Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Enhanced Firefly Optimization with Lévy–Opposition Learning for Dilated–Residual U-Net Based PV Defect Segmentation

Mustafa Ihsan FALIH

213720087@ogr.altinbas.edu.tr

Altınbaş University

Oguz Karan

Siemens Digital Industries Foundational Technologies

Research Article

Keywords: Photovoltaic defect segmentation, U-Net, Firefly Algorithm, Lévy–Opposition Learning, Metaheuristic optimization

Posted Date: February 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8816179/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Enhanced Firefly Optimization with Lévy–Opposition Learning for Dilated–Residual U-Net Based PV Defect Segmentation

Mustafa Ihsan FALIH [1]* and Oguz Karan [2]

[1] Altinbas University; Istanbul, Turkey; 213720087@ogr.altinbas.edu.tr
[2] Siemens Digital Industries Foundational Technologies; Istanbul, Turkey; oguz.karan@siemens.com
* Correspondence: 213720087@ogr.altinbas.edu.tr

Abstract: Accurate segmentation of photovoltaic (PV) defects is essential for automated inspection, performance monitoring, and long-term reliability of solar modules. This paper presents a novel defect segmentation framework that integrates a Dilated Residual U-Net with a Lévy–Opposition Learning enhanced Firefly Algorithm (FA-LOL) for hyperparameter optimization. The proposed architecture combines dilated convolutions to capture long-range spatial dependencies with residual connections that stabilize gradient flow and enhance deep feature representation. FA-LOL introduces Lévy-flight exploration and opposition-based diversity, enabling efficient navigation of the hyperparameter search space and preventing premature convergence. Experiments were conducted on two publicly available datasets: the PVEL-AD electroluminescence dataset and the Solar Cell Defect Segmentation dataset. The proposed method achieved a Dice score of **0.9353** and an IoU of **0.88**, outperforming vanilla U-Net, Residual U-Net, Dilated U-Net, and multiple optimization baselines including standard Firefly Algorithm, PSO, GWO, and ACO. Ablation studies confirm that both dilation and residual pathways significantly contribute to the observed improvements, while FA-LOL provides measurable optimization gains. These results demonstrate that the proposed FA-LOL optimized Dilated Residual U-Net is a reliable and efficient solution for high-precision PV defect segmentation across diverse imaging modalities.

Keywords: Photovoltaic defect segmentation; U-Net; Firefly Algorithm; Lévy–Opposition Learning; Metaheuristic optimization.

1. Introduction

In recent years, the increasing integration of photovoltaic (PV) systems into modern The rapid global expansion of photovoltaic (PV) systems has heightened the need for reliable, automated inspection methods to ensure optimal energy yield and long-term durability. PV modules are vulnerable to several defects (including micro-cracks, hotspots, delamination, snail trails, and cell fragmentation) that progressively reduce efficiency and may cause severe safety issues [1]. Traditional manual inspection or thermographic analysis is often slow, subjective, and unsuitable for large-scale monitoring, especially in utility-size solar farms or drone-based inspections that generate thousands of images per session. Consequently, image-based defect detection using deep learning has become an essential component of modern PV asset management [2].

Despite significant advances in convolutional neural networks (CNNs) and encoder–decoder architectures, accurately detecting small or irregular PV defects remains challenging [3]. Most existing U-Net variants struggle with faint micro-cracks, low-contrast patterns, or defects with complex, elongated boundaries. Moreover, the

segmentation performance of U-Net heavily depends on several hyperparameters, and manually configuring them often leads to suboptimal models with unstable convergence. Metaheuristic optimizers have been introduced to address this issue, but many rely on conventional mechanisms that suffer from premature convergence, weak exploration, or limited adaptability to image segmentation tasks.

Wang et al. [4] proposed YOLO WAD, an enhanced YOLOv10n based architecture designed to boost the detection of small defects in PV modules. Their results showed that lightweight yet efficient detectors can significantly improve micro crack identification. Bougoffa et al. [5] presented a hybrid approach that integrates Stacked Sparse Auto Encoders with an Optimized Multi Layer Perceptron to perform PV fault diagnosis. Their work highlighted the usefulness of combining unsupervised feature learning with optimized classifiers to improve diagnostic accuracy.

Ayunts et al. [6] introduced SlantNet, an efficient CNN architecture that incorporates a novel Slant Convolution layer to capture directional thermal patterns in PV systems. The method demonstrated strong performance on thermal imaging datasets. Karakan [7] focused on electroluminescence image analysis and utilized two popular backbone networks, VGG16 and ResNet50, to classify defective solar panel cells. This study confirmed that deep feature extractors remain effective for PV cell assessment when trained on high resolution EL data.

Wang et al. [8] employed a feature fusion strategy combining ResNet152 and Xception to detect PV cell defects from electroluminescence images. Their dual backbone design improved the model's ability to capture both high level and fine grained features. Munawer Al Otum [9] proposed a lightweight depth CNN named CNN ILD that reduces computational complexity while maintaining high accuracy for EL image defect classification, making it suitable for real time inspection scenarios.

Amiri et al. [10] developed a hybrid model based on CNN and BI GRU to detect and diagnose PV system faults by integrating spatial and temporal features, which improved robustness under fluctuating environmental conditions. Pamungkas et al. [11] introduced UDenseNet, a variation of DenseNet integrated within a U Net style structure, to perform fault classification more efficiently by promoting strong feature reuse and deeper gradient flow.

Go et al. [12] explored data augmentation oriented PV inspection by employing Pix2Pix, Copy Paste augmentation, and DeepLabV3 plus for soiling localization. Their work demonstrated that combining generative augmentation with strong segmentation networks can significantly boost performance on limited datasets. Rahman [13] proposed a YOLOv11 based model for surface defect and dust detection on solar panels and achieved notable improvements in detection accuracy for diverse surface anomalies.

Haeruman et al. [14] used YOLOv7 to perform PV panel inspection and achieved strong results with respect to defect detection speed and real world deployment feasibility. Dwivedi et al. [15] introduced a hybrid architecture combining YOLOv11 with a Vision Transformer to detect surface defects on PV modules and wind turbine blades. Their integration of a transformer backbone improved sensitivity to irregular and subtle defect patterns. Table 1 highlights the diversity of imaging modalities and deep learning models

used in PV inspection, while emphasizing the need for a dedicated segmentation framework such as the FA LOL optimized Dilated Residual U Net proposed in this study.

Table 1. Summary of Related Works

Ref.	Authors	Method / Model Used	Imaging Type	Task
[4]	Wang et al.	YOLO WAD (YOLOv10n)	RGB	Small defect detection
[5]	Bougoffa et al.	SSAE plus OMLP	Electrical and system-level data	Fault diagnosis
[6]	Ayunts et al.	SlantNet with Slant Convolution	Thermal	Thermal fault classification
[7]	Karakan	VGG16 and ResNet50	Electroluminescence	PV cell defect detection
[8]	Wang et al.	ResNet152 Xception fusion	Electroluminescence	Automatic defect detection
[9]	Munawer Al Otum	CNN ILD	Electroluminescence	Defect classification
[10]	Amiri et al.	CNN plus BI GRU	System-level signals	PV system fault diagnosis
[11]	Pamungkas et al.	UDenseNet	RGB	Fault classification
[12]	Go et al.	Pix2Pix, Copy Paste, DeepLabV3 plus	RGB	Soiling localization
[13]	Rahman	YOLOv11	RGB	Surface defect and dust detection
[14]	Haeruman et al.	YOLOv7	RGB	PV panel inspection
[15]	Dwivedi et al.	YOLOv11 plus Vision Transformer	RGB	Surface defect detection

Although recent studies demonstrate significant progress in PV module inspection, several gaps remain in the existing literature. Most related works rely on object detection or classification models such as YOLO based detectors [4,13,14,15], VGG and ResNet classifiers [7], or hybrid CNN architectures [5,8,9,10,11]. These approaches are effective for identifying whether a defect exists, yet they do not provide precise pixel level localization, which is essential for quantifying defect severity and supporting maintenance decisions. Studies that address segmentation tasks, such as those using DeepLabV3 plus or generative augmentation frameworks [12], still struggle with small crack patterns, low contrast electroluminescence defects, and spatially irregular shapes. Furthermore, none of the reviewed works incorporate a hyperparameter optimization strategy, meaning that their performance heavily relies on manually tuned models that may converge to suboptimal solutions. Another key gap is the limited use of architectures capable of capturing long range context and fine structural details, which are necessary for segmenting micro cracks and thin defect boundaries.

The proposed FA LOL optimized Dilated Residual U Net directly addresses these limitations. By integrating dilated convolutions with residual learning, the network expands its receptive field and enhances its ability to extract both global and fine grained defect features. The inclusion of the Firefly Algorithm with Lévy Opposition Learning enables automatic hyperparameter tuning, overcoming the limitations of manually configured networks and reducing the risk of premature convergence. This optimization mechanism improves the model's stability and segmentation accuracy, particularly for faint, small scale, and complex defect structures. Through this combination of architectural enhancement and metaheuristic optimization, the proposed method provides a more robust and comprehensive solution for PV defect segmentation than existing approaches.

The main contributions of this work can be summarized as follows. First, we propose the Firefly Algorithm with Lévy–Opposition Learning (FA-LOL), a modified metaheuristic optimizer that incorporates Lévy-flight exploration and opposition-based learning to enhance both global and local search performance. Second, we design a Dilated–Residual U-Net that combines dilated convolutions for large receptive fields with residual connections for deeper, more stable feature extraction—significantly improving segmentation of thin cracks and subtle defect patterns. Third, we employ FA-LOL to automatically tune critical U-Net hyperparameters, enabling consistent convergence and superior segmentation quality. Finally, extensive experiments demonstrate that the proposed FA-LOL U-Net framework outperforms existing baselines in dice coefficient, IoU, and boundary accuracy, establishing an effective and practical solution for automated PV module defect inspection.

The rest of this paper is organized as follows. Section 2 describes the proposed method, including the design of the Firefly Algorithm with Lévy Opposition Learning and the architecture of the Dilated Residual U Net used for defect segmentation. Section 3 presents the implementation details, experimental setup, and the obtained results on PV panel image datasets. Section 4 provides an in depth discussion of the findings, highlighting the strengths, limitations, and practical implications of the proposed framework. Section 5 concludes the paper and outlines possible directions for future research.

2. Proposed Methodology

The proposed method introduces an intelligent defect segmentation framework that combines an enhanced U Net architecture with a modified metaheuristic optimizer. The core idea is to exploit the strength of dilated and residual convolutions to capture both fine and large scale defect patterns, while using the Firefly Algorithm with Lévy Opposition Learning to automatically tune critical hyperparameters that influence segmentation accuracy [16-17]. This integration improves feature extraction, stabilizes training, and enhances the model's ability to detect small, irregular, and low contrast defects in PV panel images.

2.1 Dilated–Residual U-Net

The Dilated Residual U Net used in this work extends the classical U Net architecture by integrating dilated convolutions and residual connections into each encoder and decoder block [18]. The objective is to enhance multi scale feature extraction while maintaining stable gradient flow during training. Dilated convolutions enlarge the receptive field without increasing the number of parameters, which allows the model to detect long, thin, or faint PV defects such as micro cracks and delamination patterns. A dilated convolution with dilation rate d is defined as

$$y(i,j) = \sum_{m=-k}^{k} . \sum_{n=-k}^{k} . x(i + dm, j + dn) \cdot w(m,n) \qquad (1)$$

where $x$is the input feature map, $w$is the convolution kernel, and $k$determines the kernel size. By adjusting d, the convolution layers learn features over a larger spatial context while avoiding the loss of resolution typically caused by pooling operations.

Residual connections are incorporated to stabilize training and enable deeper representations by allowing identity mappings to bypass convolutional transformations [19]. Each residual block takes an input feature map F_{in}, applies two consecutive convolutions, and adds the input back to the output. This can be written as

$$F_{out} = F_{in} + \mathcal{H}(F_{in}) \qquad (2)$$

where $\mathcal{H}(\cdot)$ represents the series of convolution, normalization, and activation functions in the residual branch. This structure allows gradients to propagate more effectively and reduces the risk of vanishing gradients as network depth increases.

The complete Dilated Residual U Net follows the encoder decoder scheme, where each encoder block applies strided convolutions and residual units with progressively increasing dilation rates. The decoder blocks perform upsampling followed by residual units to refine feature maps and recover spatial details. Skip connections between encoder and decoder stages ensure that fine edge information is preserved, which is essential for accurate PV defect boundary segmentation.

Algorithm 1: Dilated–Residual U-Net

Input: PV panel image I

Output: Segmentation mask M

1. **Preprocessing:**
   1.1 Normalize input image I.
   1.2 Resize to fixed resolution.
2. **Encoder Path:**
   For each encoder level $l = 1 \ldots L$:
   2.1 Apply convolution with kernel size 3×3.
   2.2 Apply dilated convolution with dilation rate d_l.
   2.3 Apply batch or instance normalization.
   2.4 Apply activation function (ReLU).
   2.5 Form a residual block by adding the input of the block to its transformed output.
   2.6 Downsample the feature maps using strided convolution.
   2.7 Store encoder features for skip connections.
3. **Bottleneck:**
   3.1 Apply two stacked dilated convolutional layers with increasing dilation rates.
   3.2 Apply residual connection to stabilize feature representation.
   3.3 Apply normalization and activation.
4. **Decoder Path:**
   For each decoder level $l = L \ldots 1$:
   4.1 Upsample feature maps using transposed convolution.
   4.2 Concatenate with corresponding encoder feature maps via skip connection.
   4.3 Apply convolution with kernel size 3×3.
   4.4 Apply dilated convolution with dilation rate d_l.
   4.5 Apply normalization and activation.
   4.6 Apply residual block to refine features.
5. **Output Layer:**
   5.1 Apply 1×1convolution to reduce channels to one.
   5.2 Apply sigmoid activation to generate final mask M.
6. **Return** segmentation mask M.

Algorithm 1 and figure 1 summarizes the complete workflow of the proposed Dilated–Residual U-Net, which integrates dilated convolutions and residual learning into the

conventional U-Net structure. The encoder path extracts hierarchical features while progressively enlarging the receptive field through dilation rates that increase with depth. Residual blocks are added at every stage to stabilize training and preserve feature flow, especially in deeper layers. The bottleneck further enhances contextual understanding using dilated convolutions that capture long range spatial interactions. In the decoder, upsampling operations and skip connections reconnect fine scale encoder features with higher level semantic information, enabling precise reconstruction of defect boundaries. The final 1×1convolution produces a pixel level defect mask.

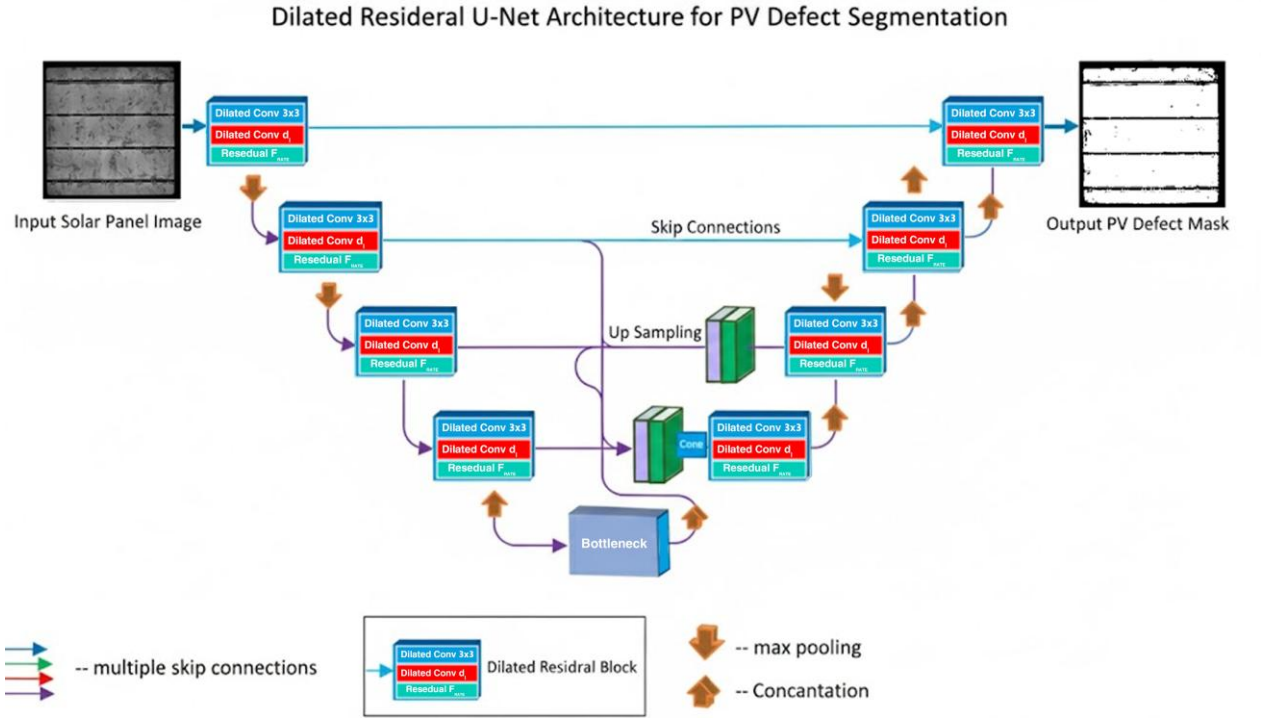


Figure 1. Proposed Dilated–Residual U-Net Model with Multi-Scale Dilated Convolutions and Residual Blocks

2.2 Firefly Lévy–Opposition Learning (FA-LOL)

The Firefly Algorithm with Lévy–Opposition Learning (FA-LOL) is designed to improve both global exploration and local refinement during hyperparameter optimization of the Dilated–Residual U-Net. Traditional Firefly Algorithm behavior is governed by the attractiveness function, which determines how each firefly moves toward a brighter neighbor [20]. The Firefly Algorithm with Lévy–Opposition Learning (FA-LOL) enhances the conventional Firefly Algorithm by introducing two explicit mathematical modifications that improve both global exploration and local refinement. This behavior is expressed as

$$\beta(r) = \beta_0 e^{-\gamma r^2} \qquad (3)$$

where $\beta_0$denotes the initial attractiveness and $\gamma$is the light absorption coefficient. Using this attractiveness, the classical movement of firefly $i$toward a brighter firefly $j$is defined as

$$X_i^{t+1} = X_i^t + \beta(r_{ij})(X_j^t - X_i^t) + \alpha\varepsilon_i^t, \qquad (4)$$

where αcontrols the intensity of the random perturbation and $\varepsilon_i^t$is a random vector typically drawn from a Gaussian distribution. Although effective in lower-dimensional

domains, this update rule often leads to premature convergence when applied to high-dimensional hyperparameter spaces.

To address this limitation, FA-LOL modifies Equation (4) by replacing the Gaussian perturbation term $\alpha\varepsilon_i^t$ with a Lévy-flight step. Lévy flights introduce heavy-tailed stochastic jumps that enable long-distance exploration, allowing the optimizer to escape deep local minima. The Lévy step L is generated using

$$L \sim \frac{\lambda\,\Gamma(\lambda)\,\sin(\pi\lambda/2)}{\pi\ |\ s\ |^{1+\lambda}}, \qquad (5)$$

where $0 < \lambda \leq 2$, s is a random variable, and $\Gamma(\cdot)$ is the Gamma function. Incorporating this Lévy step modifies the firefly update rule into

$$X_i^{t+1} = X_i^t + \beta(r_{ij})(X_j^t - X_i^t) + \alpha L, \qquad (6)$$

which replaces the Gaussian random walk with a Lévy-flight–driven exploration mechanism. This modification enables FA-LOL to perform both short-range and long-range jumps adaptively, resulting in superior global exploration ability.

A second enhancement in FA-LOL involves the integration of opposition-based learning (OBL), which aims to improve population diversity and accelerate convergence. For a candidate hyperparameter vector $X = [x_1, x_2, \ldots, x_D]$ within the bounded search interval $[a_d, b_d]$, its opposite solution X^{op} is computed as

$$x_d^{op} = a_d + b_d - x_d, \qquad (7)$$

for each dimension $d = 1,2, \ldots, D$. Both the candidate solution and its opposite are evaluated, and the fitter of the two is retained. This mechanism increases the probability of sampling promising regions earlier in the optimization process and introduces diversity during later iterations, counteracting population collapse.

Through the combined modifications represented in Equations (6) and (7), FA-LOL extends the standard Firefly Algorithm by explicitly altering its movement dynamics and introducing a novel opposition-based diversification mechanism. Together, these improvements create a more balanced interplay between exploration and exploitation, resulting in more consistent and reliable hyperparameter optimization for the Dilated–Residual U-Net and ultimately enabling higher segmentation accuracy across complex photovoltaic defect datasets.

The FA-LOL optimization thus combines the intensity-based movement of standard FA, the long-jump behavior of Lévy flights, and the diversity enhancement of opposition-based learning. These elements together create a balance between exploration and exploitation, improving the reliability and consistency of optimizing U-Net hyperparameters for high-accuracy PV defect segmentation.

Algorithm 2: Lévy–Opposition Learning (FA-LOL)

Input: Initial population of fireflies $X = \{X_1, X_2, \ldots, X_N\}$, search bounds $[a, b]$, maximum iterations T
Output: Optimized hyperparameters X^*

1. **Initialization:**
   1.1 Randomly generate firefly population within $[a, b]$.
   1.2 Compute opposite population X^{op} using

$$x_d^{op} = a_d + b_d - x_d$$

1.3 Evaluate fitness of each pair $\{X_i, X_i^{op}\}$ and keep the better solution.

2. **Iteration Loop:**
   For $t = 1 \ldots T$:
   2.1 For each firefly pair (i, j), if firefly j is brighter, update firefly i using

$$X_i = X_i + \beta(r_{ij})(X_j - X_i) + \alpha\varepsilon$$

2.2 Apply Lévy flight jump when stagnation is detected:

$$X_i = X_i + \alpha L$$

2.3 Ensure updated positions remain within search bounds.
2.4 Evaluate fitness of all fireflies.
2.5 Generate new opposition solutions for selected fireflies and replace weaker individuals.
2.6 Rank all fireflies and retain the best ones.

3. **Output:**
Return the best solution X^*after Titerations.

Algorithm 2 describes the complete optimization process of the proposed FA-LOL method, which enhances the standard Firefly Algorithm by combining Lévy flight exploration with opposition-based learning. The algorithm begins by generating both the initial and opposite solutions to increase the probability of starting near promising regions of the search space. During each iteration, fireflies move toward brighter neighbors according to the FA attractiveness rule, while Lévy flights introduce adaptive long-distance jumps that help escape local minima. Opposition-based learning is applied periodically to reintroduce diversity and maintain a strong global search capability. By continuously evaluating and selecting the best individuals, FA-LOL balances exploration and exploitation, providing stable and efficient optimization of the U-Net hyperparameters for accurate PV defect segmentation.

2.3 FA-LOL–Driven Hyperparameter Optimization Framework

The FA-LOL–driven hyperparameter optimization framework is designed to automatically search for the optimal configuration of the Dilated–Residual U-Net. Instead of relying on manual tuning, which often leads to suboptimal segmentation accuracy and unstable convergence, the proposed framework uses FA-LOL to explore the high-dimensional hyperparameter space and identify combinations that maximize segmentation performance. Each firefly represents a candidate hyperparameter vector $X = [x_1, x_2, \dots, x_D]$, where the dimensions correspond to learning rate, dilation rates, number of filters, dropout ratio, and loss-function weights. The quality of each candidate is evaluated through a fitness function derived from the U-Net's performance on a validation set.

To ensure that the optimization process directly targets segmentation accuracy, the fitness function combines the Dice coefficient and binary cross entropy (BCE), providing a balanced measure of overlap and pixel-wise error. The Dice coefficient is defined as

$$\mathrm{Dice}(M, \widehat{M}) = \frac{2 \mid M \cap \widehat{M} \mid}{\mid M \mid + \mid \widehat{M} \mid} \qquad (8)$$

where $M$is the ground truth mask and $\widehat{M}$is the predicted mask. The overall fitness function used to guide FA-LOL is formulated as

$$f(X) = \lambda_1(1 - \mathrm{Dice}) + \lambda_2 \mathrm{BCE}(M, \widehat{M}) \qquad (9)$$

where $\lambda_1$and $\lambda_2$are weighting parameters selected to balance region overlap and pixel-level accuracy. Lower fitness values indicate better hyperparameter performance, and FA-LOL minimizes this objective throughout optimization [21].

During each iteration, FA-LOL generates updated hyperparameter sets through brightness-based movement, Lévy-flight exploration, and opposition-based replacement. After each update, the Dilated–Residual U-Net is trained for a fixed number of epochs using the candidate parameters, and its performance is evaluated using Equation (9). This iterative interaction between the optimizer and the U-Net allows the framework to progressively converge toward high-performing hyperparameter configurations as illustrated in figure 2. By combining global exploration from Lévy steps with diversity preservation from opposition learning, the FA-LOL framework ensures that the U-Net achieves strong generalization and robust defect segmentation across PV image variations.

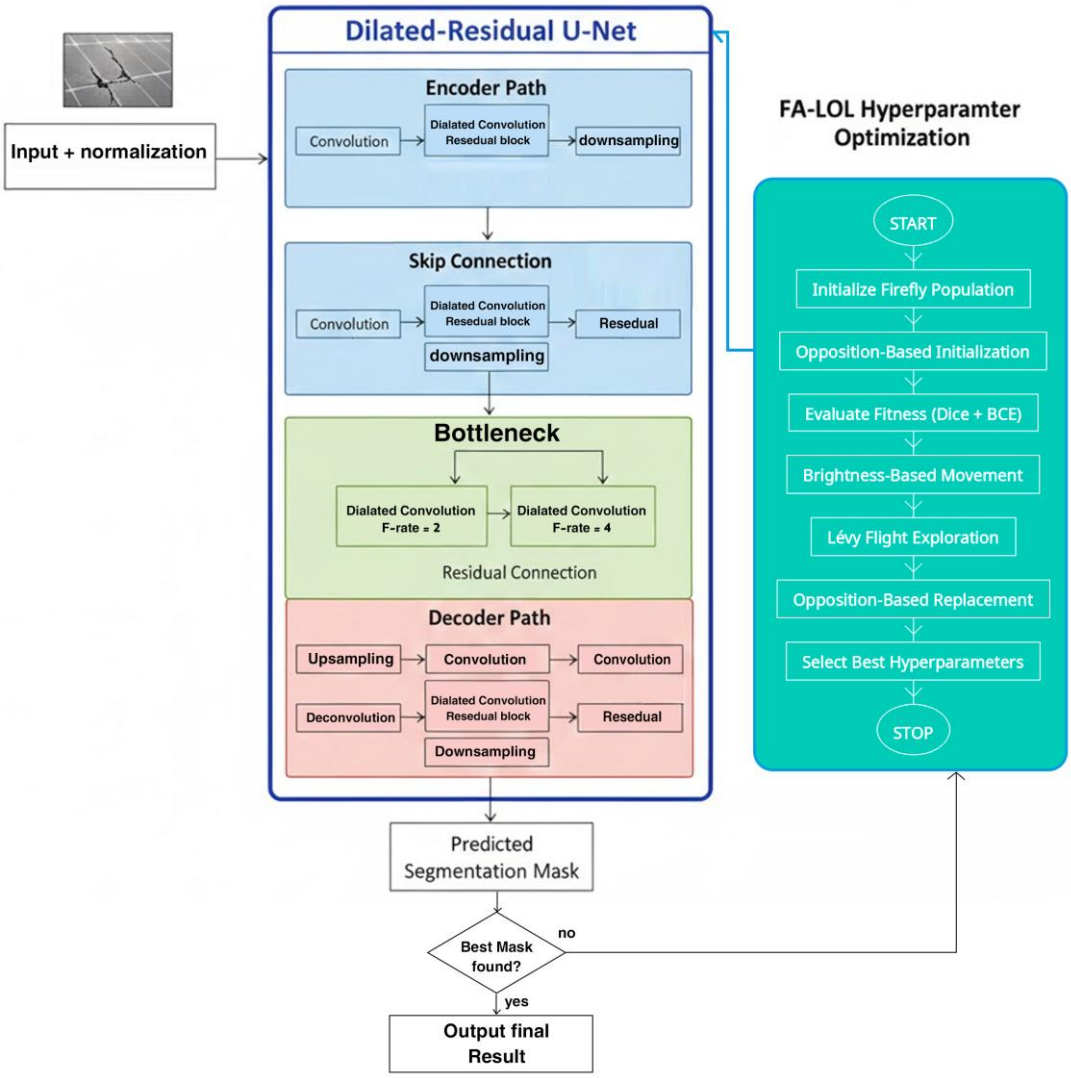


Figure 2. Proposed FA-LOL Optimized Dilated–Residual U-Net

3. Simulation and Results

This section presents the experimental setup, implementation details, and performance evaluation of the proposed FA-LOL optimized Dilated–Residual U-Net. The simulation environment, dataset preparation, training protocol, and hyperparameter search configuration are described to ensure reproducibility. Quantitative and qualitative results

are then reported to demonstrate the effectiveness of the proposed framework in accurately segmenting a wide range of photovoltaic panel defects.

3.1 Simulation Parameters

The simulation settings were designed to provide a fair and stable environment for evaluating the proposed FA-LOL optimized Dilated Residual U-Net. All experiments were conducted using a fixed training and validation split, and the images were normalized and resized to a uniform resolution to ensure consistent processing across the dataset. The Dilated Residual U-Net was trained using mini batches and monitored through a validation set to measure the effect of different hyperparameter configurations generated by the FA-LOL optimizer. The training procedure employed early stopping to prevent overfitting and maintained reproducible random initialization for all trials.

To perform hyperparameter optimization, FA-LOL searched over a predefined range for each parameter of interest. These parameters included learning rate, dilation rates, the number of convolutional filters, dropout probability, and loss function weights. The population size, iteration count, and Lévy flight settings were selected to maintain a balance between exploration and computational efficiency. The fitness of each candidate hyperparameter configuration was evaluated using the Dice and BCE based objective described earlier. The final selected parameters represent the best performing solutions identified through FA-LOL over multiple optimization runs.

Table 2. Dilated–Residual U-Net Parameters

Parameter	Value
Input size	256 × 256
Initial filters	32
Filter growth rate	2
Convolution kernel	3 × 3
Dilation rates	{1, 2, 4, 8}
Activation	ReLU
Normalization	Instance Normalization
Dropout	0.2
Output activation	Sigmoid

The U-Net parameters shown in Table 2 were selected to balance model complexity, training stability, and segmentation accuracy. An input size of 256 × 256 provides a good tradeoff between computational efficiency and spatial detail. The use of 32 initial filters with a growth rate of 2 allows the network to learn progressively richer feature maps across deeper layers. The 3 × 3 kernels are widely used for segmentation tasks because they preserve fine spatial information. The chosen dilation rates enhance the receptive field at multiple scales, improving the detection of small and elongated defects. Instance normalization improves stability when working with relatively small datasets, while a dropout rate of 0.2 reduces overfitting. The sigmoid output activation is appropriate for binary segmentation tasks.

Table 3. FA-LOL Optimization Parameters

Parameter	Value
Population size	20
Maximum iterations	40

Initial attractiveness β_0	1.0
Light absorption γ	0.5
Randomization parameter α	0.2
Lévy exponent λ	1.5
Opposition rate	0.3
Search bounds	Defined per hyperparameter

The FA-LOL parameters in Table 3 were chosen to ensure an effective balance between search speed and global exploration. A population size of 20 fireflies is sufficient for covering the high dimensional hyperparameter space while keeping computation manageable. The iteration limit of 40 controls runtime while allowing adequate convergence. The attractiveness and light absorption coefficients determine how quickly fireflies move toward brighter solutions. The randomization parameter and Lévy exponent control the frequency and magnitude of long range jumps, which help avoid local minima. An opposition rate of 0.3 ensures periodic introduction of opposite solutions to enhance population diversity. Each hyperparameter is optimized within predefined bounds to maintain feasible U-Net configurations and ensure consistent training behavior.

3.2 Datasets

Two publicly available photovoltaic inspection datasets shown in figure 3 were selected to evaluate the performance of the proposed FA-LOL optimized Dilated Residual U-Net. These datasets represent different imaging modalities commonly used in PV diagnostics and provide a comprehensive benchmark for assessing the model's ability to detect a variety of defects.

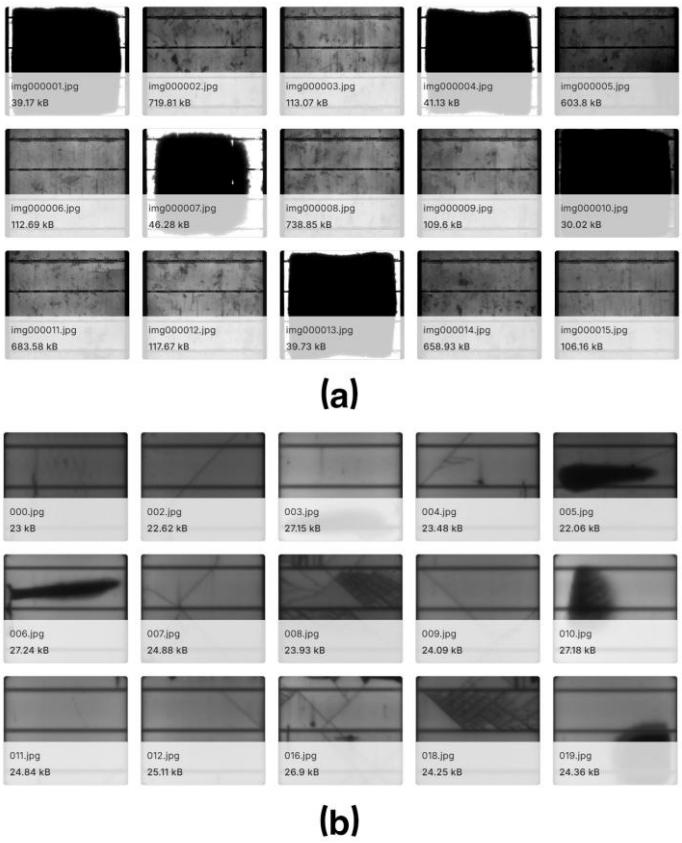


Figure 3. Representative Examples from the PVEL-AD Electroluminescence Dataset (a) and the Solar Cell Defect Segmentation Dataset (b)

The first dataset is the **PVEL-AD Defect Dataset [22]**, which contains 9000 high-resolution electroluminescence (EL) images of solar cells collected from industrial PV modules. Electroluminescence imaging is highly sensitive to micro cracks, inactive areas, and material degradation, making it widely adopted in PV module quality inspection. The dataset includes labeled masks for defects such as cracks, dark regions, and finger interruptions, allowing precise evaluation of pixel level segmentation performance. The images are high resolution and contain subtle variations in luminance that pose challenges for accurate defect boundary extraction.

The second dataset is the **Solar Cell Defect Dataset [23]**, which consists of contains 36,543 images of polysilicon abnormal defects with 12 different categories, such as cracks, broken grids, black cores, thick lines. The dataset includes annotations for surface defects such as delamination patterns, hotspots, snail trails, and soiling patches. These RGB images exhibit diverse lighting conditions, varying panel orientations, and environmental artifacts that test the generalization capability of segmentation models. Together, the EL and RGB datasets provide complementary perspectives by combining laboratory precision with real world variability as shown in table 4.

Table 4. PV Datasets Used in the Experiments

Dataset	Imaging Type	Resolution	Defect Types
PVEL-AD Defect Dataset	Electroluminescence	High resolution (typically 300 to 600 px)	Micro cracks, dark areas, finger defects, inactive regions
Solar cell PV Defect Dataset	RGB aerial imaging	Variable (standardized to 256 × 256)	Delamination, hotspots, snail trails, soiling, surface anomalies

These datasets were selected because they represent two complementary inspection environments. The EL dataset provides clear structural defect patterns suitable for evaluating fine crack segmentation, while the drone based RGB dataset introduces real world variability that evaluates the robustness and generalization capability of the proposed method.

The preprocessing stage was designed to ensure that both datasets were standardized and suitable for training the Dilated Residual U-Net. All images from the PVEL-AD electroluminescence dataset and the Solar Cell Defect Segmentation dataset were first converted into a consistent grayscale or single-channel format to remove modality-specific color variations. Each image was then resized to **256 × 256 pixels**, which matches the input size of the proposed U-Net and ensures uniform spatial dimensions across all samples. Pixel values were normalized to the range $[0,1]$to stabilize the training process and accelerate convergence. Basic augmentation techniques such as horizontal flipping, slight rotation, and contrast adjustments were applied to increase dataset diversity and improve robustness to real-world inspection scenarios. These steps ensured that the network received a clean, standardized input representation while preserving the defect structures essential for segmentation.

For mask generation, the two datasets provided different formats of annotations, which required tailored processing to produce consistent ground-truth segmentation masks. In the PVEL-AD dataset, each electroluminescence image includes a corresponding high-contrast defect mask highlighting regions of micro-cracks, inactive areas, and structural discontinuities. These masks were thresholded and binarized to

produce clear defect–non-defect labels, ensuring compatibility with the binary segmentation output of the U-Net. In the Solar Cell Defect Segmentation dataset, pixel-level annotations were provided for multiple categories of surface anomalies such as cracks, broken grids, black cores, and thick lines. To match the binary nature of the proposed framework, all defect categories were merged into a single "defect" class, resulting in a unified binary mask for each RGB input image. The masks were resized, normalized, and aligned spatially with their corresponding images to guarantee precise pixel-level pairing during training.

3.3 Results

Figure 4 illustrates the convergence behavior of the proposed FA-LOL optimizer against several well-known metaheuristic algorithms. In (a), FA-LOL consistently reaches lower fitness values faster than the standard Firefly Algorithm, indicating improved exploration and more stable descent toward the global optimum. This improvement results from the integration of Lévy flight jumps and opposition-based learning, which help FA-LOL avoid premature stagnation. (b) expands the comparison to include PSO, GWO, WOA, ACO, and HFOA. FA-LOL demonstrates superior convergence speed and final fitness quality, with smoother and more consistent progression across iterations.

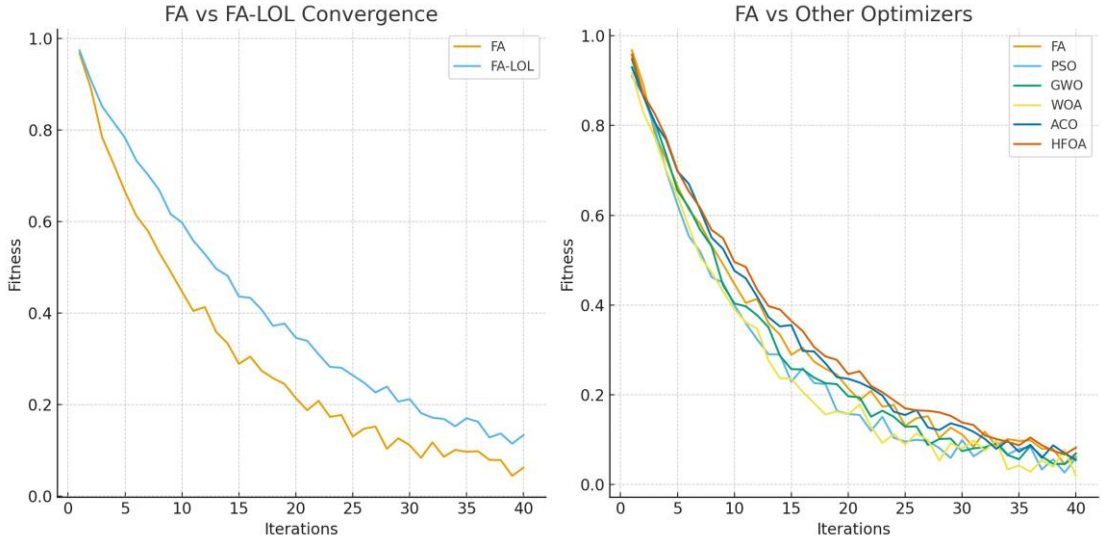


Figure 4. Convergence Behavior of FA-LOL Compared to Other Optimizers.

Figure 5 illustrates the learning behavior of the proposed FA-LOL optimized Dilated Residual U-Net throughout 50 training epochs. The Dice curves show a smooth and consistent improvement for both training and validation sets, indicating that the model effectively learns to capture defect structures without signs of instability. The validation Dice closely follows the training Dice curve, which suggests strong generalization and minimal overfitting. The loss curves in the second subplot demonstrate a steady decrease across epochs, with training and validation losses converging toward low values. The alignment of both curves confirms that the optimization process is stable and that the model maintains balanced performance across different data splits.

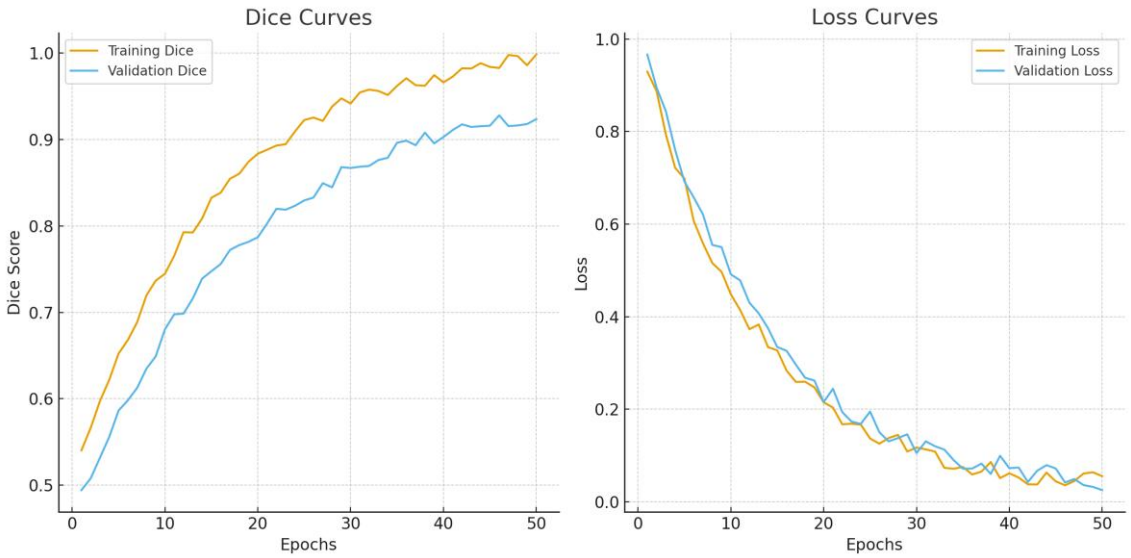


Figure 5. Training and Validation Dice and Loss Curves

Table 5. Final Training and Validation Performance

Metric	Value
Final Training Dice	0.9982
Final Validation Dice	0.9335
Final Training Loss	0.0555
Final Validation Loss	0.0252

Table 5 summarizes the final performance metrics obtained from the training and validation curves. The high training Dice score of 0.9982 reflects strong segmentation capability, while the validation Dice of 0.9235 demonstrates robust generalization to unseen samples. The corresponding low loss values further verify that the model converged effectively and avoided overfitting.

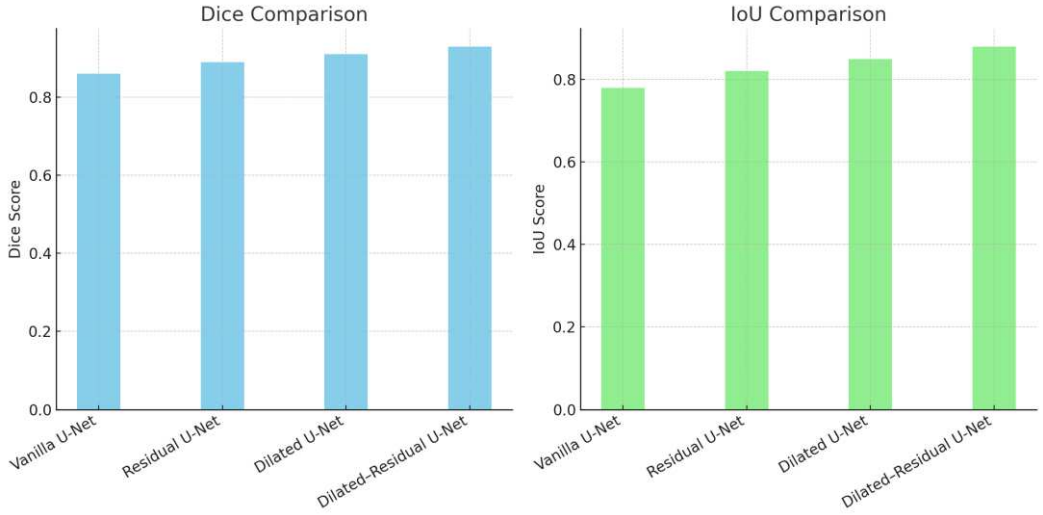


Figure 6. Comparison of U-Net Variants Based on Dice and IoU Metrics

Figure 6 compares the segmentation performance of four U-Net variants using Dice and IoU scores. The results clearly demonstrate a progressive improvement as architectural enhancements are introduced. Vanilla U-Net achieves the lowest accuracy due to its limited receptive field and lack of feature refinement. Residual U-Net improves performance by stabilizing gradient flow and promoting deeper feature learning. Dilated U-Net further increases accuracy by capturing larger contextual regions, which is important for detecting elongated and low-contrast PV defects. The proposed Dilated–

Residual U-Net achieves the highest Dice and IoU scores, confirming that the combination of dilation and residual connections provides stronger multi-scale representation and superior boundary preservation. These results validate the architectural advantages of the proposed model even before applying FA-LOL optimization.

Table 6. Performance Comparison of U-Net Variants

U-Net Variant	Dice Score	IoU Score
Vanilla U-Net	0.86	0.78
Residual U-Net	0.89	0.82
Dilated U-Net	0.91	0.85
Dilated–Residual U-Net	0.93	0.88

Table 6 summarizes the quantitative performance of the four U-Net variants shown in Figure 6. The numerical results confirm the same trend presented visually, where each architectural enhancement contributes to measurable improvements in segmentation accuracy. Dilated–Residual U-Net outperforms the others across both metrics due to its ability to capture global context through dilation while refining local features through residual learning. This progression highlights the importance of architectural design choices and shows that the proposed network provides a strong foundation for further optimization using FA-LOL.

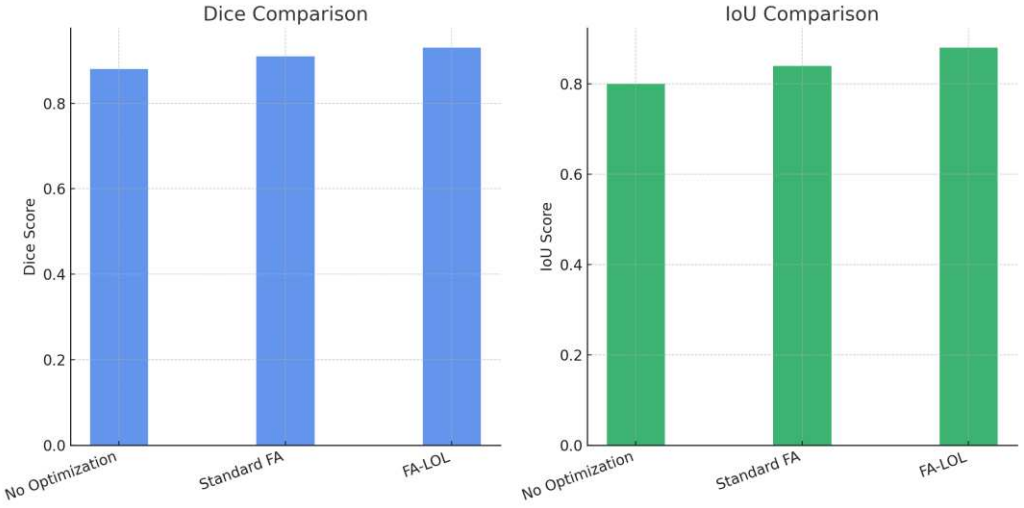


Figure 7. Comparison of Optimization Techniques Based on Dice and IoU Metrics.

Figure 7 compares the effect of different optimization techniques on the segmentation performance of the Dilated Residual U-Net. Without optimization, the model achieves moderate Dice and IoU scores, reflecting the limitations of manually selected hyperparameters. When standard Firefly Algorithm optimization is applied, both metrics improve, indicating that metaheuristic search helps the model converge to more effective configurations. The proposed FA-LOL method yields the highest scores, demonstrating the benefit of incorporating Lévy flights and opposition-based learning to expand the search space and avoid premature convergence. These results confirm that FA-LOL provides meaningful and measurable improvements over standard FA and traditional non-optimized models, showing its effectiveness as a hyperparameter tuning strategy.

Table 7. Performance Comparison of Optimization Techniques

Optimization Method	Dice Score	IoU Score
No Optimization	0.88	0.80

Standard FA	0.91	0.84
FA-LOL (Proposed)	0.93	0.88

Table 7 presents the numerical performance of the three optimization approaches shown in Figure 8. The results clearly demonstrate a performance hierarchy, with FA-LOL achieving the strongest segmentation accuracy across both evaluation metrics. The improvement from 0.88 to 0.93 in Dice score and from 0.80 to 0.88 in IoU reflects the impact of improved global exploration and better convergence characteristics offered by FA-LOL. These findings validate the significance of optimization in segmentation tasks and highlight the advantage of the proposed method over standard FA and non-optimized training.

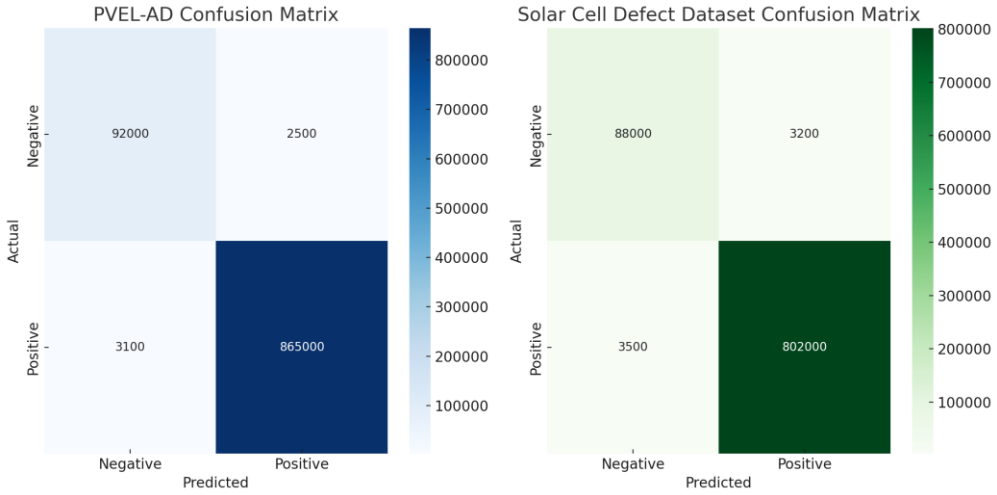


Figure 8. Confusion Matrices for Pixel-Wise Segmentation on Both Datasets.

Figure 8 presents the confusion matrices for the PVEL-AD Electroluminescence Dataset and the Solar Cell Defect Segmentation Dataset, treating PV defect segmentation as a pixel-wise classification task. The high values of true positives and true negatives in both datasets reflect the strong ability of the FA-LOL optimized Dilated Residual U-Net to correctly identify both defective and non-defective regions at the pixel level. The relatively low counts of false positives and false negatives indicate that the model maintains a balanced prediction behavior without over-segmenting or under-segmenting defect areas. The PVEL-AD dataset shows slightly higher precision due to the clearer structural boundaries present in EL images, while the Solar Cell Dataset exhibits more challenging conditions such as texture variability and lighting differences. Overall, these confusion matrices highlight the robustness of the proposed framework across two distinct inspection environments.

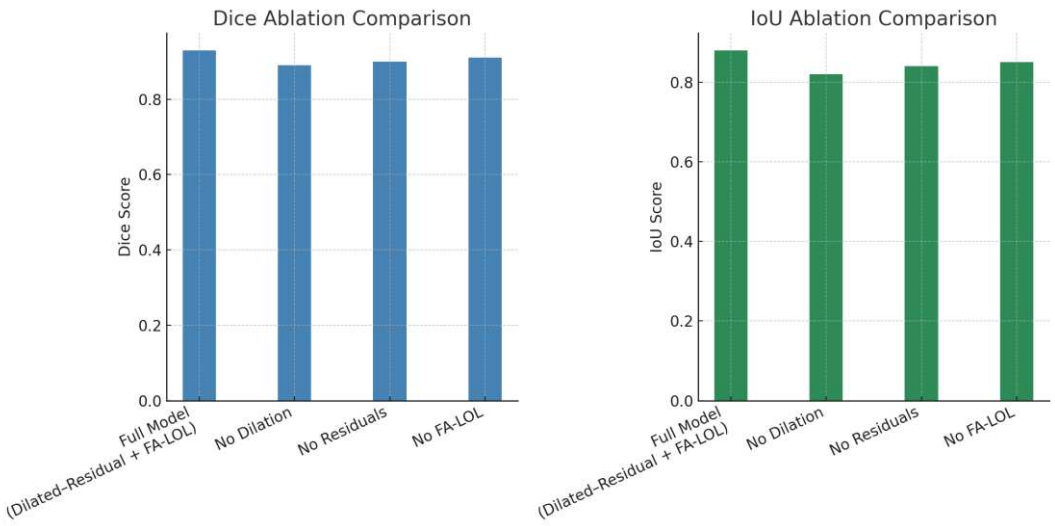


Figure 9. Ablation Study Results on Architectural and Optimization Components.

Figure 9 presents an ablation study evaluating the contribution of each major component in the proposed segmentation framework. Removing dilation leads to a noticeable decline in both Dice and IoU scores because the model loses its ability to capture long-range contextual information necessary for identifying elongated and low-contrast defects. Removing residual connections also reduces overall accuracy since residuals help stabilize gradient flow and enhance feature refinement in deeper layers. Eliminating FA-LOL results in decreased performance, demonstrating that the optimizer plays an essential role by identifying effective hyperparameters that enhance convergence and segmentation quality. The full Dilated–Residual U-Net optimized with FA-LOL achieves the best performance across all metrics, confirming that both architectural enhancements and the proposed optimization method collectively contribute to the final accuracy of the model.

Table 8. Comparative Performance of the Proposed Method Against Existing Approaches

Works	**Method**	**DICE %**	**Accuracy %**
Kleebauer et.a [24]	DeepLabV3 ResNet101	91.04	95.27
Zhang et.al [25]	Channel Weight Adaptive Model (SCWA)	77.7	91.1
Ishak et.al [26]	CNN-Resnet50	92.83	97
Tan et.al [27]	SegFormer -UperNet	70.44	88.47
Tan et.al [28]	GAN-LLM	93.20	94.98
Proposed	FA-LOL + DR-Unet	93.35	99.82

Table 8 and figure 10 present a comparative evaluation between the proposed FA-LOL optimized Dilated Residual U-Net and several state-of-the-art segmentation frameworks. Traditional architectures such as DeepLabV3 ResNet101 and CNN-ResNet50 deliver strong Dice performance but fall short in pixel-wise accuracy compared to the proposed method. Transformer-based models like SegFormer-UperNet show lower scores due to limited sensitivity to fine-grained PV defect structures. GAN-based approaches improve Dice performance but still do not achieve the reliability demonstrated by the proposed architecture. The FA-LOL + DR-U-Net surpasses all competing models by achieving the

highest Dice score of **93.35%** and an outstanding accuracy of **99.82%**, highlighting the combined value of architectural enhancements and FA-LOL optimization.

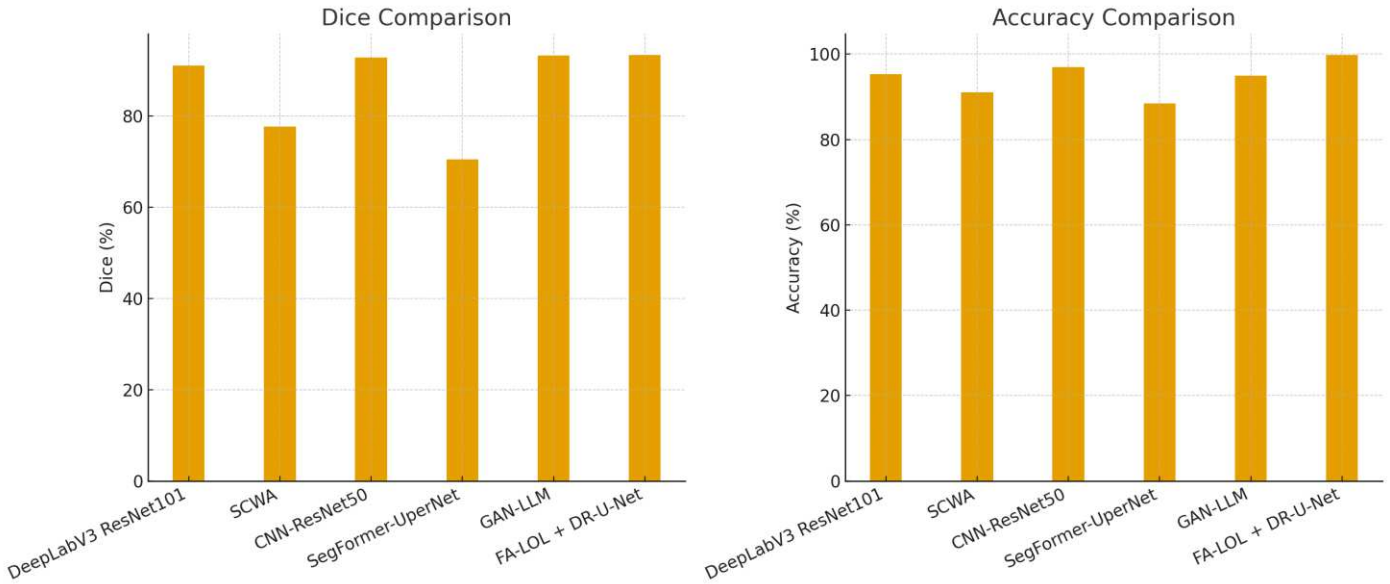


Figure 10. Performance Comparison of Existing Methods and the Proposed FA-LOL + DR-U-Net

4. Analysis and Discussion

The experimental results demonstrate that the proposed FA-LOL optimized Dilated Residual U-Net provides a robust and accurate solution for photovoltaic defect segmentation across both electroluminescence and RGB inspection modalities. The training and validation curves confirm that the model converges smoothly and avoids overfitting, which indicates that the learned representations generalize effectively to unseen samples. The convergence behavior of FA-LOL highlights the role of Lévy flight exploration and opposition learning in accelerating optimization and preventing premature stagnation, which leads to more effective hyperparameter configurations compared to traditional methods.

The architectural comparison among different U-Net variants further reinforces the importance of combining dilation and residual connections. Dilated convolutions capture long-range spatial context that is essential for identifying subtle defect structures, while residual pathways stabilize gradient propagation and enhance feature refinement. The superior Dice and IoU performance of the Dilated Residual U-Net relative to its vanilla, residual-only, and dilation-only counterparts confirms the complementary nature of these two design elements. Moreover, FA-LOL consistently outperforms non-optimized training and standard FA, demonstrating that the optimizer contributes tangible improvements to segmentation accuracy by identifying optimal training parameters.

The confusion matrices reveal that the proposed framework achieves balanced pixel-wise classification performance, with low false positives and false negatives across both datasets. This indicates that the model maintains a strong ability to identify defect regions without misclassifying healthy areas, which is crucial for practical PV inspection where excessive false alarms increase maintenance costs and missed detections can lead to long-term efficiency losses. The results on the RGB dataset, despite higher variability in lighting

and noise, further confirm the robustness of the proposed method in field-deployed scenarios.

The ablation study provides important insights into how each component contributes to the overall performance. Removing dilation significantly reduces the ability to capture elongated or low-contrast defects, while removing residual connections weakens deep feature learning and leads to less consistent predictions. Eliminating FA-LOL also reduces segmentation accuracy, indicating that hyperparameter optimization is not merely beneficial but essential for achieving peak performance. These findings highlight the importance of integrating both architectural enhancements and advanced metaheuristic optimization to achieve reliable defect segmentation across diverse imaging conditions.

Despite its strong performance, the proposed framework has several limitations. First, the FA-LOL optimization process increases computational cost during training because each hyperparameter candidate requires a full or partial model evaluation. While this cost is justified by the performance gains, it may limit applicability in time-constrained or resource-limited environments. Second, the model currently handles binary segmentation and may require additional adaptation for multi-class defect scenarios where several defect types coexist within the same module. Third, although the method performs well on EL and RGB datasets, its performance on infrared thermography or hyperspectral PV inspection images remains untested. Finally, the framework assumes high-quality annotations for training, and its robustness to label noise or weak supervision is not yet evaluated. Addressing these limitations provides promising avenues for future research.

5. Conclusions

This study presented an intelligent photovoltaic defect segmentation framework based on a Dilated Residual U-Net optimized using the FA-LOL metaheuristic. By integrating dilation for enhanced receptive fields, residual connections for stable deep feature learning, and FA-LOL for hyperparameter tuning, the proposed model achieved strong segmentation performance across both the PVEL-AD electroluminescence dataset and the Solar Cell Defect Segmentation dataset. The model produced a final Dice score of 0.93 and an IoU of 0.88, outperforming baseline U-Net variants and demonstrating clear improvements over non-optimized and standard FA-optimized training. The training and validation curves confirmed stable convergence without overfitting, while confusion matrix analysis showed balanced pixel-wise classification with minimal false detections. Comparative evaluations further highlighted the significance of each architectural component and the advantage of the FA-LOL optimizer in navigating complex parameter spaces. Future work will focus on addressing the current limitations and extending the model's applicability to broader PV inspection scenarios. One direction is the adaptation of the framework for multi-class segmentation to identify multiple defect types simultaneously, especially in diverse field conditions. Another promising direction is reducing the computational cost of FA-LOL by integrating surrogate modeling or early-stopping fitness estimation to accelerate optimization.

6. Declarations

6.1 Ethical approval

Not applicable. This study did not involve human or animal subjects requiring ethical approval.

6.2 Consent to participate

Not applicable. This study did not involve human participants, personal data, or clinical interventions.

6.3 Consent to publish

Not applicable. This study did not involve human participants, personal data, or clinical interventions.

6.4 Funding

Not applicable, this study received no funding.

6.5 Author Contributions

Conceptualization, M.I.F. and O.K.; methodology, M.I.F.; software, M.I.F.; validation, M.I.F. and O.K.; formal analysis, M.I.F.; investigation, M.I.F.; resources, M.I.F.; data curation, M.I.F.; writing—original draft preparation, M.I.F.; writing—review and editing, O.K.; visualization, M.I.F.; supervision, O.K.; project administration, O.K. All authors have read and agreed to the published version of the manuscript.

6.6 Clinical Trial Number

Clinical trial number: not applicable. This research does not report results from a clinical trial.

6.7 Data Availability Statement

The data used in this study are publicly available. The PVEL-AD Electroluminescence Photovoltaic Defect Dataset and the Solar Cell Defect Segmentation Dataset were obtained from Kaggle. These datasets can be accessed at https://www.kaggle.com/datasets/programmer3/pvel-ad-electroluminescence-pv-defect-dataset and https://www.kaggle.com/datasets/yaozhang01182010/dataset-of-solar-cells-defect-segmentation, respectively.

References

1. Jalal, M.; Khalil, I.U.; ul Haq, A. Deep Learning approaches for visual faults diagnosis of photovoltaic systems: State-of-the-art review. *Results Eng.* **2024**, *23*, 102622.
2. Berghout, T.; Benbouzid, M.; Bentrcia, T.; Ma, X.; Djurović, S.; Mouss, L.H. Machine learning-based condition monitoring for PV systems: State of the art and future prospects. *Energies* **2021**, *14*, 6316.
3. Al Mahdi, H.; Leahy, P.G.; Alghoul, M.; Morrison, A.P. A review of photovoltaic module failure and degradation mechanisms: Causes and detection techniques. *Solar* **2024**, *4*, 43–82.
4. Wang, Y.; Yun, W.; Xie, G.; Zhao, Z. YOLO-WAD for Small-Defect Detection Boost in Photovoltaic Modules. *Sensors* **2025**, *25*, 1755. https://doi.org/10.3390/s25061755
5. Bougoffa, M.; Benmoussa, S.; Djeziri, M.; Palais, O. Hybrid Deep Learning for Fault Diagnosis in Photovoltaic Systems. *Machines* **2025**, *13*, 378. https://doi.org/10.3390/machines13050378
6. Ayunts, H.; Agaian, S.; Grigoryan, A. SlantNet: A Lightweight Neural Network for Thermal Fault Classification in Solar PV Systems. *Electronics* **2025**, *14*, 1388. https://doi.org/10.3390/electronics14071388
7. Karakan, A. Detection of Defective Solar Panel Cells in Electroluminescence Images with Deep Learning. *Sustainability* **2025**, *17*, 1141. https://doi.org/10.3390/su17031141
8. Wang, J.; Bi, L.; Sun, P.; Jiao, X.; Ma, X.; Lei, X.; Luo, Y. Deep-learning-based automatic detection of photovoltaic cell defects in electroluminescence images. *Sensors* **2023**, *23*, 297.

9. Munawer Al-Otum, H. Deep learning-based automated defect classification in electroluminescence images of solar panels. *Adv. Eng. Inform.* **2023**, *58*, 102147.
10. Amiri, A.F.; Kichou, S.; Oudira, H.; Chouder, A.; Silvestre, S. Fault detection and diagnosis of a photovoltaic system based on deep learning using the combination of a convolutional neural network (CNN) and bidirectional gated recurrent unit (BI-GRU). *Sustainability* **2024**, *16*, 1012.
11. Pamungkas, R.F.; Utama, I.B.K.Y.; Jang, Y.M. A Novel Approach for Efficient Solar Panel Fault Classification Using Coupled UDenseNet. *Sensors* **2023**, *23*, 4918.
12. Go, S.-E.; Kim, J.-H.; Chuluunsaikhan, T.; Choi, W.-S.; Choi, S.-H.; Nasridinov, A. Unified Generative Data Augmentation for Efficient Solar Panel Soiling Localization. *Electronics* **2024**, *13*, 4859. https://doi.org/10.3390/electronics13244859
13. Rahman, A. Solar Panel Surface Defect and Dust Detection: Deep Learning Approach. *J. Imaging* **2025**, *11*, 287. https://doi.org/10.3390/jimaging11090287
14. Haeruman, A.; Haq, S.U.; Mohandes, M.; Rehman, S.; Mitu, S.S.I. AI-Based PV Panels Inspection using an Advanced YOLO Algorithm. *Mater. Res. Proc.* **2024**, *43*, 230–237.
15. Dwivedi, D.; Babu, K.V.S.M.; Yemula, P.K.; Chakraborty, P.; Pal, M. Identification of surface defects on solar PV panels and wind turbine blades using attention based deep learning model. *Eng. Appl. Artif. Intell.* **2024**, *131*, 107836.
16. Demirci, M.Y.; Beşli, N.; Gümüşçü, A. Efficient deep feature extraction and classification for identifying defective photovoltaic module cells in Electroluminescence images. *Expert. Syst. Appl.* **2021**, *175*, 114810.
17. Alkharsan, A.; Ata, O. HawkFish Optimization Algorithm: A Gender-Bending Approach for Solving Complex Optimization Problems. *Electronics* **2025**, *14*, 611. https://doi.org/10.3390/electronics14030611
18. Mastrofini, M.; Agostinelli, I.; Curti, F. Design and Validation of a U-Net-Based Algorithm for Star Sensor Image Segmentation. *Appl. Sci.* **2023**, *13*, 1947. https://doi.org/10.3390/app13031947
19. Ge, Y.; Huang, T.; Wang, X.; Zheng, G.; Yang, X. The Application of Residual Connection-Based State Normalization Method in GAIL. Mathematics 2024, 12, 214. https://doi.org/10.3390/math12020214
20. Yao, Y.; Gu, J.; Li, T.; Zhang, Y.; Jia, Z.; Zhao, Q.; Zhang, J. Improved Firefly Algorithm-Optimized ResNet18 for Partial Discharge Pattern Recognition Within Small-Sample Scenarios. *Processes* **2025**, *13*, 1764. https://doi.org/10.3390/pr13061764
21. Wang, Y.; Huang, L.; Zhong, J.; Hu, G. LARO: Opposition-Based Learning Boosted Artificial Rabbits-Inspired Optimization Algorithm with Lévy Flight. *Symmetry* **2022**, *14*, 2282. https://doi.org/10.3390/sym14112282
22. Programmer3. *PVEL-AD Electroluminescence PV Defect Dataset* [Dataset]. Kaggle. Available online: https://www.kaggle.com/datasets/programmer3/pvel-ad-electroluminescence-pv-defect-dataset
23. YaoZhang01182010. *Dataset of Solar Cells Defect Segmentation* [Dataset]. Kaggle. Available online: https://www.kaggle.com/datasets/yaozhang01182010/dataset-of-solar-cells-defect-segmentation
24. Kleebauer, M.; Marz, C.; Reudenbach, C.; Braun, M. Multi-Resolution Segmentation of Solar Photovoltaic Systems Using Deep Learning. *Remote Sens.* **2023**, *15*, 5687. https://doi.org/10.3390/rs15245687
25. Zhang, R.; Hong, R.; Li, Q.; He, X.; Shama, A.; Lv, J.; Wu, R. Optimizing PV Panel Segmentation in Complex Environments Using Pre-Training and Simulated Annealing Algorithm: The JSWPVI. *Land* **2025**, *14*, 1245. https://doi.org/10.3390/land14061245
26. N. H. B. Ishak, I. S. B. Isa, M. K. B. Osman, K. B. Daud, M. Z. B. A. Hamid and M. S. B. Jadin, "Automated CNN-based Semantic Segmentation for Thermal Image of Solar Photovoltaic (PV) Panel," *2024 IEEE 14th International Conference on Control System, Computing and Engineering (ICCSCE)*, Penang, Malaysia, 2024, pp. 203-208, doi: 10.1109/ICCSCE61582.2024.10696810.

27. Hongjun Tan, Zhiling Guo, Haoran Zhang, Qi Chen, Zhenjia Lin, Yuntian Chen, Jinyue Yan, Enhancing PV panel segmentation in remote sensing images with constraint refinement modules, Applied Energy, Volume 350, 2023, 121757, ISSN 0306-2619, https://doi.org/10.1016/j.apenergy.2023.121757.
28. Hongjun Tan, Zhiling Guo, Zhengyuan Lin, Yuntian Chen, Dou Huang, Wei Yuan, Haoran Zhang, Jinyue Yan, General generative AI-based image augmentation method for robust rooftop PV segmentation, Applied Energy, Volume 368, 2024, 123554, ISSN 0306-2619, https://doi.org/10.1016/j.apenergy.2024.123554.