

AI-Driven Predictive Maintenance for Solar Photovoltaic Systems: A Machine Learning Framework for Sustainable Energy Infrastructure

Ananya Roy

Department of Electrical Engineering,
National Institute of Technology, Jaipur, Rajasthan, India.
E-mail: ananya.roy@example-nitj.ac.in
ORCID: 0000-0002-1111-2222

Vikram Deshmukh

Department of Computer Science and Engineering,
Indian Institute of Information Technology, Pune, Maharashtra, India.
E-mail: vikram.deshmukh@example-iiitp.ac.in
ORCID: 0000-0003-3333-4444

Upasana Gupta

Centre for Sustainable Systems Research,
BallPen Research Foundation, New Delhi, India.
Corresponding author: u.gupta@example-ballpen.org
ORCID: 0000-0001-5555-6666

(Received on 03 August 2026; Revised on 21 August 2026; Accepted on 02 September 2026)

Abstract

Unplanned downtime in solar photovoltaic (PV) installations remains a major barrier to reliable renewable energy supply, particularly for small and mid-scale plants that lack continuous manual inspection. This paper proposes a machine learning framework that combines sensor-derived time-series features (irradiance, module temperature, current-voltage ratios) with a gradient-boosted ensemble classifier to predict panel degradation and impending faults up to 72 hours before failure. The model was trained and validated on a twelve-month operational dataset from a 500 kW grid-connected plant in western India, supplemented by a synthetic fault-injection dataset to balance rare failure classes. The proposed framework achieved a fault-prediction accuracy of 94.2 percent and reduced false-negative alerts by 31 percent compared with a threshold-based baseline. Results indicate that predictive maintenance driven by interpretable machine learning can meaningfully lower operations and maintenance costs while improving energy yield, offering a scalable pathway for AI-enabled sustainability in distributed solar infrastructure. Limitations related to dataset generalizability and the need for explainability in safety-critical deployment are discussed, along with directions for future multi-site validation.

Keywords— predictive maintenance; solar photovoltaic systems; machine learning; sustainable energy; fault detection; gradient boosting.

1. Introduction

Solar photovoltaic (PV) capacity has expanded rapidly over the past decade as national grids pursue decarbonization targets, yet operational efficiency has not kept pace with installed capacity. Undetected faults such as micro-cracking, hotspot formation, and inverter degradation can reduce energy yield by 10 to 25 percent annually, and routine manual inspection is costly and difficult to scale across distributed sites (IEA, 2024). Artificial intelligence offers a route to condition-based maintenance that shifts inspection from fixed schedules to data-driven predictions, reducing both downtime and unnecessary maintenance visits.

This study addresses two open problems in the predictive-maintenance literature for PV systems: first, most published models are validated on simulated or single-season data, limiting confidence in year-round deployment; second, few studies report the trade-off between prediction accuracy and false-negative rate, which is the operationally relevant metric for maintenance planning. The present work develops and validates a gradient-boosted ensemble model on twelve months of real plant telemetry, explicitly optimizing for reduced false negatives given the asymmetric cost of a missed fault.

The remainder of this paper is organized as follows. Section 2 reviews related work on AI-based condition monitoring for renewable infrastructure. Section 3 describes the dataset, feature engineering, and model architecture. Section 4 presents experimental results and comparative analysis. Section 5 discusses implications, limitations, and directions for future work, followed by conclusions.

2. Related Work

2.1 Condition Monitoring in Renewable Infrastructure

Early condition-monitoring systems for PV plants relied on rule-based thresholds applied to inverter and string-level telemetry (Patel & Rao, 2021). While computationally inexpensive, threshold methods generalize poorly across sites with differing irradiance profiles and degrade in accuracy as panels age. Subsequent work introduced statistical process control and IoT-enabled sensor networks to improve detection latency, but these approaches remained reactive rather than predictive.

2.2 Machine Learning for Fault Prediction

Supervised learning methods, particularly random forests and gradient boosting, have shown improved fault-classification performance over rule-based baselines (Kumar, Singh, & Mehta, 2023; Alam & Sharma, 2023). Deep learning architectures, including convolutional and recurrent networks, have been applied to thermal imagery and time-series telemetry with reported accuracies above 90 percent under controlled conditions (Chatterjee, Kumar, & Das, 2022). However, Singh, Bansal, and Iyer (2024) note that black-box deep models face adoption barriers in operations teams that require interpretable justification before scheduling maintenance, motivating continued interest in tree-based ensembles that support feature-importance analysis.

Gupta and Verma (2024) further argue that predictive analytics for distributed renewable assets in emerging-economy contexts must account for constrained maintenance budgets typical of small and medium enterprises operating captive solar installations, a consideration that directly informed the cost-sensitive design of the present model.

3. Methodology

3.1 Dataset and Feature Engineering

Operational data were collected from a 500 kW grid-connected solar plant over a twelve-month period at fifteen-minute resolution, comprising irradiance, module temperature, string current, string voltage, and inverter efficiency. A synthetic fault-injection procedure was applied to the minority fault classes (hotspot, string disconnection, and inverter derating) to address class imbalance, following established data-augmentation practice for rare-event prediction.

Twelve engineered features were derived, including rolling seven-day efficiency deviation, current-voltage ratio drift, and temperature-normalized power output. Features were standardized and split into training (70

percent), validation (15 percent), and test (15 percent) sets using a temporal split to prevent information leakage across time.

3.2 Model Architecture

A gradient-boosted decision tree ensemble (XGBoost) was selected for its balance of predictive accuracy and interpretability through feature-importance scoring. Hyperparameters (maximum depth, learning rate, number of estimators) were tuned via five-fold cross-validation on the training set, with the class-weighted log-loss objective adjusted to penalize false negatives more heavily than false positives, reflecting the higher operational cost of an undetected fault relative to an unnecessary inspection.

$$F(x) = \sum_{m=1}^M \gamma_m h_m(x), \text{ subject to } L(y, F) = \sum_i w_i \ell(y_i, F(x_i)) \quad (1)$$

where $h_m(x)$ denotes the m -th regression tree, γ_m its learned weight, and w_i the class-specific penalty weight applied to observation i , with w_i set higher for fault-class instances during training.

4. Results and Analysis

Table 1 summarizes model performance against a threshold-based baseline and a standard (unweighted) gradient-boosted classifier. The cost-sensitive model achieved the highest overall accuracy and the lowest false-negative rate, confirming that reweighting the loss function toward fault classes materially improves early-warning reliability without a large precision penalty.

Model	Accuracy (%)	Precision (%)	Recall (%)	False Negative Rate (%)
Threshold-based baseline	78.4	71.2	64.9	35.1
Unweighted XGBoost	90.6	88.3	85.7	14.3
Cost-sensitive XGBoost (proposed)	94.2	91.5	93.1	6.9

Table 1. Comparative performance of fault-prediction models on the held-out test set.

Feature-importance analysis identified rolling efficiency deviation and current-voltage ratio drift as the two strongest predictors, jointly accounting for 47 percent of total gain, consistent with prior findings linking these features to early-stage panel degradation (Kumar, Singh, & Mehta, 2023). The model's 72-hour prediction horizon allowed maintenance scheduling that reduced average unplanned downtime from 4.8 to 3.1 days per fault event over the evaluation period, an estimated 35 percent reduction.

5. Discussion

The results support the case that cost-sensitive, tree-based ensembles offer a practical middle ground between the poor generalization of threshold methods and the interpretability limitations of deep learning approaches, particularly for maintenance teams operating under budget constraints typical of distributed and small-scale solar assets. Nonetheless, several limitations warrant caution before broader deployment.

First, the model was validated on a single plant and climate zone; performance under different irradiance regimes and panel technologies remains untested and should be confirmed through multi-site trials. Second, while feature importance provides partial interpretability, it does not constitute a causal explanation suitable for safety-critical maintenance decisions, and future work should incorporate formal explainability methods such as SHAP value analysis. Third, the synthetic fault-injection procedure, while necessary to address

class imbalance, may not fully capture the statistical characteristics of naturally occurring rare faults, and results should be re-validated as more organic fault data accumulate.

6. Conclusion

This paper presented a cost-sensitive gradient-boosted machine learning framework for predictive maintenance of grid-connected solar PV systems, achieving 94.2 percent fault-prediction accuracy and a 31 percent reduction in false negatives relative to a threshold-based baseline, with an estimated 35 percent reduction in unplanned downtime in the evaluation deployment. The findings support predictive, AI-enabled maintenance as a scalable lever for improving the reliability and cost-efficiency of distributed renewable energy infrastructure. Future work will extend validation to multiple sites and panel technologies and incorporate formal explainability methods to support adoption in safety-critical operational settings.

Conflict of Interest

The authors confirm that there is no conflict of interest to declare for this publication.

Acknowledgments

The authors thank the plant operations team for providing anonymized telemetry data and acknowledge institutional computing support from their respective departments. No external funding was received for this study.

AI Disclosure

No generative AI tools were used to produce the research results, data analysis, or conclusions reported in this manuscript. Generative AI was used solely for language editing of the draft, and all such output was reviewed and verified by the authors, who take full responsibility for the content of the publication.

References

- Alam, M., & Sharma, R. (2023). Machine learning approaches for fault diagnosis in photovoltaic systems: A review. *Renewable and Sustainable Energy Reviews*, 178, 113245. <https://doi.org/10.1016/j.rser.2023.113245>
- Chatterjee, S., Kumar, P., & Das, A. (2022). Deep learning-based anomaly detection for grid-connected solar plants. *IEEE Transactions on Sustainable Energy*, 13(4), 2011–2022. <https://doi.org/10.1109/TSTE.2022.3167890>
- Gupta, U., & Verma, N. (2024). Predictive analytics for renewable infrastructure in emerging economies: An MSME perspective. *Journal of Cleaner Production*, 434, 139872. <https://doi.org/10.1016/j.jclepro.2023.139872>
- IEA. (2024). *Renewables 2024: Analysis and forecast to 2030*. International Energy Agency.
- Kumar, A., Singh, R., & Mehta, S. (2023). Random forest and gradient boosting for solar panel degradation forecasting. *Solar Energy*, 251, 112–124. <https://doi.org/10.1016/j.solener.2023.01.012>
- Patel, D., & Rao, K. (2021). IoT-enabled condition monitoring of photovoltaic modules: Architecture and case study. *Sustainable Energy Technologies and Assessments*, 45, 101112. <https://doi.org/10.1016/j.seta.2021.101112>
- Singh, H., Bansal, M., & Iyer, V. (2024). Explainable AI for predictive maintenance in smart grids. *Applied Energy*, 356, 122341. <https://doi.org/10.1016/j.apenergy.2023.122341>
- World Bank. (2023). *Tracking SDG 7: The energy progress report*. World Bank Group.