University of the Built Environment
EST. 1919

Patron: His Majesty King Charles III

Horizons, 60 Queen's Road, Reading, RG1 4BS
ube.ac.uk | +44(0)118 921 4696 | enquiries@ube.ac.uk

University of the Built Environment
Institutional Research Repository

Title	An Industry-Ready Deep Learning Framework for Power Loss Prediction in Half-Cut Photovoltaic Modules
Author(s)	mahmoud Dhimish and Angela Lee
ORCID	https://orcid.org/; https://orcid.org/0000-0003-0769-5215
Type	Article
Publication title	Solar Compass
Publisher	Elsevier
ISSN/ISBN	2772-9400
Date	10 November 2026
Version	
DOI	
Repository link	https://ube.repository.guildhe.ac.uk/id/eprint/310/
Link to publication	https://www.sciencedirect.com/journal/solar-compas...
Notes	This publication version may differ from the final version

Copyright:
The University aims to make research outputs available to a broader audience via its digital Repository. Where copyright permits, full text material held in the Repository is made freely available.
URLs from GuildHE Research Repositories maybe freely distributed and linked to. Please refer to each manuscript for any further copyright restrictions.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page, and the content is not changed in any way.

An Industry-Ready Deep Learning Framework for Power Loss Prediction in Half-Cut Photovoltaic Modules

Mahmoud Dhimish[1,2] and Angela Lee[2]

[1] Department of Electrical and Photonics Engineering, Technical University of Denmark, 4000 Roskilde, Sjælland, Denmark

[2] University of the Built Environment, Horizons, 60 Queen's Road, Reading, RG1 4BS, UK

Corresponding author, Mahmoud Dhimish, m.dhimish@ube.ac.uk or mahdh@dtu.dk

Abstract

Electroluminescence (EL) imaging is widely used for photovoltaic (PV) module inspection, but in industrial practice it remains largely qualitative and operator dependent. This limitation is particularly relevant for half-cut PV modules, where complex current redistribution means that defect morphology does not directly indicate electrical severity. This study proposes a hybrid two-stage deep learning and physics-based framework for quantitative power-loss prediction from EL images. The first stage segments cracks and low-luminescence regions, while the second translates the extracted spatial defect features into module-level power loss. The predicted degradation is then coupled with a single-diode PV model to reconstruct the corresponding current–voltage characteristics under standard test conditions. Using a dataset of 500 modules, the segmentation model achieved a mean Dice coefficient of 0.86, while the power-loss model achieved an R^2 of 0.93 and a mean absolute error of 0.0183. Reconstructed current–voltage curves showed maximum-power deviations below 5.5% in representative cases. The framework therefore links EL-visible defect morphology directly to electrical severity, supporting objective pre-installation assessment of half-cut PV modules.

Keywords: Deep learning; Defect segmentation; Electroluminescence imaging; Photovoltaic modules.

1. Introduction

The global transition toward low-carbon buildings has accelerated the integration of energy-generating materials within the building envelope. Among these, building-integrated photovoltaic (BIPV) systems [1,2], shown in Fig. 1(a), represent a class of multifunctional façade elements that simultaneously serve structural, environmental, and energy-generating roles. Unlike conventional rooftop photovoltaic (PV) installations, BIPV modules are embedded directly into façades, curtain walls, skylights, and shading systems, thereby acting as laminated composite building materials rather than external add-ons [3]. Their performance and durability therefore influence not only energy yield but also the serviceability, safety, and sustainability of the built environment.

Modern PV façade modules are multilayer laminated composites composed of tempered glass, encapsulant polymers (e.g., ethylene-vinyl acetate [4]), crystalline silicon cells, back sheet materials, and metallic interconnections, as presented in Fig. 1(b). These layers are thermally bonded to form mechanically stable and environmentally resistant panels. Recent advancements in PV architecture, particularly the widespread adoption of half-cut cell designs (Fig. 1(c)), have improved module efficiency and shade tolerance by electrically dividing full cells into two independent sub-cells connected in parallel strings [5]. This configuration reduces resistive losses and enhances performance under partial shading; however, it also modifies current pathways and alters crack sensitivity behavior within the silicon layer. As a result, damage mechanisms in half-cut modules cannot be directly extrapolated from traditional full-cell systems [6].

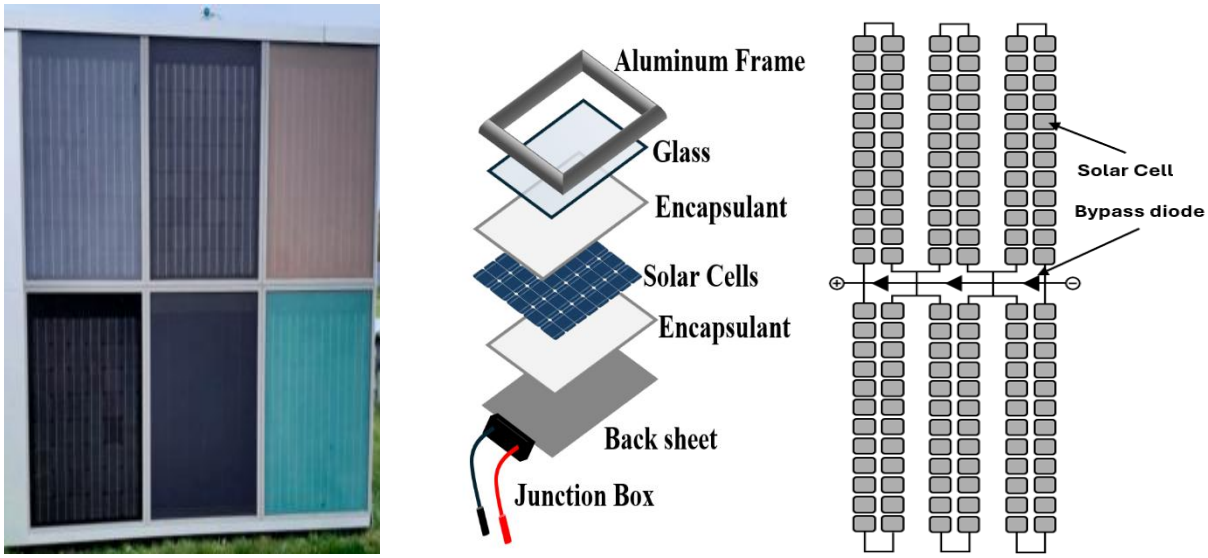

(a) (b) (c)

Figure 1. (a) BIPV system at the Technical University of Denmark; (b) Structure of a crystalline silicon PV module; (c) Moden half-cut PV modules architecture.

From a materials engineering perspective, microcracks in the silicon layer represent a critical damage mechanism in laminated PV composites [7]. These cracks may originate during wafer sawing, cell handling, lamination, transportation, or installation. While some microcracks remain electrically inactive, others propagate under thermo-mechanical loading induced by wind pressure [8], façade deflection, thermal cycling, or mechanical stresses during installation [9]. Electrically active cracks disrupt current continuity, create localized recombination zones, and may accelerate degradation through hotspot formation. In façade-integrated systems, such degradation has broader implications: reduced energy performance, increased maintenance frequency, and premature PV module replacement [10].

The sustainability dimension of BIPV materials therefore extends beyond operational energy generation. The embodied energy and carbon footprint associated with silicon purification, cell fabrication, glass production, and lamination are substantial. Premature failure or early replacement of defective modules undermines lifecycle carbon reduction targets and contributes to material waste [11]. Consequently, robust pre-installation quality control and early-stage damage assessment are essential to ensure long-term façade reliability and to maximize lifecycle sustainability benefits.

In laminated BIPV façade modules, damage within the silicon cell layer is often electrically critical but mechanically subtle. Microcracks, inactive cell regions, and interconnect disruptions may not be detectable through standard visual inspection or surface-based assessment. Electroluminescence (EL) imaging has therefore emerged as a specialized functional non-destructive testing technique for PV modules [12-15]. Under forward electrical bias, intact silicon junctions emit near-infrared radiation, whereas cracked or electrically isolated regions exhibit reduced or absent luminescent response [16], as demonstrated in Fig. 2. EL imaging thus enables high-resolution mapping of electrically active and inactive regions within the encapsulated cell layer without physically separating the laminate. Importantly, the technique assesses functional integrity rather than solely mechanical damage.

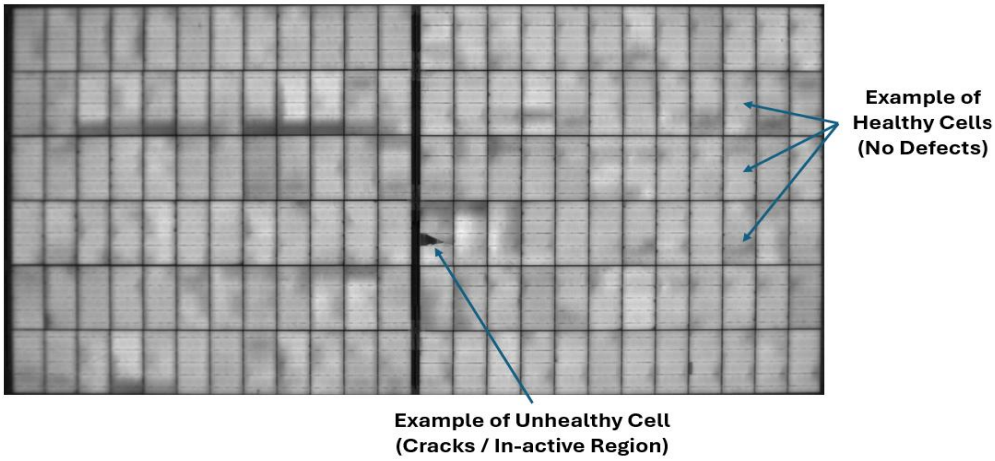

Figure 2. Typical EL image of half-cut PV module, with example markers on healthy (non-defective) and unhealthy (defective) solar cells.

Despite its sensitivity, EL-based inspection remains predominantly qualitative in industrial practice. Damage interpretation is often manual, reliant on operator expertise, and limited to visual grading schemes that do not systematically quantify defect morphology or its impact on module performance. Furthermore, conventional assessment approaches rarely consider the influence of evolving PV architectures, such as half-cut cell configurations, on defect criticality [17]. As laminated PV elements become increasingly integrated into building envelopes, there is a growing need to transform EL imaging from a visual diagnostic tool into a quantitatively rigorous non-destructive testing framework capable of supporting objective grading, certification, and sustainability-driven decision-making.

Although EL imaging has become a standard inspection tool in PV module manufacturing and laboratory diagnostics, its practical implementation remains largely qualitative. In industrial environments, EL images are typically assessed visually to identify macroscopic cracks, inactive regions, or dark areas indicating electrically isolated cell segments [18]. Such visual grading procedures are often based on heuristic criteria or simplified defect categorization schemes. While these approaches can identify severe defects, they lack reproducibility, quantitative consistency, and objective thresholds for performance-based decision-making [19].

Recent advances have introduced machine learning and deep learning techniques for automated defect detection in EL images. Convolutional neural networks (CNNs) and segmentation architectures have demonstrated high accuracy in identifying cracks or classifying defect types [20,21]. However, most existing studies focus primarily on pixel-level segmentation performance or classification accuracy [22-24], rather than linking defect morphology to functional severity or certification thresholds. In other words, the majority of current artificial intelligence (AI)-based approaches answer the question: “Where is the crack?” but do not address the more critical engineering question: “How significant is this defect in terms of functional degradation and serviceability?” This distinction is particularly important when PV modules are used as laminated façade materials (i.e., in large-scale BIPV systems), where installation decisions and long-term sustainability depend on defect criticality rather than mere presence.

On the other hand, the widespread adoption of half-cut PV cell architecture introduces additional complexity that is insufficiently addressed in current literature [25]. In half-cut configurations, cells are electrically divided and connected in parallel substrings, altering current distribution paths and modifying crack sensitivity behavior. Certain crack orientations or positions may have limited electrical impact due to redundancy in current pathways, whereas others can isolate substantial active areas and induce disproportionate performance losses. Conventional crack detection frameworks generally treat defects as homogeneous features [26] without incorporating architectural topology into the assessment process.

This distinction is increasingly important as half-cut cell architectures are becoming widely adopted in contemporary crystalline-silicon PV modules, progressively replacing conventional full-cell designs in many applications. In a conventional full-cell module, crack-induced inactive areas directly influence the current-carrying capability of the affected cell and string. In half-cut modules, however, the divided-cell configuration and parallel substring arrangement create different current redistribution pathways. Consequently, defects with similar EL morphology or inactive areas may produce substantially different electrical losses depending on their location, orientation, and interaction with the module topology. Power-loss relationships developed for conventional full-cell modules therefore cannot necessarily be transferred directly to half-cut architectures, motivating a dedicated defect-to-performance prediction framework for this increasingly deployed module technology.

Therefore, to support sustainable deployment of laminated BIPV façade elements (particularly those with half-cut PV modules designs), EL inspections must evolve from qualitative toward a quantitative, architecture-aware evaluation framework. Such a framework should (i) extract reproducible morphological descriptors from EL images, (ii) relate these descriptors to functional degradation (i.e., accurate power loss estimation) behavior considering module topology, and (iii) provide objective grading metrics suitable for pre-installation certification and lifecycle management decisions. Achieving this transition requires the integration of advanced data-driven modelling with material- and architecture-specific knowledge of laminated composite behavior. Addressing this need forms the central motivation of the present study.

2. Research Questions and Paper Contributions

Based on the knowledge gap identified earlier, namely the absence of a quantitative and functionally grounded relationship between EL-derived defect morphology and electrical degradation in half-cut PV modules, the present study addresses the following research questions:

1) Can spatial defect information extracted from electroluminescence images be translated into accurate module-level power loss under standard test conditions?
2) Is it possible to establish a physically consistent mapping between image-derived defect patterns and electrical performance without directly using electrical parameters during inference?
3) Can a cascaded learning framework provide sufficient predictive reliability for industrial pre-installation screening of laminated half-cut PV façade modules?

To answer these questions, this work develops a physics-based, two-stage deep learning framework that transforms EL imaging from qualitative PV modules defect visualization into quantitative electrical performance prediction. In the first stage, a convolutional segmentation network identifies cracks and low-luminescence regions with high spatial

fidelity, producing structured defect maps that preserve morphological and contextual information across the module surface. In the second stage, these spatial representations are encoded into compact latent descriptors and used to predict module-level power loss through supervised regression trained against measured current–voltage (I–V) characteristics. The predicted degradation is subsequently translated into reconstructed I–V curves using a physics-based model.

Overall, the principal novelty of the proposed framework lies not solely in defect segmentation or electrical regression individually, but in their integration within a half-cut-module-specific workflow that links spatial EL defect morphology to quantitative power loss, physics-based I–V reconstruction, and decision-level module screening.

3. Materials and Methods

3.1. Laminated Half-Cut PV Modules

The experimental investigation was conducted on commercially available laminated half-cut PV modules representative of contemporary façade-integrated crystalline silicon systems. The modules employ monocrystalline p-type passivated emitter and rear contact (PERC) solar cells arranged in a 144-half-cell configuration. The modules consist of a multilayer laminated composite structure comprising tempered front glass (3.2mm thickness), ethylene-vinyl acetate (EVA) [27] encapsulant layers, silicon half-cells interconnected through metallic busbars and ribbons, a polymeric back sheet, and an aluminum frame. The electrical specifications of the PV modules are summarized in Table 1, and Fig. 3 shows the corresponding I–V curve at standard test conditions (STC), where irradiance 1000 W/m^2, cell temperature 25 °C, at air mass 1.5 global (AM1.5G).

Table 1. Electrical specifications of the PV modules used in this work; parameters under STC conditions.

Parameter	Value	Unit
Solar cell technology	Monocrystalline p-type PERC cells	n/a
Cell count	144	n/a
Junction box	IP68	n/a
Maximum power (P_{mp})	390	W
Open circuit voltage (V_{oc})	48.91	V
Maximum power voltage (V_{mp})	40.55	V
Short circuit current (I_{sc})	10.16	A
Maximum power current (I_{mp})	9.62	A
Temperature coefficient of I_{sc}	+0.051	%/°C
Temperature coefficient of V_{oc}	−0.289	%/°C
Maximum efficiency	20.4	%

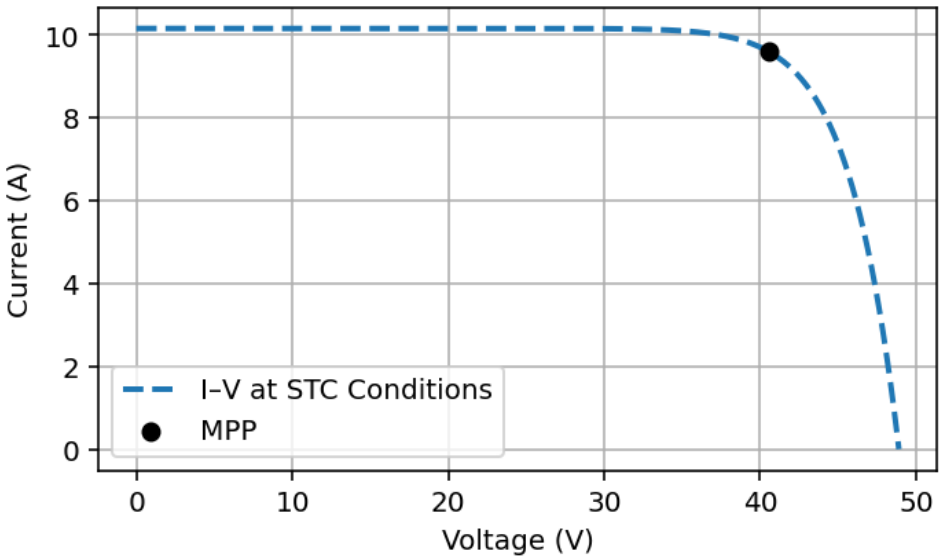

Figure 3. Theoretical I–V curve of the tested PV module at STC conditions.

3.2. EL Imaging Procedure

EL imaging was employed as a functional non-destructive testing technique to assess electrically active and inactive regions within the laminated half-cut PV modules described in Section A. EL imaging was conducted under forward-bias conditions, whereby a direct current (at $I_{sc} = 10.16\ A$) was injected into the module to stimulate radiative recombination in intact silicon junctions. Electrically continuous regions emit near-infrared radiation, whereas microcracks, inactive areas, and disrupted interconnections exhibit reduced or absent luminescent response [28].

EL images were acquired using a digital single-lens reflex (DSLR) camera (Nikon D3500, complementary metal–oxide–semiconductor sensor) with a native resolution of 6000 × 4000 pixels. The camera was mounted on a rigid tripod system and aligned normally to the plane of the module to ensure uniform spatial sampling and minimize perspective distortion. The distance between the camera and the module surface was fixed at 4.5 m for all experiments. All imaging was performed in a fully darkened laboratory environment to eliminate ambient light interference and maximize signal-to-noise ratio in the near-infrared emission range.

Camera acquisition parameters were standardized for all measurements to ensure reproducibility. The ISO sensitivity was set to 400, the aperture was fixed at f/2.5, and the exposure time was maintained at 5 seconds. These settings were selected to balance luminescent signal intensity with image sharpness while avoiding saturation of highly emissive regions. Long exposure times were necessary due to the relatively low photon emission intensity characteristic of EL imaging in modern half-cut crystalline silicon modules. The fixed acquisition protocol ensures that variations in image intensity arise from material and electrical behavior rather than changes in camera configuration.

3.3. Deep Learning-Based Defect Detection and Power Loss Prediction Framework

The overall methodological framework developed in this work is illustrated in Fig. 4. The proposed approach is designed as a hybrid two-stage deep learning and physics-based pipeline linking EL imaging with quantitative estimation of module-level power loss under STC conditions. The deep learning component performs spatial defect segmentation and subsequent power-loss regression, while physical knowledge is incorporated through the single-diode PV model used to reconstruct the predicted I–V characteristic from the estimated degradation. Thus, no governing equation or physics-based constraint is imposed directly within the neural-network loss functions; instead, the physics-based component operates as the final electrical reconstruction stage, ensuring that the predicted degradation is translated into an electrical response consistent with PV operating principles.

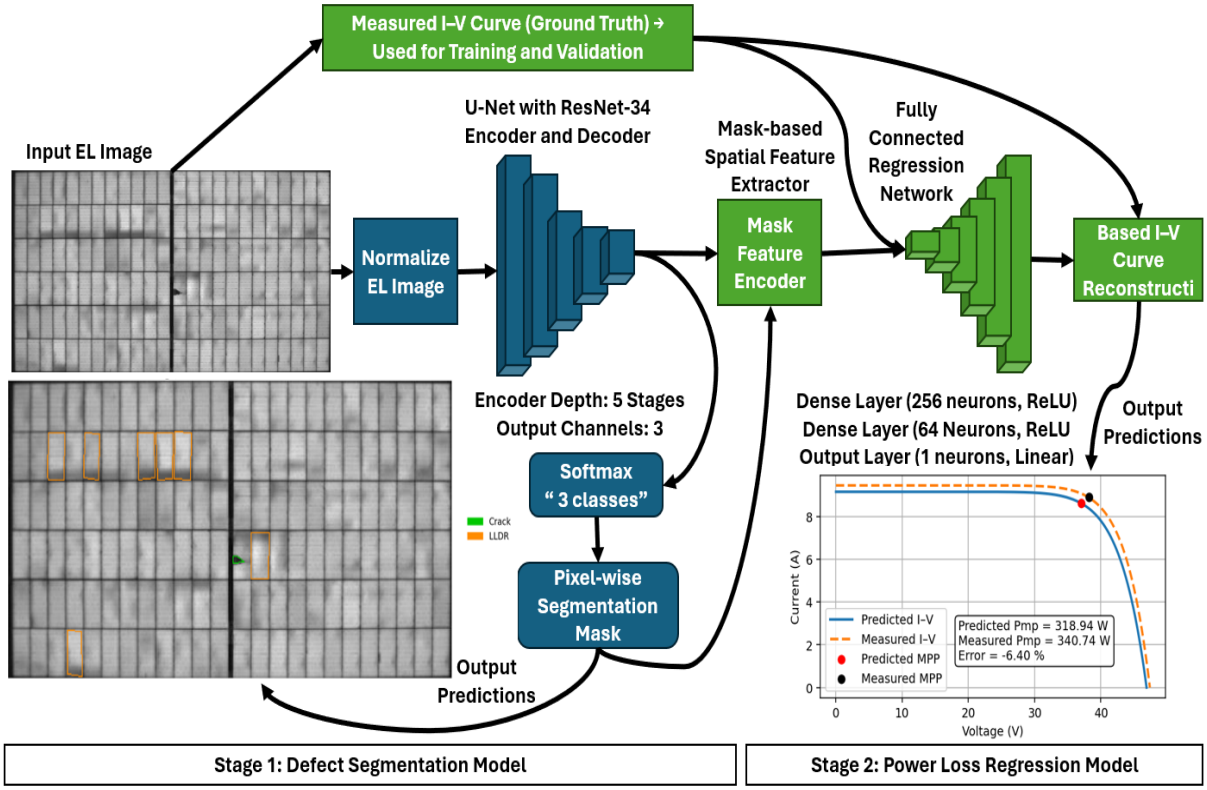

Figure 4. Cascaded two-stage framework for defect detection and power loss prediction from EL images of PV modules.

A central methodological feature of the proposed approach is the separation between defect detection (stage 1) and power loss prediction (stage 2), while maintaining a consistent information flow between the two stages. The EL image serves as the primary input to the model. Measured I–V curves obtained under STC are used exclusively to compute the true power loss for training supervision and for benchmarking during validation. At no stage is electrical information injected into the inference pathway measured, ensuring physically consistent and non-leaking architecture.

During model development, several alternative network configurations and single-stage approaches were explored through preliminary iterative trials. These trials indicated that the adopted two-stage architecture provided the most suitable balance between predictive performance and interpretability by maintaining a clear separation between defect localization and electrical performance estimation. The U-Net with ResNet-34 encoder was consequently selected for Stage 1 because its multi-scale feature extraction and skip connections are well suited to resolving both narrow crack features and extended low-luminescence regions. These preliminary architecture-selection trials formed part of model development rather than a formal benchmarking study.

3.4. Defect Segmentation Model: Stage 1

The first stage of the proposed framework, illustrated in Fig. 4 (Stage 1), is responsible for extracting spatially resolved defect information from EL images of laminated half-cut PV modules. The objective of this stage is not to estimate electrical performance directly, but to generate a structured representation of spatial defect morphology that can later be translated into module-level power loss.

Let $I_{EL} \in \mathbb{R}^{H \times W}$ denote a grayscale EL image of spatial resolution $H \times W$, where each pixel intensity corresponds to near-infrared radiative recombination emitted under forward-bias excitation. Prior to network processing, the image is normalized to ensure consistent intensity scaling across the dataset, using (1).

$$I_n(x, y) = \frac{I_{EL}(x,y) - I_{min}}{I_{max} - I_{min}} \quad (1)$$

where I_{min} and I_{max} represent the minimum and maximum intensity values within the module region, and $I_n(x, y) \in [0,1]$ is the normalized EL intensity at pixel location (x, y). This step removes global exposure variation and ensures that the network focuses on relative spatial contrast patterns rather than absolute brightness levels.

The normalized image I_n is then provided as input to a U-Net architecture with a ResNet-34 encoder backbone [29], as shown in Fig. 4. The encoder consists of five hierarchical stages, each composed of convolutional layers, batch normalization, and rectified linear unit (ReLU) activations [30]. At each layer l, feature maps are computed using (2).

$$F_l = \sigma(W_l * F_{l-1} + b_l) \quad (2)$$

where F_{l-1} is the input feature map to layer l, W_l is the learnable convolutional kernel tensor, $*$ is the convolution operation, b_l is the bias vector, $\sigma(\cdot)$ is the ReLU activation function, and F_l is the output feature map at layer l.

This hierarchical encoding enables the network to learn multi-scale spatial representations, ranging from fine crack edges to broader low-luminescence regions. Down sampling operations within the encoder increase the receptive field, allowing the

model to capture contextual interactions across half-cell regions. The decoder path of the U-Net progressively restores spatial resolution through up sampling layers and skip connections that concatenate encoder feature maps with decoder representations. These skip connections preserve fine spatial detail [31], which is essential for accurately delineating narrow crack structures that may span only a few pixels in width.

The final decoder layer produces a three-channel output tensor as (3), corresponding to three semantic classes, (i) Class 0: Background, (ii) Class 1: Crack, and (iii) Class 2: Low-Luminescence Region (LLDR).

$$Z(x, y) \in \mathbb{R}^3 \tag{3}$$

To convert the raw logits $Z_c(x, y)$ into class probabilities, a softmax activation is applied using (4). Here $p_c(x, y)$ is the predicted probability that pixel (x, y) belongs to class c, $Z_c(x, y)$ is the logit for class c.

$$p_c(x, y) = \frac{e^{Z_c(x,y)}}{\sum_{k=0}^{2} e^{Z_k(x,y)}} \tag{4}$$

Then the predicted pixel-wise segmentation mask $\widehat{M}(x, y)$ is then obtained using (5), which assigns each pixel to the class with the highest predicted probability.

$$\widehat{M}(x, y) = \arg \max_{c \in \{0,1,2\}} p_c(x, y) \tag{5}$$

Because crack regions typically occupy a small fraction of the total module area, class imbalance is addressed using a hybrid segmentation loss. The total segmentation loss is defined as (6).

$$L_{seg} = \lambda_1 L_{Dice} + \lambda_2 L_{CE} \tag{6}$$

where L_{seg} is the overall segmentation loss, L_{Dice} is the Dice loss, L_{CE} is the categorical cross-entropy loss, λ_1 and λ_2 are weighting coefficients controlling the relative contribution of each term. The Dice loss is defined as (7), where $M_c(x, y)$ is the ground-truth binary mask for class c. This formulation emphasizes overlap accuracy between predicted and true defect regions and improves sensitivity to small crack structures.

$$L_{Dice} = 1 - \frac{2 \sum_{x,y} p_c(x,y) M_c(x,y)}{\sum_{x,y} p_c(x,y) + \sum_{x,y} M_c(x,y)} \tag{7}$$

Finally, at stage 1, the categorical cross-entropy term is defined as (8).

$$L_{CE} = - \sum_{x,y} \sum_{c=0}^{2} M_c\,(x, y) \log p_c(x, y) \tag{8}$$

where $M_c(x, y)$ is the one-hot encoded ground-truth label. This component stabilizes multi-class classification and penalizes incorrect pixel assignments.

In this work, equal weighting ($\lambda_1 = \lambda_2 = 0.5$) was adopted to balance region-overlap optimization with probabilistic classification stability. The output of Stage 1 is therefore a

spatially structured segmentation mask $\widehat{M}$ identifying crack and low-luminescence regions across the module surface. Importantly, this stage does not estimate electrical performance directly. Instead, it produces a physically interpretable defect map that serves as the structured input to Stage 2 of the framework, where the electrical consequence of these detected defects is quantified in terms of module power loss.

3.5. Power Loss Regression and Physics-Based I–V Reconstruction: Stage 2

The second stage of the proposed framework, illustrated in Fig. 4 (Stage 2), translates the spatial defect information extracted in Stage 1 into a quantitative estimate of the module-level power loss under STC conditions. This stage establishes the functional link between detected cracks and low-luminescence regions and the corresponding reduction in maximum power output. Each module in the dataset is electrically characterized using a calibrated I–V tracer under STC. From the measured I–V curve, the maximum power point $MPP = P_{mp}$ is extracted. The true power loss is defined as (9).

$$P_{loss} = \frac{P_{ref} - P_{mp}}{P_{ref}} \tag{9}$$

Where P_{ref} is the manufacturer-rated nominal maximum power (W), P_{mp} is the measured maximum power under STC (W), and P_{loss} is the dimensionless fractional power loss. This value serves as the supervisory signal for training the regression network. Importantly, the measured I–V curve is not used as an input feature during inference; it is used only to compute the ground-truth P_{loss} during training and to benchmark predicted I–V curves during validation. The pixel-wise segmentation mask $\widehat{M}(x, y)$ generated in Stage 1 provides a structured representation of defects locations. However, the regression network requires a compact feature representation rather than a full-resolution mask. To obtain this representation, a mask-based spatial feature extractor is applied. Let $\widehat{M}_c(x, y)$ denote the binary mask for class $c \in \{1,2\}$, corresponding to crack and LLDR regions, respectively. The relative area fraction for each defect class is computed as (10).

$$A_c = \frac{1}{HW} \sum_{x=1}^{H} \sum_{y=1}^{W} \widehat{M}_c\,(x, y) \tag{10}$$

where H and W are the image height and width, and A_c represents the normalized spatial coverage of class c.

In addition to area fractions, spatial embeddings derived from the encoder feature maps are aggregated using global average pooling [32]. For a feature map $F_k(x, y)$ with channel index k, the pooled feature is defined as (11). Where z_k is the aggregated value of channel k, and $F_k(x, y)$ is the value at pixel (x, y).

$$z_k = \frac{1}{HW} \sum_{x=1}^{H} \sum_{y=1}^{W} F_k\,(x, y) \tag{11}$$

Therefore, the final latent defect representation vector is denoted as (12). Where d is the dimensionality of the aggregated feature space. This vector encapsulates both defect extent and spatial context in a compact numerical form suitable for regression. The latent feature vector $\mathbf{z}$ is provided as input to a fully connected regression network, as shown in Fig. 4. The network consists of two hidden dense layers followed by a single linear output neuron. The first hidden layer computes as (13).

$$\mathrm{z} \in \mathbb{R}^{d} \tag{12}$$

$$\mathrm{h}_1 = \sigma(\mathrm{W}_1\mathrm{z} + \mathrm{b}_1) \tag{13}$$

where $\mathrm{W}_1 \in \mathbb{R}^{256\times d}$ is the weight matrix, $\mathrm{b}_1 \in \mathbb{R}^{256}$ is the bias vector, $\sigma(\cdot)$ is the ReLU activation function, and $\mathrm{h}_1 \in \mathbb{R}^{256}$ is the hidden representation. The second hidden layer follows the same principle and is computed as in (14), where $\mathrm{W}_2 \in \mathbb{R}^{64\times 256}$, $\mathrm{b}_2 \in \mathbb{R}^{64}$, and $\mathrm{h}_2 \in \mathbb{R}^{64}$.

$$\mathrm{h}_2 = \sigma(\mathrm{W}_2\mathrm{h}_1 + \mathrm{b}_2) \tag{14}$$

The final predicted power loss is, therefore, obtained using a linear output layer (15), where $\mathrm{W}_r \in \mathbb{R}^{1\times 64}$, $b_r \in \mathbb{R}$, and $\hat{P}_{loss}$ is the predicted fractional power loss. At this stage, the linear activation is used at the output to allow continuous, unconstrained prediction.

$$\hat{P}_{loss} = \mathrm{W}_r\mathrm{h}_2 + b_r \tag{15}$$

The regression network is trained using mean squared error (MSE), as in (16), where N is the number of training samples, $P_{loss}^{(i)}$ is the true power loss derived from measured I–V curves, and $\hat{P}_{loss}^{(i)}$ is the predicted value.

$$L_{reg} = \frac{1}{N}\sum_{i=1}^{N}\left(P_{loss}^{(i)} - \hat{P}_{loss}^{(i)}\right)^2 \tag{16}$$

This loss penalizes deviations between predicted and measured power loss, ensuring that the network learns the functional relationship between spatial defect patterns and electrical performance reduction. The predicted power loss is subsequently used to reconstruct the module's I–V curve. The predicted maximum power is computed as (17), where P_{ref} is the nominal rated power.

$$\hat{P}_{mp} = (1 - \hat{P}_{loss})P_{ref} \tag{17}$$

The reconstructed I–V curve was generated using the single-diode model under STC, parameterized using manufacturer datasheet values. This step ensures that the final output is not merely a scalar estimate but a physically interpretable electrical characteristic curve. The reconstructed I–V curve is then compared against the measured I–V curve for benchmarking and validation, as illustrated in Fig. 4.

In conclusion, Stage 2 transforms spatial defect detection into an electrically meaningful performance metric. The regression network first estimates module-level power loss from the extracted defect representation. This data-driven prediction is then coupled

with the single-diode PV model, which constitutes the physics-based component of the framework, to reconstruct the corresponding I–V characteristic under STC. The physical model is therefore applied as an electrical consistency and reconstruction layer rather than as a governing-equation constraint within the neural-network training process.

3.6. Implementation and Training Configuration

A total of 1000 full-module EL images were used in this study, all acquired from the same 144-half-cell monocrystalline p-type PERC module technology described in Section 3.1 and using the standardized EL imaging protocol presented in Section 3.2. The dataset was divided into two completely independent groups. The first 500 modules constituted the model-development dataset and were used for 70% training, 15% validation, and 15% internal testing. The remaining 500 modules were kept entirely separate from model development and were used exclusively for final independent evaluation. None of these 500 modules or their EL images were used during training, validation, model selection, or hyperparameter optimization. Corresponding measured I–V characteristics were available for the modules, providing the electrical ground truth for power-loss prediction and subsequent validation.

To improve generalization and robustness to imaging variability, data augmentation was applied online during training. Augmentation operations included random horizontal flipping, random rotation within ±10°, brightness scaling within ±10%, and addition of Gaussian noise ($\sigma = 0.01$). Although the dataset comprised 500 unique EL images, the stochastic augmentation strategy effectively expanded the diversity of training samples to 1500; and same strategy when training the model was developed, 70% of the dataset used for training, 15% for validation and 15% for testing. All models were implemented using the PyTorch deep learning framework and trained using Google Colab Pro with NVIDIA graphics processing unit (GPU) acceleration (Tesla T4 or A100, depending on availability). This configuration provided sufficient computational capacity for training the ResNet-34-based U-Net segmentation backbone and the fully connected regression network.

The segmentation model (Stage 1) comprises approximately 24–26 million trainable parameters, primarily due to the ResNet-34 encoder and U-Net decoder structure. The regression network (Stage 2), consisting of two dense layers (256 and 64 neurons) and a single linear output neuron, contains approximately 0.1–0.3 million trainable parameters depending on the dimensionality of the aggregated defect feature vector. Optimization was performed using the Adam optimizer with an initial learning rate of 1×10^{-4}.

A learning rate scheduler (Reduce-on-Plateau [33]) was applied, reducing the learning rate by a factor of 0.5 if validation loss did not improve over 10 consecutive epochs. This adaptive strategy stabilized convergence and prevented oscillatory training behavior. The models were trained for a maximum of 120 epochs with a batch size of 8. Early stopping

was employed with a patience of 15 epochs, monitoring validation loss to prevent overfitting. Weight initialization followed standard practice. The ResNet-34 encoder was initialized with ImageNet-pretrained weights [34], and the fully connected regression layers were initialized using He (Kaiming [35]) initialization. Dropout regularization was applied to reduce overfitting. Dropout rate of 0.3 in the segmentation decoder bottleneck, and dropout rates of 0.3 and 0.2 after the first and second fully connected layers in the regression network, respectively.

Segmentation and regression were trained sequentially in a cascaded manner. First, the segmentation network was trained to convergence. Subsequently, the regression network was trained using the predicted segmentation masks and corresponding ground-truth power loss values. During inference, the two stages operate consecutively as illustrated in Fig. 4 and detailed in Fig. 5. For segmentation, pixel-level accuracy and Dice similarity coefficient were computed. For regression, performance was assessed using mean absolute error (MAE), root mean squared error (RMSE), and coefficient of determination (R^2). The predicted power loss $\hat{P}_{loss}$ was further used to reconstruct the predicted I–V curve.

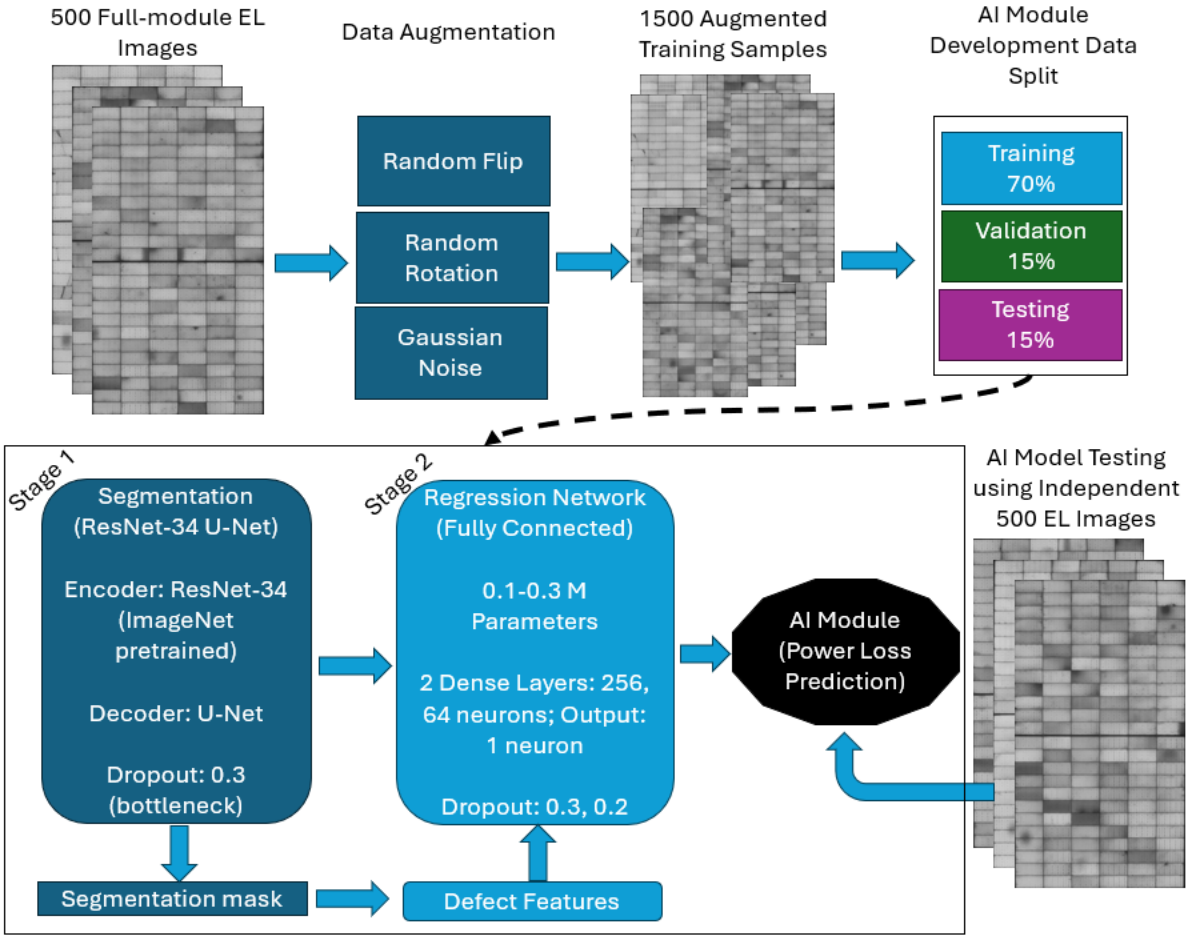

Figure 5. Detailed architectural and training configuration of the proposed two-stage deep learning framework.

Model development used the stratified training, validation, and internal test subsets described above. The validation subset was used during model development and early stopping, while the internal test subset was retained for evaluation of segmentation performance. Following completion of model development, the power-loss regression framework was further evaluated using the separate independent dataset of 500 previously unseen modules. This independent dataset was not used for training, validation, model selection, or hyperparameter optimization. K-fold cross-validation was not employed.

4. Results and Discussion

4.1. Stage 1: Defect Segmentation Performance

The performance of the proposed U-Net with ResNet-34 encoder for pixel-wise defect segmentation was first evaluated in terms of training convergence and segmentation accuracy. Fig. 6 illustrates the evolution of training and validation accuracy across 120 epochs. The model exhibits stable and monotonic convergence, with both curves increasing steadily during the initial learning phase (epochs 1–40) before gradually reaching a plateau.

As shown in Fig. 6, validation accuracy closely follows the training curve throughout the optimization process, indicating limited overfitting and effective regularization through dropout, data augmentation, and learning rate scheduling. The final validation accuracy stabilizes at approximately 0.91–0.92, while training accuracy converges near 0.94. The relatively small generalization gap (<3%) demonstrates that the model learns defect-relevant spatial features rather than memorizing training samples.

The rapid convergence observed within the first 40 epochs suggests that the pretrained ResNet-34 encoder effectively transfers hierarchical spatial representations from ImageNet initialization to the EL imaging domain. After this stage, performance improvements become incremental, reflecting refinement of crack boundary delineation and low-luminescence region detection.

While pixel-level accuracy provides an overall measure of segmentation correctness, it can be biased by background dominance due to class imbalance. Therefore, class-wise Dice similarity coefficients were computed on the held-out validation and test sets to provide a more robust evaluation of segmentation performance, particularly for the crack class, which occupies a relatively small spatial fraction of the module surface. The Dice similarity coefficient (DSC) is defined as (18). where P denotes the predicted defect region and G represents the corresponding ground-truth mask.

$$DSC = \frac{2|P \cap G|}{|P| + |G|} \tag{18}$$

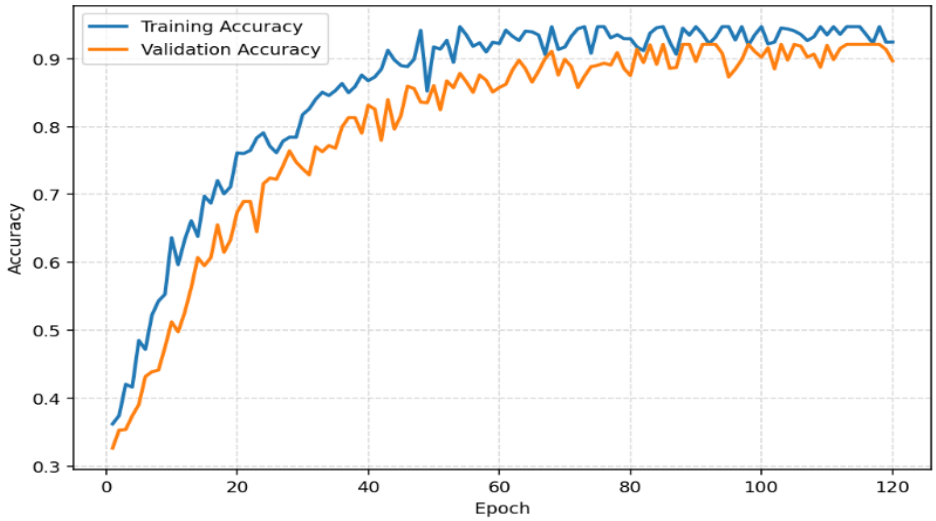

Figure 6. Training and validation accuracy of the defect segmentation model.

The quantitative segmentation performance is summarized in Table 2. The results demonstrate that the model achieves strong overlap accuracy across defect classes. In particular, the Dice score for low-luminescence regions exceeds 0.89 on the test set, reflecting the model's ability to capture extended regions of reduced radiative recombination with high spatial consistency. For crack detection, the Dice score of 0.82 indicates reliable identification of thin and spatially sparse crack structures, which are inherently more challenging due to their narrow morphology and lower pixel representation. The slightly lower Dice score for cracks compared to LLDR class is expected, given the fine-scale nature of crack patterns and their sensitivity to thresholding and boundary precision.

Table 2. Quantitative Segmentation Performance of the Stage 1 Model

Class	Dice (Validation)	Dice (Test)
Crack	0.84	0.82
LLDR	0.91	0.89
Mean Dice	0.88	0.86

Representative qualitative segmentation results are presented in Fig. 7(a) and Fig. 7(b), illustrating the predicted defect masks overlaid on two example EL images from the test dataset. The contour-based visualization highlights the spatial extent of crack regions (yellow) and low-luminescence regions (purple) as identified by the proposed segmentation model. As observed, the predicted crack contours closely follow the irregular morphology of fracture patterns, while LLDR regions are consistently localized in areas exhibiting reduced EL intensity. The results demonstrate accurate delineation of both thin crack structures and broader electrically inactive zones, confirming that the model preserves fine spatial detail while maintaining robustness across different defect severities and distributions.

A representative example of a severe electrical fault is illustrated in Fig. 7(c), where a defective bypass diode results in the deactivation of an entire substring within the PV module. In this case, the affected substring exhibits near-zero electroluminescent emission and is consistently identified by the proposed segmentation model as a continuous LLDR. Unlike localized crack-induced degradation, this failure mode produces a large-scale electrically inactive area spanning multiple cells, which the model successfully captures as a unified defect region. The accurate delineation of the full substring demonstrates that the network does not merely detect local intensity variations but generalizes structurally coherent electrical failures. This example highlights the model's ability to identify both fine-scale crack morphology and large-area performance-limiting faults, confirming the robustness and physical consistency of the proposed segmentation framework.

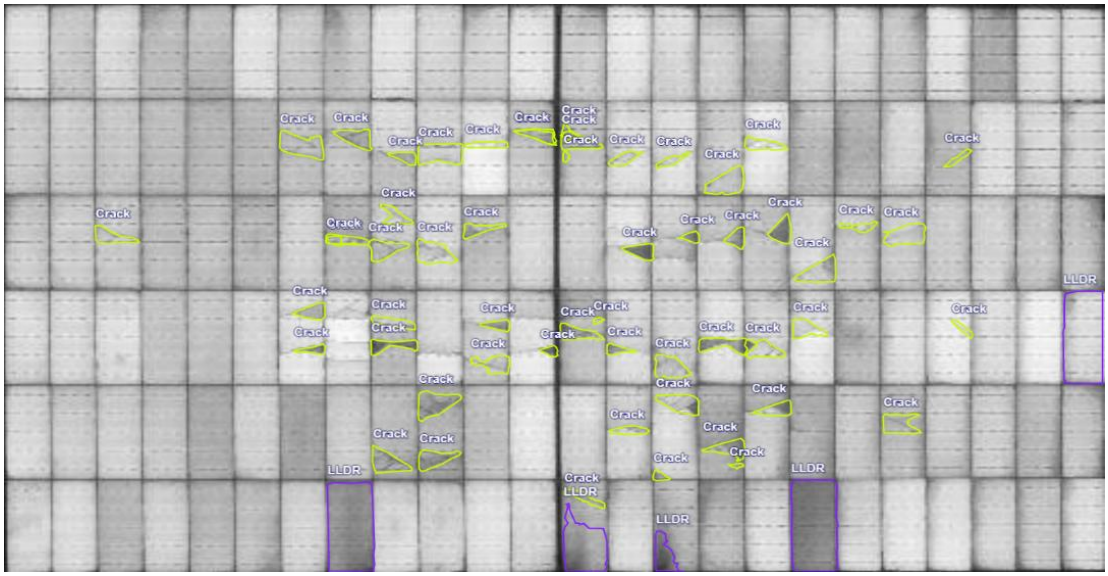

(a)

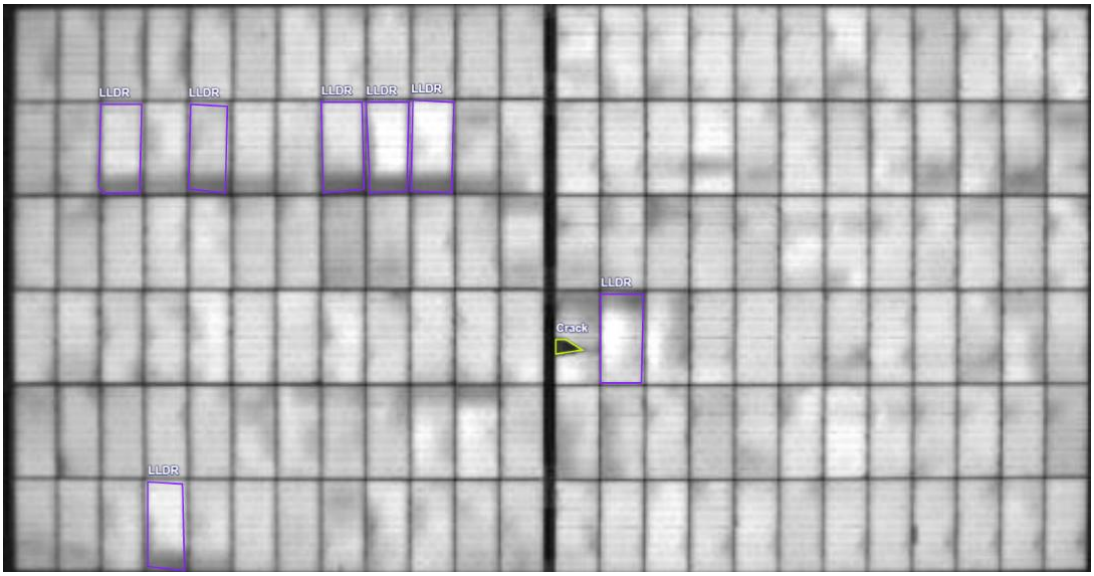

(b)

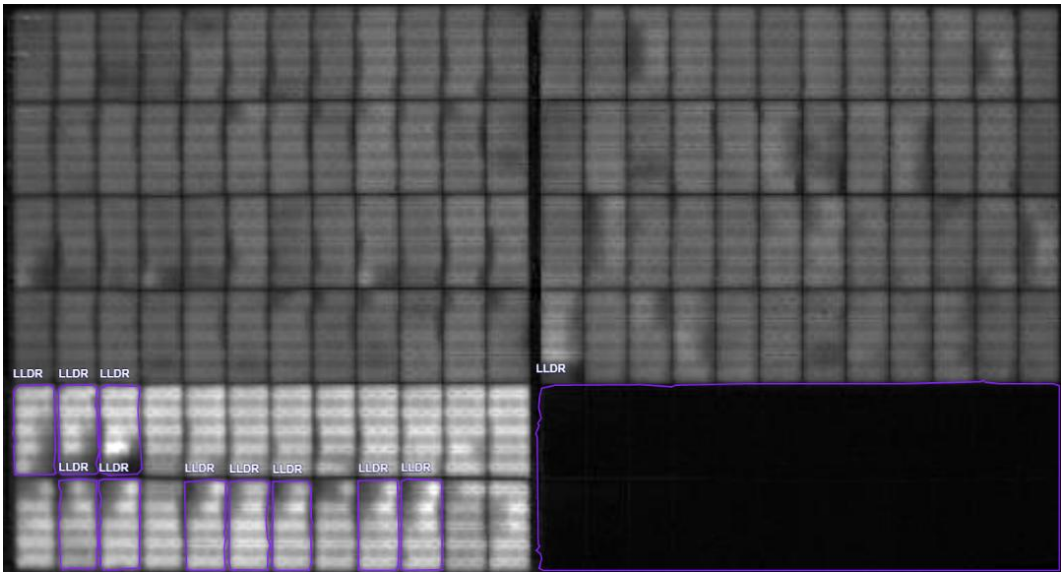

(c)

Figure 7. Segmentation results for three representative EL module images from the test dataset. (a) example with distributed crack patterns and localized LLDR regions, (b) example illustrating clustered crack formations and extended low-luminescence areas, (c) example of PV module exhibiting bypass diode failure. The affected substring is entirely classified as LLDR, indicating large-scale electrical inactivity.

4.2. Stage 2: Power Loss Regression and Physics-Based I–V Reconstruction

Following spatial defect segmentation, the second stage of the proposed framework quantifies the electrical impact of detected defects by predicting module-level power loss under STC conditions. The regression network was trained using the latent defect representation derived from Stage 1 and evaluated on an independent held-out test set comprising 500 module samples.

Fig. 8 illustrates the relationship between predicted power loss, obtained from calibrated I–V curve measurements. The dashed line represents ideal 1:1 agreement. Shaded regions indicate ±2% and ±5% absolute error tolerance bands relative to the identity line, while data points are color-coded according to absolute prediction error (<2%, 2–5%, and >5%). Each point corresponds to an individual PV module.

The predicted values closely follow the ideal agreement line across the full degradation range (0–30%). The regression model achieves MAE = 0.0183, RMSE = 0.0243, and R^2 = 0.9287. The high coefficient of determination ($R^2 \approx 0.93$) indicates that the learned mapping from spatial defect morphology to electrical degradation captures the dominant variance in module-level performance reduction. The relatively low MAE and RMSE values further confirm that prediction deviations remain within narrow tolerance bounds for the majority of samples, with most modules falling inside the ±5% error band.

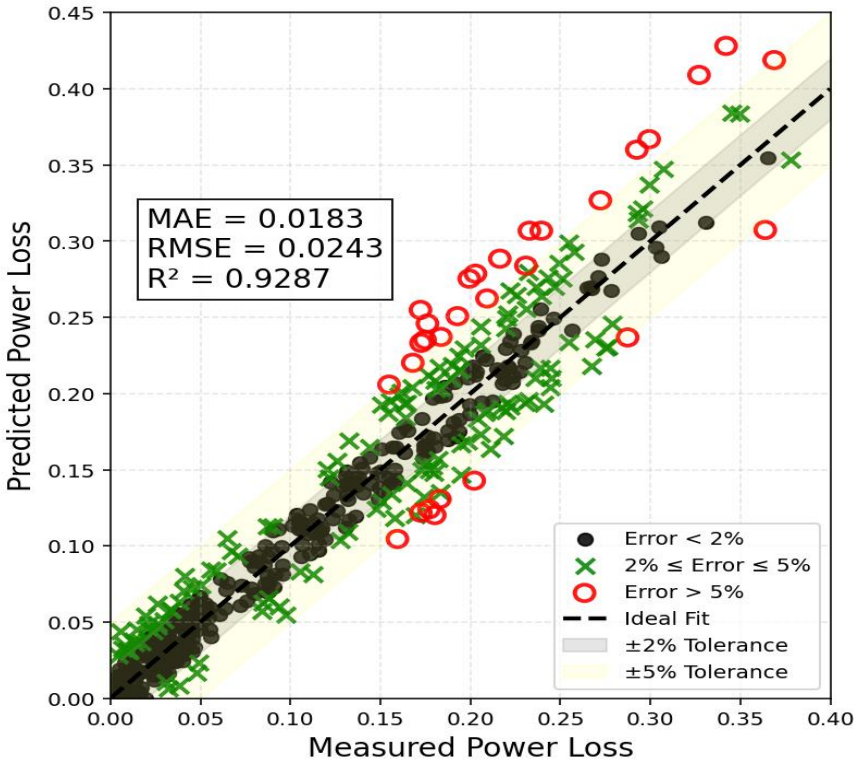

Figure 8. Predicted versus measured PV module-level power loss for the test dataset (n = 500). The dashed line represents the ideal 1:1 agreement. Shaded regions represent ±2% and ±5% error tolerance bands relative to the ideal fit. Data points are color-coded by absolute prediction error: black (<2%), green (2–5%), and red (>5%).

A moderate increase in residual dispersion is observed in the higher power-loss region (>15%), reflecting the increased electrical complexity associated with severe crack propagation, partial cell isolation, non-uniform current redistribution, and possible bypass diode activation. Such behavior is physically consistent with nonlinear defect–performance interactions in half-cut PV architectures, where localized mechanical damage can induce disproportionate electrical losses. In contrast, low-degradation modules (<5%) exhibit tight clustering around the identity line, demonstrating robust sensitivity to minor defect-induced performance reductions.

To further evaluate statistical reliability, residual analysis was performed on the test dataset, as presented in Fig. 9. The mean bias error (MBE) is 0.0041, corresponding to an average overestimation of 0.41%, indicating negligible systematic prediction bias. The residual distribution is approximately symmetric and centered near zero, with the majority of errors confined within ±3% absolute deviation. A small number of higher-magnitude residuals are observed in the upper degradation range, reflecting increased electrical nonlinearity associated with severe crack propagation and substring-level failures. This mild increase in variance at higher degradation levels is physically consistent with current redistribution and bypass diode activation effects in half-cut PV architectures.

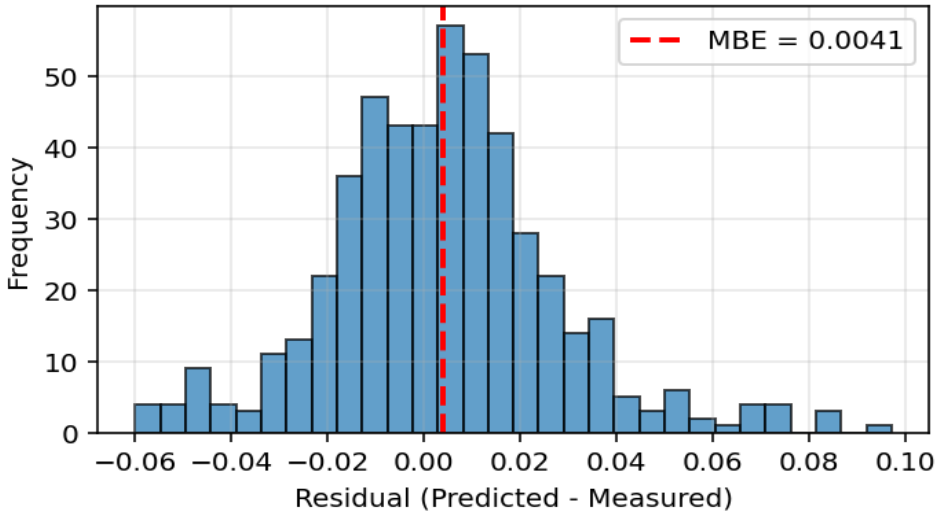

Figure 9. Residual distribution of predicted module-level power loss on the held-out test dataset.

To further assess the physical consistency of the predicted power loss values, the reconstructed I–V curves were compared against experimentally measured I–V characteristics for three representative modules, shown in Fig. 10(a–c). These modules correspond to the same samples previously presented in Fig. 7(a-c), respectively, thereby linking spatial defect morphology to electrical performance behavior.

In Fig. 10(a), corresponding to the module with distributed crack patterns (Fig. 7(a)), the predicted maximum power is 348.48 W compared to a measured value of 362.39 W, yielding a relative deviation of −3.84%. The predicted I–V curve closely follows the measured characteristic in both the constant-current region and near the MPP, with minor divergence observed in the voltage roll-off region. This deviation is consistent with moderate series resistance effects that are not fully linear.

Fig. 10(b), associated with the module shown in Fig. 7(b), demonstrates improved agreement, with a predicted maximum power of 373.83 W versus a measured value of 379.98 W, corresponding to an error of −1.62%. The near-overlapping I–V curves indicate accurate estimation of both current and voltage behavior, confirming that spatially localized crack formations translate into proportionally moderate power loss without large-scale distortion of the curve shape. In contrast, Fig. 10(c) presents the module previously identified in Fig. 7(c) as exhibiting substring-level LLDR behavior due to bypass diode malfunction. In this case, the predicted maximum power is 300.28 W compared to a measured 317.56 W, resulting in a −5.44% deviation. A larger error is expected given the nonlinear electrical behavior associated with substring deactivation and bypass diode conduction, which introduces abrupt voltage reduction and complex current redistribution effects. Despite this complexity, the reconstructed curve preserves the overall trend and correctly captures the substantial reduction in maximum power.

Across all three cases, the predicted I–V characteristics maintain strong alignment with measured curves (error band between −1.62% to −5.44%), particularly around the maximum power region, which is most critical for energy yield estimation, demonstrating that the regression-derived power loss values translate into physically meaningful and quantitatively reliable electrical reconstructions.

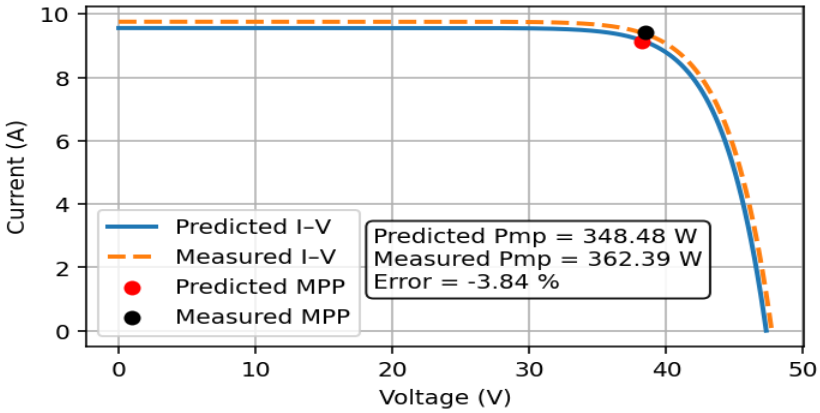

(a)

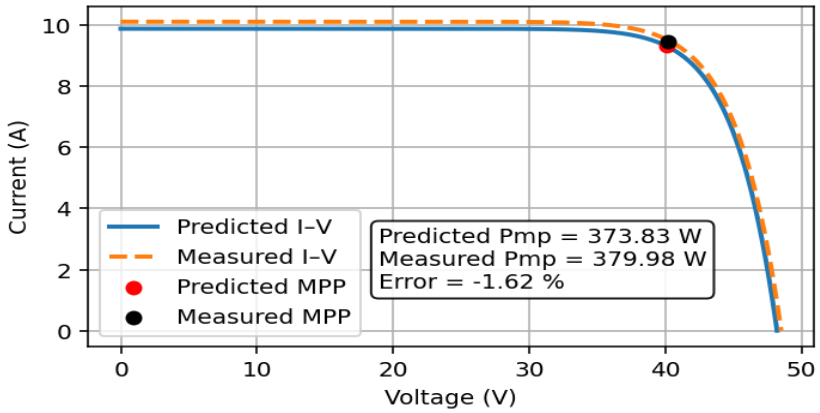

(b)

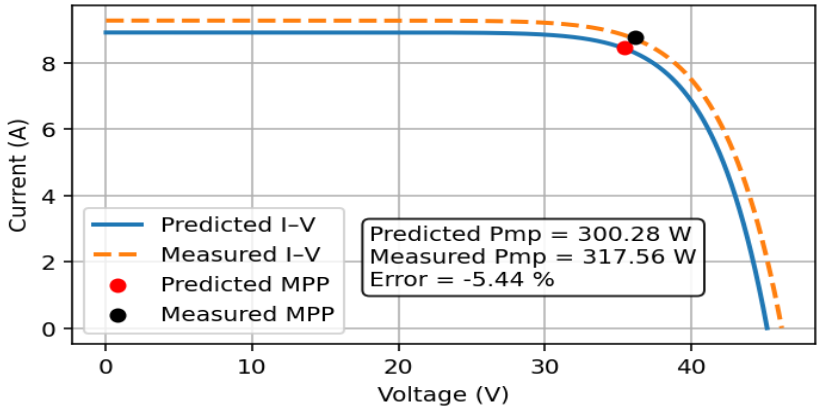

(c)

Figure 10. Comparison of reconstructed and measured I–V curves for three representative modules. (a) module with distributed crack patterns (corresponding to Fig. 7(a)); (b) module with clustered crack formations (Fig. 7(b)), (c) module exhibiting substring-level low-luminescence behavior due to bypass diode malfunction (Fig. 7(c)).

4.3. Industrial Decision-Level Evaluation

While the regression framework predicts continuous fractional power loss, industrial deployment requires discrete acceptance decisions. In practical PV module screening, degradation severity is commonly categorized into three bands: Accept (<3%), Monitor (3–15%), and Reject (>15%). To evaluate the suitability of the proposed framework for industrial decision support, predicted and measured power loss values were mapped into these three severity classes. Let the true fractional power loss be denoted as P_{true} and the predicted value as P_{pred}. The industrial classification function $C(P)$ is defined as (19).

$$C(P) = \begin{cases} 0, & P < 0.03 \text{(Accept)} \\ 1, & 0.03 \leq P \leq 0.15 \text{(Monitor)} \\ 2, & P > 0.15 \text{(Reject)} \end{cases} \tag{19}$$

The confusion matrix for the held-out test dataset is shown on Fig. 11. The overall decision accuracy is 83.4%. Importantly, no severely degraded modules (>15% power loss) were misclassified as Accept, indicating zero false-accept cases for critical defects. The recall for the Reject category reaches 88.1%, demonstrating reliable identification of performance-limiting modules.

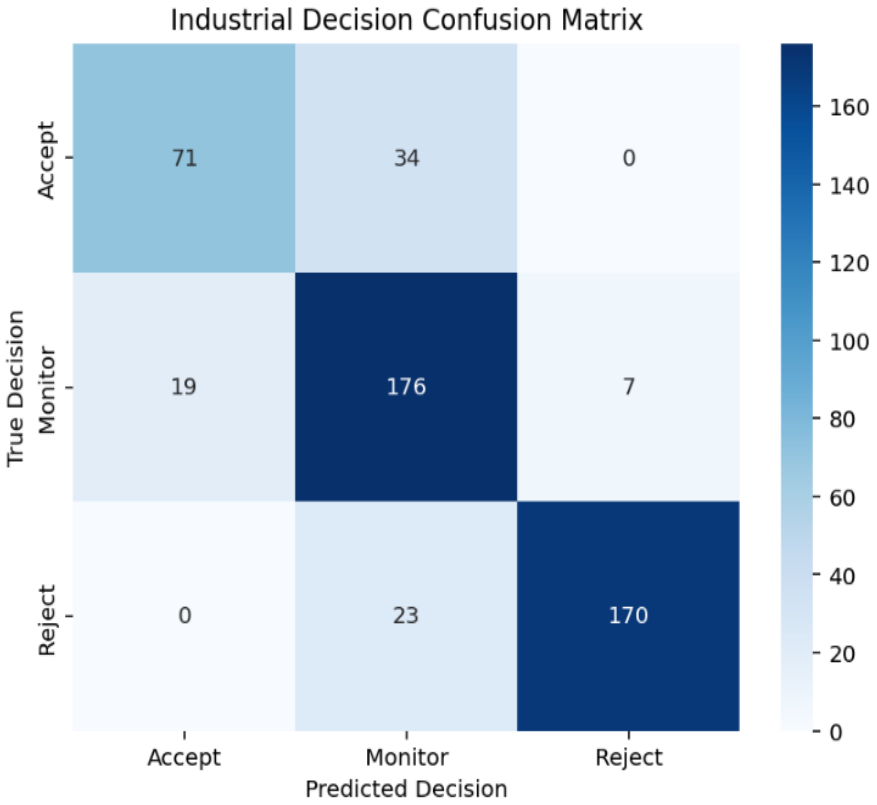

Figure 11. Confusion matrix for industrial decision-level classification based on predicted and measured power loss. Modules are categorized into Accept (<3%), Monitor (3–15%), and Reject (>15%) classes. No severely degraded modules were misclassified as Accept, demonstrating risk-aware industrial screening capability.

Misclassifications occur primarily near threshold boundaries between adjacent classes (Accept–Monitor and Monitor–Reject), which is expected given the continuous nature of degradation progression. From an industrial perspective, conservative misclassification toward intermediate categories (e.g., Reject predicted as Monitor) is operationally less critical than false acceptance of severely degraded modules. The results therefore confirm that the proposed framework provides not only accurate continuous prediction but also risk-aware and practically deployable decision support for pre-installation PV module certification.

To further quantify decision reliability, class-wise precision, recall, and F1-score metrics are summarized in Fig. 12. The Reject category achieves a precision of 0.96 and an F1-score of 0.92, confirming that modules classified as Reject are almost always truly performance-critical. This high precision is particularly important in industrial screening, where erroneous rejection decisions may incur economic cost, while false acceptance of severely degraded modules poses operational risk.

The Monitor category exhibits balanced performance (precision = 0.76, recall = 0.87), indicating strong sensitivity to moderate degradation while maintaining reasonable specificity. The Accept category shows comparatively lower recall (0.68), reflecting conservative boundary behavior near the 3% threshold, where small regression deviations may shift modules into the Monitor class. From an industrial perspective, such conservative misclassification is preferable to underestimating degradation severity.

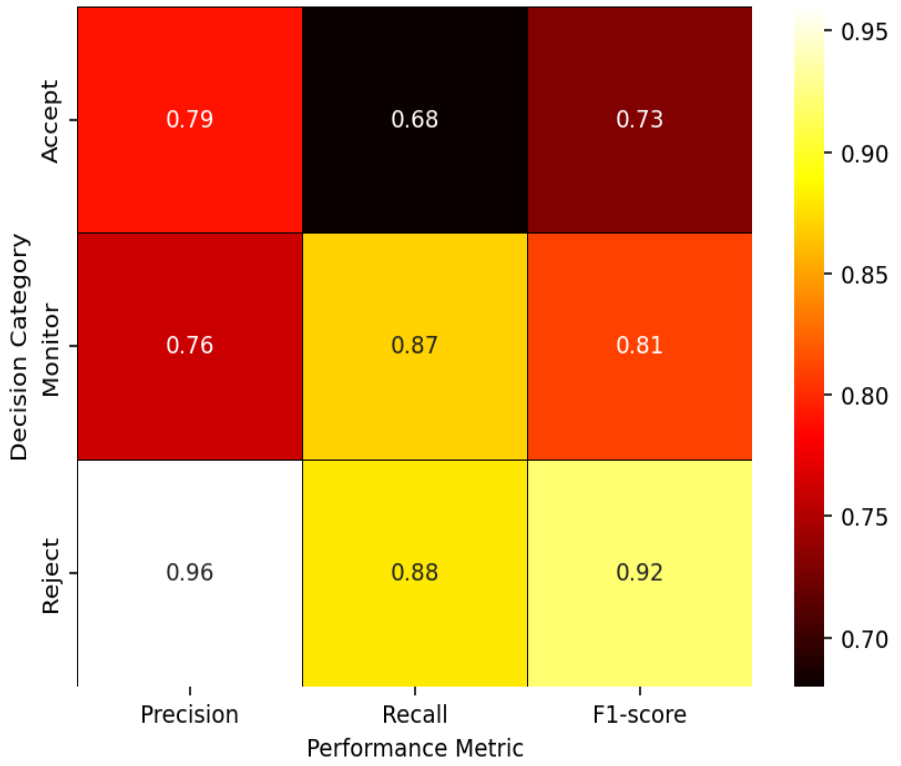

Figure 12. Precision, recall, and F1-score for industrial decision-level classification of PV modules into Accept (<3%), Monitor (3–15%), and Reject (>15%) categories.

Additionally, from an integration perspective, the proposed framework can complement existing PV inspection and asset-monitoring workflows rather than replace conventional electrical monitoring. Following EL image acquisition, the processing sequence comprises image preprocessing, defect segmentation, power-loss estimation, and decision-level classification. The resulting power-loss prediction and Accept/Monitor/Reject assessment can therefore be incorporated alongside existing I–V, thermography, or asset-management records. The present study demonstrates the diagnostic workflow and independent predictive performance; direct integration within a deployed monitoring platform remains a future implementation step.

From a deployment perspective, the proposed framework is computationally practical for integration into existing EL inspection workflows. The complete inference pipeline, including image preprocessing, defect segmentation, power-loss prediction, and decision-level classification, requires approximately 2 seconds per module. Although GPU acceleration was used during model development and training, inference can be implemented on a laptop-class computer connected directly to the EL imaging system, without requiring dedicated high-performance computing infrastructure. In practice, an acquired EL image can therefore be transferred directly to the prediction software and automatically processed to provide the estimated power loss and corresponding Accept/Monitor/Reject classification. The approximately 2 seconds processing time per module enables sequential screening of large module batches and makes the approach readily scalable to manufacturing, pre-installation quality-control, and periodic inspection workflows. Furthermore, the independent evaluation on 500 previously unseen modules demonstrates that the trained framework can be applied to new modules without model retraining for each individual inspection.

4.4. Comparative Analysis

Table 3 summarizes representative EL-based electrical prediction frameworks reported in the literature [36-40] and contrasts them with the proposed model. As shown in Table 3, existing approaches fall into two main categories:

1) Defect-centric models, which use EL imaging for classification or intermediate feature extraction before regression.
2) Direct EL to I–V regression models, which map images to electrical parameters or full I–V curves using CNN or transformer architectures.

While these studies [36-40] demonstrate promising electrical prediction accuracy, they share a common limitation: they are architecture-agnostic. None of the existing works explicitly incorporate module electrical topology into the modeling strategy. In particular, modern half-cut PV module designs, which exhibit parallel substring configurations, altered current pathways, and distinct mismatch propagation behavior, have not been explicitly modeled in prior EL-based regression frameworks. The proposed work

addresses this gap by introducing an architecture-aware EL-to-electrical pipeline, specifically designed for half-cut PV modules. Rather than treating the module as a uniform electrical entity, the model accounts for structural segmentation and topology-consistent electrical behavior. This improves prediction stability, particularly in the MPP region where substring mismatch effects are amplified.

The present framework was developed for controlled EL inspection and STC-referenced power-loss prediction; therefore, weather variables such as irradiance, ambient temperature, and cloud conditions are not direct model inputs. Although the framework is independently evaluated using previously unseen modules, deployment under outdoor or daylight EL conditions may introduce additional domain shifts through changes in background illumination, module temperature, image contrast, and measurement noise. Future adaptation could therefore include field-specific fine-tuning, environmental augmentation, domain-adaptation techniques, and incorporation of environmental measurements where required.

Despite the independent evaluation performed in this study, several limitations remain. The framework was developed primarily for a specific half-cut monocrystalline PERC architecture and controlled EL imaging conditions; therefore, its generalization to other module technologies, manufacturers, and outdoor/daylight EL acquisition requires further investigation. In addition, the regression stage depends on the quality of the preceding defect segmentation. Future work will therefore extend the framework across broader PV technologies and field conditions, investigate domain-adaptation strategies, and incorporate uncertainty quantification for confidence-aware decision support.

Table 3. Comparative summary of EL-based electrical prediction frameworks for PV-module-level available in the literature [36-40] compared with our proposed model*.

Ref.	Year	Method	Main Findings	Main Limitations
[36]	2020	Pixel-intensity EL features + CNN cell classification + polynomial regression	Achieved adjusted R^2 = 0.88 for power prediction; CNN corrosion classification accuracy ≈95%; demonstrated Rs prediction (adjusted R^2 ≈0.73)	Requires normalized baseline data; limited to specific degradation modes (e.g., busbar corrosion); Not for half-cut PV topology modelling
[37]	2023	EL-based power regression using random forest, gradient boosting, RNN, CNN; ResNet50 feature extraction; I–V based ground truth	Best model: Gradient Boosting + ResNet50 (MAE ≈ 0.0341, MSE ≈ 0.0021); demonstrated direct cell-level power prediction from EL; synthetic balancing had limited impact	Limited dataset (603 real cells); laboratory-controlled single-cell setup; not validated on full modules or field conditions
[38]	2024	Two-stage deep learning pipeline: YOLOv8 (object detection) + MobileNetV2 (classification) + CNN regression	Fault detection: accuracy ≈0.93 (black-core), 0.868 (crack), 0.95 (edge); Power estimation nMAE ≈0.035, nRMSE ≈0.049; Series resistance nMAE ≈0.045	Small dataset (29 modules); laboratory EL only; limited fault types (3 classes); moderate object detection mAP (~0.46)
[39]	2025	Swin Transformer (SWinV2-B) to regress full normalized I–V curve (400 points) from EL images	Predicted MPP, Isc, Voc with <1% mean error; SDM Rs error ≈5%; demonstrated curve-level regression and k-fold validation (R^2 ≈0.99)	Limited to mildly degraded c-Si modules; poor generalization to unseen datasets (MPP error ≈13.9%)
[40]	2026	CNN regression mapping paired EL20+EL80 images to 200-point normalized I–V curve with smoothness regularization	MAE ≈1–2% for Isc, Pmpp, Impp, Vmpp; Voc error ≈0%; validation accuracy ≈95%; inference <0.5 s/module; single-EL image viable	Sensitive to EL image quality and cropping; normalized I–V only (no absolute ratings); limited to c-Si modules (Sandia dataset)
This Work	2026	Two-stage architecture-aware pipeline: cell-level segmentation + defect-aware regression with residual electrical consistency analysis	Architecture-aware I–V /power estimation; robust across heterogeneous degradation; improved stability in MPP region; industrial decision-level accuracy; interpretable residual diagnostics	Requires accurate segmentation stage; evaluated primarily on controlled EL datasets; field daylight EL validation ongoing

* Note: table contains the following abbreviations: RNN, recurrent neural network; nMAE, normalized mean absolute error; nRMSE, normalized root mean squared error; mAP, mean average precision; SDM, single-diode model.

5. Conclusion

This study addressed the critical limitation of qualitative EL-based inspection in laminated half-cut PV façade modules by developing a physics-based, two-stage deep learning framework for quantitative functional assessment. Unlike conventional defect detection approaches that focus solely on crack localization, the proposed method establishes a direct and measurable link between spatial defect morphology and module-level electrical degradation under standard test conditions.

The segmentation stage achieved a mean Dice coefficient of 0.86 in identifying crack and low-luminescence regions, demonstrating reliable spatial defect extraction across diverse degradation patterns. The subsequent regression stage predicted fractional power loss with a mean absolute error of 0.0183, a root mean squared error of 0.0243, and a coefficient of determination of 0.93. Residual analysis confirmed negligible systematic bias (MBE = 0.0041) and statistically stable prediction behavior. Furthermore, predicted degradation values enabled reconstruction of physically consistent I–V curves, with maximum power deviations remaining below 5.5% in representative cases.

Beyond continuous prediction accuracy, the framework was evaluated at the industrial decision level using Accept (<3%), Monitor (3–15%), and Reject (>15%) thresholds. The overall classification accuracy reached 83.4%, with zero false-accept cases for severely degraded modules and a precision of 0.96 for the Reject category. These results demonstrate risk-aware decision capability suitable for pre-installation certification workflows. The findings confirm that electroluminescence imaging can be transformed from a qualitative diagnostic tool into a quantitative, performance-oriented screening methodology for façade-integrated PV systems. By enabling objective defect-to-performance translation, the proposed approach supports improved reliability management, reduced premature module replacement, and enhanced lifecycle sustainability in building-integrated PV applications.

Future work will focus on expanding the dataset to include field-aged modules under varying environmental conditions, incorporating uncertainty quantification for probabilistic decision support, and extending the framework toward real-time deployment within automated drone-based EL inspection platforms.

Funding: No external funding received for the present research.

Availability of data and materials: The data will be made available on reasonable request to the corresponding author of the paper.

Code availability: Not applicable.

Competing interests: The authors declare that they have no competing interests.

References

[1] M. Sechilariu, B. Wang, and F. Locment, "Building integrated photovoltaic system with energy storage and smart grid communication," *IEEE Trans. Ind. Electron.*, vol. 60, no. 4, pp. 1607–1618, Apr. 2013.

[2] B. Liu, S. Duan, and T. Cai, "Photovoltaic DC-building-module-based BIPV system—Concept and design considerations," *IEEE Trans. Power Electron.*, vol. 26, no. 5, pp. 1418–1429, May 2011.

[3] Z. R. K. Abojela, M. K. M. Desa, and A. H. Sabry, "Current prospects of building-integrated solar PV systems and the application of bifacial PVs," *Front. Energy Res.*, vol. 11, Art. no. 1164494, 2023.

[4] G. Säckl, G. M. Wallner, J. Duchoslav, M. Tiefenthaler, and D. Stifter, "Advanced analysis of ethylene vinyl acetate copolymer materials for photovoltaic modules," *Polym. Test.*, vol. 132, Art. no. 108381, 2024.

[5] N. Bosco, "Turn your half-cut cells for a stronger module," *IEEE J. Photovolt.*, vol. 12, no. 5, pp. 1149–1153, Sept. 2022.

[6] S. Hwang and Y. Kang, "Reliability study on the half-cutting PERC solar cell and module," *Energy Rep.*, vol. 10, pp. 678–685, 2023.

[7] G. Goudelis, P. I. Lazaridis, and M. Dhimish, "A review of models for photovoltaic crack and hotspot prediction," *Energies*, vol. 15, no. 12, Art. no. 4303, 2022.

[8] A. R. Wittwer *et al.*, "Wind loading and its effects on photovoltaic modules: An experimental–computational study to assess the stress on structures," *Sol. Energy*, vol. 240, pp. 315–328, 2022.

[9] M. Dhimish, A. Ahmad, and A. M. Tyrrell, "Inequalities in photovoltaics modules reliability: From packaging to PV installation site," *Renew. Energy*, vol. 192, pp. 805–814, 2022.

[10] K. A. Kim, G. S. Seo, B. H. Cho, and P. T. Krein, "Photovoltaic hot-spot detection for solar panel substrings using AC parameter characterization," *IEEE Trans. Power Electron.*, vol. 31, no. 2, pp. 1121–1130, Feb. 2016.

[11] G. Badran and V. K. Lazarov, "From waste to resource: Exploring the current challenges and future directions of photovoltaic solar cell recycling," *Solar*, vol. 5, no. 1, Art. no. 4, Feb. 2025.

[12] A. M. Karimi, J. S. Fada, N. A. Parrilla, B. G. Pierce, M. Koyutürk, R. H. French, and J. L. Braid, "Generalized and mechanistic PV module performance prediction from computer vision and machine learning on electroluminescence images," *IEEE J. Photovolt.*, vol. 10, no. 3, pp. 878–887, May 2020.

[13] G. A. dos Reis Benatto, C. Mantel, S. Spataru, A. A. S. Lancia, N. Riedel, S. Thorsteinsson, P. B. Poulsen, H. Parikh, S. Forchhammer, and D. Sera, "Drone-based daylight electroluminescence imaging of PV modules," *IEEE J. Photovolt.*, vol. 10, no. 3, pp. 872–877, May 2020.

[14] K. G. Bedrich, W. Luo, M. Pravettoni, D. Chen, Y. Chen, Z. Wang, P. J. Verlinden, P. Hacke, Z. Feng, J. Chai, and Y. Wang, "Quantitative electroluminescence imaging analysis for performance estimation of PID-influenced PV modules," *IEEE J. Photovolt.*, vol. 8, no. 5, pp. 1281–1288, Sept. 2018.

[15] R. Yang, B. Du, P. Duan, Y. He, H. Wang, Y. He, and K. Zhang, "Electromagnetic induction heating and image fusion of silicon photovoltaic cell electrothermography and electroluminescence," *IEEE Trans. Ind. Informat.*, vol. 16, no. 7, pp. 4413–4422, July 2020.

[16] F. M. A. Mazen, R. A. A. Seoud, and Y. O. Shaker, "Deep learning for automatic defect detection in PV modules using electroluminescence images," *IEEE Access*, vol. 11, pp. 57783–57795, 2023.

[17] M. W. Akram, G. Li, Y. Jin, C. Zhu, A. Javaid, M. Z. Akram, and M. U. Khan, "Study of manufacturing and hotspot formation in cut cell and full cell PV modules," *Sol. Energy*, vol. 203, pp. 247–259, 2020.

[18] M. Dhimish and P. Mather, "Ultrafast high-resolution solar cell cracks detection process," *IEEE Trans. Ind. Informat.*, vol. 16, no. 7, pp. 4769–4777, July 2020.

[19] D. P. Winston, M. S. Murugan, R. M. Elavarasan, R. Pugazhendhi, O. J. Singh, P. Murugesan, M. Gurudhachanamoorthy, and E. Hossain, "Solar PV's micro crack and hotspots detection technique using NN and SVM," *IEEE Access*, vol. 9, pp. 127259–127269, 2021.

[20] Q. Liu, M. Liu, C. Wang, and Q. J. Wu, "An efficient CNN-based detector for photovoltaic module cells defect detection in electroluminescence images," *Sol. Energy*, vol. 267, Art. no. 112245, 2024.

[21] S. Hassan and M. Dhimish, "A survey of CNN-based approaches for crack detection in solar PV modules: Current trends and future directions," *Solar*, vol. 3, no. 4, pp. 663–683, Dec. 2023.

[22] C. Bu, R. Shen, W. Bai, P. Chen, R. Li, R. Zhou, J. Li, and Q. Tang, "CNN-based defect detection and classification of PV cells by infrared thermography method," *Nondestruct. Test. Eval.*, vol. 40, no. 5, pp. 1752–1769, 2025.

[23] E. Arnaudo, G. Blanco, A. Monti, G. Bianco, C. Monaco, P. Pasquali, and F. Dominici, "A comparative evaluation of deep learning techniques for photovoltaic panel detection from aerial images," *IEEE Access*, vol. 11, pp. 47579–47594, 2023.

[24] S. Hassan and M. Dhimish, "Enhancing solar photovoltaic modules quality assurance through convolutional neural network-aided automated defect detection," *Renew. Energy*, vol. 219, Art. no. 119389, 2023.

[25] M. K. Basher, M. Nur-E-Alam, M. M. Rahman, K. Alameh, and S. Hinckley, "Aesthetically appealing building integrated photovoltaic systems for net-zero energy buildings: Current status, challenges, and future developments—A review," *Buildings*, vol. 13, no. 4, Art. no. 863, 2023.

[26] Z. Wang, H. Chen, and Z. Cao, "CPDD: A cross-scenario photovoltaic defect detector based on fine-grained feature autoencoding and pseudo-box contrastive learning," *IEEE Trans. Semicond. Manuf.*, vol. 38, no. 3, pp. 612–623, Aug. 2025.

[27] C. Hirschl, M. Biebl-Rydlo, M. DeBiasio, W. Mühleisen, L. Neumaier, W. Scherf, G. Oreski, G. Eder, B. Chernev, W. Schwab, and M. Kraft, "Determining the degree of crosslinking of ethylene vinyl acetate photovoltaic module encapsulants—A comparative study," *Sol. Energy Mater. Sol. Cells*, vol. 116, pp. 203–218, 2013.

[28] X. Hu, T. Chen, J. Xue, G. Weng, S. Chen, H. Akiyama, and Z. Zhu, “Absolute electroluminescence imaging diagnosis of GaAs thin-film solar cells,” *IEEE Photon. J.*, vol. 9, no. 5, pp. 1–9, Oct. 2017.

[29] M. Abdelsattar, A. AbdelMoety, and A. Emad-Eldeen, “ResNet-based image processing approach for precise detection of cracks in photovoltaic panels,” *Sci. Rep.*, vol. 15, no. 1, Art. no. 24356, 2025.

[30] Y. Sun, G. Szűcs, and A. R. Brandt, “Solar PV output prediction from video streams using convolutional neural networks,” *Energy Environ. Sci.*, vol. 11, no. 7, pp. 1811–1818, 2018.

[31] L. Zhuang, Z. Zhang, and L. Wang, “The automatic segmentation of residential solar panels based on satellite images: A cross learning driven U-Net method,” *Appl. Soft Comput.*, vol. 92, Art. no. 106283, 2020.

[32] S. Hassanss and M. Dhimish, “Dual spin max pooling convolutional neural network for solar cell crack detection,” *Sci. Rep.*, vol. 13, no. 1, Art. no. 11099, 2023.

[33] G. Ruan, X. Chen, Y. Li, E. G. Lim, L. Fang, L. Jiang, Y. Du, and F. Wang, “SkyNet: A deep learning architecture for intra-hour multimodal solar forecasting with ground-based sky images,” *Renew. Energy*, Art. no. 124354, 2025.

[34] D. Marmanis, M. Datcu, T. Esch, and U. Stilla, “Deep learning earth observation classification using ImageNet pretrained networks,” *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 1, pp. 105–109, Jan. 2016.

[35] J. Liu, Y. Liu, and Q. Zhang, “A weight initialization method based on neural network with asymmetric activation function,” *Neurocomputing*, vol. 483, pp. 171–182, 2022.

[36] A. M. Karimi, J. S. Fada, N. A. Parrilla, B. G. Pierce, M. Koyutürk, R. H. French, and J. L. Braid, “Generalized and mechanistic PV module performance prediction from computer vision and machine learning on electroluminescence images,” *IEEE J. Photovolt.*, vol. 10, no. 3, pp. 878–887, May 2020.

[37] H. F. M. Romero, L. Hernández-Callejo, M. Á. G. Rebollo, V. Cardenoso-Payo, V. A. Gomez, J. I. M. Aragonés, and R. T. Moyo, “Optimized estimator of the output power of PV cells using EL images and I–V curves,” *Sol. Energy*, vol. 265, Art. no. 112089, 2023.

[38] E. Ozturk, E. Ogliari, M. Sakwa, A. Dolara, N. Blasuttigh, and A. M. Pavan, “Photovoltaic modules fault detection, power output, and parameter estimation: A deep learning approach based on electroluminescence images,” *Energy Convers. Manage.*, vol. 319, Art. no. 118866, 2024.

[39] B. K. Byford, L. E. Boucheron, B. H. King, and J. L. Braid, “Advanced photovoltaic module characterization: Using image transformers for current–voltage curve prediction from electroluminescence images,” *IEEE J. Photovolt.*, vol. 15, no. 4, pp. 557–565, July 2025.

[40] M. Dhimish, G. A. dos Reis Benatto, R. G. Vieira, and P. B. Poulsen, “EL-to-IV: Deep learning-based prediction of photovoltaic current-voltage curves from electroluminescence imaging,” *Energies*, vol. 19, no. 4, Art. no. 876, 2026.