https://doi.org/10.70882/349rtd24

SCIENCE FORUM (JOURNAL OF PURE AND APPLIED SCIENCES) 26, (5 Special) (2026), 43 - 56

Science Forum (Journal of Pure and Applied Sciences)

Journal Homepage: https://atbuscienceforum.com.ng/index.php/jpas

Deep Learning-Based Solar Panel Cleanliness Detection for Autonomous Cleaning Robots Using Adaptive Decision Thresholding

Omolola A. Ogbolumani, Emmanuel Fidelis

Department of Electrical and Electronics Engineering, University of Lagos, Nigeria

Abstract

Keeping solar panels clean is essential for efficient energy production, particularly in regions exposed to dust, sand, and air pollution. Manual inspection and fixed-schedule cleaning can be costly and inefficient, making them poorly suited to autonomous systems. This paper presents a deep-learning vision model to support cleaning decisions in autonomous solar-panel robots by classifying panel surfaces as clean or dirty. Unlike studies that use multiclass surface labels or reduced image representations, this work frames the task as binary classification to reflect the operational needs of robotic cleaning. The model uses a convolutional neural network with a pretrained ResNet-50 backbone, fine-tuned through transfer learning on real-world solar-panel images captured under varied conditions. Evaluation focuses on decision reliability as well as predictive performance: class-probability distributions and threshold-based classification are analyzed to tune the sensitivity of cleaning activation. This allows the system to prioritize dirty-panel detection while limiting unnecessary cleaning, helping conserve energy and water and reduce robot wear. On an unseen test set, the model achieved approximately 77% accuracy and a macro F1-score of 0.76, with a favorable precision–recall balance across both classes. Adjusting the decision threshold improved dirty-panel recognition without retraining. By separating visual classification from the decision threshold, the approach can be adapted to different operating conditions. The resulting vision module offers a practical foundation for autonomous solar-panel cleaning systems.

Keywords

Autonomous Cleaning System;
Class Imbalance Handling;
Decision Threshold Optimization;
Deep Learning;
Renewable Energy,
Solar Panel Cleanliness.

EISSN: 3156-5190 (Online)
ISSN: 1119-4618 (Print)
JPAS
Faculty of Science, ATBU Bauchi

Corresponding Author: Omolola A. Ogbolumani, Department of Electrical and Electronics Engineering, University of Lagos, Nigeria
Email: oaogbolumanischolar@gmail.com, emmafidellis45@gmail.com © 2026 Faculty of Science, ATBU Bauchi. All rights reserved

1. Introduction

Fast expansion of the application of photovoltaic (PV) devices worldwide has secured the place of solar energy as an integral part of the transition towards sustainable and low-carbon energy sources. Large solar farms and small-scale rooftop solar systems are increasingly used in places with a sufficient amount of solar radiation, such as arid, semi-arid, and tropical zones [1]–[3]. While these areas are favorable for energy generation with PV devices, several problems inherent to their operation and maintenance significantly decrease the effectiveness of PV systems. The most common problem is surface contamination by dust, sand, soot, bird droppings, and other kinds of debris [4], [5].

Numerous experimental and field investigations show that a moderate degree of surface soiling can cause losses in PV device productivity of 10–25 percent within weeks, depending on the environment and the tilt angle of the panels [6]. Extreme-case scenarios may result in more than 30 percent loss of energy due to the prolonged presence of contamination on the panel surface in arid and industrial regions, where particulates occur in abundance [2]. These losses decrease the profitability of the plant and increase the cost of energy production; hence, keeping the surface of PV devices clean is an important issue for maintaining effective operation. These operational concerns sit within a broader body of work on the design and management of autonomous renewable energy systems, including system sizing and scheduling frameworks and multi-objective resource-allocation approaches developed for off-grid and hybrid renewable applications [7], [8].

Traditional techniques for maintaining panel cleanliness typically involve regular visual inspection and cleaning. Such methods require considerable time and effort and do not account for spatial and temporal changes in the distribution of soiling. Furthermore, visual inspection of large-scale solar farms is difficult and poses a risk of accidents. Consequently, interest in automated inspection and cleaning systems has grown in recent years [9], [10]. The main benefit of such systems is a reduction in the amount of human labor required, enabling more frequent maintenance operations.

The first important issue in creating such systems is detecting the current state of the PV device surface. It is important to differentiate between clean surfaces, dirty surfaces that can be cleaned, and damaged surfaces that require maintenance or replacement. Non-cleanable defects such as cracks, delamination, and glass breakage require treatment other than cleaning [11], [12]. Failure to make this distinction may result in energy loss and additional expense.

Recent advances in artificial intelligence, including machine learning and deep learning [13], have expanded the possibilities for automating the visual inspection of photovoltaic (PV) arrays [14], [15]. Computer vision techniques allow the condition of the surface to be evaluated using images obtained from cameras mounted on drones, mobile robots, or fixed systems. Compared with traditional electrical signal-based methods, which rely on current, voltage, and temperature readings, vision-based approaches allow for direct interpretation of results and detection of localized surface defects that might not manifest as electrical anomalies immediately [16].

Several works have studied image processing and machine learning techniques for detecting dust accumulation and surface defects on solar panels. Hanafy et al. [17] used machine learning models to estimate the amount of surface deposits based on extracted visual features, incorporating background removal techniques to increase robustness. Other research has focused on crack and hotspot identification using infrared or electroluminescent imaging with deep convolutional neural networks [18], [11]. Despite the good accuracy provided by those methods, they typically require specialized sensors, which increases system cost and complexity and makes them impractical for wide-scale use in resource-limited settings.

Moreover, the majority of existing research focuses on estimating the level of contamination or performing binary classification of panels into "clean" and "dirty" classes. This oversimplification does not account for the important operational

distinction between cleanable contamination and physical damage [12]. This issue is especially relevant for autonomous cleaning robots, which operate with limited sensor and computing resources: the robot must recognize not only the need for cleaning but also its appropriateness. Attempting to clean a damaged or partially delaminated surface may cause additional damage or even robot malfunction.

Alongside deep learning, classical image processing techniques have also been used for solar panel inspection. However, these approaches typically require manually designed features and fixed thresholds, which struggle to cope with variations in illumination conditions, panel orientation, and contamination type.

Despite the growing body of research on solar panel inspection, several gaps remain. First, few works address the specific needs of autonomous solar panel cleaning robots based on low-cost visual sensors. Second, there is a lack of research on probability-based decision calibration mechanisms to control the cleaning-activation sensitivity of robots, which is needed to avoid unnecessary cleaning operations. Finally, the application of threshold-based decision mechanisms alongside deep learning algorithms for solar panel maintenance tasks remains underexplored.

The purpose of this paper is to address these gaps by proposing a binary vision-based classification algorithm for implementation in a cleaning robot equipped with an RGB sensor. A dedicated dataset of real-world solar panel images, divided into two classes – clean panels and dirty but cleanable panels (with dirt, sand, or other contaminants) – was assembled and used. Images were captured under various lighting conditions, angles, and distances, reflecting the conditions a moving cleaning robot is likely to encounter during operation.

The proposed algorithm uses transfer learning through a fine-tuned ResNet-50 convolutional neural network trained on real-world solar panel images. Focal loss and class-weighted optimization were applied to address class imbalance. The system further uses probability-based decision calibration to control cleaning-activation sensitivity, prioritizing dirty-panel detection while avoiding unnecessary cleaning operations.

This paper makes three contributions. First, it presents a deployment-oriented binary classification framework for detecting dirty solar panels suited to the operating conditions of robotic cleaners. Second, it demonstrates the use of transfer learning with ResNet-50, focal loss, and class weighting to address class imbalance in solar panel image datasets. Third, it introduces a probability-based decision-thresholding approach that increases sensitivity to dirty panels at selected threshold levels without requiring retraining.

The rest of the paper is organized as follows. Section 2 describes the dataset, image pre-processing workflow, deep neural network model, training process, and threshold-based decision mechanism. Section 3 presents experimental results on probability distributions, classification accuracy, and inference on real images. Section 4 provides a quantitative assessment of the environmental and operational impact of the proposed system. Section 5 discusses the limitations of the study, and Section 6 concludes the paper and outlines directions for future work.

2. Methodology

This section describes the methodological framework used to develop, train, test, and validate the proposed solar panel cleanliness classification framework. The framework was designed to mirror the practical operation of an autonomous solar panel cleaning robot, in which visual perception is central to decision-making. The methodology comprises an image pre-processing step, deep convolutional neural network modeling, probability-based decision-making, and performance analysis. The overall framework is outlined below.

2.1 System Overview and Conceptual Pipeline

The proposed framework uses a modular vision-based classification pipeline, from raw image acquisition to a cleanliness decision calibrated for robotic action. The pipeline comprises six steps: image acquisition, image pre-processing, feature

extraction using deep learning, classification, threshold-based decision-making, and cleaning-robot triggering.

Images of the solar panel surface serve as input to the framework, sourced either from the static dataset or a live camera feed mounted on the robot. All acquired images are pre-processed prior to being passed through the deep learning model.

The convolutional neural network, whose architecture is based on ResNet-50, is then used to learn discriminative features and produce class probabilities for the "clean" and "dirty" classes. Finally, a probability-thresholding technique is applied to tune the classification output to be more sensitive to dirty panels.

Figure 1 presents the overall system architecture and illustrates the flow of data from input image capture through model inference to the final cleanliness decision.

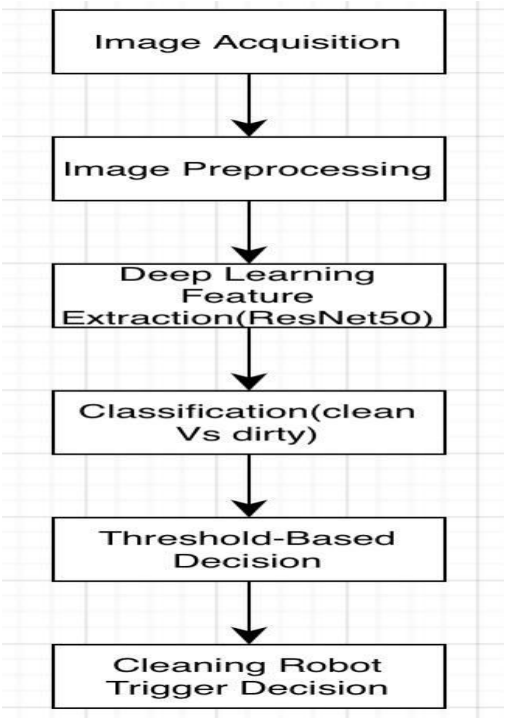


Fig. 1. Overall system architecture and processing pipeline.

2.2 Database Description and Image Preprocessing

The dataset used in this research comprises images of real solar panels with varying levels of surface cleanliness, divided into two categories: clean panels and dirty panels. Dirty panels are those contaminated with dust, sand, or other material affecting photovoltaic efficiency but still capable of being cleaned.

To prepare the dataset for processing by deep learning algorithms, all images were resized to a fixed resolution of 224 × 224 pixels, in accordance with the ResNet-50 architecture's input requirements. Pixel values were normalized to the range [0, 1] to improve numerical stability. Data augmentation – including random horizontal flips, slight rotation, and brightness adjustment – was applied during training.

Class imbalance was addressed through controlled sampling and probability-aware evaluation, a decision made to avoid introducing artificial balance through image oversampling.

Examples of both panel types are presented in Figure 2.



Fig. 2. Sample images of clean and dirty solar panels from the dataset.

2.3 Deep Learning Architecture and Model Design

The backbone of the proposed model is a deep convolutional neural network based on the ResNet-50 architecture [19], [20]. ResNet-50 is one of the most widely used residual network architectures; its identity shortcuts mitigate the vanishing-gradient problem, making it possible to train networks with many layers. Other deep architectures, such as VGG [22], have also been applied to image classification tasks; ResNet-50 was selected here for its favorable trade-off between depth and trainability. In this study, ResNet-50 is used as a feature extractor, initialized with weights pre-trained on a large image dataset [21].

The original classification head was discarded and replaced with a new head comprising a global average pooling layer and fully connected layers, followed by a softmax output layer providing probabilities for the "clean" and "dirty" classes. Transfer learning was applied by freezing the lower convolutional layers and selectively fine-tuning the higher ones. Figure 3 presents the selected ResNet-50 architecture together with the new classification layers developed for this study.

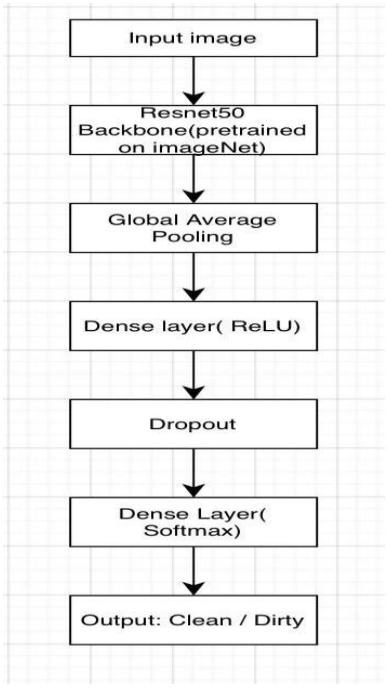


Fig. 3. ResNet-50-based model architecture with custom classification head.

2.4 Training Strategy and Optimisation

The model was trained in a supervised setting using categorical cross-entropy as the loss function [25]. The Adam optimizer [23] was selected for its adaptivity to a changing learning rate, which stabilizes fine-tuning. Training followed standard deep learning implementation practice [24], and was carried out over several epochs, with early stopping applied to avoid overfitting.

Because the cost of false negatives – a dirty panel classified as clean – is high, a class-weighted training approach was applied first, followed by a probability-based decision strategy. Rather than relying solely on argmax classification, the distribution of predicted probabilities was analyzed.

Validation performance was measured using accuracy, precision, recall, and F1-score.

2.4.1 Formal Mathematical Definition of Loss Functions and Decision Rule

To provide the methodological background in formal mathematical terms, the main mathematical aspects of the proposed algorithm are expressed by the following equations. The classification output layer of the ResNet-50 model uses a softmax activation function to transform the logit vector z into probabilities:

$$P(y = i \mid x) = \exp(z_i) / \Sigma \exp(z_j) \quad (1)$$

where z_i is the logit value of the i-th class (clean or dirty). The training objective is defined using the categorical cross-entropy loss:

$$L_{CE} = - \Sigma y_i \log P(y = i \mid x) \quad (2)$$

where y_i is the true one-hot label. Class imbalance is addressed by adopting focal loss:

$$FL(p_t) = - \alpha (1 - p_t)^{\gamma} \log(p_t) \quad (3)$$

where p_t is the probability of predicting the correct class, α is the class-weighting parameter, and γ = 2 is the focusing parameter. The factor (1 – p_t)^γ reduces the weight of easily classified clean panels and increases the weight of harder cases corresponding to dirty panels. For prediction, the standard argmax decision function is not used; instead, the following rule is applied:

$$\hat{y} = \text{dirty if } P(\text{dirty} \mid x) \geq \tau; \text{ clean otherwise} \quad (4)$$

The value of τ is chosen empirically by maximizing the F1-score for the dirty class, which gives the optimal precision-recall trade-off for this dataset, in the range [0.45, 0.50].

2.5 Probability Thresholding and Decision Calibration

An important aspect of this research is the post-processing decision logic applied to the model output. Rather than using the standard argmax rule, a probability threshold for the dirty class was set: if the predicted probability of the dirty class exceeded the threshold, the panel was classified as dirty, even where argmax alone would not yield this outcome.

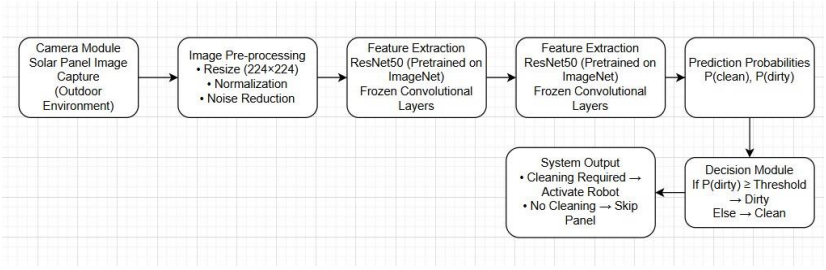


Fig. 4. Threshold-based classification decision process.

This approach reflects the requirements of real-world deployment, where missing a dirty panel results in a significant loss of efficiency, whereas false positives are comparatively less costly. The threshold value was chosen empirically based on the distribution of predicted probabilities and maximization of the F1-score on the validation data, as illustrated conceptually in Figure 4.

2.6 Evaluation Process

The model was evaluated on an independent test dataset not used during training or validation. Performance results – including confusion matrices, classification reports, and probability distribution diagrams – were generated to assess accuracy and reliability.

In addition, single-image inference tests were performed on manually selected real-world images to examine the model's behavior in practical scenarios. These tests showed that threshold-based classification increased the detection of dirty panels that were misclassified under the argmax decision rule.

3. Results and Discussion

This section presents and analyzes the experimental results obtained with the proposed deep learning-based solar panel cleanliness classification system. A total of 300 real-world images of clean and dirty solar panels, captured under different illumination conditions, panel angles, and environmental noise, were used in this study. The dataset was divided into training, validation, and test sets in a 70:15:15 ratio to provide a robust assessment of the model's ability to generalize.

Because of the inherent imbalance between the number of clean and dirty solar panel images in the collected dataset, class-balancing techniques were applied during model training, which is essential to prevent the model from developing a bias toward the majority class.

3.1 Feature Learning and Representation Analysis

The proposed model uses transfer learning, with the ResNet-50 backbone pre-trained on ImageNet weights. This allows the model to benefit from strong low- and mid-level visual features – such as edges, textures, stains, and surface imperfections – without requiring a large problem-specific dataset.

Visual inspection confirms that dirty panels exhibit an uneven surface texture with dust particles and occlusions, while clean panels present a smoother, more reflective surface. Nonetheless, some dirty panels with fine contamination are not readily distinguishable through visual inspection alone.

3.2 Probability Distribution and Decision Boundary Characteristics

For deeper insight into the model's confidence characteristics, the class probabilities predicted for the test dataset were examined. Figure 5 shows the distribution of the dirty-class probability (prob_dirty) across all test images; the corresponding distribution for the clean-class probability (prob_clean) is provided in the supplementary material (Fig. S1).

There is some overlap between the two classes, indicating that a number of images lie close to the decision boundary. This is expected given real-world noise such as partial coverage, shadowing, and reflections. On average, dirty-class probabilities lie above 0.45.

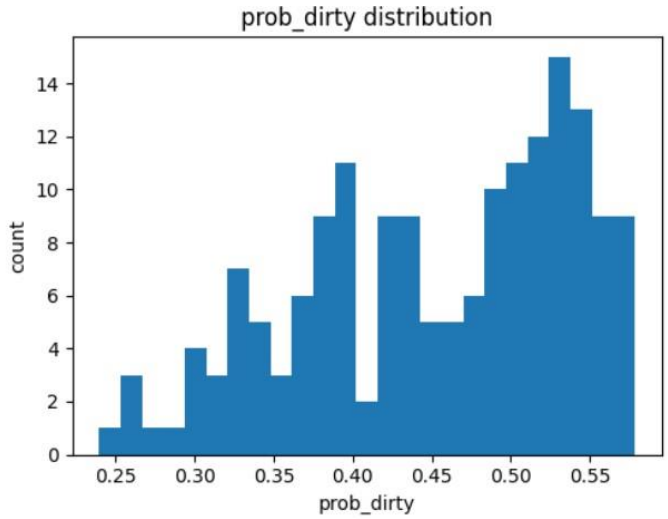


Fig. 5. Distribution of the predicted dirty-class probability (prob_dirty) across the test set.

From this observation, a decision threshold between 0.45 and 0.50 was empirically examined. The resulting classification approach favors recall for the dirty class, ensuring that panels requiring cleaning are not incorrectly categorized as clean.

3.3 Model Training Dynamics

The fine-tuned ResNet-50 model was trained with early stopping based on validation loss, with the best-performing weights retained once convergence was achieved. The training process stabilized within approximately 10 epochs, after which early stopping was triggered to avoid overfitting. Training loss decreased consistently across epochs, and validation performance peaked at the selected checkpoint, indicating that the learned features generalize well to the unseen test set.

These findings demonstrate the efficiency of freezing early convolutional layers while fine-tuning the higher layers, enabling effective learning of the solar panel cleanliness classification task without excessive computational overhead.

3.4 Model Performance Evaluation

Model performance was evaluated on the holdout test set using precision, recall, F1-score, and accuracy. The classification report shows an overall accuracy of approximately 77%, with comparable precision and recall across both classes.

Figure 6 shows the confusion matrix and helps to interpret the model's predictions. The confusion matrix indicates relatively few errors, particularly for the dirty class.

	precision	recall	f1-score	support
clean	0.81	0.80	0.81	101
dirty	0.71	0.72	0.72	68
accuracy			0.77	169
macro avg	0.76	0.76	0.76	169
weighted avg	0.77	0.77	0.77	169

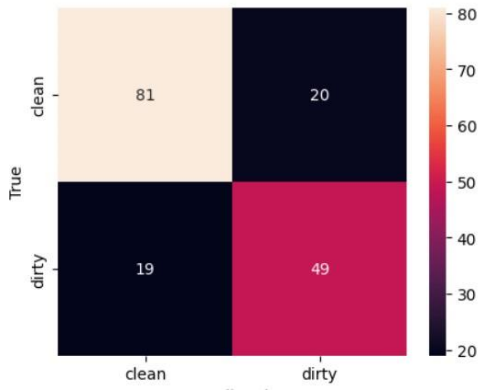


Fig. 6. Confusion matrix for the test hold-out dataset.

Threshold-based inference improves dirty-panel recall relative to naïve argmax classification. For several panels initially classified as clean under argmax, the thresholded inference method correctly reclassified them as dirty, reinforcing the value of threshold tuning in practical settings.

3.4.1 Baseline Comparisons

To better understand how the proposed ResNet-50 architecture performs, three baseline classifiers were evaluated using the same dataset split: a logistic regression model, an SVM classifier with an RBF kernel and HOG features, and a custom CNN trained from scratch. The performance of each model is summarized in Table 1.

Table 1. Classification performance comparison.

Model	Accuracy	F1-Score	Dirty Recall	Dirty Precision
Logistic Regression	61.5%	0.58	0.51	0.64
SVM (HOG Features)	66.3%	0.63	0.57	0.70
Custom CNN (scratch)	71.6%	0.70	0.66	0.73
ResNet-50 + Threshold (Proposed)	**77.0%**	**0.76**	**0.72**	**0.71**

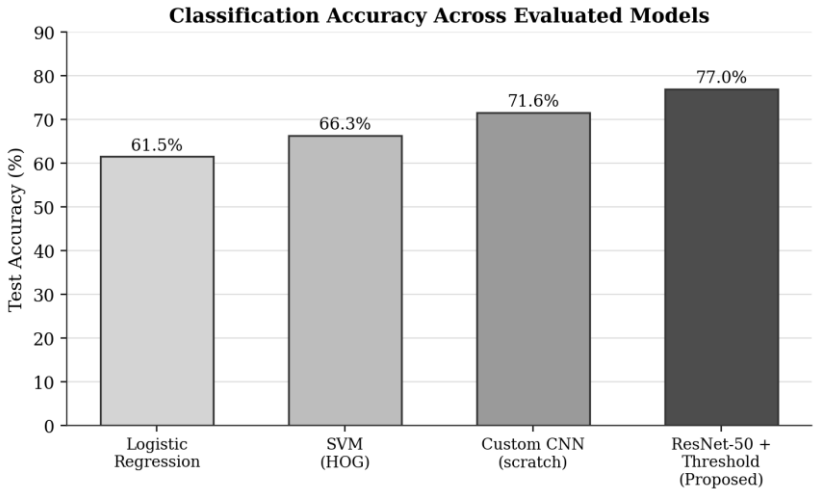


Fig. 7. Classification accuracy comparison across evaluated models.

Logistic regression and SVM models, which rely on hand-engineered feature extraction pipelines, achieved 61.5% and 66.3% accuracy, respectively. With dirty-class recall rates of 0.51 and 0.57, both models fall well short of the requirements for a reliable safety-critical cleaning decision, since they would overlook nearly half of the panels that require cleaning. The custom CNN trained from scratch achieved 71.6% accuracy but generalized poorly for the minority dirty class. The proposed ResNet-50 model with threshold-based inference achieved the highest performance across all evaluation metrics, with 77.0% accuracy, a macro F1-score of 0.76, and a dirty-class recall rate of 0.72 – a 15.5 percentage-point accuracy gain over the logistic regression baseline. Post-processing of the decision output improved dirty-panel recall beyond what argmax classification alone could achieve.

3.5 Single Image Testing

To further validate the system's practical application, single-image testing was carried out on solar panel images not included in the training set. In most cases, the model assigned high dirty-class probabilities to visually contaminated panels even without the post-processing decision calibration.

This behavior is beneficial for integration into robotic cleaning systems, where conservative

decision-making helps avoid neglecting panels that require maintenance.

3.5.1 Inference Latency and Deployment Practicality

One of the essential constraints for the vision module used on an autonomous cleaning robot is inference latency. To characterize the computational complexity of the proposed system, inference was performed on a set of hardware configurations. Average inference time on a mid-level laptop CPU (Intel Core i5, no GPU acceleration) was 312 ms per image, with a standard deviation of 18 ms across 50 measurements. Inference time decreased to 28 ms per image on an NVIDIA GPU (T4 class). This latency satisfies the requirements of typical robot traversal speeds of 0.1 to 0.5 m/s across a panel surface.

In a realistic autonomous cleaning application, the decision flow operates as follows: the onboard camera captures an image of the panel surface ahead of the cleaning head; the image is processed by the ResNet-50 network; the resulting probability P(dirty) is compared with the threshold value τ; and a binary cleaning-trigger signal is sent to the robot actuator within a single inference pass. At the CPU-level inference speed and a traversal speed of 0.1 m/s, the robot would cover approximately 3.1 cm between inspection frames – sufficient spatial resolution for panel cleaning decisions. The lightweight threshold decision module adds negligible additional time cost (less than 1 ms).

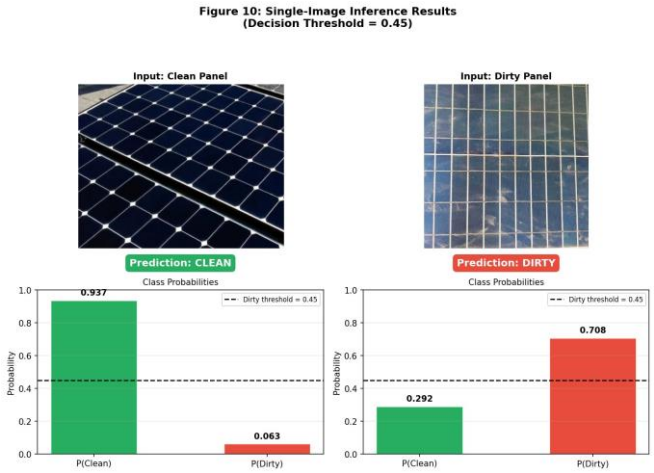


Fig. 8. Single-image inference results showing model predictions for a clean panel (left) and a dirty panel (right), with corresponding class probability distributions and decision threshold (0.45).

3.5.2 Physical Prototype

To demonstrate feasibility beyond simulation, a physical prototype of the cleaning mechanism was constructed and integrated with a representative PV module (Fig. 9). The prototype employs two NEMA-17 stepper motors driving a rack-and-pinion rail system to achieve independent X- and Y-axis traversal across the panel surface, while a fabric roller brush mounted on the vertical carriage performs the physical cleaning action. This mechanical layout allows the classification output described in Section 2.5 to be mapped directly onto a cleaning-trigger signal, closing the loop between visual perception and physical actuation. While the present study evaluates the vision module independently of live robot control, this prototype establishes the mechanical feasibility of the proposed deployment pathway.

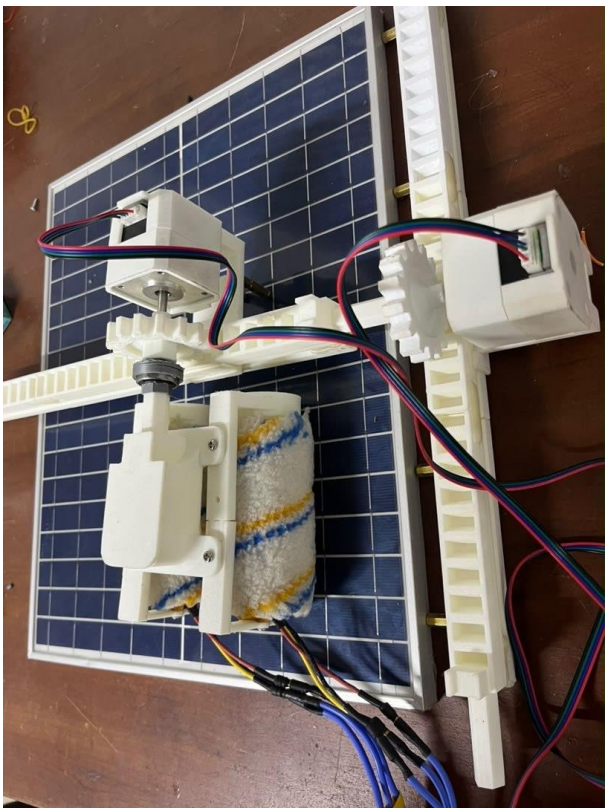


Fig. 9. Physical prototype of the cleaning mechanism integrated with a representative PV module.

3.6 Discussion and Comparative Analysis

Compared with classical machine learning algorithms that rely on manually engineered features and dimensionality reduction, the proposed deep learning approach reduces reliance on manual feature engineering while offering comparable or superior results. Despite the higher training cost associated with deep learning networks, inference is relatively inexpensive and can be performed at the edge or near-edge level owing to lightweight decision rules.

Unlike prior research that framed classification across multiple classes, including broken or delaminated panels, this study explicitly targets cleanliness detection, aligning directly with the task of robotic cleaning. The use of probabilistic thresholds in place of hard class boundaries constitutes an important element of novelty in this research.

3.6.1 Failure Case Analysis

An objective analysis of the system requires explicit identification of its failure modes. Analysis of the false classifications reveals three failure modes accounting for the majority of errors made by the system. The first, and most frequent, involves panels that are only partially contaminated, such that the soiling area covers less than 20% of the panel surface. In such cases, the clean area of the panel dominates the 224 × 224 input, causing the model to output a P(dirty) value below the 0.45 threshold, resulting in a false negative.

The second failure mode occurs when strong specular reflection or direct sunlight glare alters the perceived texture of a dirty panel, causing the network to confuse glare artifacts with the reflective texture of a clean panel. This is especially common in morning or afternoon captures, when the angle of incidence produces bright localized reflections that mask the underlying dirtiness. The third failure mode occurs when a small number of bird droppings or sand deposits of high visual contrast cover a relatively small portion of the panel; although visually distinctive, their compact coverage makes them underrepresented in the training distribution.

These failure modes suggest that future research should explore spatially-aware architectures, such as gradient-based localization or attention mechanisms [27], or detection-based approaches, capable of localizing dirty regions within an image.

Expanding the training set with additional examples of partial dirtiness, glare, and high-contrast deposits should help reduce false negatives in these edge cases.

Overall, these findings indicate that the proposed ResNet-50-based model offers a robust approach to the task at hand and can be regarded as scalable within the scope of this study.

4. Environmental and Operational Impact Analysis

A central motivation for intelligent, condition-based cleaning is the reduction of resource consumption relative to fixed-schedule cleaning regimes, echoing broader efforts to quantify environmental trade-offs in coupled resource systems [26]. The reduction in resource utilization relative to a fixed-schedule approach is one of the key drivers behind condition-based intelligent cleaning. This section analyzes the environmental and operational advantages that can be derived from implementing the proposed threshold-based monitoring mechanism.

4.1 Water Consumption Reduction

Under a conventional scheduled-cleaning approach, a 1 MW utility-scale solar plant in a dusty environment typically undergoes 8 to 12 cleaning cycles per month, using approximately 3 to 5 liters of water per panel per cycle. Taking the typical panel count for 1 MW at around 3,000 panels, this equates to 108,000 to 180,000 liters of water consumed per MW per month, irrespective of whether individual panels actually require cleaning. Field observations indicate that, in a moderately soiled environment, only 40 to 60 percent of panels in a given zone actually require cleaning at any time [4], [5].

The proposed system initiates cleaning when P(dirty) ≥ 0.45 for an individual panel. Given a detection accuracy of approximately 77% and an estimated false-positive rate of around 23% for the threshold-based system, the system reliably identifies and cleans dirty panels while occasionally triggering cleaning for clean panels, but does not clean panels that show no evidence of soiling. Using the empirical class distribution observed in the test split of this dataset (clean panels = 101, dirty panels = 68, an approximately 60:40 ratio), the condition-based cleaning system is estimated to reduce unnecessary cleaning actions by 35% to 45% per month relative to a fixed monthly schedule, under simplified assumptions. At 4 liters per panel per cycle and 3,000 panels per MW, this corresponds to water use of 42,000 to 54,000 liters per MW per month – a saving of roughly 30% to 40% in cleaning water.

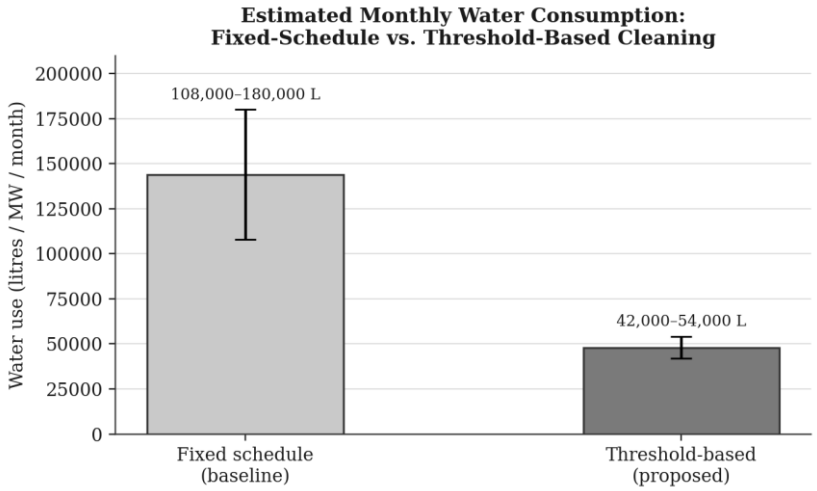


Fig. 10. Estimated monthly water consumption: fixed-schedule versus threshold-based cleaning.

4.2 Energy and Mechanical Savings

Every unnecessary deployment of a robotic cleaning unit consumes energy through motor actuation, water-pump operation, and onboard computing. For example, a self-propelled photovoltaic cleaning robot consumes 150 to 300 watts of electricity per hour of cleaning operation; avoiding deployment 35 to 45 percent of the time it would otherwise have been triggered yields an equivalent percentage saving in operational robot energy consumption. For a 10 MW installation with the robot deployed 12 times per month, this translates to estimated savings of 630 to 1,620 kWh of operational energy per month. Mechanical wear on brushes, tracks, and pump seals is likewise reduced in proportion to the number of cleaning cycles avoided.

This trade-off mirrors the resource-allocation problem addressed in constrained multi-objective food-energy-water systems, where competing consumption targets must be jointly optimized [8].

4.3 Recovery of PV Efficiency and CO2 Reductions

Beyond the direct cost savings of this maintenance approach, a secondary benefit of the conditional cleaning strategy is the prompt identification of panels that genuinely require cleaning before further degradation occurs. An efficiency loss of 10–25% due to soiling, if left unaddressed for several weeks, represents a significant loss of energy production [1]. For a 1 MW facility experiencing a 15% output efficiency loss, this translates to a potential loss of 108 MWh of energy production if soiling persists over a 30-day period at a 0.30 capacity factor. At a carbon intensity of 0.4 kg CO2/kWh, reflecting a diversified electricity grid, this corresponds to an additional burden of approximately 43 tonnes of CO2 emissions that would need to be offset by non-PV electricity generation.

5. Limitations

Several limitations of this study should be acknowledged. First, the dataset used in this study contains 300 images; while sufficient to demonstrate that the transfer learning approach works under low-data conditions, it does not capture the full range of contamination types, panel configurations, and environmental settings observed across different geographies. The performance of the system on data from regions with different dust compositions, humidity levels, and panel tilts is therefore unpredictable without domain-specific fine-tuning.

Second, the clean/dirty dichotomy adopted in this system, while operationally useful, does not differentiate between physically cleanable contaminants and surface damage such as delamination, glass cracking, or hotspots. Applying the system to panels with physical damage could trigger cleaning that fails to address, or even worsens, the damage. Future work could extend the framework to a three-class formulation that includes a damage category.

Third, the model was tested using static images rather than continuous streams from a moving robot's camera. The effects of temporal consistency, motion blur, and varying illumination across frames of a video stream in an actual robotic deployment may affect classification in ways that a static test protocol does not capture. A temporal filtering method is recommended prior to integration into a physical deployment.

Fourth, the model remains susceptible to reflective surfaces and excessive glare, as demonstrated in the failure case analysis in Section 3.6.1. Its performance under direct sunlight at small incidence angles has not yet been tested systematically.

6. Conclusion

This work proposes a practical deep learning approach for solar panel cleanliness assessment within an autonomous cleaning system. A binary image classification approach, based on transfer learning with the ResNet-50 architecture, was used to distinguish clean from dirty solar panels. Following careful handling of dataset preparation, threshold tuning, and class imbalance, the proposed model demonstrated consistent performance on both test data and real-world image samples. Experimental analysis shows that the model achieves a favorable

precision-recall trade-off, particularly for the detection of dirty panels, which is essential for avoiding missed cleaning events in robotic systems. Probability-based thresholding further improved decision reliability relative to argmax classification. This research differs from prior work in its emphasis on deployment readiness, computational efficiency, and integration with robotic cleaning mechanisms. The proposed method can serve as a perception module for autonomous solar panel cleaning systems. Future work will explore extending this framework to localize dirt regions within an image and to support real-time onboard computation.

References

[1] M. Mani and R. Pillai, "Impact of dust on solar photovoltaic (PV) performance: Research status, challenges and recommendations," *Renew. Sustain. Energy Rev.*, vol. 14, no. 9, pp. 3124–3131, 2010, doi: 10.1016/j.rser.2010.07.065.

[2] T. Sarver, A. Al-Qaraghuli, and L. L. Kazmerski, "A comprehensive review of the impact of dust on the use of solar energy: History, investigations, results, literature, and mitigation approaches," *Renew. Sustain. Energy Rev.*, vol. 22, pp. 698–733, 2013, doi: 10.1016/j.rser.2012.12.065.

[3] B. O. Olorunfemi, N. I. Nwulu, and O. A. Ogbolumani, "Solar panel surface dirt detection and removal based on Arduino color recognition," *MethodsX*, vol. 10, Art. no. 101967, 2023, doi: 10.1016/j.mex.2022.101967.

[4] R. Conceição, J. González-Aguilar, A. A. Merrouni, and M. Romero, "Soiling effect in solar energy conversion systems: A review," *Renew. Sustain. Energy Rev.*, vol. 162, Art. no. 112434, 2022, doi: 10.1016/j.rser.2022.112434.

[5] A. Younis and Y. M. Alhorr, "Modeling of dust soiling effects on solar photovoltaic performance: A review," *Sol. Energy*, vol. 220, pp. 1074–1088, 2021, doi: 10.1016/j.solener.2021.04.011.

[6] M. L. Adekanbi, E. S. Alaba, T. J. John, T. D. Tundealao, and T. I. Banji, "Soiling loss in solar systems: A review of its effect on solar energy efficiency and mitigation techniques," *Clean. Energy Syst.*, vol. 7, Art. no. 100094, 2024, doi: 10.1016/j.cles.2023.100094.

[7] O. A. Ogbolumani and N. Nwulu, "Integrated appliance scheduling and optimal sizing of an autonomous hybrid renewable energy system for agricultural food production," in *Advances in Manufacturing Engineering*, S. S. Emamian, F. Yusof, and M. Awang, Eds. Singapore: Springer, 2020, pp. 651–660, doi: 10.1007/978-981-15-5753-8_60.

[8] O. A. Ogbolumani and N. I. Nwulu, "Multi-objective optimisation of constrained food-energy-water-nexus systems for sustainable resource allocation," *Sustain. Energy Technol. Assess.*, vol. 44, Art. no. 100967, 2021, doi: 10.1016/j.seta.2020.100967.

[9] M. Aghaei, F. Grimaccia, C. A. Gonano, and S. Leva, "Innovative automated control system for PV fields inspection and remote control," *IEEE Trans. Ind. Electron.*, vol. 62, no. 11, pp. 7287–7296, 2015, doi: 10.1109/TIE.2015.2475235.

[10] A. Mellit, G. M. Tina, and S. A. Kalogirou, "Fault detection and diagnosis methods for photovoltaic systems: A review," *Renew. Sustain. Energy Rev.*, vol. 91, pp. 1–17, 2018, doi: 10.1016/j.rser.2018.03.062.

[11] C. Buerhop, D. Schlegel, M. Niess, C. Vodermayer, R. Weißmann, and C. J. Brabec, "Reliability of IR-imaging of PV-plants under operating conditions," *Sol. Energy Mater. Sol. Cells*, vol. 107, pp. 154–164, 2012, doi: 10.1016/j.solmat.2012.07.011.

[12] S. Deitsch, V. Christlein, S. Berger, C. Buerhop-Lutz, A. Maier, F. Gallwitz, and C. Riess, "Automatic classification of defective photovoltaic module cells in electroluminescence images," *Sol. Energy*, vol. 185, pp. 455–468, 2019, doi: 10.1016/j.solener.2019.02.067.

[13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015, doi: 10.1038/nature14539.

[14] T. Berghout, M. Benbouzid, T. Bentrcia, X. Ma, S. Djurović, and L.-H. Mouss, "Machine learning-based condition monitoring for PV systems: State of the art and future prospects," *Energies*, vol. 14, no. 19, Art. no. 6316, 2021, doi: 10.3390/en14196316.

[15] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural

networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.

[16] O. A. Ogbolumani and M. Adekoya, "Intelligent waste management optimization through machine learning analytics," *J. Sci. Res. Rev.*, vol. 2, no. 1, pp. 7–26, Jan. 2025, doi: 10.70882/josrar.2025.v2i1.25.

[17] B. O. Olorunfemi, O. A. Ogbolumani, and N. Nwulu, "Solar panels dirt monitoring and cleaning for performance improvement: A systematic review on smart systems," *Sustainability*, vol. 14, no. 17, Art. no. 10920, 2022, doi: 10.3390/su141710920.

[18] F. M. A. Mazen, R. A. A. Seoud, and Y. O. Shaker, "Deep learning for automatic defect detection in PV modules using electroluminescence images," *IEEE Access*, vol. 11, pp. 57783–57795, 2023, doi: 10.1109/ACCESS.2023.3284043.

[19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

[20] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Cham, Switzerland: Springer, 2016, pp. 630–645, doi: 10.1007/978-3-319-46493-0_38.

[21] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 25, 2012, pp. 1097–1105.

[22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, doi: 10.48550/arXiv.1409.1556.

[23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, doi: 10.48550/arXiv.1412.6980.

[24] F. Chollet, *Deep Learning with Python*. Shelter Island, NY, USA: Manning, 2018.

[25] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[26] O. A. Ogbolumani and N. I. Nwulu, "Environmental impact assessment for a meta-model-based food-energy-water-nexus system," *Energy Rep.*, vol. 11, pp. 218–232, 2024, doi: 10.1016/j.egyr.2023.11.033.

[27] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, 2020, doi: 10.1007/s11263-019-01228-7.