

Bridging the Data Gap: Spatially Conditioned Diffusion Model for Anomaly Generation in Photovoltaic Electroluminescence Images

Shiva Hanifi¹, Sasan Jafarnejad¹, Raphael Frank¹

¹Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Luxembourg
{shiva.hanifi, sasan.jafarnejad, raphael.frank}@uni.lu

Abstract

Reliable anomaly detection (AD) in photovoltaic (PV) modules is critical for maintaining solar energy efficiency. However, the development of robust computer vision models for PV inspection is limited by the scarcity of large-scale, diverse, and balanced datasets. In this study, we introduce a spatially conditioned denoising diffusion probabilistic model, PV-DDPM, designed to generate anomalous electroluminescence (EL) images across four PV cell types: multi-crystalline silicon (multi-c-Si), mono-crystalline silicon (mono-c-Si), half-cut multi-c-Si, and interdigitated back contact (IBC) with dogbone interconnect. PV-DDPM enables the controlled synthesis of both multi-defect and single-defect scenarios by conditioning on binary masks that represent structural features and defect positions. To the best of our knowledge, this is the first framework that jointly models multiple PV cell types while allowing simultaneous generation of diverse anomaly types. We also introduce E-SCDD dataset, an enhanced version of the SCDD dataset, comprising 1,000 pixel-wise annotated EL images spanning 30 semantic classes, and 1,768 unlabeled synthetic samples. Quantitative evaluation shows our generated images achieve a Fréchet Inception Distance (FID) of 4.10 and Kernel Inception Distance (KID) of 0.0023 ± 0.0007 across all categories. Furthermore, compared to training on the original SCDD dataset, training the vision-language anomaly detection model AA-CLIP on E-SCDD improves pixel-level AUC and average precision by 1.70 and 8.34 points, respectively. Code and data are publicly released to ensure reproducibility.

Code — <https://github.com/sntubix/pv-ddpm>

Dataset — <https://huggingface.co/datasets/shivahanifi/extended-scdd>

1 Introduction

Solar energy plays a key role in the global transition toward clean and sustainable energy, driving continuous advancements in photovoltaic (PV) technology. The rapid expansion of large-scale solar farms, comprising thousands of PV modules, has increased the demand for efficient cleaning, monitoring, and maintenance. Although PV panels have an expected lifespan of up to 25 years, they degrade at an average annual rate of 0.5%, which can increase significantly due to

environmental stress and material aging (Jordan et al. 2016). To address these degradation challenges, advanced monitoring methods, particularly imaging-based techniques such as electroluminescence (EL), have become essential for detecting and localizing anomalies (Jahn et al. 2018; Köntges et al. 2014; Maziuk et al. 2023).

Advances in robotics and artificial intelligence (AI), especially computer vision (CV), now enable automated PV inspections (Iqbal et al. 2019). CV-based defect detection is emerging as a promising approach for predictive maintenance in industrial PV applications. However, these methods are constrained by the scarcity of large-scale, diverse, and well-annotated EL datasets that represent a wide range of defect types and cell architectures (Deutsch et al. 2019; Fioresi et al. 2022; Pratt, Govender, and Klein 2021).

To address these challenges, we propose a diffusion-based, mask-conditioned framework for generating synthetic EL images of defective PV modules, offering precise spatial control over defect placement. This approach enhances dataset diversity and supports the development of robust anomaly detection (AD) methods for automated PV inspection applications, including autonomous robotic inspection systems. To the best of our knowledge, this is the first work that explicitly separates different PV cell architectures and enables the generation of multiple defect types within a single sample.

The main contributions of this work are as follows:

- We propose PV-DDPM, a denoising diffusion probabilistic model (DDPM) (Ho, Jain, and Abbeel 2020) conditioned on binary masks of PV defects and features. It is capable of synthesizing realistic EL images for four PV cell types and architectures: multi-crystalline silicon (multi-c-Si), mono-crystalline silicon (mono-c-Si), half-cut multi-c-Si, and interdigitated back contact (IBC) cells with dogbone interconnects.
- Our method enables the generation of both single and multiple defect types in individual samples, increasing variability and realism in training data.
- We introduce the E-SCDD dataset, an extension of the SCDD dataset (Pratt, Mattheus, and Klein 2023), comprising a total of 2,786 EL images. It includes 1,000 pixel-wise annotated images spanning 30 semantic classes (14 intrinsic features and 16 defect types), and

1,786 additional unlabeled synthetic images.

- We train and evaluate AA-CLIP (Ma et al. 2025), vision-language AD model, on SCDD and E-SCDD, achieving a pixel-level AUC of 94.96 and AP of 55.47, exceeding SCDD by 1.70 and 8.34 points, respectively, demonstrating E-SCDD’s effectiveness for robust, fine-grained AD.

The rest of this paper is organized as follows: Section 2 reviews the state-of-the-art in synthetic EL image generation and PV defect detection. Section 3 introduces the E-SCDD dataset and Section 4 details the proposed methodology. Section 5 presents a comprehensive evaluation of our method and dataset. Finally, Section 6 summarizes the key findings and explores promising directions for future research.

2 Related Work

EL Datasets

Modern PV defect detection employs various imaging modalities such as EL, photoluminescence (PL), infrared (IR), ultraviolet (UV), and visual RGB imaging. Each modality differs in imaging conditions, sensitivity, and suitability for specific cell architectures and defect types (Maziuk et al. 2023; Rico Espinosa, Bressan, and Giraldo 2020; Herrmann and Jahn 2021). Among these, EL imaging is the most widely used, with numerous datasets varying in scale, annotation detail, and cell-type diversity. Table 1 summarizes the most relevant publicly available datasets, providing information on the number of labeled and unlabeled images, annotation formats, class coverage, and supported cell architectures. While datasets such as ELPV and SDLE provide image-level labels for a limited number of classes, datasets like PVEL-AD offer bounding box annotations for larger collections. The SCDD dataset is one of the most comprehensive, providing pixel-level segmentation masks with detailed labeling across 29 classes, 16 defect types and 13 intrinsic features. Considering SCDD’s extensive coverage of cell architectures and fine-grained annotations, we adopt it as our baseline. We further extend it by adding a new feature class and incorporating additional real and synthetic EL images, resulting in the proposed E-SCDD dataset (see Section 3).

Synthetic Data Generation

Generative models have become central to addressing the data imbalance in EL datasets. Among these, generative adversarial networks (GANs), notably deep convolutional GANs (DCGANs) and their variants, are widely adopted for synthesizing realistic EL images (Luo, Cheng, and Zheng 2019; Romero et al. 2023; Liu et al. 2024; Guan et al. 2025). SIGAN (Su et al. 2021) introduces a spatial identity loss to preserve background structures while generating realistic defect patterns, ensuring consistency in non-defective regions. Designed primarily for data augmentation in semantic segmentation, SIGAN was released with the EL-2019 dataset and demonstrated notable improvements in segmentation performance. In a related approach, attention-enhanced GANs (Zhou et al. 2023) improve visual realism

by incorporating spatial attention mechanisms to better capture complex interactions between defects and background textures.

Complementary to GAN-based methods, physics-driven approaches synthesize EL images based on solar cell behaviors, such as current-voltage (I-V) curve modeling (Fioresi et al. 2022). These methods offer physically grounded samples that enhance model robustness (Li et al. 2021). Although less common, they offer the advantage of embedding domain knowledge and addressing class imbalance.

Hybrid strategies that combine GANs with domain adaptation techniques have also been explored to reduce the distribution shift between synthetic and real data, showing promise in enhancing generalization for real-world defect detection (Shou et al. 2020).

Recently, diffusion-based generative models have emerged as a robust alternative to GANs. Unlike GANs, which often suffer from mode collapse or training instability, diffusion models, such as denoising diffusion probabilistic models (DDPMs) (Ho, Jain, and Abbeel 2020), generate images through a gradual denoising process, enabling more stable training and improved diversity in the synthesized samples. The Photovoltaic Defect Image Generator (PDIG) based on Stable Diffusion (Li et al. 2025) is a recent work applying Stable Diffusion with domain adaptation modules to synthesize PV defect images. However, it generates only a single defect type per sample and does not account for diverse cell architectures.

Despite significant advancements, challenges remain. Existing approaches predominantly target a limited number of common defect types (e.g., cracks, shunts, finger interruptions) and often generate only a single defect type per image, limiting diversity and realism. Additionally, the generation of synthetic EL images across diverse PV cell types, such as mono-/multi-c-Si, half-cut mono-/multi-c-Si, and IBC dog-bone designs, has received no attention, leaving a gap in the comprehensive modeling of real-world module diversity.

In this work, we address these limitations using a mask-conditioned DDPM (Dorjsembe, Pao, and Xiao 2024), which enables explicit control over cell type, defect type, and the number of defects per image.

Anomaly Detection

Recent research on AD in EL images of PV cells spans supervised segmentation and detection models as well as unsupervised and self-supervised approaches developed to mitigate dataset scarcity. Fully supervised models achieve precise localization and classification when sufficient annotations are available. SEiPV-Net segments 24 classes using class-weighted loss functions to address label imbalance (Joe et al. 2023). Similarly, (Hanifi et al. 2026) extend this approach to 29 classes by employing a combination of different encoder-decoder architectures. For instance-level localization, BAF-Detector employs a bidirectional attention feature pyramid to capture multiscale defect representations (Su, Chen, and Zhou 2022). To reduce annotation costs, several unsupervised and semi-supervised methods adapt general-purpose AD frameworks to EL imagery. Wave-Flow (Yang, Zhang, and Ma 2025) models nonsta-

Dataset	Labeled	Unlabeled	Annotation	Classes (def./feat.)	Cell types
ELPV (Buerhop-Lutz et al. 2018)	2,624	–	image-level	4 / –	mono-/multi-c-Si (earlier generation)
E-ELPV (Grisanti et al. 2024)	2,624	–	image-level	6 / –	mono-/multi-c-Si (earlier generation)
SDLE (French, Karimi, and Braid 2019)	1,028	–	image-level	3 / –	multi-c-Si
PVEL-AD (Su, Zhou, and Chen 2023)	36,543	–	bounding box	13 / –	multi-c-Si
UCF (Fioresi et al. 2022)	11,851	17,000	region-level	10 / –	mono-/multi-c-Si
EL-2019 (Su et al. 2021)	540	–	pixel-level	3 / –	multi-c-Si
SCDD (Pratt, Mattheus, and Klein 2023)	695	150,000	pixel-level	16 / 13	mono-/multi-c-Si, half-cut, IBC dogbone
E-SCDD (Ours)	1,000	151,700	pixel-level	16 / 14	mono-/multi-c-Si, half-cut, IBC dogbone

Table 1: Summary of publicly available EL datasets, including the number of labeled and unlabeled images, annotation type, number of annotated classes, and PV cell types represented.

tionary multi-c-Si textures using wavelet-domain normalizing flows, while hybrid inpainting models inspired by DRAEM (Tan, Shi, and Chen 2023) leverage context attention and masking for defect restoration. Recent frameworks such as MLA-SDAL and ISLAD (Chang et al. 2024; Han et al. 2025) incorporate attention-based scale alignment to improve robustness to illumination and manufacturing variability. Although these approaches improve generalization compared to fully supervised networks, their scalability to unseen defect types and diverse cell architectures remains limited.

Vision–language models (VLMs) such as CLIP (Radford et al. 2021), WinCLIP (Jeong et al. 2023), CLIP-AD (Chen et al. 2025), and AA-CLIP (Ma et al. 2025) enable zero- and few-shot AD by aligning image and text embeddings. In this paradigm, textual prompts describing defect concepts are compared with image embeddings to produce similarity-based anomaly maps. This approach supports open-vocabulary detection, weak localization, and efficient bootstrapping of labels for active learning. Nevertheless, VLMs have not yet been systematically evaluated on PV EL datasets.

AA-CLIP (Ma et al. 2025), with its alignment-aware attention pooling and adaptive fine-tuning mechanisms that enhance cross-modal consistency and anomaly sensitivity, is particularly well suited to the fine-grained, spatially localized defect signatures in EL images; therefore, this work adopts and evaluates it for PV AD.

3 E-SCDD Dataset

We introduce the E-SCDD dataset, an extended version of the original SCDD dataset (Pratt, Mattheus, and Klein 2023), designed to improve class diversity, semantic granularity, and representation across multiple PV cell architectures.

The original SCDD dataset comprises 695 annotated EL images with detailed pixel-wise masks spanning 29 semantic classes, including 16 defect types and 13 intrinsic features. Figure 1 illustrates representative annotations, highlighting three feature classes (*ribbons*, *sp mono*, *sp multi*) and three defect classes (*corrosion cell*, *inactive*, *crack*). De-

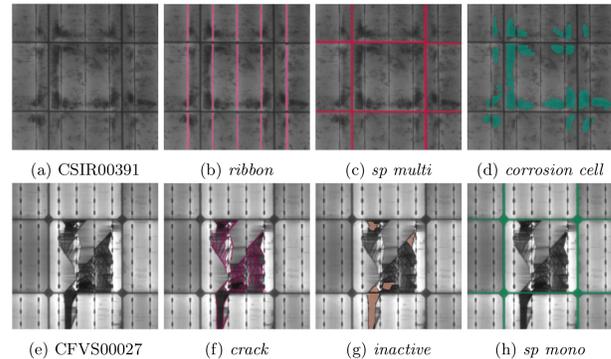


Figure 1: Representative ground truth annotations. The first row shows (a) EL image CSIR00391 with annotated features ((b)*ribbon*, (c)*sp multi*) and defect ((d)*corrosion cell*). The second row shows (e) EL image CFVS00027 with defects ((f)*crack*, (g)*inactive*) and feature ((h)*sp mono*).

tailed class explanations are provided in the supplementary materials.

With the E-SCDD dataset, we introduce a novel intrinsic feature class, *sp multi halfcut*, increasing the total number of semantic classes to 30. This addition allows a more precise characterization of the distinctive features of half-cut multi-c-Si cells. Furthermore, we refined the annotations corresponding to the half-cut multi-c-Si cells, relabeling the spacing as *sp multi halfcut* instead of original *sp multi* labels.

Additionally, to address class imbalance and further diversify the dataset, we annotated 110 additional real EL images from the original SCDD’s unlabeled pool. Images were chosen to ensure representation across four cell types: mono-c-Si, multi-c-Si, half-cut multi-c-Si, and IBC dogbone cells. Samples for each cell type were selected randomly and initially screened to ensure sufficient image quality prior to annotation. These new annotations include 40 images of multi-c-Si cells, 43 images of half-cut multi-c-Si cells, 14 images of mono-c-Si cells, and 13 images of IBC dogbone cells. It is important to note that the fifth cell type, half-cut mono-c-Si cells, remains underrepresented due to limited availability

in the original dataset and the absence of samples in the unlabeled pool.

To ensure annotation consistency, we developed an interactive annotation assistant tool that provides concise class definitions and visual examples, along with image-based and class-based search capabilities. This tool is publicly available in the project’s GitHub repository.

The annotation process was carried out on the Roboflow platform¹ following a structured three-stage protocol:

- First, an annotator labeled selected samples, using Roboflow’s Smart Polygon tool (a Segment Anything Model (Kirillov et al. 2023) we fine-tuned on SCDD dataset).
- Second, the initial annotations were independently reviewed by a CV domain specialist with expertise in PV systems to ensure semantic correctness and consistency.
- Finally, a senior PV technology expert carried out a thorough quality assurance review prior to the integration of the annotated samples into the E-SCDD dataset.

Consequently, the total number of annotated real EL images increased to 805.

In addition to the real images, we generated 195 synthetic EL images, annotated using the same procedure. This synthetic set covers four PV cell categories: 80 IBC dogbone, 65 mono-c-Si, 30 half-cut multi-c-Si, and 20 multi-c-Si cell images. Furthermore, 1,768 unlabeled synthetic EL images are included to facilitate additional studies. The methodology for synthetic data generation and the selection criteria are detailed in Sections 4 and 5.

The proposed E-SCDD dataset consists of 1,000 annotated EL images of PV modules, comprising 695 original images, 110 newly annotated images from SCDD, and 195 annotated synthetic images, with pixel-wise ground truth masks across 30 semantic classes (16 defect and 14 intrinsic feature classes), along with 1,768 unlabeled synthetic images.

4 Methods

PV-DDPM: Photovoltaic Denoising Diffusion Probabilistic Model for Conditional EL Image Generation

To synthesize EL images of PV modules, we adopted a mask-conditioned image generation framework based on a denoising diffusion probabilistic model (DDPM) (Ho, Jain, and Abbeel 2020). Our approach, inspired by prior work in medical image synthesis (Dorjsembe, Pao, and Xiao 2024), extends conditional diffusion models to the PV domain. Figure 2 provides an overview of the proposed PV-DDPM architecture and its conditioning masks.

Conditioning Masks To accommodate different cell types, defect types, and spatial distribution of defects in the generated EL images, PV-DDPM is conditioned on binary masks that explicitly encode this information.

To create these informative conditioning masks, we first categorized the EL images and their corresponding ground

truth annotations into four groups based on the spacing class annotations: mono-c-Si (*sp mono*), multi-c-Si (*sp multi*), half-cut multi-c-Si (*sp multi halfcut*), and IBC dogbone (*sp dogbone*).

Each ground truth annotation was then used to generate binary segmentation masks in two forms:

1. Background mask: highlights all pixels belonging to the background (*bckgrnd*) class in white, with all other pixels in black. This mask captures the complete set of features and defects present in the EL image.
2. Feature-defect mask: identifies all feature classes in the EL image, rendering their pixels in white. For each defect class, it then creates a binary mask that combines the relevant feature regions with that specific defect.

This subclass-conditioned formulation offers two key advantages. First, it increases the number of training image-mask pairs. Second, it enables the model to learn category-specific structural and defect patterns, thereby enhancing the fidelity and diversity of the synthesized EL images. Figure 3 depicts four samples of the EL images and their corresponding background and feature-defect masks.

Conditional Diffusion Framework In the PV-DDPM framework, a forward process progressively corrupts an image with Gaussian noise, while a learned reverse process reconstructs it. Given an input image x_0 , the forward process defines a Markov chain $q(x_t | x_{t-1})$ with additive Gaussian noise over T timesteps. The noisy image at timestep t is given by:

$$x_t = \sqrt{\bar{\alpha}_t} \cdot x_0 + \sqrt{1 - \bar{\alpha}_t} \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (1)$$

A cosine noise schedule proposed in (Nichol and Dhariwal 2021) is adopted to compute the variance schedule $\bar{\alpha}_t$, which stabilizes training by avoiding abrupt transitions near $t = 0$.

The reverse process employs a U-Net denoiser to predict the added noise. The model is conditioned on binary masks that highlight regions of interest (both defects and structural features), which are concatenated channel-wise with the noisy input image. This conditioning strategy enables targeted reconstruction of cell types and defect patterns.

Training Procedure We trained the model for 1,000 diffusion steps using the Adam optimizer (learning rate 5×10^{-4} , batch size 8), with gradient accumulation for memory efficiency and an exponential moving average (EMA) to stabilize weights. An L_2 reconstruction loss guides training, with early stopping based on validation loss (patience of 10 epochs) to mitigate overfitting.

The model is initially trained on the complete dataset encompassing all PV cell categories. To specialize generation for specific cell types, we then fine-tuned on category-specific subsets. While the mono-c-Si and multi-c-Si categories provide numerous high-quality samples, the remaining categories have fewer and often lower-quality examples, which affects the quality of the generated samples.

¹<https://roboflow.com>

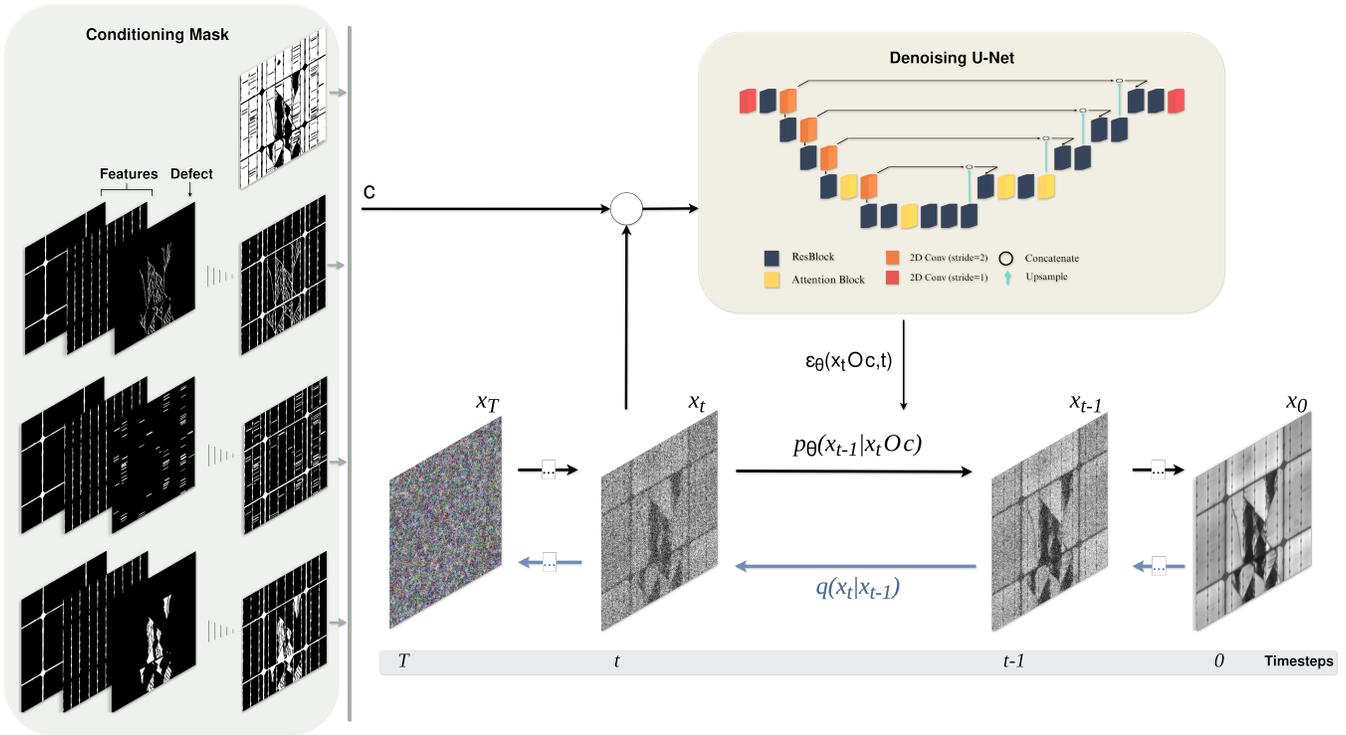


Figure 2: PV-DDPM architecture: illustrating the forward diffusion process and the U-Net-based denoising network. Binary conditioning masks, full background masks or specific feature-defect mask, are combined with the original EL image to guide the generation process.

Anomaly Detection Models

To validate the quality of our synthetic data and demonstrate the dataset’s applicability to robotic maintenance tasks, we trained AA-CLIP (Ma et al. 2025) model using both the SCDD and E-SCDD datasets. We modified the groundtruth annotations by mapping the feature classes into the background and preserving only anomalous areas. To avoid data leakage from synthetic images (which closely resemble specific real images), we ensured that each synthetic image was placed in the same data split as its corresponding real image. Training was performed in both few-shot (2-shot, 16-shot, 32-shot) and full-shot settings.

Performance of the model for each of the datasets and settings has been evaluated using the image-level and pixel-level Area Under the Receiver Operating Characteristic Curve (AUROC), and average precision (AP) metrics.

5 Results

Using the PV-DDPM model conditioned on binary masks, we generated 10 samples per conditioning mask, yielding a total of 1,788 images: 70 half-cut multi-c-Si, 180 IBC dog-bone, 802 mono-c-Si, and 736 multi-c-Si cells. Figure 3 illustrates PV-DDPM’s conditional synthesis capabilities, presenting four PV cell types, the original real EL images alongside their corresponding binary masks. When conditioned on the background masks, the PV-DDPM generates samples that realistically reproduce multiple simultaneous

defects. Conversely, conditioning on feature-defect masks enables the targeted synthesis of cells exhibiting only a single defect type superimposed on the actual cell structure. This demonstrates the model’s flexibility in producing both complex multi-defect scenarios and controlled single-defect instances.

Next, we evaluated the quality of the generated samples and assessed the performance of the semi-supervised method on the proposed E-SCDD dataset.

Synthetic Data Evaluation

The quality of the synthetic EL images generated by the PV-DDPM model was evaluated using both qualitative and quantitative methods. Domain experts in PV systems conducted detailed visual inspections to assess various aspects of the generated images. In parallel, quantitative evaluation was carried out using customized implementations of the Fréchet Inception Distance (FID) (Heusel et al. 2017) and Kernel Inception Distance (KID) (Bí et al. 2018) metrics, adapted to capture domain-specific characteristics.

Qualitative Evaluation For each category, a subset of the 20 most visually realistic samples was selected for expert evaluation. A domain expert assessed the samples based on five key criteria: format consistency, feature and defect alignment, texture similarity, visual irregularities, and overall quality. This expert feedback was iteratively incorporated to refine the PV-DDPM model and its outputs. Specifically,

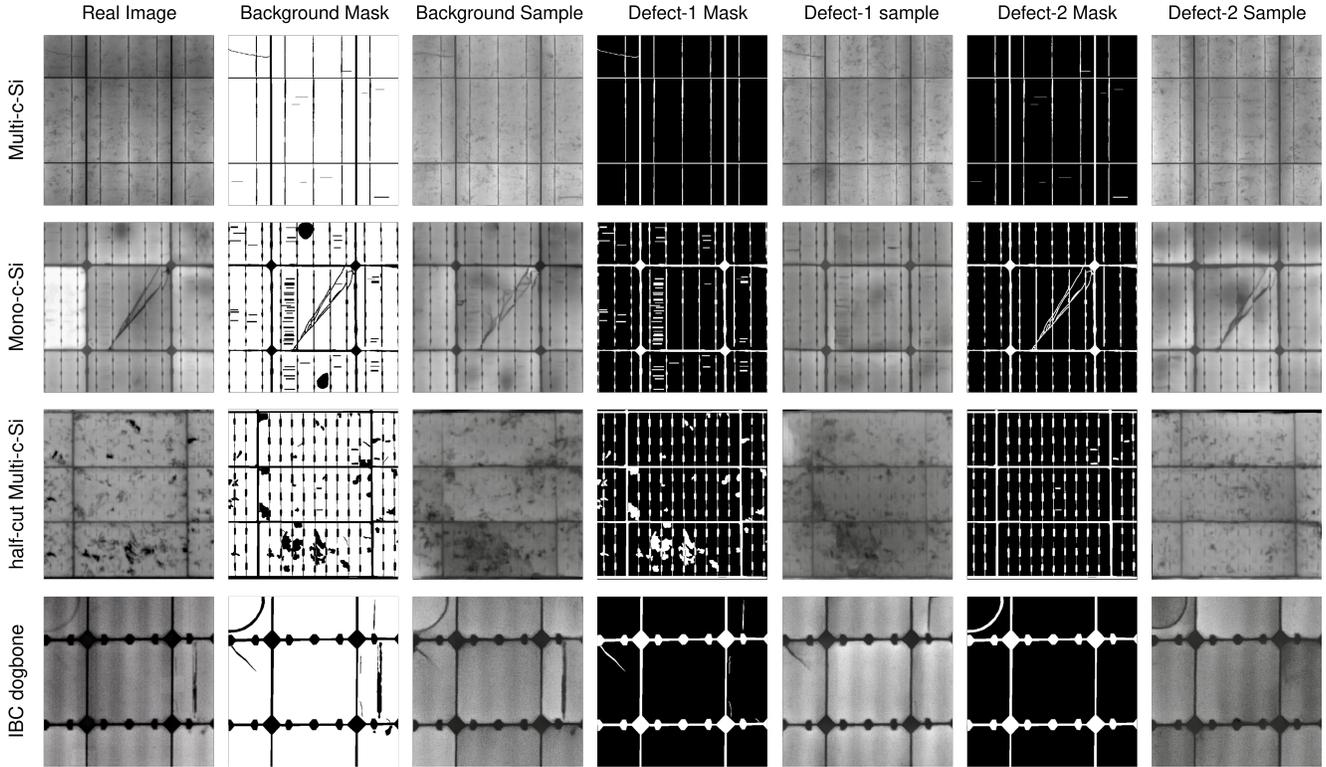


Figure 3: Real EL images of four PV cell types and their corresponding PV-DDPM-generated images. Binary masks derived from the real images represent either all features and defects (background masks) or all features combined with a specific defect (feature-defect masks). Conditioning on background masks produces multi-defect samples, while feature-defect masks enable controlled synthesis of single-defect instances across the different cell types.

informed by detailed evaluations of the generated images across cell types, we trained dedicated models for each cell type and eliminated padding to prevent dark edge artifacts.

Overall, synthetic samples of the multi-c-Si category were considered the most realistic samples across all the evaluation criteria, followed by the images generated for mono-c-Si category. This trend aligns with expectations, given the relatively larger number of training samples available for these categories and the higher quality of the available samples. The lower FID and KID scores for these two categories presented in Table 2 further align with the visual assessment.

To visually assess the distribution overlap between real and generated images of each cell type and also inter-class overlaps, we employed pairwise controlled manifold approximation (PaCMAP) for dimensionality reduction (Wang et al. 2021). To extract the features from real and generated images, we employed a DeepLabv3 model (Chen et al. 2017) with a ResNet101 (He et al. 2016) backbone pre-trained on ImageNet (Deng et al. 2009) and fine-tuned on the E-SCDD dataset without the inclusion of synthetic samples.

The extracted features were projected to a 2D space using PaCMAP, with hyperparameters set to 17 neighbors, a feature preponderance ratio (FP) of 2, and a mid-neighbor ratio (MN) of 1. Figure 4 shows the resulting scatter plot where

real and generated samples are differentiated by shape (circles and squares, respectively), while consistent color tones are used for each PV cell type, with real samples rendered in darker tones and generated ones in lighter shades to enhance interpretability. The clear separation between the four cell types confirms the quality of the generated images. Additionally, the tight clustering of generated samples with their real counterparts within each category supports the fidelity of the synthesis process.

Quantitative Evaluation To quantitatively evaluate the quality of the generated EL images, we implemented custom versions of FID and KID metrics. These metrics are widely used to assess the quality of generative models by comparing the distribution of real and generated images in a learned feature space.

FID computes the distance between two multivariate Gaussian distributions fitted to feature representations of real and synthetic images, capturing both mean and covariance differences. KID, on the other hand, estimates the squared Maximum Mean Discrepancy (MMD) between feature distributions using a polynomial kernel, and has the advantage of being unbiased even with smaller sample sizes (as is the case in our evaluation).

Unlike standard implementation of these metrics, which uses features extracted from the InceptionV3 net-

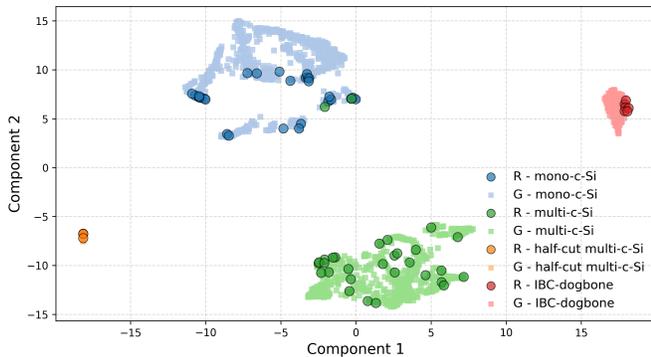


Figure 4: PaCMAP 2D embedding of feature representations for real (R) and generated (G) EL images across four PV cell types.

Category	FID↓	KID↓
Multi-c-Si	2.4551	0.0008 ± 0.0001
Mono-c-Si	3.2420	0.0031 ± 0.0009
Half-cut multi-c-Si	4.2725	0.0009 ± 0.0004
IBC dogbone	6.4163	0.0044 ± 0.0014
Average	4.0965	0.0023 ± 0.0007

Table 2: Category-wise quantitative evaluation of generated EL images using FID and KID metrics. Lower values indicate better performance.

work (Szegedy et al. 2016) pre-trained on the ImageNet dataset, we employ a domain-specific feature extractor. In particular, we use the fine-tuned DeepLabv3-ResNet101 network (from our PaCMAP analysis) to extract features. This adaptation ensures that the feature space is tailored to the domain of EL images, and results in more reliable FID and KID scores.

For each of the four categories of the PV modules, we computed the FID and KID metrics on the images of the original test set and the synthetic images of the same category. Table 2 reports the results, with the multi-c-Si category achieving the lowest (best) scores and more realistic samples. The other categories show slightly higher values, but still demonstrate good fidelity relative to real images.

Anomaly Detection

Table 3 reports the performance of the AA-CLIP model on the SCDD and E-SCDD datasets across varying data levels. Image-level AUROC and AP remain consistently high across all settings, demonstrating AA-CLIP’s strong ability to distinguish defective from non-defective cells even under minimal supervision. In contrast, pixel-level performance improves gradually with increased data, reflecting the greater complexity of precise defect localization. Notably, in the full-shot setting, AA-CLIP achieves a pixel-level AUC of 94.96 and AP of 55.47 on E-SCDD, surpassing the corresponding SCDD results by 1.70 and 8.34 points, respectively. These gains highlight the effectiveness of E-SCDD, whose expanded diversity of synthetic and real

Dataset	Data Level	Pixel-Level		Image-Level	
		AUC↑	AP↑	AUC↑	AP↑
SCDD	2-shot	73.25	11.74	91.55	99.88
	16-shot	89.01	39.25	97.18	99.96
	32-shot	91.12	36.66	92.96	99.90
	full-shot	93.26	47.13	52.11	99.11
E-SCDD	2-shot	79.64	17.79	98.36	99.97
	16-shot	81.68	19.69	84.43	99.73
	32-shot	87.04	25.69	86.89	99.77
	full-shot	94.96	55.47	79.60	99.10

Table 3: Full-shot and Few-shot performance of AA-CLIP on SCDD and E-SCDD datasets. Pixel-level and image-level metrics are reported as AUC and AP.

samples enables more robust learning and generalization for fine-grained AD. Qualitative results of these experiments are provided in the supplementary materials.

6 Conclusion

This study introduced a mask-conditioned diffusion-based framework (PV-DDPM) for generating realistic EL images of PV modules, effectively modeling multiple cell architectures and diverse defect types. By leveraging binary masks to guide the generation process, our method enables precise control over defect configurations and allows creation of both multi-defect and single-defect scenarios. We also developed the E-SCDD dataset, which expands existing PV defect datasets in terms of both semantic richness and cell type diversity. Comprehensive evaluations confirmed the visual and statistical quality of the synthetic images. Moreover, downstream experiments demonstrated that incorporating this dataset improves the performance of the AA-CLIP model, achieving higher pixel-level and image-level AUC and AP metrics. Our results highlight the significant potential of combining advanced generative models with enriched datasets to support predictive maintenance and autonomous robotic inspection of PV systems. Future work will focus on improving the technical quality of the generated images through close collaboration with PV experts, expanding the dataset’s coverage of underrepresented classes, and exploring multi-modal generative approaches to further enhance dataset variability and broaden its application scope.

7 Acknowledgments

The authors would like to thank Marc Köntges, Andrej Wentnagel, and Andreas Kokkas for their expert support and valuable contributions to this research. The authors also gratefully acknowledge the financial support provided by SolarCleanso S.A. (<https://solarcleanso.com/en/>)

References

- Bí, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. DEMYSTIFYING MMD GANS.
- Buerhop-Lutz, C.; Deitsch, S.; Maier, A.; Gallwitz, F.; Berger, S.; Doll, B.; Hauch, J.; Camus, C.; and Brabec, C. J.

2018. A benchmark for visual identification of defective solar cells in electroluminescence imagery. In *35th European PV Solar Energy Conference and Exhibition*, 1287–1289.
- Chang, Z.; Zhang, A. J.; Wang, H.; Xu, J.; and Han, T. 2024. Photovoltaic Cell Anomaly Detection Enabled by Scale Distribution Alignment Learning and Multiscale Linear Attention Framework. *IEEE Internet of Things Journal*, 11(16): 27816–27827.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017. Rethinking Atrous Convolution for Semantic Image Segmentation.
- Chen, X.; Zhang, J.; Tian, G.; He, H.; Zhang, W.; Wang, Y.; Wang, C.; and Liu, Y. 2025. CLIP-AD: A Language-Guided Staged Dual-Path Model for Zero-Shot Anomaly Detection. *Communications in Computer and Information Science*, 2201 CCIS: 17–33.
- Deutsch, S.; Christlein, V.; Berger, S.; Buerhop-Lutz, C.; Maier, A.; Gallwitz, F.; and Riess, C. 2019. Automatic classification of defective photovoltaic module cells in electroluminescence images. *Solar Energy*, 185: 455–468.
- Deng, J.; Dong, W.; Socher, R.; Li, L. J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. *2009 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009*, 248–255.
- Dorjsembe, Z.; Pao, H.-K.; and Xiao, F. 2024. Polyp-DDPM: Diffusion-Based Semantic Polyp Synthesis for Enhanced Segmentation.
- Fioresi, J.; Colvin, D. J.; Frota, R.; Gupta, R.; Li, M.; Seigneur, H. P.; Vyas, S.; Oliveira, S.; Shah, M.; and Davis, K. O. 2022. Automated Defect Detection and Localization in Photovoltaic Cells Using Semantic Segmentation of Electroluminescence Images. *IEEE Journal of Photovoltaics*, 12(1): 53–61.
- French, R. H.; Karimi, A. M.; and Braid, J. L. 2019. Electroluminescent (EL) Image Dataset of PV Module Under Step-wise Damp Heat Exposures.
- Grisanti, M.; Spatafora, M. A. N.; Ortis, A.; and Battiato, S. 2024. E-ELPV: Extended ELPV Dataset for Accurate Solar Cells Defect Classification. *Lecture Notes in Networks and Systems*, 822: 837–848.
- Guan, Y.; Liu, Y.; Wang, J.; Wang, T.; Yi, Q.; Jiang, W.; Gu, X.; Zhang, Y.; Zhang, L.; Han, T.; Huang, B.; and Hu, L. 2025. Novel Multi-scale Attention Generative Adversarial Network for Photovoltaic Solar Cell Defect Inspection Using Electroluminescence Images. *IEEE Access*.
- Han, T.; Wang, X.; Guo, J.; Chang, Z.; and Chen, Y. 2025. Health-Aware Joint Learning of Scale Distribution and Compact Representation for Unsupervised Anomaly Detection in Photovoltaic Systems. *IEEE Transactions on Instrumentation and Measurement*, 74.
- Hanifi, S.; Jafarnejad, S.; Cormier, M.; and Frank, R. 2026. Multi-class Semantic Segmentation of Photovoltaic Module Defects and Features: Towards Industrial Robotic Applications. *Lecture Notes in Computer Science*, 15707 LNAI: 37–48.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition.
- Herrmann, W.; and Jahn, U. 2021. Qualification of PV Power Plants using Mobile Test Equipment Fact sheet 2 0 2 1.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models.
- Iqbal, J.; Al-Zahrani, A.; Alharbi, S. A.; and Hashmi, A. 2019. Robotics Inspired Renewable Energy Developments: Prospective Opportunities and Challenges. *IEEE Access*, 7: 174898–174923.
- Jahn, U.; Herz, M.; Köntges, M.; Parlevliet, D.; Paggi, M.; Tsanakas, I.; Stein, J.; Berger, K. A.; Ranta, S.; French, R. H.; Richter, M.; and Tanahashi, T. 2018. Review on infrared and electroluminescence imaging for PV field applications : International Energy Agency Photovoltaic Power Systems Programme. *IEA PVPS Task 13, Subtask 3.3 : Photovoltaic Power Systems Programme. Report IEA-PVPS T13-12:2018*, 1–94.
- Jeong, J.; Zou, Y.; Kim, T.; Zhang, D.; Ravichandran, A.; and Dabeer, O. 2023. WinCLIP: Zero-/Few-Shot Anomaly Classification and Segmentation.
- Joe, S. .; Rehman, M. U. .; Jang, Y. .; Chong, K. T.; Eesaar, H.; Joe, S.; Rehman, M. U.; Jang, Y.; and Chong, K. T. 2023. SEiPV-Net: An Efficient Deep Learning Framework for Autonomous Multi-Defect Segmentation in Electroluminescence Images of Solar Photovoltaic Modules. *Energies 2023, Vol. 16, Page 7726*, 16(23): 7726.
- Jordan, D. C.; Kurtz, S. R.; VanSant, K.; and Newmiller, J. 2016. Compendium of photovoltaic degradation rates. *Progress in Photovoltaics: Research and Applications*, 24(7): 978–989.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W. Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *Proceedings of the IEEE International Conference on Computer Vision*, 3992–4003.
- Köntges, M.; Kurtz, S.; Packard, C.; Jahn, U.; Berger, K. A.; Kato, K.; Friesen, T.; Liu, H.; and Van Iseghem, M. 2014. *Review of Failures of Photovoltaic Modules*. ISBN 9783906042169.
- Li, D.; Su, B.; Zhang, H.; Li, Y.; and Chen, H. 2025. Photovoltaic Defect Image Generator with Boundary Alignment Smoothing Constraint for Domain Shift Mitigation.
- Li, M.; Fioresi, J.; Pannaci, N. K.; and Davis, K. O. 2021. Improving the Accuracy of Pixel-level Defect Classification of Electroluminescence Images. *Conference Record of the IEEE Photovoltaic Specialists Conference*, 1458–1459.
- Liu, X.; Jiang, S.; Wu, M.; and Xing, W. 2024. Defect detection of solar cells based on deep convolutional generative adversarial networks. *2024 5th International Conference on Computer Vision, Image and Deep Learning, CVIDL 2024*, 1502–1506.

- Luo, Z.; Cheng, S. Y.; and Zheng, Q. Y. 2019. GAN-Based Augmentation for Improving CNN Performance of Classification of Defective Photovoltaic Module Cells in Electroluminescence Images. *IOP Conference Series: Earth and Environmental Science*, 354(1): 012106.
- Ma, W.; Zhang, X.; Yao, Q.; Tang, F.; Wu, C.; Li, Y.; Yan, R.; Jiang, Z.; and Zhou, S. K. 2025. AA-CLIP: Enhancing Zero-Shot Anomaly Detection via Anomaly-Aware CLIP. *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- Maziuk, M.; Jasińska, L.; Domaradzki, J.; and Chodasewicz, P. 2023. IMAGING METHODS OF DETECTING DEFECTS IN PHOTOVOLTAIC SOLAR CELLS AND MODULES: A SURVEY. *Metrology and Measurement Systems*, 30(3): 381–401.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved Denoising Diffusion Probabilistic Models.
- Pratt, L.; Govender, D.; and Klein, R. 2021. Defect detection and quantification in electroluminescence images of solar PV modules using U-net semantic segmentation. *Renewable Energy*, 178: 1211–1222.
- Pratt, L.; Mattheus, J.; and Klein, R. 2023. A benchmark dataset for defect detection and classification in electroluminescence images of PV modules using semantic segmentation. *Systems and Soft Computing*, 5.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *Proceedings of Machine Learning Research*, 139: 8748–8763.
- Rico Espinosa, A.; Bressan, M.; and Giraldo, L. F. 2020. Failure signature classification in solar photovoltaic plants using RGB images and convolutional neural networks. *Renewable Energy*, 162: 249–256.
- Romero, H. F. M.; Hernández-Callejo, L.; Rebollo, M. G.; Cardeñoso-Payo, V.; Gómez, V. A.; Bello, H. J.; Moyo, R. T.; and Aragonés, J. I. M. 2023. Synthetic Dataset of Electroluminescence Images of Photovoltaic Cells by Deep Convolutional Generative Adversarial Networks. *Sustainability 2023, Vol. 15, Page 7175*, 15(9): 7175.
- Shou, C.; Hong, L.; Ding, W.; Shen, Q.; Zhou, W.; Jiang, Y.; and Zhao, C. 2020. Defect detection with generative adversarial networks for electroluminescence images of solar cells. *Proceedings - 2020 35th Youth Academic Annual Conference of Chinese Association of Automation, YAC 2020*, 312–317.
- Su, B.; Chen, H.; and Zhou, Z. 2022. BAF-Detector: An Efficient CNN-Based Detector for Photovoltaic Cell Defect Detection. *IEEE Transactions on Industrial Electronics*, 69(3): 3161–3171.
- Su, B.; Zhou, Z.; and Chen, H. 2023. PVEL-AD: A Large-Scale Open-World Dataset for Photovoltaic Cell Anomaly Detection. *IEEE Transactions on Industrial Informatics*, 19(1): 404–413.
- Su, B.; Zhou, Z.; Chen, H.; and Cao, X. 2021. SIGAN: A Novel Image Generation Method for Solar Cell Defect Segmentation and Augmentation.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the Inception Architecture for Computer Vision.
- Tan, J.; Shi, S.; and Chen, H. 2023. Anomaly Detection for Photovoltaic Modules Based on Image Inpainting. *Chinese Control Conference, CCC, 2023-July*: 6897–6902.
- Wang, Y.; Huang, H.; Rudin, C.; and Shaposhnik, Y. 2021. Understanding How Dimension Reduction Tools Work: An Empirical Approach to Deciphering t-SNE, UMAP, TriMap, and PaCMAP for Data Visualization. *Journal of Machine Learning Research*, 22(201): 1–73.
- Yang, B.; Zhang, Z.; and Ma, J. 2025. Wavelet-Based Normalized Flow for Anomaly Detection in Photovoltaic Electroluminescent with Nonstationary Textures. *IEEE Sensors Journal*, 25(1): 891–903.
- Zhou, Y.; Pei, S.; Chen, H.; and Liu, K. 2023. Generation of Solar Cell Defect Images Based on Multi-Perceptual Fields and Attention Mechanism. *Chinese Control Conference, CCC, 2023-July*: 7484–7489.