ministry of Education and Science of Ukraine. This research was funded by the European Union's external assistance instrument under the Method Program for Research and Innovation "Horizon 2020." The work was conducted as part of the project titled "Intelligent System for Recognizing Defects in Green Energy Facilities Using UAVs" (state registration No. 0124U004665).

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

### References

1. Tsanakas, J., Ha, L., & Buerhop, C. Faults and infrared thermographic diagnosis in operating c-Si photovoltaic modules: a review of research and future challenges. *Renewable and Sustainable Energy Reviews*, 2016, vol. 62, pp. 695–709. DOI: 10.1016/j.rser.2016.04.079.
2. Grimaccia, F., Leva, S., Dolara, A., & Aghaei, M. Survey on PV modules common faults after an O&M flight extensive campaign over different plants in Italy. *IEEE Journal of Photovoltaics*, 2017, vol. 7, no. 3, pp. 810–816. DOI: 10.1109/JPHOTOV.2017.2674977.
3. Kandeal, A., Elkadeem, M., Thakur, A., Abdelaziz, G., Sathyamurthy, R., Kabeel, A., Yang, N., & Sharshir, S. Infrared thermography-based condition monitoring of solar photovoltaic systems: a mini review of recent advances. *Solar Energy*, 2021, vol. 223, pp. 33–43. DOI: 10.1016/j.solener.2021.05.032.
4. Gallardo-Saavedra, S., Hernández-Callejo, L., & Duque-Pérez, O. Technological review of the instrumentation used in aerial thermographic inspection of photovoltaic plants. *Renewable and Sustainable Energy Reviews*, 2018, vol. 93, pp. 566–579. DOI: 10.1016/j.rser.2018.05.027.
5. Zefri, Y., ElKettani, A., Sebari, I., & Ait Lamallam, S. Thermal infrared and visual inspection of photovoltaic installations by UAV photogrammetry—application case: Morocco. *Drones*, 2018, vol. 2, no. 4, p. 41. DOI: 10.3390/drones2040041.
6. Melnychenko, O., Scislo, L., Savenko, O., Sachenko, A., & Radiuk, P. Intelligent integrated system for fruit detection using multi-UAV imaging and deep learning. *Sensors*, 2024, vol. 24, no. 6, p. 1913. DOI: 10.3390/s24061913.
7. Michail, A., Livera, A., Tziolis, G., Georghiou, G., Panayiotou, C., Loutsiou, A., & Lilli, C. A comprehensive review of unmanned aerial vehicle-based approaches to support photovoltaic plant diagnosis. *Heliyon*, 2024, vol. 10, no. 1, p. e023983. DOI: 10.1016/j.heliyon.2024.e023983.
8. Lee, D., & Park, J. Developing inspection methodology of solar energy plants by thermal infrared sensor on board unmanned aerial vehicles. *Energies*, 2019, vol. 12, no. 15, article no. 2928. DOI: 10.3390/en12152928.
9. Aghaei, M., Dolara, A., Leva, S., & Grimaccia, F. Image resolution and defects detection in PV inspection by unmanned technologies. In: *Proceedings of the 2016 IEEE Power and Energy Society General Meeting (PESGM)*. Boston, MA, USA, IEEE, 2016, pp. 1–5. DOI: 10.1109/PESGM.2016.7741605.
10. Buerhop-Lutz, C., Bommes, L., Schlipf, J., Pickel, T., Fladung, A., & Peters, I. Infrared imaging of photovoltaic modules: A review of the state of the art and future challenges facing gigawatt-scale PV power stations. *Progress in Energy*, 2022, vol. 4, no. 4, article no. 042010. DOI: 10.1088/2516-1083/ac890b.
11. Phoolwani, U., Sharma, T., Singh, A., & Gawre, S. IoT based solar panel analysis using thermal imaging. In: *Proceedings of the 2020 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)*. Bhopal, India, IEEE, 2020, pp. 1–6. DOI: 10.1109/SCEECS48394.2020.9114.
12. Lysyi, A., Sachenko, A., Radiuk, P., Lysyi, M., Melnychenko, O., Ishchuk, O., & Savenko, O. Enhanced fire hazard detection in solar power plants: An integrated UAV, AI, and SCADA-based approach. *Radioelectronic and Computer Systems*, 2025, vol. 2025, no. 2, pp. 99–117. DOI: 10.32620/reks.2025.2.06.
13. Morando, L., Recchiuto, C., Calla, J., Scuteri, P., & Sgorbissa, A. Thermal and visual tracking of photovoltaic plants for autonomous UAV inspection. *Drones*, 2022, vol. 6, no. 11, article no. 347. DOI: 10.3390/drones6110347.
14. Svystun, S., Scislo, L., Pawlik, M., Melnychenko, O., Radiuk, P., Savenko, O., & Sachenko, A. DyTAM: Accelerating wind turbine inspections with dynamic UAV trajectory adaptation. *Energies*, 2025, vol. 18, no. 7, article no. 1823. DOI: 10.3390/en18071823.
15. Vlaminck, M., Heidbuchel, R., Philips, W., & Luong, H. Region-based CNN for anomaly detection in

PV power plants using aerial imagery. *Sensors*, 2022, vol. 22, no. 3, article no. 1244. DOI: 10.3390/s22031244.

16. Di Tommaso, A., Genduso, F., Miceli, R., & Galluzzo, G. A multi-stage model based on YOLOv3 for defect detection in PV panels using IR and visible imaging by UAV. *Renewable Energy*, 2022, vol. 193, pp. 941–962. DOI: 10.1016/j.renene.2022.04.046.

17. Su, Y., Tao, F., Jin, J., & Zhang, C. Automated overheated region object detection of photovoltaic module with thermography image. *IEEE Journal of Photovoltaics*, 2021, vol. 11, no. 2, pp. 535–544. DOI: 10.1109/JPHOTOV.2020.3045680.

18. Duranay, Z. B. Fault detection in solar energy systems: a deep learning perspective. *Electronics*, 2023, vol. 12, no. 21, article no. 4397. DOI: 10.3390/electronics12214397.

19. Meng, S., Yue, Y., & Xu, T. Enhanced YOLOv11 framework for accurate multi-fault detection in UAV photovoltaic inspection. *Sensors*, 2025, vol. 25, no. 17, article no. 5311. DOI: 10.3390/s25175311.

20. Tan, M., Pang, R., & Le, Q. EfficientDet: scalable and efficient object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Seattle, WA, USA, IEEE, 2020, pp. 10778–10787. DOI: 10.1109/CVPR42600.2020.01079.

21. Akram, M., Li, G., Jin, Y., Chen, X., Zhu, H., Ahmad, A., & Riaz, F. Automatic detection of photovoltaic module defects in infrared images with deep learning. *Solar Energy*, 2020, vol. 198, pp. 175–186. DOI: 10.1016/j.solener.2020.01.055.

22. Zefri, Y., Sebari, I., Hajji, H., & Aniba, G. Developing a deep learning-based layer-3 solution for thermal infrared large-scale photovoltaic module inspection from orthorectified big UAV imagery data. *International Journal of Applied Earth Observation and Geoinformation*, 2022, vol. 106, article no. 102652. DOI: 10.1016/j.jag.2021.102652.

23. Jia, Y., Chen, G., & Zhao, L. Defect detection of photovoltaic modules based on improved VarifocalNet. *Scientific Reports*, 2024, vol. 14, article no. 15170. DOI: 10.1038/s41598-024-66234-3.

24. Wang, B., Chen, Q., Wang, M., Chen, Y., Zhang, Z., Liu, X., Gao, W., Zhang, Y., & Zhang, H. PVF-10: a high-resolution unmanned aerial vehicle thermal infrared image dataset for fine-grained photovoltaic fault classification. *Applied Energy*, 2024, vol. 376, article no. 124187. DOI: 10.1016/j.apenergy.2024.124187.

25. Alfaro-Mejía, E., Loaiza-Correa, H., Franco-Mejía, E., Restrepo-Girón, A., & Nope-Rodríguez, S. Dataset for recognition of snail trails and hot spot failures in monocrystalline Si solar panels. *Data in Brief*, 2019, vol. 26, article no. 104441. DOI: 10.1016/j.dib.2019.104441.

26. Qureshi, U., Rashid, A., Altini, N., Bevilacqua, V., & La Scala, M. Explainable intelligent inspection of solar photovoltaic systems with deep transfer learning: considering warmer weather effects using aerial radiometric infrared thermography. *Electronics*, 2025, vol. 14, no. 4, article no. 755. DOI: 10.3390/electronics14040755.

27. Oulefki, A., Agaian, S., El Afou, Y., Djahel, S., Zenkouar, K., & Taleb-Ahmed, A. Detection and analysis of deteriorated areas in solar PV modules using unsupervised sensing algorithms and 3D augmented reality. *Heliyon*, 2024, vol. 10, no. 6, article no. e27973. DOI: 10.1016/j.heliyon.2024.e27973.

28. Liao, K.-C., & Lu, J.-H. Using UAV to detect solar module fault conditions of a solar power farm with IR and visual image analysis. *Applied Sciences*, 2021, vol. 11, no. 4, article no. 1835. DOI: 10.3390/app11041835.

29. Svystun, S., Melnychenko, O., Radiuk, P., Savenko, O., Sachenko, A., & Lysyi, A. Thermal and RGB images work better together in wind turbine damage detection. *International Journal of Computing*, 2024, vol. 23, no. 4, pp. 526–535. DOI: 10.47839/ijc.23.4.3752.

30. Rohith, G., Rajalakshmi, R., Manish, D., & Narasimhan, R. Fusion-Solar-Net for solar panel fault detection. *Results in Engineering*, 2025, vol. 27, article no. 106513. DOI: 10.1016/j.rineng.2025.106513.

31. Lai, Y.-S., Hsieh, C.-C., Liao, T.-W., Huang, C.-Y., Yeh, C.-H., & Chen, W.-H. Deep learning-based automatic defect detection of photovoltaic modules in infrared, electroluminescence, and red–green–blue images. *Energy Conversion and Management*, 2025, vol. 332, article no. 119783. DOI: 10.1016/j.enconman.2025.119783.

32. Niccolai, A., Grimaccia, F., & Leva, S. Advanced asset management tools in photovoltaic plant monitoring: UAV-based digital mapping. *Energies*, 2019, vol. 12, no. 24, article no. 4736. DOI: 10.3390/en12244736.

33. Bommes, L., Pickel, T., Buerhop-Lutz, C., Hauch, J., Brabec, C., & Peters, I. Computer vision tool for detection, mapping, and fault classification of photovoltaics modules in aerial IR videos. *Progress in Photovoltaics: Research and Applications*, 2021, vol. 29, no. 12, pp. 1236–1251. DOI: 10.1002/pip.3448.

34. Kolahi, M., Esmailifar, S. M., Moradi Sizkouhi, A. M., & Aghaei, M. Digital-PV: a digital twin-based platform for autonomous aerial monitoring of large-scale photovoltaic power plants. *Energy Conversion and Management*, 2024, vol. 321, article no. 118963. DOI: 10.1016/j.enconman.2024.118963.

35. Kishor, I., Mamodiya, U., Patil, V., Naik, N., Kumar, S., & Yadav, R. AI-integrated autonomous robotics for solar panel cleaning and predictive maintenance using drone and ground-based systems. *Scientific Reports*, 2025, vol. 15, article no. 32187. DOI: 10.1038/s41598-025-17313-6.

36. DJI. *Matrice 300 RTK* [Computer program], 2020. Available at: <https://www.dji.com/>

global/support/product/matrice-300 (accessed 3 September 2025).

37. DJI. *Zenmuse H20T* [Computer program], 2020. Available at: <https://enterprise.dji.com/zenmuse-h20-series> (accessed 1 September 2025).

38. NVIDIA Corporation. *Jetson AGX Orin 32GB Module* [Computer program], 2022. Available at: <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/> (accessed 1 September 2025).

39. Canonical Ltd. *Ubuntu* [Computer program]. Version 20.04.6 LTS (Focal Fossa), 2022. Available at: <https://releases.ubuntu.com/focal/> (accessed 4 September 2025).

40. Open Robotics. *ROS Noetic Ninjemys* [Computer program], 2020. Available at: <https://www.openrobotics.org/> (accessed 3 September 2025).

41. DJI. *DJI Onboard SDK* [Computer program]. Version 4.1.0, 2021. Available at: <https://github.com/dji-sdk/Onboard-SDK> (accessed 30 September 2025).

42. DJI. *DJI Payload SDK* [Computer program]. Version 3.12.0, 2025. Available at: <https://developer.dji.com/payload-sdk/> (accessed 3 September 2025).

43. Python Software Foundation. *Python 3 Programming Language* [Computer program]. Version 3.12.0, 2023. Available at: <https://www.python.org/> (accessed 22 September 2025).

44. OpenCV Development Team. *OpenCV* [Computer program]. Version 4.9.0, 2023. Available at: <https://opencv.org/> (accessed 30 September 2025).

45. NumPy Developers. *NumPy* [Computer program]. Version 2.3.0, 2025. Available at: <https://numpy.org/> (accessed 2 September 2025).

46. PyTorch Foundation. *PyTorch* [Computer program]. Version 2.0, 2023. Available at: <https://pytorch.org/> (accessed 1 September 2025).

47. Ultralytics. *YOLOv11m-seg Architecture* [Computer program], 2024. Available at: <https://docs.ultralytics.com/models/yolo11/> (accessed 31 September 2025).

48. ZeroMQ Project. *ZeroMQ* [Computer program]. Version 4.3.5, 2023. Available at: <https://zeromq.org/> (accessed 1 September 2025).

49. Microsoft Corporation. *Microsoft Azure Cloud Platform* [Computer program], 2010. Available at: <https://azure.microsoft.com/> (accessed 1 September 2025).

50. ZigBee Alliance. *ZigBee Specification* [Standard]. Document No. 053474r20, 2005. Available at: <https://zigbeealliance.org/> (accessed 3 September 2025).

51. Rainio, O., Teuho, J., & Klén, R. Evaluation metrics and statistical tests for machine learning. *Scientific Reports*, 2024, vol. 14, no. 1, article no. 6086. DOI: 10.1038/s41598-024-56706-x.

Received 17.08.2025, Received in revised form 15.10.2025

Accepted date 17.11.2025, Published date 08.12.2025

## МЕТОД ІНСПЕКТУВАННЯ ФОТОЕЛЕКТРИЧНИХ МОДУЛІВ З БПЛА ІЗ ВИКОРИСТАННЯМ ЗЛІТТЯ ТЕПЛОВІЗІЙНИХ ТА RGB-ДАНИХ

**A. M. Лисий, A. O. Саченко, P. M. Радюк, M. I. Лисий,  
O. B. Мельниченко, D. I. Загородня**

Предметом статті є проєктування та експериментальне оцінювання інтелектуальної гранично-хмарної кіберфізичної системи автоматизованого інспектування фотоелектричних модулів (ФЕМ) на промислових сонячних електростанціях на основі мультипалігрових тепловізійних інфрачервоних та RGB-зображень, отриманих за допомогою безпілотних літальних апаратів (БПЛА). Метою цього дослідження є підвищення точності виявлення дефектів ФЕМ та операційної ефективності через розроблення нового методу, який пітвторює дані сенсорів із перекріттям на компактний, геоприв'язаний реєстр дефектів, що підлягають усуненню. Мету дослідження досягнуто через отримане підвищення середньої точності виявлення (mAP) та зниження частоти хибних спрацьовувань, що спричинені надмірністю даних, з одночасною мінімізацією використання пропускної здатності каналу зв'язку. Завдання, які розв'язано в роботі: розроблено інваріантне до палітри теплове подання, що усуває залежність від псевдокольорової візуалізації та внутрішніх параметрів камери; об'єднано стійкий тепловий потік із контрастно покращеними RGB-даними за допомогою адаптивного механізму шлюзованого злиття, здатного знижувати вагу ненадійних модальностей; реалізовано бортовий контур активного сприйняття, який переорієнтує підвіс БПЛА для отримання додаткових ракурсів неоднозначних аномалій малої площини; розроблено модуль геопросторової кластеризації та дедублікації, що об'єднує повторні виявлення в унікальні події несправностей, придатні для інтеграції з інструментами SCADA та ГІС; кількісно оцінено переваги запропонованої архітектури на публічних наборах даних та в реальних польових випробуваннях у порівнянні з одномодальними базовими підходами. Використані в дослідженнях методи та алгоритми дозволяють отримати високоточні результати інспектування ФЕМ з використанням дешевих та надійних тепловізійних та RGB-камер, що може зробити цю технологію доступною для широкого використання в енергетичній сфері.

ліджені методи включають глибокі згорткові нейронні мережі на основі магістралі семантичної сегментації YOLOv11m-seg, навченої із членом регуляризації консистентності палітри, шлізоване злиття теплових та RGB-ембедингів на рівні ознак, розрахунок коригувальних поворотів підвісу на основі формули Родрігеса для адаптивного повторного захоплення, щільнісну просторову кластеризацію з метрикою гаверсинуса для географічної дедублікації виявлень, а також статистичний аналіз ефективності з використанням mAP, макро-усередненої F1-міри, повноти та індикатора хибних спрацьовувань, що спричинені дублюванням. Отримано такі результати: на наборі даних PVF-10 запропонована система досягає mAP@0.5 = 0.903, перевищуючи лише теплові та лише RGB-детектори на 12–16 відсоткових пунктів; на наборі даних STHS-277 вона досягає mAP@0.5 = 0.887; інваріантне до палітри навчання разом з адаптивним повторним захопленням підвищують повноту виявлення малих цілей до 0.86; геопросторова кластеризація знижує частоту хибних спрацьовувань через дублювання на 12–15 відсоткових пунктів; польова валідація на дахових та наземних електростанціях підтверджує 96 % повноти з низьким середньоквадратичним відхиленням між автоматичним та ручним підрахунком дефектів; а передача виключно релевантної телеметрії зменшує обсяг переданих із борту даних на 60–67 % при збереженні діагностичної точності. У підсумку, наукова новизна отриманих результатів полягає в уніфікованій, інваріантній до палітри, мультимодальній гранично-хмарній кіберфізичній архітектурі, яка поєднує сенсорику БПЛА, активне сприйняття, геопросторовий аналіз та передачу даних з урахуванням пропускної здатності в єдиний операційний метод інспектування фотоелектричних модулів, створюючи масштабовану основу для обслуговування великих сонячних електростанцій за технічним статном.

**Ключові слова:** фотоелектричні модулі; інспектування з БПЛА; виявлення дефектів; тепловізійна інфрачервона зйомка; RGB-зображення; глибоке навчання; мультимодальне злиття; інваріантність до палітри; геопросторова кластеризація.

**Лисий Андрій Миколайович** – асп. каф. комп’ютерної інженерії та інформаційних систем, Хмельницький національний університет, Хмельницький, Україна.

**Саченко Анатолій Олексійович** – д-р техн. наук, проф., директор Науково-дослідного інституту інтелектуальних комп’ютерних систем, Західноукраїнський національний університет, Тернопіль, Україна; Радомський університет імені Казімежа Пуласького, Радом, Польща.

**Радюк Павло Михайлович** – д-р філос., доц. каф. комп’ютерних наук, Хмельницький національний університет, Хмельницький, Україна.

**Лисий Микола Іванович** – д-р техн. наук, проф., Національна академія Державної прикордонної служби України імені Богдана Хмельницького, Хмельницький, Україна.

**Мельниченко Олександр Вікторович** – д-р філос., старш. викл. каф. комп’ютерної інженерії та інформаційних систем, Хмельницький національний університет, Хмельницький, Україна.

**Загородня Діана Іванівна** – канд. техн. наук, доц., доц. каф. інформаційно-обчислювальних систем і управління, Західноукраїнський національний університет, Тернопіль, Україна.

**Andrii Lysyi** – PhD Student, the Department of Computer Engineering and Information Systems, Khmelnytskyi National University, Khmelnytskyi, Ukraine,  
e-mail: Andrii.Lysyi1@gmail.com, ORCID: 0009-0001-0065-9740, Scopus Author ID: 59475081500.

**Anatoliy Sachenko** – Doctor of Technical Sciences, Professor, Director of the Research Institute for Intelligent Computer Systems, West Ukrainian National University, Ternopil, Ukraine; Kazimierz Pulaski University of Radom, Radom, Poland,  
e-mail: as@wunu.edu.ua, ORCID: 0000-0002-0907-3682, Scopus Author ID: 35518445600.

**Pavlo Radiuk** – PhD in Computer Science, Associate Professor, the Department of Computer Science, Khmelnytskyi National University, Khmelnytskyi, Ukraine,  
e-mail: radiukp@khmnu.edu.ua, ORCID: 0000-0003-3609-112X, Scopus Author ID: 57216894492.

**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: 0000-0002-9858-706X, Scopus Author ID: 59482244100.

**Oleksandr Melnychenko** – PhD in Computer Science, Senior Lecturer, the Department of Computer Engineering and Information Systems, Khmelnytskyi National University, Khmelnytskyi, Ukraine,  
e-mail: melnychenko@khmnu.edu.ua, ORCID: 0000-0001-8565-7092, Scopus Author ID: 45961382900.

**Diana Zahorodnia** – Candidate of Technical Sciences, Associate Professor, Associate Professor at the Department of Information and Computing Systems and Control, West Ukrainian National University, Ternopil, Ukraine,  
e-mail: dza@wunu.edu.ua, ORCID: 0000-0002-9764-3672, Scopus Author ID: 54421527700.