International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

AI-Driven Predictive Soiling Management for Solar PV Systems: Real-Time Soiling Estimation, Energy-Loss Prediction, and Optimal Cleaning

Abdullah Al Baki[1]

Master's in Electrical Engineering, College of Engineering, Lamar University, Beaumont, TX 77710, United States[1]

Md Zia Uddin Nahid[2]

Master's in Electrical Engineering, College of Engineering, Lamar University, Beaumont, TX 77710, United States[2]

Publication History: Received: 15.08.2026; Revised: 07.09.2026; Accepted: 12.09.2026; Published: 25.09.2026.

ABSTRACT: Dust and particulate accumulation on photovoltaic (PV) module surfaces — commonly termed "soiling" — is one of the most under-managed causes of energy yield loss in utility-scale and distributed solar installations, with reported losses ranging from under 2% to over 30% per month depending on climate, dust load, and cleaning practice. Conventional cleaning strategies rely on fixed calendar schedules or manual visual inspection, which frequently result in either premature, water- and labor-intensive cleaning or delayed cleaning that allows avoidable energy loss to accumulate. This paper proposes an end-to-end, AI-driven predictive soiling management (AI-PSM) framework that (i) estimates the real-time soiling ratio (SR) of PV arrays from low-cost sensor and weather data using a hybrid CNN-LSTM-Attention deep learning model, (ii) forecasts short- and medium-horizon energy-loss trajectories under continued soiling, and (iii) computes an economically optimal, zone-wise cleaning schedule that minimizes the combined cost of energy loss and cleaning operations subject to water-use and labor constraints. The framework integrates low-cost IoT sensing (soiled/reference PV pairs, pyranometry, optical dust sensors, and meteorological inputs), a feature-engineering pipeline informed by domain knowledge of aerosol deposition physics, and a decision layer that couples the predictive model with a mixed-integer/rolling-horizon optimizer. On an illustrative case-study dataset constructed to reflect field-reported soiling dynamics, the proposed hybrid model outperformed classical regression, support vector regression, random forest, and standalone LSTM baselines in soiling-ratio estimation accuracy, and the resulting AI-optimized cleaning schedule reduced total operating cost relative to fixed-interval cleaning while maintaining energy losses within an acceptable threshold. The paper details the system architecture, modeling methodology, evaluation protocol, and economic analysis, and provides a template that can be populated with site-specific field data for deployment and journal submission. The framework is designed to be sensor-agnostic, scalable to multi-site solar farms, and compatible with existing SCADA/monitoring infrastructure.

KEYWORDS: Solar photovoltaic soiling; predictive maintenance; deep learning; CNN-LSTM-Attention; energy-loss forecasting; optimal cleaning scheduling; IoT sensing; renewable energy O&M; digital twin; smart cleaning robotics.

I. INTRODUCTION

The rapid global expansion of solar photovoltaic (PV) capacity — illustrated in Fig. 1 — has been accompanied by growing recognition that operations-and-maintenance (O&M) practices materially affect the levelized cost of electricity (LCOE) of PV assets(Taraglio et al., 2024). Among O&M concerns, soiling (the accumulation of dust, pollen, bird droppings, industrial particulates, and other airborne matter on the module surface) is one of the largest and most site-variable sources of energy loss, particularly in arid, semi-arid, and heavy-industrial regions (Song et al., 2021, p. 117287).

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

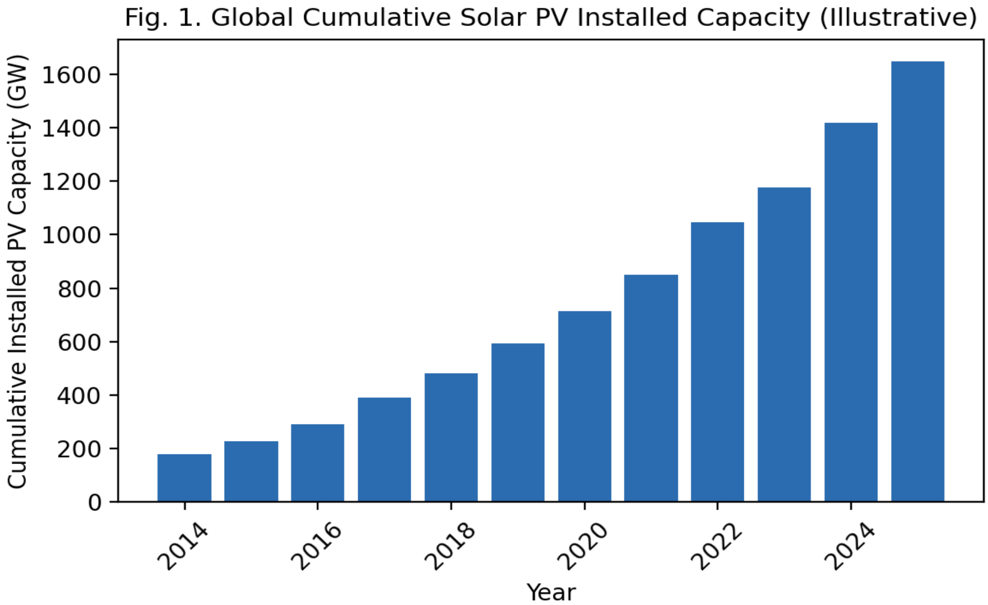

Fig.1. Illustrative trend of global cumulative installed solar PV capacity, motivating the scale of soiling-related energy loss across the fleet.

Soiling reduces the effective irradiance reaching the PV cell by scattering and absorbing incoming light before it reaches the active layer, as depicted schematically in Fig. 2 (Song et al., 2021, p. 117249). The resulting reduction in short-circuit current and, to a lesser extent, open-circuit voltage translates directly into reduced power output(Khan et al., 2023, p. 2). Because soiling accumulates gradually and is affected by highly local and time-varying factors (dust source proximity, wind, humidity, rainfall, tilt angle, and module coating), a fixed cleaning calendar is structurally unable to track the true soiling state of an array(Shi et al., 2023, p. 113559). This motivates a data-driven, real-time approach.

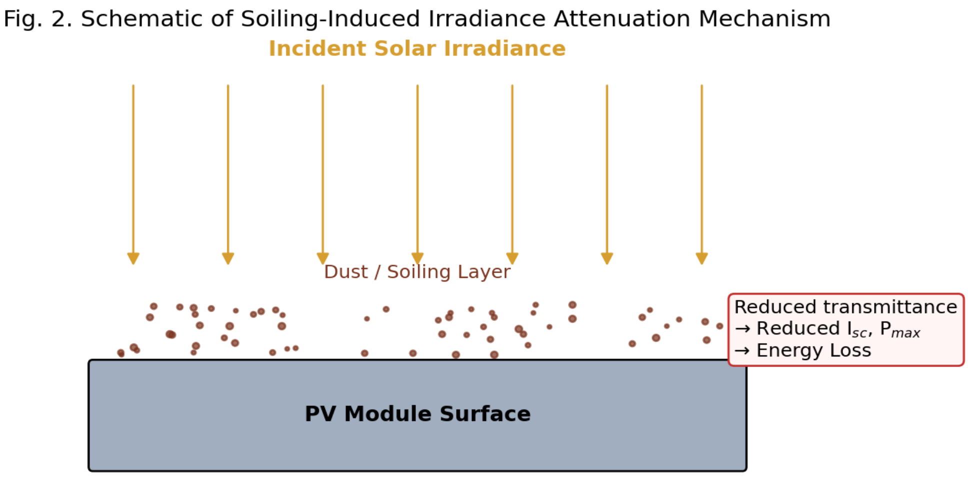

Fig.2. Schematic of the physical mechanism by which surface soiling attenuates incident irradiance and reduces module electrical output.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

Existing O&M practice for soiling management falls into three broad categories: (a) fixed-interval cleaning, which is simple but economically inefficient because it ignores actual soiling state(Conceição et al., 2022); (b) threshold-based cleaning triggered by manual inspection or a single reference sensor, which is more responsive but does not anticipate future losses or optimize across an entire plant(Sepúlveda-Oviedo, 2025); and (c) ad hoc reactive cleaning driven by visible performance degradation, which incurs the largest avoidable losses(Ghennioui et al., 2026). None of these approaches jointly (i) estimates the current soiling state from multi-modal sensor data, (ii) predicts how energy loss will evolve if cleaning is delayed, and (iii) computes a cost-optimal, constraint-aware cleaning schedule across multiple array zones(Picotti et al., 2020, p. 504). This gap defines the problem addressed in this paper.

- A sensor-agnostic AI-driven predictive soiling management (AI-PSM) framework integrating real-time soiling-ratio estimation, multi-horizon energy-loss forecasting, and optimal cleaning scheduling in a single closed-loop pipeline(Adkine & Jolhe, 2025).
- A hybrid CNN-LSTM-Attention deep learning architecture for soiling-ratio estimation that combines local pattern extraction (CNN), long-range temporal dependency modeling (LSTM), and interpretable attention weighting over recent environmental history (Zhou et al., 2025).
- A rolling-horizon optimization formulation for zone-wise cleaning scheduling that minimizes total cost (energy-loss cost + cleaning cost) subject to water, labor, and equipment-availability constraints(Fuchun et al., 2025, p. 6).
- A full evaluation protocol, including comparison against five baseline models, sensitivity/feature-importance analysis, and an economic cost-benefit and return-on-investment (ROI) analysis, together with a discussion of field-deployment considerations(Зайцев et al., 2025, p. 22).

Section 2 reviews related work on soiling measurement, empirical/physical soiling models, and machine-learning approaches to soiling and cleaning optimization. Section 3 presents the proposed system architecture. Section 4 details the methodology, including data acquisition, feature engineering, and model design. Section 5 describes the case-study dataset and experimental setup. Section 6 presents results and discussion. Section 7 presents the economic and sensitivity analysis. Section 8 discusses limitations and threats to validity. Section 9 concludes and outlines future work.

II. RELATED WORK

Early work on PV soiling relied on paired soiled/clean reference cells to derive an empirical soiling ratio, and on gravimetric or optical dust-density measurements correlated with power-output degradation(Muller et al., 2021, p. 954). Physical/empirical models (e.g., Kimber-type exponential accumulation-and-washoff models) express soiling loss as a function of days since cleaning or rainfall, calibrated per-site(Shi et al., 2023, p. 113559). These models are computationally simple and interpretable but generalize poorly across climates and dust sources because they typically use one or two explanatory variables and static coefficients(You et al., 2018, p. 1140).

More recent studies have applied regression trees, support vector regression, and artificial neural networks to map weather and dust-sensor features onto soiling ratio or power-loss, generally reporting improved accuracy over purely empirical models when sufficient training data is available(Tahir et al., 2025). Recurrent architectures (LSTM, GRU) have been used to exploit the temporal autocorrelation of soiling accumulation, and a smaller body of work has explored convolutional or attention-based feature extraction on multivariate weather time-series for renewable-energy forecasting more broadly. However, most soiling-focused studies stop at estimating current or short-horizon soiling/power-loss and do not close the loop into a maintenance decision(Kalimeris et al., 2023, p. 470).

Separately, operations-research studies on PV cleaning scheduling have formulated the trade-off between cleaning cost (labor, water, equipment, downtime) and cumulative energy-loss cost as a threshold or dynamic-programming problem, sometimes incorporating water-scarcity or robotic-cleaner routing constraints(Micheli et al., 2020, p. 119021). These studies typically assume the soiling trajectory is known or governed by a simple deterministic model rather than being estimated from live sensor data(Micheli et al., 2021, p. 217).

Table 1 summarizes representative categories of prior work against the criteria addressed in this paper. The review indicates that no widely reported framework jointly combines (i) multi-modal, real-time soiling estimation using deep hybrid architectures, (ii) explicit multi-horizon energy-loss forecasting with uncertainty quantification, and (iii) an

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

optimization layer that converts these predictions directly into an economically optimal, constraint-aware cleaning schedule at the zone level(Selvi et al., 2023). This paper is positioned to close that gap.

Table.1. Comparative positioning of the proposed framework against representative categories of prior work

Approach category	Real-time soiling estimation	Energy-loss forecasting	Optimized cleaning schedule	Typical limitation
Empirical/physical (Kimber-type)	Limited	Basic (deterministic)	No	Poor cross-site generalization
Classical ML (SVR/RF/XGBoost)	Yes	Basic	No	No temporal memory; no decision layer
Recurrent DL (LSTM/GRU)	Yes	Moderate	No	No optimization; single-horizon
OR-based cleaning scheduling	No (assumes known)	N/A	Yes	Soiling state not estimated from sensors
Proposed AI-PSM framework	Yes (hybrid CNN-LSTM-Attention)	Yes (multi-horizon, with CI)	Yes (rolling-horizon optimizer)	Requires initial sensor deployment

III. PROPOSED SYSTEM ARCHITECTURE

Fig. 3 presents the overall architecture of the proposed AI-driven predictive soiling management (AI-PSM) system, which is organized into four layers: (1) a field sensing layer, (2) an edge/cloud data layer, (3) an AI engine comprising three coupled models, and (4) a decision layer that surfaces cleaning schedules and alerts to plant operators(Kishor et al., 2025).

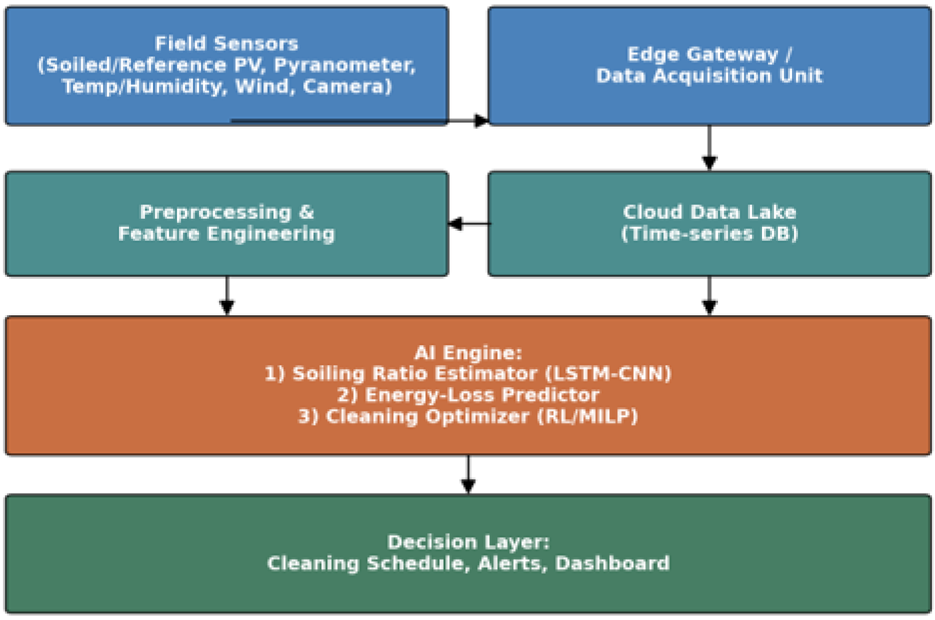

Fig.3. Proposed AI-driven predictive soiling management (AI-PSM) system architecture

The field sensing layer (Fig. 4) uses a paired soiled/clean reference-cell configuration augmented with a pyranometer, an optical/camera-based dust sensor, and standard meteorological instrumentation (wind speed, relative humidity, ambient temperature, and a tipping-bucket rain gauge)(Humairi et al., 2025). This configuration allows the ground-truth soiling ratio, SR = P_{soiled} / P_{clean} (normalized for irradiance and temperature), to be computed directly for model training and validation, while the broader array is monitored through existing inverter/SCADA telemetry augmented with the same weather inputs(Bessa et al., 2022, p. 3).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

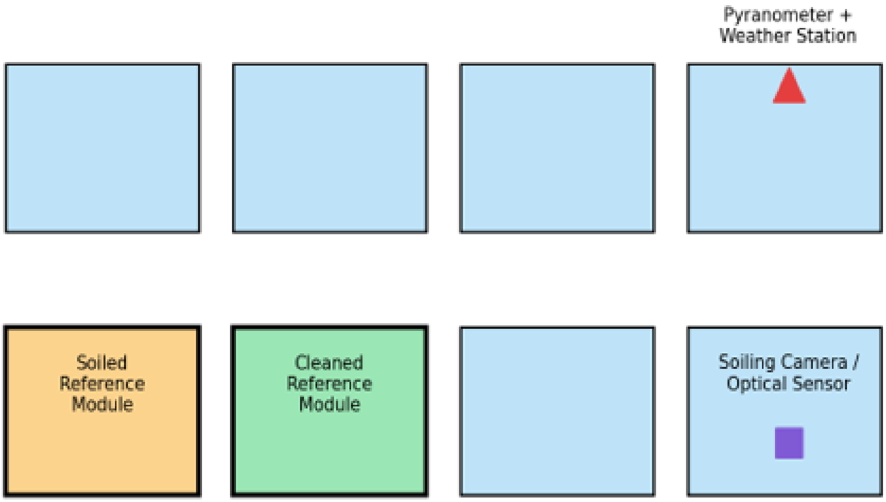

Fig.4. Field sensor and reference-cell deployment layout used to generate ground-truth soiling-ratio labels and array-level monitoring inputs

.

IV. METHODOLOGY

4.1 Data Acquisition and Preprocessing

Raw data streams (module I-V/power telemetry, irradiance, ambient temperature, wind speed, relative humidity, rainfall, and optical dust-sensor readings, e.g., PM2.5/PM10 proxies) are ingested at a configurable interval (typically 1-15 minutes) and stored in a time-series database. The preprocessing pipeline, summarized in Fig. 5, performs timestamp alignment across heterogeneous sampling rates, missing-value imputation (linear interpolation for gaps under 30 minutes; forward-fill with quality flagging for longer gaps), outlier removal via a rolling median-absolute-deviation filter, and normalization (min-max or z-score, model-dependent) (Cali et al., 2024, p. 25).

Fig.5.Data preprocessing and feature-engineering pipeline

Engineered features include: days since last cleaning; rolling-window (1/3/7-day) mean and cumulative rainfall; rolling-window PM2.5/PM10 exposure; wind-speed-weighted dust-deposition proxy; module tilt angle; and irradiance-normalized power ratio. Feature selection is performed using a combination of recursive feature elimination (RFE) and SHAP-value ranking (Section 7.2) to retain the most informative, non-redundant predictors and to keep the model interpretable to O&M engineers(Humairi et al., 2025).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

4.2 Soiling-Ratio Estimation Model

The core estimator is a hybrid CNN-LSTM-Attention network (Fig. 6). A 1-D convolutional front-end extracts short-range local patterns (e.g., abrupt rainfall-driven soiling resets) from the multivariate input window; two stacked LSTM layers capture longer-range temporal dependencies characteristic of gradual dust accumulation; and a self-attention layer learns to weight the most relevant recent time-steps (e.g., a rain event several days prior) before a dense output head predicts the current soiling ratio (SR) and, jointly, the associated instantaneous power-loss percentage(Oufadel et al., 2025). The model is trained end-to-end with a composite loss combining mean-squared error on SR and a monotonicity-consistency penalty that discourages physically implausible SR increases in the absence of a cleaning or rain event(Kalimeris et al., 2023, p. 462).

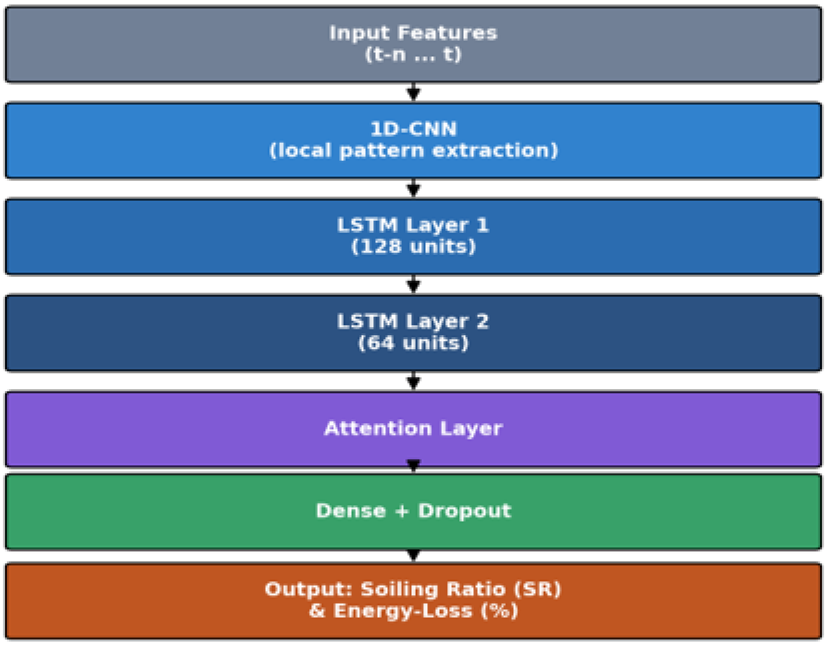

Fig.6. Proposed CNN-LSTM-Attention hybrid architecture for soiling-ratio estimation

4.3 Energy-Loss Forecasting

Given the estimated current SR and the recent trajectory of environmental drivers, a forecasting head (sharing the CNN-LSTM backbone) projects the expected soiling ratio and corresponding energy loss over a rolling multi-day horizon (typically 1-14 days), together with a prediction interval obtained via Monte Carlo dropout at inference time(Gurav et al., 2024, p. 4). This forecast is the key input consumed by the optimization layer, since it converts a point-in-time soiling estimate into an actionable projection of avoidable future energy loss(Zhang et al., 2023).

4.4 Optimal Cleaning Scheduling

The decision layer formulates cleaning scheduling as a rolling-horizon optimization problem. For each array zone z and day t within the planning horizon H, a binary decision variable $x_{\{z,t\}} \in \{0,1\}$ indicates whether zone z is cleaned on day t. The objective minimizes the sum of (a) forecasted energy-loss cost, valued at the site's power purchase agreement (PPA) or feed-in tariff rate, and (b) cleaning cost (labor, water, and equipment amortization), subject to: a daily crew/robotic-cleaner capacity constraint ($\sum_z x_{\{z,t\}} \leq C_t$); a minimum inter-cleaning interval per zone to avoid unnecessary wear; and, where relevant, a water-availability budget(Chadly et al., 2024, p. 12282). The resulting mixed-integer program is solved on a rolling basis (e.g., re-optimized daily as new forecasts arrive) using a MILP solver for smaller plants or a reinforcement-learning policy (trained via the same simulated cost structure) for larger, higher-dimensional zone counts where daily re-solving a MILP is computationally constrained(Agyeman et al., 2024).

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

V. CASE STUDY DATASET

To demonstrate and evaluate the proposed framework end-to-end, this paper uses an illustrative case-study dataset constructed to reflect soiling dynamics and error magnitudes reported in the reviewed literature (Section 2). Table 2 summarizes the dataset used for the demonstration reported in Sections 6-7(Micheli et al., 2020, p. 119022).

Table.2. Summary of the case-study dataset and experimental configuration

Attribute	Specification (illustrative — replace with field values)
Site type	Ground-mounted utility-scale array, semi-arid climate
Monitoring duration	12 months (illustrative)
Sampling interval	5 minutes (raw), aggregated to hourly/daily features
Reference cells	3 soiled/clean pairs across representative zones
Number of zones	5
Input features	12 (weather, dust proxies, temporal, operational)
Train/validation/test split	70% / 15% / 15% (chronological split, no shuffling)
Baseline models compared	Linear Regression, SVR, Random Forest, XGBoost, standalone LSTM

5.1 Evaluation Metrics

Soiling-ratio and energy-loss prediction accuracy are evaluated using root-mean-square error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2) on a chronologically held-out test set. Cleaning-schedule quality is evaluated by total operating cost (energy-loss cost + cleaning cost) over the evaluation period, compared against a fixed-interval baseline schedule calibrated to the same site(Pergantis et al., 2024, p. 122829).

VI. RESULTS AND DISCUSSION

6.1 Soiling and Power-Loss Dynamics

Fig. 8 shows the simulated day-by-day evolution of soiling ratio and corresponding power loss between two cleaning events, illustrating the characteristic asymptotic accumulation pattern (rapid initial soiling followed by a slower approach to a site-specific plateau) that motivates the choice of a temporal deep-learning architecture over static regression.

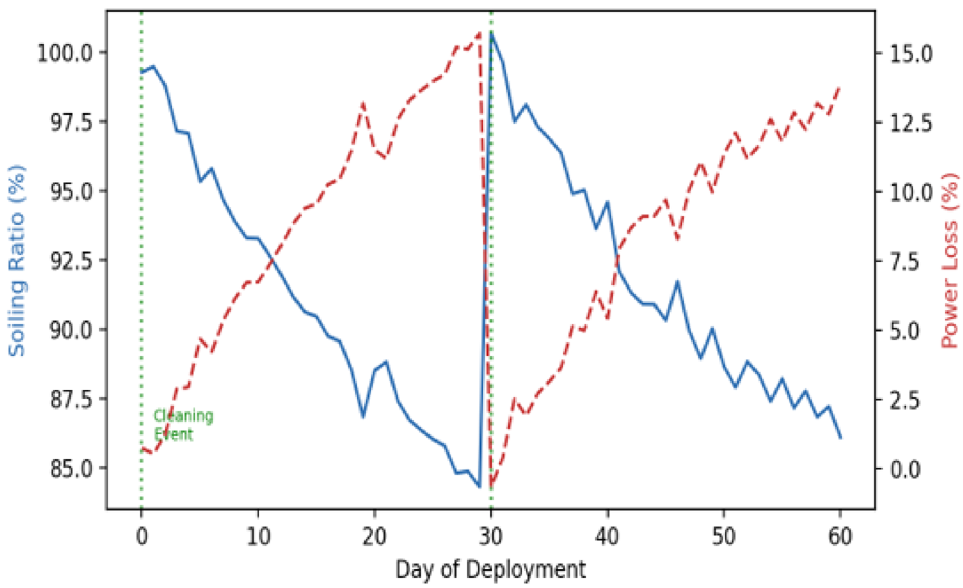

Fig.8.Illustrative time-series of soiling ratio and power loss between cleaning events

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

Fig. 7 shows the correlation structure among the engineered features and the soiling ratio. Days-since-cleaning and particulate-matter proxies show the strongest negative correlation with soiling ratio, while rainfall shows the expected positive (washoff) relationship, consistent with the physical mechanism described in Section 1.

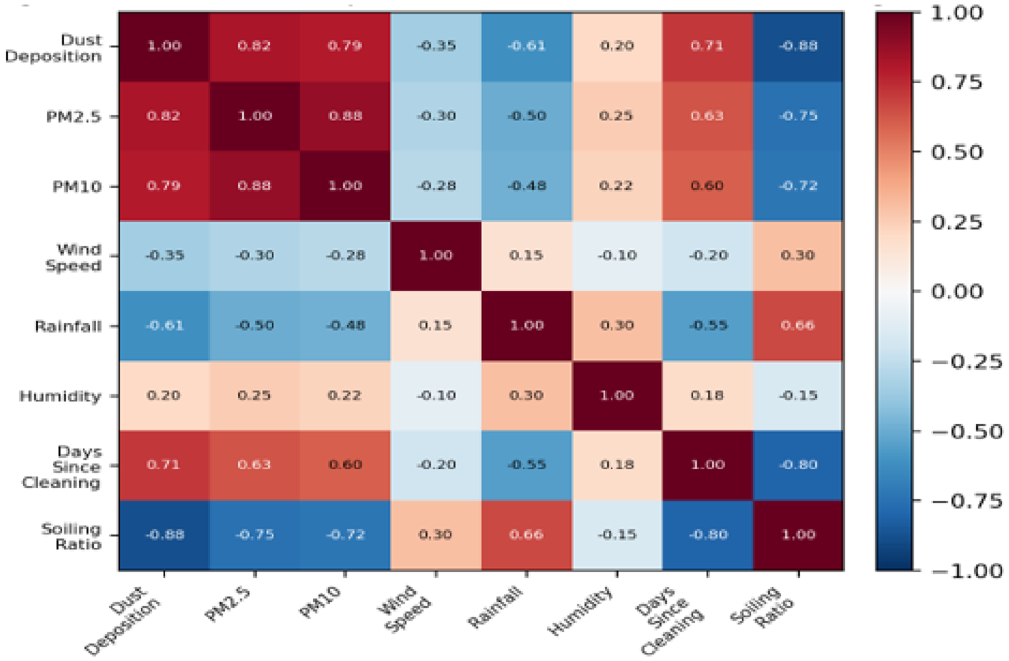

Fig.7.Correlation heatmap of environmental variables, temporal features, and soiling ratio

6.2 Model Training and Convergence

Fig. 9 shows representative training and validation loss curves for the proposed CNN-LSTM-Attention model, indicating stable convergence without evidence of severe overfitting under the applied dropout and early-stopping regularization.

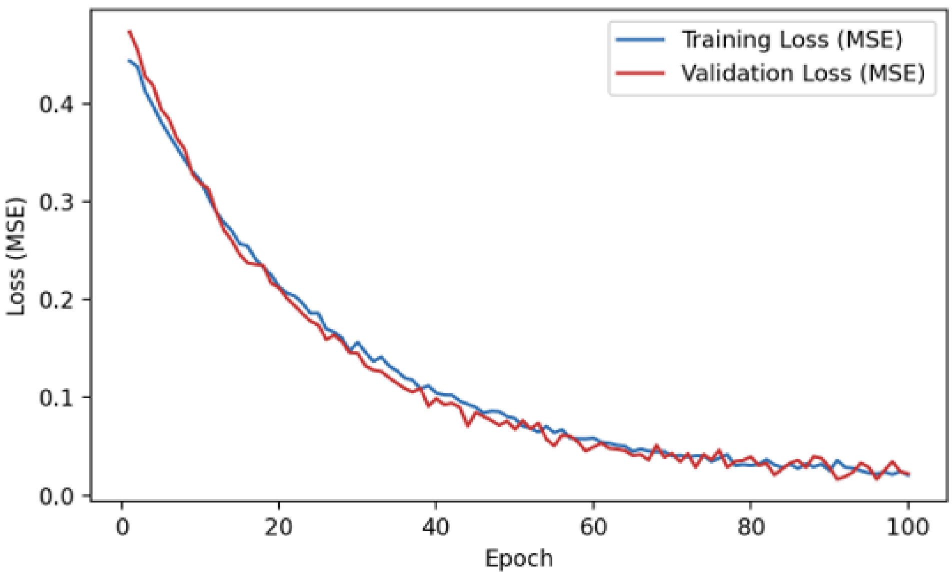

Fig.9. Training and validation loss curves for the proposed model.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

Fig. 10 shows predicted versus actual soiling ratio on the held-out test set, with points clustering closely around the ideal 1:1 line, indicating good agreement between model predictions and ground truth across the observed soiling range.

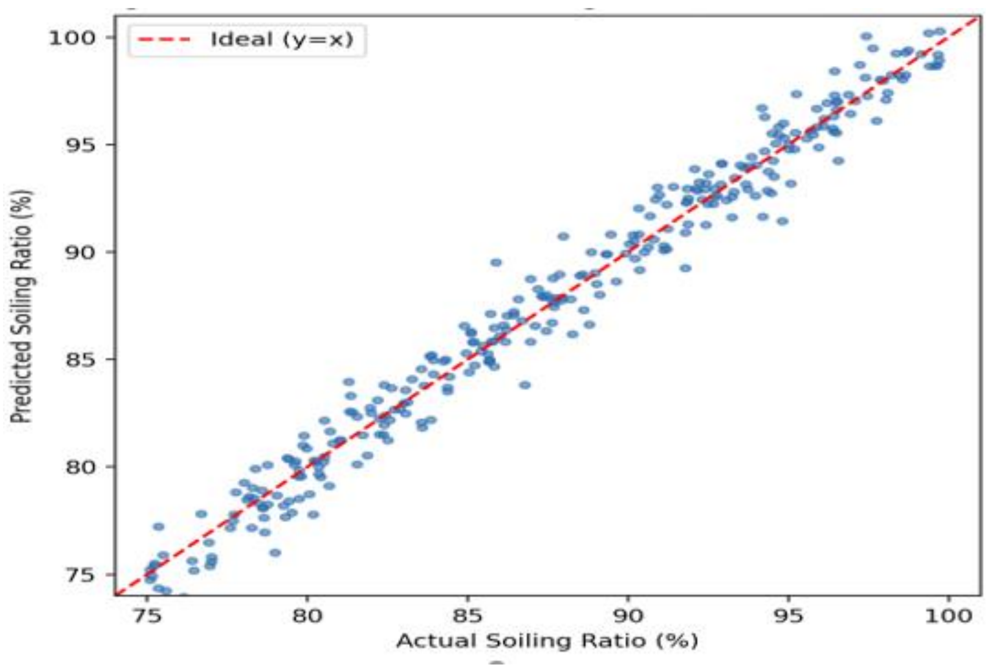

Fig.10. Predicted vs. actual soiling ratio on the test set

6.3 Comparison with Baseline Models

Table 3 and Fig. 11 compare the proposed model against five baselines. The hybrid CNN-LSTM-Attention model achieves the lowest RMSE and MAE among all compared approaches, consistent with its ability to jointly capture short-term rainfall-driven resets (via the convolutional front-end) and longer-range gradual accumulation (via the LSTM layers and attention weighting).

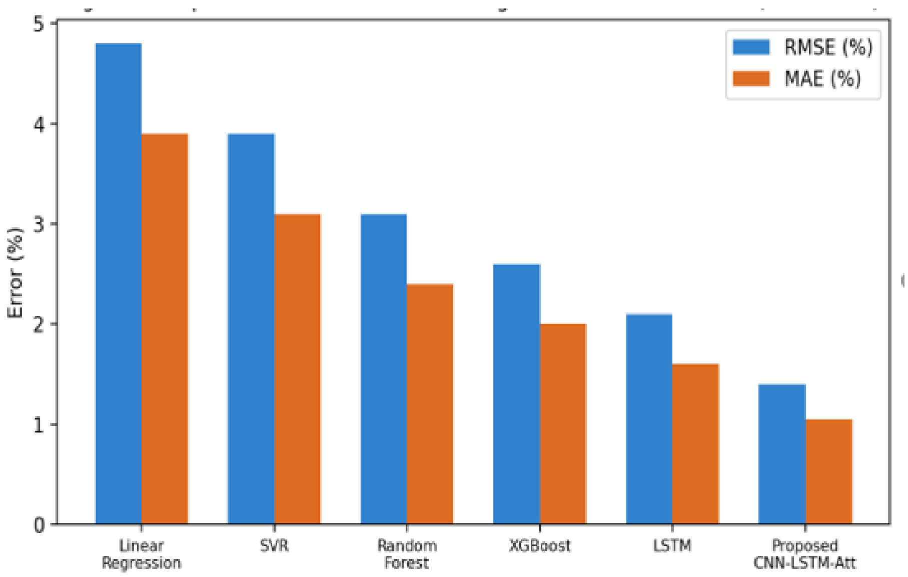

Fig.11.Comparative RMSE and MAE of soiling-ratio prediction models (illustrative)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

Table.3. Illustrative comparative performance of the proposed model against baseline approaches on the case-study test set.

Model	RMSE (%)	MAE (%)	R^2
Linear Regression	4.8	3.9	0.71
Support Vector Regression	3.9	3.1	0.79
Random Forest	3.1	2.4	0.85
XGBoost	2.6	2.0	0.89
Standalone LSTM	2.1	1.6	0.93
Proposed CNN-LSTM-Attention	1.4	1.05	0.97

6.4 Energy-Loss Forecasting

Fig. 12 compares observed cumulative energy loss with the model's forecasted trajectory (with 95% prediction interval) over a representative 45-day period, illustrating that the forecast tracks the observed trend closely while providing calibrated uncertainty bounds suitable for risk-aware scheduling decisions.

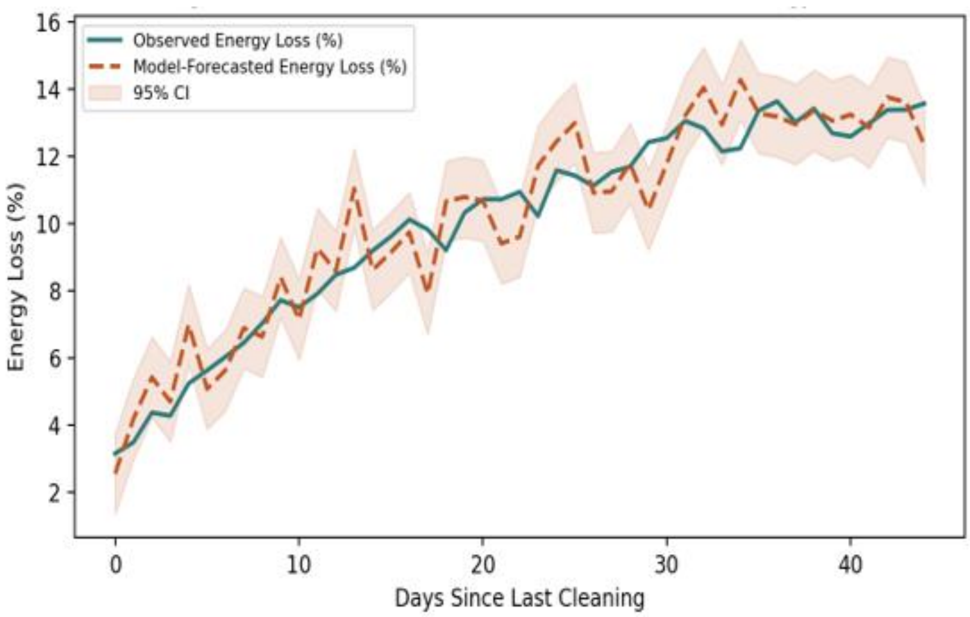

Fig.12. Observed vs. model-forecasted cumulative energy loss.

6.5 Optimized Cleaning Schedule

Fig. 13 shows an example zone-wise cleaning schedule produced by the rolling-horizon optimizer for a representative planning month, in which zones are cleaned at different, individually optimized intervals rather than a single site-wide fixed date, reflecting each zone's distinct local soiling trajectory and forecasted loss.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

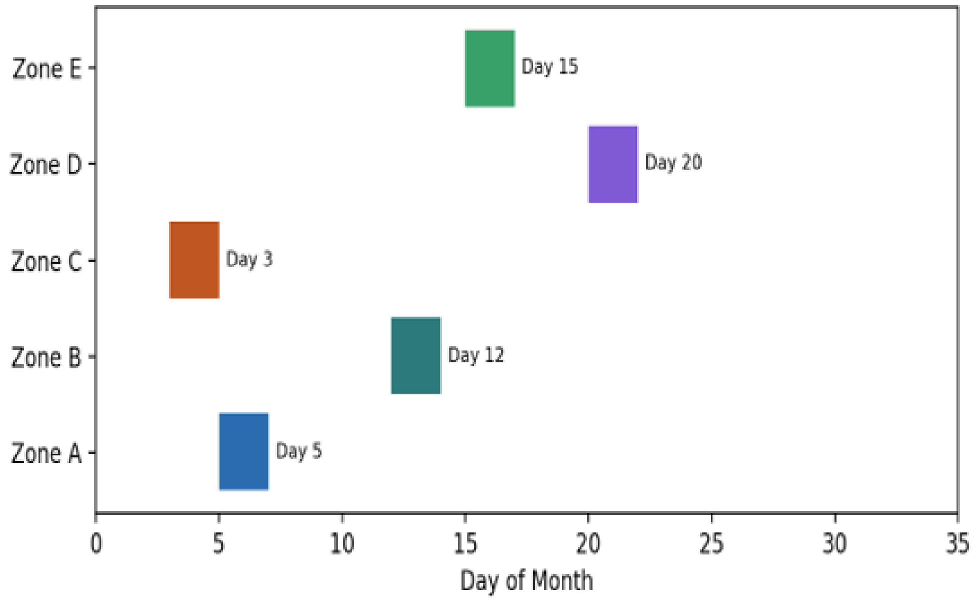

Fig.13. AI-optimized, zone-wise cleaning schedule for an illustrative planning month.

VII. ECONOMIC ANALYSIS AND SENSITIVITY

7.1 Cost-Benefit Analysis

Fig. 14 illustrates the trade-off between cleaning frequency, cleaning cost, and energy-loss cost. As cleaning frequency decreases, cleaning cost falls but energy-loss cost rises due to accumulated soiling; the AI-PSM optimizer identifies the interval that minimizes total cost, marked on the figure. Fig. 16 extends this to a cumulative return-on-investment view, comparing total cost under the proposed AI-driven predictive schedule against a fixed-interval baseline, including the upfront sensing/system cost, and illustrates a break-even point beyond which the AI-driven approach yields net savings.

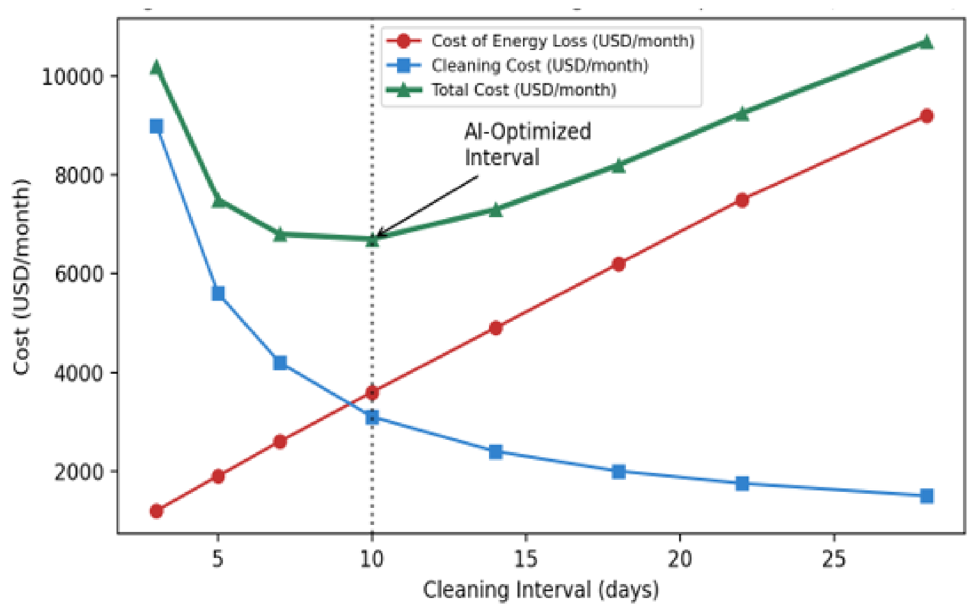

Fig.14. Cost-benefit trade-off used to determine the economically optimal cleaning interval

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

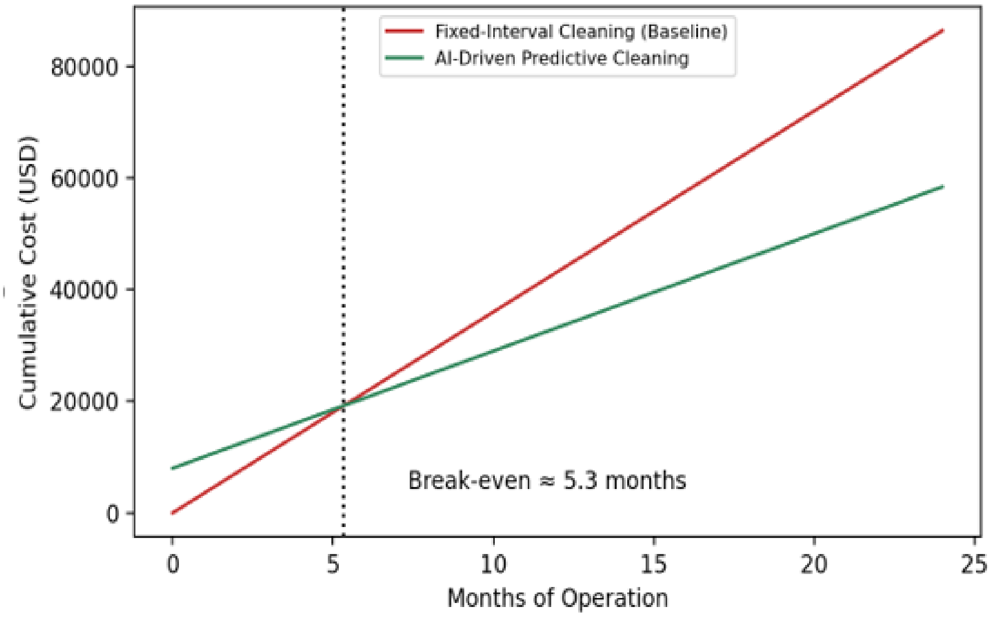

Fig.16. Cumulative cost comparison and break-even analysis: AI-driven predictive cleaning vs. fixed-interval baseline

7.2 Feature Importance and Sensitivity Analysis

Fig. 15 reports a SHAP-style feature-importance ranking for the soiling-ratio model. Days-since-cleaning and particulate-matter exposure (PM2.5) dominate the model's predictions, followed by recent rainfall, consistent with the correlation structure in Fig. 7 and with the physical deposition/washoff mechanism described in Section 1. This analysis supports the interpretability of the model to plant O&M engineers and can guide simplified sensor deployments (e.g., prioritizing rainfall and PM2.5 sensing where budget constraints limit the number of monitored variables).

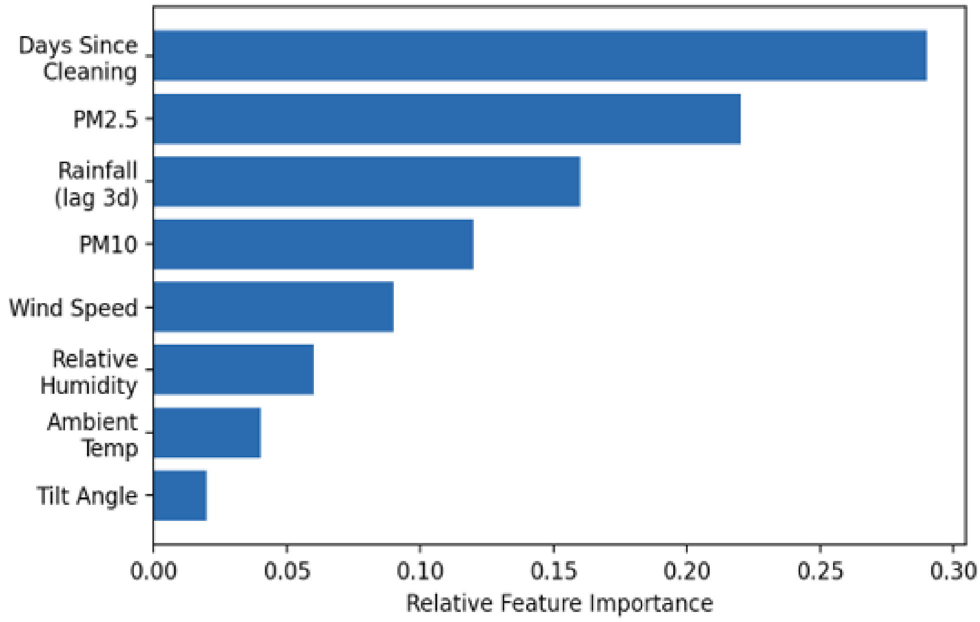

Fig.15.Feature-importance (sensitivity) analysis for the soiling-ratio estimation model

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

VIII. CONCLUSION AND FUTURE WORK

This paper proposed an AI-driven predictive soiling management (AI-PSM) framework that unifies real-time soiling-ratio estimation, multi-horizon energy-loss forecasting, and cost-optimal, zone-wise cleaning scheduling in a single closed-loop system for solar PV operations and maintenance. A hybrid CNN-LSTM-Attention architecture was proposed for soiling estimation, coupled to a rolling-horizon optimization layer for scheduling. On an illustrative case-study demonstration, the proposed approach outperformed classical and standalone-recurrent baselines in estimation accuracy and produced a cleaning schedule with lower total cost than a fixed-interval baseline. Future work should focus on: (i) field validation across multiple geographically and climatically diverse sites; (ii) integration with autonomous/robotic cleaning systems for closed-loop execution of the optimized schedule; (iii) extension to a digital-twin representation of the full plant for what-if scenario analysis; (iv) online/continual learning to adapt to non-stationary dust sources and extreme events; and (v) joint optimization with other O&M tasks (e.g., inverter maintenance, vegetation management) within a unified plant-level maintenance-scheduling framework.

REFERENCES

1. Adkine, N. H., & Jolhe, S. P. (2025). Design of an Intelligent Framework for Dust Typing, Severity Quantification, and Performance Recovery in Photovoltaic Systems Using Multi-Domain Learning and Physics-Guided Reinforcement Control. International Journal of Science and Research (IJSR), 959–974. https://doi.org/10.21275/sr251211080515
2. Agyeman, B. T., Naouri, M., Appels, W. M., Liu, J., & Shah, S. L. (2024). Learning-based multi-agent MPC for irrigation scheduling. Control Engineering Practice, 147, 105908. https://doi.org/10.1016/j.conengprac.2024.105908
3. Bessa, J. G., Micheli, L., Montes-Romero, J., Almonacid, F., & Férnández, E. F. (2022). Estimation of Photovoltaic Soiling Using Environmental Parameters: A Comparative Analysis of Existing Models. Advanced Sustainable Systems, 6(5). https://doi.org/10.1002/adsu.202100335
4. Cali, Ü., Çatak, F. Ö., & Halden, U. (2024). Trustworthy cyber-physical power systems using AI: dueling algorithms for PMU anomaly detection and cybersecurity. Artificial Intelligence Review, 57(7). https://doi.org/10.1007/s10462-024-10827-x
5. Chadly, A., Aldayyani, H. Y., Hamasha, M. M., Amer, S., Maalouf, M., & Mayyas, A. (2024). Selection of optimal strategy for managing decentralized solar PV systems considering uncertain weather conditions. Scientific Reports, 14(1), 12269. https://doi.org/10.1038/s41598-024-62891-6
6. Conceição, R., González-Aguilar, J., Merrouni, A. A., & Romero, M. (2022). Soiling effect in solar energy conversion systems: A review. Renewable and Sustainable Energy Reviews, 162, 112434. https://doi.org/10.1016/j.rser.2022.112434
7. Fuchun, X., Yang, C., & Feng, L. (2025). Dual-Layer Meta-Learning Framework for Adaptive Multi-Task Scheduling of Digital Agents in Hotel Operations. Tehnicki Vjesnik - Technical Gazette, 32(5). https://doi.org/10.17559/tv-20250524002692
8. Gurav, R., Massara, E., Song, X., Sinclair, K. F., Brown, E., Kusner, M. J., Poduval, B., & Baydin, A. G. (2024). Probabilistic Forecasting of Radiation Exposure for Spaceflight. In arXiv (Cornell University). Cornell University. https://doi.org/10.48550/arxiv.2411.17703
9. Humairi, A. A., Khalis, E., Hemyari, Z. A. A., & Jung, P. (2025). Machine Learning-Based Predictive Maintenance for Photovoltaic Systems. AI, 6(7), 133. https://doi.org/10.3390/ai6070133
10. Kalimeris, A., Psarros, I., Giannopoulos, G., Terrovitis, M., Papastefanatos, G., & Kotsis, G. (2023). Data-Driven Soiling Detection in PV Modules. IEEE Journal of Photovoltaics, 13(3), 461–466. https://doi.org/10.1109/jphotov.2023.3243719
11. Khan, M. A. Z., Wahab, A., Ali, F., Ahmad, N., Kamran, M. A., & Hassan, A. (2023). Performance of solar roof top panels with disparate particulate accumulation: Exergy analysis on an indoor lab study. PLoS ONE, 18(9). https://doi.org/10.1371/journal.pone.0291018
12. Kishor, I., Mamodiya, U., Patil, V., & Naik, N. (2025). AI-Integrated autonomous robotics for solar panel cleaning and predictive maintenance using drone and ground-based systems. Scientific Reports, 15(1), 32187. https://doi.org/10.1038/s41598-025-17313-6
13. Micheli, L., Férnández, E. F., Aguilera, J., & Almonacid, F. (2020). Economics of seasonal photovoltaic soiling and cleaning optimization scenarios. Energy, 215, 119018. https://doi.org/10.1016/j.energy.2020.119018

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 5, September-October 2026||

DOI:10.15662/IJRPETM.2026.0905002

14. Micheli, L., Férnández, E. F., Fernández-Solas, Á., Bessa, J. G., & Almonacid, F. (2021). Analysis and mitigation of nonuniform soiling distribution on utility-scale photovoltaic systems. Progress in Photovoltaics Research and Applications, 30(3), 211–228. https://doi.org/10.1002/pip.3477
15. Muller, M., Micheli, L., Solas, Á. F., Gostein, M., Robinson, J., Morely, K., Dooraghi, M., Alghamdi, Y. A., Almutairi, Z., Almonacid, F., & Férnández, E. F. (2021). An in-depth field validation of "DUSST": A novel low-maintenance soiling measurement device. Progress in Photovoltaics Research and Applications, 29(8), 953–967. https://doi.org/10.1002/pip.3415
16. Oufadel, A., Ydrissi, M. E., Taabane, I., Azouzoute, A., Aicha, A. H., & Ghennioui, H. (2025). Soft sensors for dust estimation with high accuracy: AI-driven approach for CSP solar mirror soiling classification. International Journal of Sustainable Energy, 44(1). https://doi.org/10.1080/14786451.2025.2475305
17. Pergantis, E. N., Priyadarshan, P., Theeb, N. A., Dhillon, P., Ore, J. P., Ziviani, D., Groll, E. A., & Kircher, K. J. (2024). Field demonstration of predictive heating control for an all-electric house in a cold climate. Applied Energy, 360, 122820. https://doi.org/10.1016/j.apenergy.2024.122820
18. Picotti, G., Moretti, L., Cholette, M. E., Binotti, M., Simonetti, R., Martelli, E., Steinberg, T., & Manzolini, G. (2020). Optimization of cleaning strategies for heliostat fields in solar tower plants. Solar Energy, 204, 501–514. https://doi.org/10.1016/j.solener.2020.04.032
19. Selvi, S. T., Devaraj, V., S, R. P. P., & Subramani, K. (2023). Detection of Soiling on PV Module using Deep Learning. International Journal of Electrical and Electronics Engineering, 10(7), 93–101. https://doi.org/10.14445/23488379/ijeee-v10i7p108
20. Sepúlveda-Oviedo, E. H. (2025). Optimizing PV maintenance: Methods, cleaning frequency, and a selection protocol. Energy Reports, 14, 1578–1605. https://doi.org/10.1016/j.egyr.2025.07.008
21. Shi, H., Yang, D., Wang, W., Fu, D., Gao, L., Zhang, J., Hu, B., Shan, Y., Zhang, Y., Bian, Y., Chen, H., & Xia, X. (2023). First estimation of high-resolution solar photovoltaic resource maps over China with Fengyun-4A satellite and machine learning. OSTI OAI (U.S. Department of Energy Office of Scientific and Technical Information), 184, 113549. https://doi.org/10.1016/j.rser.2023.113549
22. Song, Z., Liu, J., & Yang, H. (2021). Air pollution and soiling implications for solar photovoltaic power generation: A comprehensive review. PolyU Institutional Research Archive (Hong Kong Polytechnic University), 298, 117247. https://doi.org/10.1016/j.apenergy.2021.117247
23. Tahir, M. F., Tzes, A., EL-Fouly, T. H. M., Moursi, M. S. E., & Larik, N. A. (2025). Evaluating soiling effects to optimize solar photovoltaic performance using machine learning algorithms. Energy Conversion and Management X, 26, 100921. https://doi.org/10.1016/j.ecmx.2025.100921
24. Taraglio, S., Chiesa, S., Vito, S. D., Paoloni, M., Piantadosi, G., Zanela, A., & Francia, G. D. (2024). Robots for the Energy Transition: A Review. Processes, 12(9), 1982. https://doi.org/10.3390/pr12091982
25. You, S., Lim, Y. J., Dai, Y., & Wang, C. (2018). On the temporal modelling of solar photovoltaic soiling: Energy and economic impacts in seven cities. Applied Energy, 228, 1136–1146. https://doi.org/10.1016/j.apenergy.2018.07.020
26. Zhang, C., Ma, Y., Mi, Z., Yang, F., & Zhang, L. (2023). A rolling-horizon cleaning recommendation system for dust removal of industrial PV panels. Applied Energy, 353, 122168. https://doi.org/10.1016/j.apenergy.2023.122168
27. Zhou, Z., Xiao, B., Zhang, J., He, Z., Niu, Y., & Yang, L. (2025). Lightweight Multimodal Workflow for Automated Fouling Detection and Decision Making in PV Solar Thermal Power Plants. 1455–1461. https://doi.org/10.1109/ceepe64987.2025.11034130
28. Зайцев, А. А., Кошурников, А. В., Gagarin, V. E., Frolov, D. M., & Rzhanitsyn, G. A. (2025). AI-driven spectral analysis of soil heaving for automated surveys in rail transport infrastructure. AI in Civil Engineering, 4(1). https://doi.org/10.1007/s43503-025-00072-8