PROJECT ABSTRACT

Solar Panel Defect Detection from EL Images using CNN

A convolutional neural network that classifies electroluminescence (EL) images of solar photovoltaic cells into healthy, micro-crack, hot spot and snail trail categories. Trained on the public ELPV dataset with transfer learning from an EfficientNetV2-S backbone, the system includes a web demo with Grad-CAM defect localization, a batch QC review queue, and an evaluation report with confusion matrix and per-class metrics.

Project type: Software

Availability: Built to order

Suitable branches: AI & Machine Learning, Electrical, Electronics

Technologies: Python 3, TensorFlow, Keras, EfficientNetV2-S, ELPV dataset, Grad-CAM, OpenCV, Jupyter

Price: Request a quotation

Prepared by Projectech · projectech.in — for academic project reference.

1. Abstract

Photovoltaic modules lose power silently: micro-cracks from transport and hail, hot spots from cracked interconnections, and snail trails from silver nanoparticle migration all appear long before they are visible to the naked eye. Electroluminescence (EL) imaging makes these defects visible — a forward-biased cell glows under an EL camera, and any dark region marks an inactive area. This project replaces manual inspection of EL cell images with a trained convolutional neural network that classifies each cell as healthy, micro-crack, hot spot or snail trail, with confidence scores and a Grad-CAM heatmap showing which region drove the decision. The model is trained on the public ELPV dataset (2,624 EL images, Deitsch et al.) using transfer learning from an EfficientNetV2-S backbone, and ships with a working web demo, a batch QC review queue, and a complete viva kit.

2. Problem Statement

EL inspection stations today rely on humans paging through hundreds of cell images per module, which is slow, fatiguing and inconsistent — subtle micro-cracks and early snail trails are routinely missed. Missed defects become field failures: cracked cells overheat, create hot spots and drag down the output of an entire string. Module manufacturers and solar maintenance teams need an automated first-pass classifier that triages EL captures reliably and explains each verdict visually, so reviewers spend their time only on flagged cells.

This project provides that classifier: a CNN trained on real EL imagery, packaged with the preprocessing pipeline, the trained weights, and a review application that mirrors how a production-line inspection station works.

3. Objectives

- Classify single-cell EL images into healthy, micro-crack, hot spot and snail trail.
- Train on the public ELPV dataset using transfer learning, with a fully documented pipeline.
- Make every prediction explainable with a Grad-CAM defect-localization heatmap.
- Deliver a web demo that classifies uploaded EL captures with class probabilities.
- Provide a batch QC mode that triages many EL captures and flags defective cells.
- Evaluate honestly on a held-out test split: confusion matrix and per-class metrics.
- Ship a complete viva kit: report, presentation and Q&A; preparation document.

4. System Architecture

The system is organized in three stages. The **data stage** loads the ELPV dataset, re-organizes expert labels into the four working classes, and applies CLAHE contrast enhancement, resizing to 224×224 and normalization, with augmentation (rotations, flips, brightness jitter, Gaussian noise) that simulates EL camera variance. The **model stage** fine-tunes an EfficientNetV2-S backbone (top blocks trainable) with a custom head — global average pooling, dropout 0.35, dense 256 with Swish, and a 4-way softmax — optimized with class-weighted categorical cross-entropy and early stopping on validation F1. The **application stage** is a single-file web demo: uploaded EL images pass through the identical preprocessing pipeline, receive a class prediction with probabilities, a Grad-CAM heatmap overlay, a written finding and a recommended maintenance action, and can be queued in batch QC mode for line-speed triage.

5. Technologies Used

- Python 3
- TensorFlow / Keras
- EfficientNetV2-S (transfer learning)
- ELPV electroluminescence dataset
- Grad-CAM explainability

- OpenCV (CLAHE preprocessing)
- NumPy, Matplotlib, scikit-learn
- Jupyter Notebook
- HTML5 + JavaScript demo app

6. Working Principle

- EL images are loaded from the ELPV dataset and re-organized into healthy, micro-crack, hot spot and snail trail classes.
- CLAHE contrast enhancement, resizing and normalization prepare each image; augmentation simulates EL camera variance.
- The EfficientNetV2-S backbone with a custom classification head is fine-tuned on the training split with class-weighted loss.
- Early stopping on validation F1 prevents overfitting; the best checkpoint is kept.
- The held-out test split is evaluated once to produce the confusion matrix and per-class precision/recall/F1.
- Grad-CAM maps the predicted class back to the last convolutional layer for the defect heatmap.
- The demo applies the identical preprocessing to uploaded EL captures and returns the verdict with findings and a recommended action.
- Batch QC mode queues many captures and flags defective cells for manual review.

7. Key Features

- 4-class EL defect classifier (healthy, micro-crack, hot spot, snail trail) on a fine-tuned EfficientNetV2-S backbone.
- Grad-CAM heatmap overlay on every prediction, so the verdict is explainable.
- Web demo: upload EL captures or browse the sample gallery, with class probabilities and written findings.
- Batch QC review queue that triages production-line EL captures and flags defects.
- Trained on the real, public ELPV dataset — no synthetic data.
- Reproducible training notebook with augmentation, class weights and early stopping.
- Honest evaluation report: confusion matrix, per-class metrics and training curves from the order's run.
- Complete viva kit: report, PPT and Q&A; preparation document.

8. Dataset & Model Details

- Dataset: ELPV (Electroluminescence PV) by Deitsch et al. — 2,624 EL images of mono- and polycrystalline silicon cells with expert defect annotations, published as a public benchmark; re-organized into the four working classes.

- Task: 4-class image classification; input 224×224×3 preprocessed EL image, output probability distribution over four defect classes.
- Model: EfficientNetV2-S backbone (blocks 1–4 frozen, 5–6 fine-tuned) + global average pooling → Dropout(0.35) → Dense(256, Swish) → Dense(4, softmax); ~1.2M trainable parameters (design target).
- Split: stratified 70/15/15 train/validation/test, cell-disjoint to prevent leakage.
- Metrics: macro-F1 92.4% is the design target for the built-to-order training run; per-class precision/recall and confusion matrix are measured on the held-out test split. No metric is claimed as measured until the training run is executed for the order — the delivered report documents the actually achieved figures.
- Explainability: Grad-CAM from the last convolutional block is sanity-reviewed in the report to confirm the model attends to defect regions.

9. Limitations

- 92.4% macro-F1 is a design target, stated honestly — the report documents the actual achieved figure after the order's training run.
- Trained on single-cell EL crops; full-module EL images must be cropped to cells first (utility included, not fully automatic).
- ELPV labels carry known label noise — some cells have multiple defects — which affects measured metrics; discussed in the report.
- Very faint early-stage snail trails may be missed; per-class recall is quantified honestly.
- The model is an educational QC build, not a certified inspection system; it does not replace professional module certification.
- Performance depends on EL capture quality — uneven bias current or stray light changes the image distribution.

10. Conclusion

Solar Panel Defect Detection from EL Images using CNN gives students a complete, honest deep-learning project on a real industrial inspection problem: real electroluminescence data, transfer learning on a modern backbone, explainable predictions through Grad-CAM, and a working deployment demo with a batch QC mode that mirrors production-line practice. The deliverables — training notebook, trained weights, demo app, evaluation report and viva kit — are prepared fresh per order, and the evaluation report documents the actually measured figures from that order's training run rather than any borrowed claim. For the viva, the student can walk through the EL physics, the dataset, the architecture choices, the Grad-CAM explanations and the per-class failure modes with genuine understanding.