Journal of Operation and Automation in Power Engineering

Vol. 15, No. 1, Jan. 2027, Pages: 83-100

http://joape.uma.ac.ir

Research Paper

Hybrid AI-Based Multimodal Soiling-State Estimation and Revenue-Aware Cleaning Optimization for PV Plant Maintenance

Hussain Ali Mutar [1], *Ahmed S. Al-Akayshee* [2], *Ibtihal Razaq Niama Al-Rubeei* [3, *], and *Haider TH. Salim ALRikabi* [3]

[1]*Computer Department, College of Computer Science and Information Technology, Wasit University, Wasit, Iraq.*
[2]*Department of Medical Instrumentation Techniques Engineering, College of Technical Engineering, The Islamic University of Najaf, Iraq.*
[2]*Jabir Ibn Hayyan University for Medical and Pharmaceutical Sciences, Najaf, Iraq.*
[3]*Electrical Engineering Department, College of Engineering, Wasit University, Wasit, Iraq.*

Abstract— *Accumulation of soiling on PV plant surfaces leads to considerable energy losses, especially in arid and semi-arid regions where soiling variability and precipitation affect the cleaning strategies. The present research proposes a hybrid and combined approach that integrates artificial intelligence designed to link the multimodal estimation of soiling state with revenue-grounded and uncertain cleaning optimization. Temporal parameters (SCADA and meteorological variables) are merged with sparse image-based indicators via a TCN architecture and XGBoost, passing the uncertainty of the predictions made by the model to the scheduler created by Mixed-Integer Linear Programming. The scheduler takes into account anticipated energy recovery, electricity prices, precipitation, cleaning expenditures, availability of labor and water, and feasibility of operations. The performance of the developed framework was tested for a full-year horizon of 8760 hours using four types of simulated PV plants, 20 Monte Carlo seeds for testing, as well as a site-specific and independent testing set of data. The estimator delivered the corresponding MAE value of 0.89%, ROC-AUC index of 0.983, and F1-score of 0.949. Operationally, the proposed strategy achieved an annual net value of approximately USD 3.09 million, compared to USD 3.01 million for the baseline model (LSTM + rules), while reducing the number of cleaning events from 66 to 34. Additionally, it reduced annual water consumption to approximately 10.8 $\times 10^3 m^3$, compared to 18.4 $\times 10^3 m^3$ for a fixed 14-day cleaning policy. These results indicate that the framework improves the timing and resource-use efficiency of cleaning decisions, while providing a unified pathway from heterogeneous photovoltaic observations to economic planning.*

Keywords—Photovoltaic soiling, PV cleaning optimization, hybrid artificial intelligence, multimodal data fusion, temporal convolutional network, XGBoost, uncertainty-aware prediction, predictive maintenance, revenue-aware scheduling.

NOMENCLATURE

Abbreviations

AI	Artificial Intelligence
F1	F1-Score
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MILP	Mixed-Integer Linear Programming
O&M	Operation and Maintenance
PV	Photovoltaic
RF	Random Forest
RMSE	Root Mean Squared Error
ROC-AUC	Receiver Operating Characteristic–Area Under the Curve
SCADA	Supervisory Control and Data Acquisition
SHAP	Shapley Additive Explanations
SR	Soiling Ratio
TCN	Temporal Convolutional Network
UAV	Unmanned Aerial Vehicle
XGBoost	Extreme Gradient Boosting

Parameters

λ	Uncertainty penalty coefficient
C_i^{clean}	Cleaning cost per operation
H	Number of scheduling intervals
N	Number of PV array groups
Δ_t	Decision interval
η_i	Cleaning effectiveness for array group
π_t	Electricity selling price
L_i	Labor requirement for one cleaning cycle
R_t	Rainfall intensity
R_{th}	Rainfall threshold for natural cleaning
W_i	Water required for one cleaning cycle

Set and Indices

A	Set of independently schedulable PV array groups
i	Index of the PV array group
T	Set of discrete scheduling intervals
t	Time index

Variables

$E_{i,t}^{rec}$	Recoverable energy after cleaning

Received: 23 Aug. 2026
Revised: 10 Oct. 2026
Accepted: 16 Oct. 2026
*Corresponding author:
E-mail: ibtihal.razaq@uowasit.edu.iq (I.R.N. Al-Rubeei)
DOI: 10.22098/joape.2026.20691.2606

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.
Copyright © 2027 University of Mohaghegh Ardabili.

$P_{i,t}^{clean}$ Expected clean-array power output
$P_{i,t}^{soil}$ Power output under soiling conditions
$x_{i,t}^{SCADA}$ SCADA measurement vector
$x_{i,t}^{Vision}$ Image-feature vector
$x_{i,t}^{Weather}$ Meteorological feature vector
$\hat{s}_{i,t}$ Estimated soiling state
$\sigma_{i,t}$ Prediction uncertainty (standard deviation)
$s_{i,t}$ True normalized soiling state
$u_{i,t}$ Cleaning decision variable
$x_{i,t}$ Multimodal observation vector

1. INTRODUCTION

Photovoltaic (PV) plants are increasingly deployed as a major source of low-carbon electricity, making reliable and economically efficient operation and maintenance essential for sustaining their energy yield. Surface soiling is considered one of the major problems that negatively affect the efficiency of PV systems at different stages. This is due to the fact that soiling builds up and diminishes depending on the climatic conditions and therefore influences the amount of electricity generated. Efficient soiling management involves proper evaluation of plant condition as well as prompt decision-making regarding maintenance solutions while respecting operating and financial constraints.

1.1. Motivation and background

Surface soiling acts as an unyielding factor that disrupts the overall performance of photovoltaic (PV) systems, especially in areas characterized as arid and semi-arid where airborne particles are generated rapidly, and there are limited opportunities for natural washing due to rainwater. Accumulation of dust and particles reduces the optical performance of the module surface and subsequently lowers the output of PV systems. Studies done in the field show that the intensity of the soiling effect varies significantly from place to place; this difference is mostly due to local climate, conditions in the atmosphere, and even the type and nature of particles in the air [1]. Moreover, the presence of various environmental conditions, atmospheric factors, characteristics of particles, and the state of humidity and winds affect the processes of deposition and removal of pollutants from PV modules [2]. These factors have stimulated considerable research on the frequency of cleaning photovoltaic panels. Recent research has addressed the issue of pollution buildup in PV systems and the price of electricity and expenses paid for cleaning, in order to understand how often cleaning should occur [3]. It has also been recommended to use numerical approaches in order to compare the results of different frequencies of cleaning PV panels and find the economic result [4]. In addition to obtaining energy, cleaning processes take into account the expenses connected with operation of PV plants, including the expense for water and labor [5, 6]. At the same time, although operating data on PV systems is available more widely nowadays, effective monitoring and forecasting systems are becoming more prevalent due to the growth of artificial intelligence technologies [7, 8]. Data-driven approaches are one of the most useful techniques that allow the understanding of the loss of electrical energy from PV systems [9]. Picture analysis is another important direction for working with soiling and consists of various methods that help identify visible pollution [10, 11]. Picture analysis can also be applied in combination with electricity measurement to find the relationship between the state of the surface and potential damage to the energy output of PV systems [12].

Thus, the research on the soiling process has developed into other important areas that can be characterized as visual inspection-oriented, providing a foundation for advanced management of surface soiling in PV systems.

1.2. Literature review

Based on these studies, the literature on PV soiling management can be approached from a few similarly related points of view. To date, most research works have focused on characterization and monitoring of soiling, on the application of AI for detection and forecasting, and finally on the optimization of cleaning processes. Taking these areas of research as interrelated will help in gaining a better perspective on how the existing solutions deal with the physical, predictive, and operational aspects of PV soiling management. Recent studies have strengthened the use of machine learning for quantitative soiling estimation. An optimized machine-learning model based on LSTM was developed to detect, quantify, and predict PV soiling and associated electrical degradation [9]. The model linked accumulated dust mass to optical and power losses and further considered maintenance recommendations and cleaning-frequency optimization. Experimental research has also investigated the effect of different cleaning frequencies using environmental measurements and several machine-learning models, including a stacking ensemble [13]. The results demonstrated strong predictive performance for soiling loss under the tested conditions while illustrating the relationship between cleaning frequency and PV performance. These studies demonstrate the increasing capability of data-driven models to estimate soiling severity from environmental and operational information; however, their primary emphasis remains on soiling-loss prediction rather than integrating probabilistic state estimation with a plant-level scheduling problem.

A parallel research direction has focused on image-based and vision-driven soiling assessment. A visual deep-learning framework was proposed for detecting PV surface anomalies and validating the detected conditions against inverter-level energy production, while also considering sustainability-oriented cleaning decisions [14]. AI-based monitoring has likewise been investigated in desert environments using an SVM classifier to distinguish dusty and clean panels, with a reported ROC-AUC of 0.790 [15]. These approaches demonstrate the value of visual information for identifying surface conditions that may not be fully represented by electrical measurements alone. Nevertheless, their principal inputs remain image- or condition-classification oriented, whereas the present framework combines sparse visual descriptors with continuously available SCADA and weather information to estimate a continuous soiling state. The shift from the identification of soiling to maintenance occurrence has gained consideration as well. A data-driven system [16] automated PV cleaning detection through the optimization of a stochastic rate recovery model by using the F1-score and considering rainfall threshold sensitivity. The performance of the system was tested in four PV sites in India and the USA and showed 78% recognition rate and 87.4% F1-score during dry time, which proves that automated event detection can be used for the quantification of soiling losses in a larger scale. Although investigative works have been conducted on predictive maintenance utilizing two-step models, which separate the clean condition assessment from the cleaning classification process. The comparative model [17] made use of precise operation measurements in combination with the regression-based simulation data to create the clean condition power baselines and perform assessment of the cleaning needs. The mean F1-score of 0.9792 obtained by comparing the results provided by real and simulated data demonstrates that the data-driven simulation may be utilized for predictive maintenance in a scalable manner. Hybrid AI technologies have also begun connecting detection and decision-making in maintenance processes. For example, a hybrid framework [18] assisted in combining traditional image processing, machine learning, and deep learning for risk-aware PV soiling detection and cleaning decision-making by taking into account variability of datasets and domain changes. This is a strong step forward from the approach based on just classification to decision-based management of soiling. However, the decision-making is still based on intelligent detection and risk-aware cleaning instead

of using the rolling-horizon optimization approach, which includes electricity price changes, rainfall forecasts, water availability, cleaning expenses, and scheduling. The presented framework enhances this decision-making aspect by treating the anticipated soiling condition and its predictions as input data for constrained economic optimization.

The literature related to scheduling has witnessed a movement toward more adaptive decision mechanisms. For instance, a dynamic cleaning scheduling system that employs reinforcement learning and utilizes two methods of decision-making (PPO and SAC) has been introduced to carry out cleaning in uncertain environmental settings [19]. In a case study carried out in Abu Dhabi, the PPO system was able to achieve cost savings of up to 13% in comparison with the other techniques applied. Moreover, an approach that adopted a cost-based cleaning optimization practice for dry environments [20] showed the application of the Grey Wolf optimizer for defining cleaning techniques. In terms of recent investigations conducted using machine learning algorithms for scheduling in Riyadh, it has been shown that environmental data and air quality indicators, along with XGBoost, Random Forest, and LSTM algorithms, were employed for establishing pollution-related energy losses [21]. In particular, XGBoost has shown a value of R^2 equal to 0.93. Furthermore, the cost-benefit analysis conducted has shown that the biweekly cleaning was able to reduce operational costs by up to 15% compared to methods with fixed frequency. All these studies indicate the significance of using adaptive scheduling and show the importance of linking predictions of the environment into cleaning decisions.

The latest studies suggest a possible physical and economic relationship between soiling and cleaning. One of the methodologies that would allow for obtaining physical relationships between soiling and cleaning was created [22]. The executed experiments have indicated that dust ingredients, cementation caused by humidity, characteristics of the solar panel, and cleaning conditions may play a significant role in both optical/electrical losses and cleaning efficiency. The examination of urban solar power plants showed the combination of the model related to the dry deposition of dirt with the energy criterion, which made it possible to determine the specific practices of doing manual cleaning [23]. The research indicated a high spatial difference between the degrees of soiling and corresponding cleaning frequencies, i.e., different regions require different cleaning practices. Finally, the research on automatic systems of cleaning has undergone development on the way to real-world implementation of maintenance. This direction is complementary to optimization-based approaches because it links estimated panel condition with automated physical intervention. The proposed framework operates at a higher decision-management level by determining when cleaning should be economically initiated, while allowing the physical cleaning mechanism to be treated as an operational resource within the scheduling formulation.

Taken together, the recent literature shows a clear progression from soiling-loss prediction and visual detection toward hybrid decision support and adaptive cleaning optimization. Recent studies have advanced temporal prediction [9, 13], vision-based assessment [14, 15], automated event detection [16], predictive maintenance using simulated operational data [17], hybrid risk-aware detection [18], adaptive scheduling [19–21], robotic cleaning optimization [24], and energy-economic cleaning analysis [22, 23]. The proposed framework builds on these developments by integrating these dimensions into a single pipeline in which multimodal measurements are used to estimate the evolving soiling state, predictive uncertainty is preserved, and the resulting information is passed to a revenue-aware MILP scheduler operating under rainfall, electricity-price, water, cleaning-cost, and operational constraints. This integrated perspective forms the basis for the challenge-gap analysis presented in Section 1.3. A comparative summary of the reviewed studies and the distinguishing characteristics of the proposed framework is provided in Table 1.

1.3. Research gap

Despite the progress made in the area of soiling assessment technology, predictive maintenance, and optimization techniques of cleaning process, none of the research directions were consolidated into a common operational decision-making procedure. Even though machine learning techniques or image processing technologies may estimate the degree of soiling or find the signs of visible contamination, the estimated state is seldom directly utilized in any subsequent maintenance decision-making processes. The complexity of the situation originates from the fact that the available information is different facts. Generally, electrical measurements are continuously available, while SCADA information can be confounded by additional sources of losses such as irradiance. At the same time, visual measurements can provide supplementary surface information, yet are also affected by other parameters.

Another gap identified is the transition from soiling estimation to making cleaning decisions that can be justified economically. Available maintenance strategies increasingly utilize electricity prices, cleaning costs, weather conditions, as well as resource constraints to implement the scheduling procedures. In spite of the fact that adaptive maintenance strategies are delivered by learning-based approaches, cleaning decisions are mostly derived from deterministic estimates or pre-set thresholds for deterministic models. In addition, cleaning of PV systems represents a complex problem that needs to balance energy recovery and cleaning costs, as well as the necessity for water and labor, rainfall forecasts, variations in the price of energy, and available service capacity. Taking into account the specificity of such tasks, the utilization of fixed time intervals is hardly possible.

Finally, the different components have been used in the literature but usually not in a sufficient way, i.e., the changes in the way the facts are analyzed have not been mentioned adequately.

1.4. Novelty, main contributions, and limitations

The proposed research addresses identified gaps by developing an integrated framework that links multimodal soiling state estimation with revenue-oriented, uncertainty-aware cleaning planning. Unlike approaches that treat soiling estimation and maintenance planning as separate tasks, the proposed framework transfers both the estimated soiling state and the associated model-derived uncertainty to the decision-making layer. Consequently, the resulting framework connects predictive information with cleaning decisions subject to economic and operational constraints. The main contributions of this study are summarized below:

1. Hybrid and multimodal soiling state estimation: A hybrid TCN-XGBoost architecture is developed to leverage complementary information from temporal electrical measurements, meteorological variables, and sparse image-derived descriptors. The TCN captures temporal dependencies, while XGBoost represents non-linear relationships between contemporaneous and aggregated features. This multimodal formulation provides a unified estimation mechanism, rather than relying on a single sensing modality or predictive model.

2. The inclusion of forecast uncertainty in maintenance decisions: This approach transfers the predictive uncertainty drawn from the model, which concerns the determined condition of dirtiness, to the level of planning. Rather than changing the condition of dirtiness itself, uncertainty is entered as a risk penalty in the economic objective. The method directly includes prediction confidence in the evaluation of the cleaning alternatives, based on the level of uncertainty measured according to the adjusted simulation method.

3. Cleaning optimization in terms of resources and revenue: A planning approach based on Mixed-Integer Linear Programming (MILP) takes into account the expected energy recovery benefits and costs, electricity rates, precipitation, cleaning costs, availability of labor and water, and predictive uncertainty at the same time. The receding horizon method enables the execution of several planning iterations using newly available information, in contrast to

Table 1. Comparative synthesis of recent PV soiling assessment, prediction, and cleaning-scheduling studies.

Ref./year	Primary focus	Main inputs/data	Core method	Main output/decision	Key capability	Relation to proposed method
[9] (2026)	Soiling-loss prediction	Environmental/operational data; dust accumulation	DLNN, FNN, LSTM, autoencoder	Dust accumulation, light-transmission loss, electrical degradation, power-loss prediction	Quantifies soiling progression and PV electrical losses	Prediction-focused; no explicit uncertainty-aware plant-level scheduling
[13] (2026)	Experimental soiling-loss prediction	Environmental measurements under different cleaning frequencies	Multiple ML models, including stacking ensemble	Soiling loss and PV performance	Connects cleaning frequency with measured PV performance	Predictive/experimental focus; no coupled multimodal scheduling optimization
[14] (2026)	Vision-based soiling detection	PV images + inverter-level energy data	Visual deep learning	Surface anomalies, energy validation, cleaning support	Links visual detection with actual inverter energy production	Primarily vision-oriented; limited continuous SCADA/weather fusion and explicit constrained optimization
[15] (2026)	Desert dust monitoring	PV-panel visual condition data	SVM classification	Dusty/clean panel classification; ROC-AUC = 0.790	Addresses dust identification in desert PV environments	Classification-oriented rather than continuous state estimation and economic scheduling
[16] (2026)	Cleaning-event detection	PV performance + rainfall/weather data; multiple installations	Automated stochastic rate-and-recovery model; F1-based parameter selection	Cleaning-event detection and soiling-rate estimation	78% recall and 87.4% F1 during dry-period utility-scale testing	Event detection rather than joint future-action optimization
[17] (2026)	Two-stage predictive maintenance	High-resolution operational + regression-based simulated data	Clean-condition regression baseline + ML cleaning classification	Clean-condition power baseline and binary cleaning decision	Mean F1 = 0.9792 between real- and simulated-data decisions	Binary threshold decision; no probabilistic multimodal state passed to constrained scheduling
[18] (2026)	Risk-aware soiling detection	PV images / visual features	Hybrid image processing, machine learning, and deep learning; ensemble fusion	Soiling probability/index and cleaning categories	Improves robustness to domain shifts; up to 85.93% accuracy on unseen data	Detection/risk support; lacks explicit revenue-aware MILP scheduling
[19] (2026)	Dynamic cleaning scheduling	Environmental conditions and cleaning-operation information	PPO and SAC reinforcement learning	Adaptive cleaning intervals/schedules	PPO reported up to 13% cost savings	Direct scheduling, but no multimodal uncertainty-aware soiling estimator or MILP formulation
[20] (2026)	Cost-effective cleaning optimization	Soiling-related losses and economic/operational information	Grey Wolf Optimizer	Cost-oriented cleaning strategy	Targets economic cleaning in arid PV environments	Optimization is not coupled with multimodal AI state estimation and uncertainty propagation
[21] (2026)	ML-based dynamic scheduling	NASA, PVGIS, and local air-quality/environmental data	XGBoost, Random Forest, LSTM	Pollution-loss prediction and cleaning schedule	XGBoost: R^2 = 0.93; up to 15% lower cost with biweekly cleaning	Strong prediction/scheduling linkage, but no explicit uncertainty-aware multimodal fusion or MILP resource constraints
[24] (2026)	Robotic cleaning schedule optimization	Five years of meteorological and dust observations from a utility-scale PV plant	XGBoost and Random Forest; constrained ML scheduling	365-day season-adaptive robotic schedule	32–33 annual events; ∼411.6 MWh gain and ∼51% lower cleaning cost in base case	Focused on robotic cleaning and site-specific assumptions; lacks broader multimodal uncertainty-aware scheduling
[22] (2026)	Soiling/cleaning dynamics	Controlled dust samples; module and cleaning conditions	Controlled experimental methodology	Soiling accumulation, optical/electrical losses, cleaning effectiveness	Reveals effects of dust composition, humidity, module properties, and cleaning conditions	Physical characterization rather than predictive decision-making or scheduling
[23] (2026)	Energy-economic cleaning analysis	Urban PV performance/environmental and cost conditions	Dry-deposition soiling model + energy-economic criterion	Site-specific soiling rates and manual cleaning strategies	Captures spatial variation and economic cleaning frequency	Manual/site-specific strategy; no AI multimodal state estimation or uncertainty-aware scheduling
[25] (2026)	Condition-based automated cleaning	PV panel condition measurements	Intelligent condition-based automated cleaning system	Condition-triggered physical cleaning	Moves maintenance toward automated condition response	Emphasizes physical automation rather than revenue-aware optimization of cleaning initiation
Proposed method (2026)	**Integrated multimodal estimation + economic scheduling**	**SCADA + weather + sparse visual descriptors**	**Temporal convolution + gradient boosting + uncertainty-weighted fusion; revenue-aware MILP**	**Probabilistic soiling state → array-level cleaning schedule**	**Couples estimation, uncertainty, revenue, rainfall, electricity price, water, labor, cost, explainability, and sensitivity analysis**	**Integrated end-to-end decision pipeline**

the static methods of cleaning execution based on fixed thresholds or intervals.

4. Integrated decision-making and prediction evaluation: The framework has been evaluated in four cases of simulated photovoltaic plants taking into account the stochastic nature of the environmental and operational conditions using site validation, independent test sets, ablation analysis, sensitivity analysis, and stability assessment, among others. Thus, the evaluation not only includes the prediction accuracy but also the economic and resource-related consequences of cleaning decisions made. The procedure of ablation of selected parts of the approach allows checking the influence of the multimodal inputs, tariff information, and planning strategy.

5. The complete link between estimation and planning: The important difference of the framework is the clear relation between predictions and optimization techniques. The predicted dirtiness level is the basis for evaluation of benefits of cleaning, while the uncertainty related to dirtiness is treated as a risk factor in the planning goal.

There are still many limitations that are available. The present evaluation depends on simulation; there has been no quantitative calibration of predictive uncertainty against real prediction errors and no validation of ground-truth labels and environmental processes derived from the simulation through continuous measurements in the field. Moreover, detailed cleaning logistics, continuous streaming of images, and component-level preservation were not included in this study. The above-discussed limitations, together with the needs of field validation and practical implementation, will be presented in detail in Section 4.5.2.

1.5. Paper layout

The following sections explain the structure of the research paper. Section 2 describes the issue of PV soiling management and its mathematical definition. Section 3 provides details of the hybrid methodology used in this research. It therefore explains how multimodal soiling prediction, uncertainty management, and income-oriented cleaning optimization work together. Section 4 presents the simulation framework, its performance, and the results of simulations conducted during research. At the end, Section 5 concludes the paper by summarizing the discussion held in previous sections.

2. PROBLEM DESCRIPTION AND MATHEMATICAL FORMULATION

This segment provides a rigorous mathematical underpinning for the proposed hybrid technique for decision-making in the soiling management of photovoltaic systems. The problem is formulated as making sequential decisions for estimating the state of soiling in photovoltaic systems based on multimodal observations and using the estimation for choosing optimum cleaning actions that are both economically and technically feasible. The proposed system takes into account SCADA data, climatic information, and additional features from images, whereby the state estimation includes uncertainty analysis and optimization of the cleaning process towards the profit-making goals.

2.1. System description

The system suggested works on the principle that each group of photovoltaic modules is an independently maintainable unit. A finite-time framework of this kind allows the processes of the accumulation of dirt and environmental state to be represented. The relevant terminologies used are detailed in the Nomenclature part.

The following assumptions define the scope of the proposed formulation:

1. All PV array groups are treated as independent maintenance units, and electrical interaction with neighboring arrays is not considered.
2. The electrical infrastructure and communication systems are assumed to operate properly throughout the scheduling horizon.
3. Cleaning can be performed only during approved periods and, in such a way, will influence the soiling condition of the array being worked on.
4. Meteorological observations, SCADA measurements, and image-based inspection features are synchronized to the same decision intervals and are therefore treated as temporally aligned observations.
5. Natural cleaning caused by rainfall and artificial cleaning performed by maintenance crews are considered the only mechanisms capable of reducing accumulated soiling.

Let

$$\mathcal{A} = 1, 2, \ldots, N \tag{1}$$

denote the set of independently schedulable PV array groups, where (N) represents the total number of maintenance units within the solar farm. Likewise, the discrete scheduling horizon is defined as:

$$\mathcal{T} = 1, 2, \ldots, H \tag{2}$$

where (H) denotes the total number of decision intervals considered by the optimization model. For every array group $i \in \mathcal{A}$ and every decision instant $t \in \mathcal{T}$, a multimodal observation vector is constructed by integrating heterogeneous information acquired from the plant monitoring infrastructure. The complete observation vector is defined as:

$$\left[\mathbf{x}_{i,t}^{SCADA}, ; \mathbf{x}_{i,t}^{\text{Weather}}, ; \mathbf{x}_{i,t}^{\text{Vision}}\right] \tag{3}$$

The superscripts identify the source of each feature group. The SCADA component contains electrical and operating measurements; the weather component provides lightweight descriptors obtained from periodic UAV or ground inspections. These modalities jointly characterize the electrical, environmental, temporal, and visible aspects of PV soiling.

The physical situation of the whole array group is expressed by a hidden state of soiling that cannot be recognized directly during the operation of the machine. The soiling state can only be calculated through the use of multimodal observations provided in Eq. (3). The normalized soiling state can be written as follows [26]:

$$s_{i,t} \in [0, 1] \tag{4}$$

where $s_{i,t} = 0$ represents a perfectly clean PV surface and $s_{i,t} = 1$ denotes the maximum practical soiling level considered by the proposed model. The temporal evolution of this hidden state is governed by environmental deposition, rainfall events, and maintenance actions, whose mathematical relationships are developed in the following subsections.

The operational objective is to determine an optimal cleaning policy for every array group throughout the scheduling horizon. Accordingly, the binary maintenance decision variable is defined as:

$$u_{i,t} = \left\{ \begin{array}{ll} 1, & \text{if array group } i \text{ is scheduled for cleaning at time } t \\ 0 & otherwise \end{array} \right\} \tag{5}$$

The decision variable establishes the link between the state estimation module and the optimization framework. The following sections describe how the hidden dirt state defined in Eq. (4) is first estimated using multimodal AI models, and the optimal maintenance decisions defined in Eq. (5) are derived through a revenue-aware optimization model considering operational costs, expected power recovery, resource constraints, and predictive uncertainty at the same time.

2.2. Mathematical representation of multimodal state variables

Soiling management in photovoltaic systems is formulated as a partially observable dynamic system, given that the actual state of the surface cannot be measured continuously during normal operation. Consequently, the framework distinguishes between the latent physical state, the available observations, and the estimated state generated by the hybrid AI model [27]. For each array group $i \in \mathcal{A}$ at time $t \in \mathcal{T}$, the complete system state is represented by the vector:

$$\left[s_{i,t}, P_{i,t}^{\text{loss}}, r_{i,t}, c_{i,t}\right]^{\top} \tag{6}$$

where $s_{i,t}$ denotes the normalized soiling ratio introduced in Eq. (4), $P_{i,t}^{\text{loss}}$ represents the corresponding power loss caused by soiling, $r_{i,t}$ represents the accumulated rainfall-related cleaning effect, and $c_{i,t}$ represents the recent cleaning history of the array group. This state vector provides a compact representation of the physical and operational conditions affecting the future cleaning requirement.

The available measurements described by Eq. (3) are related to the hidden system state through a nonlinear observation function:

$$h(\mathbf{z}_{i,t}) + \varepsilon_{i,t} \tag{7}$$

where $h(\cdot)$ represents the unknown nonlinear relationship between the physical condition of the PV array and the observed measurements, and $\varepsilon_{i,t}$ denotes measurement uncertainty caused by sensor noise, communication errors, and environmental disturbances.

Because the true soiling state cannot be directly obtained from operational measurements, the hybrid estimation module generates an approximation of the latent state as:

$$\hat{s}_{i,t} = f\theta\left(\mathbf{x}_{i,t}\right) \tag{8}$$

where $f_\theta(\cdot)$ denotes the hybrid TCN-XGBoost estimator with uncertainty-aware fusion. The estimated state $\hat{s}_{i,t}$ is subsequently used by the optimization layer to determine whether cleaning provides sufficient economic benefit.

In addition to the estimated soiling severity, the proposed framework explicitly considers prediction uncertainty. Therefore, each estimation output is accompanied by an uncertainty indicator expressed as:

$$\mathrm{Var}\left(\hat{s}_{i,t}\right) \tag{9}$$

where $\sigma^2_{i,t}$ quantifies the confidence level associated with the predicted soiling state. A larger uncertainty value indicates lower confidence in the estimation process and allows the optimization layer to adopt more conservative decisions when the available information is insufficient.

The impact of soiling on PV generation is represented by a power degradation relationship. The clean-reference power output of each array group under the prevailing environmental condition is defined as follows:

$$\eta_i G_{i,t} A_i \phi(T_{i,t}) \tag{10}$$

where $G_{i,t}$ denotes solar irradiance, A_i represents the effective PV area, η_i is the conversion efficiency, and $\phi(T_{i,t})$ describes the temperature-dependent performance correction.

The actual power output under soiling conditions is then expressed as:

$$P^{clean}_{i,t}(1 - \alpha_s s_{i,t}) \tag{11}$$

where α_s is the soiling sensitivity coefficient. Consequently, the relative generation loss caused by accumulated soiling can be obtained as:

$$\frac{P^{clean}_{i,t} - P^{soil}_{i,t}}{P^{clean}_{i,t}} \tag{12}$$

which links the estimated soiling state to the corresponding generation loss [28].

2.3. Soiling dynamics model

The process of flux deposition, natural cleaning processes, and upkeep efforts affect the development of soiling on photovoltaic systems over time. Unlike static threshold methods, this approach views soiling as a dynamic state variable, which, depending on the weather, site features, and management steps, is adopted in subsequent steps. Soiling in each array in the future is thus evaluated with the help of a state-transition theory.

Soiling between decision points can mainly be attributed to airborne particle deposition. Dry accumulation is seen as a function of environmental conditions, including wind velocity, humidity, weather, and the intrinsic dust properties of the installation site. Consequently, the hourly deposition is presented as follows [29]:

$$f_{acc}(v^{wind}_{i,t}, RH_{i,t}, D_i, \chi_i) \tag{13}$$

where $v^{wind}_{i,t}$ represents wind speed, $RH_{i,t}$ denotes relative humidity, D_i represents the site-dependent dust availability factor, and χ_i includes location-specific characteristics affecting particle adhesion. The function $f_{acc}(\cdot)$ can be calibrated using historical measurements or replaced by a learned model when sufficient field data become available.

The method used to estimate the natural reduction of soiling due to rainfall is obtained separately, since rainfall does not necessarily remove all deposited particles. In fact, the effectiveness of this rainfall on cleaning depends on three parameters: precipitation intensity, precipitation duration, and contamination level.

$$\kappa_r g(R_{i,t}) s_{i,t} \tag{14}$$

$$\begin{cases} 0, & R_{i,t} < R^{th}, \\ 1 - \exp(-\mu R_{i,t}), & R_{i,t} \geq R^{th} \end{cases} \tag{15}$$

Rainfall-induced soiling removal is modeled separately because precipitation does not necessarily remove all deposited particles. Its effectiveness depends on rainfall intensity and the accumulated contamination level, as represented by Eq. (14). To exclude negligible precipitation events, the activation function in Eq. (15) introduces a minimum effective rainfall threshold and a saturation parameter. Thus, rainfall below the threshold produces no modeled cleaning effect, whereas higher precipitation progressively increases natural cleaning.

When a manual cleaning operation is performed, the accumulated soiling level is reduced according to the effectiveness of the cleaning technology. The post-cleaning state is therefore represented as:

$$(1 - \eta_c) s_{i,t} \tag{16}$$

where η_c denotes the cleaning efficiency. Values less than unity reflect residual contamination resulting from imperfect cleaning, limited water availability, or inaccessible surface areas.

By combining dry deposition, rainfall removal, and manual cleaning effects, the complete state-transition Equation of the proposed soiling model is formulated as [30]:

$$\Delta s^{rain}_{i,t}(1 - u_{i,t}\eta_c) \tag{17}$$

where $u_{i,t}$ is the binary cleaning decision variable defined in Eq. (5). Eq. (17) provides the fundamental dynamic relationship between environmental conditions and maintenance actions. When $u_{i,t} = 0$, the future soiling state is determined only by deposition and rainfall effects. Conversely, when $u_{i,t} = 1$, the accumulated soiling is reduced according to the cleaning efficiency.

For numerical stability and physical consistency, the predicted soiling state is bounded within the feasible interval:

$$0 \leq s_{i,t} \leq 1, \qquad \forall i \in \mathcal{A}, t \in \mathcal{T} \tag{18}$$

The formulation described above is capable of providing a flexible description of the amount of dust while remaining in agreement with the optimization framework. In this context, the transition model defined by Eqs. (13)-(18) allows the scheduler to analyze cleaning decisions' economic implications based on future states instead of taking into account the instantaneous soiling level only. The next subsection describes the method of estimating the expected energy recovery due to cleaning actions, which is introduced in the revenue optimization objective.

2.4. Energy recovery and economic benefit model

The estimated soiling state is converted into recoverable energy and monetary value to determine whether cleaning is economically justified. For each array group, the clean-reference output, soiling-related loss, recoverable energy, revenue, and cleaning cost are defined below.

$$\eta_i G_{i,t} A_i \left(1 - \beta_T (T_{m,i,t} - T_{ref})\right) \tag{19}$$

where η_i denotes the nominal conversion efficiency of array group (i), $G_{i,t}$ represents the incident solar irradiance, A_i is the effective PV area, β_T is the temperature coefficient, $T_{m,i,t}$ is the module temperature, and T_{ref} is the reference temperature.

The actual power output under the current soiling condition is obtained by applying a soiling reduction factor:

$$P_{i,t}^{clean}(1 - \alpha_s s_{i,t}) \tag{20}$$

where α_s represents the sensitivity of power generation to the normalized soiling state. The instantaneous power loss caused by soiling is therefore calculated as:

$$P_{i,t}^{clean} P_{i,t}^{soil} \tag{21}$$

which represents the generation that can potentially be recovered through effective cleaning.

Since cleaning decisions are made over discrete scheduling intervals, the recoverable energy associated with cleaning array group (i) at time (t) is calculated as:

$$\sum_{\tau=t}^{t+H_c} \Delta P_{i,\tau}^{soil} \Delta t \tag{22}$$

where H_c denotes the recovery period considered after cleaning and Δt is the decision interval. Eq. (22) accounts for the energy recovered over the post-cleaning recovery period rather than only at the cleaning instant.

However, not all recovered energy has the same economic value because electricity prices may vary over time. Therefore, the monetary value of the recovered energy is determined by:

$$\sum_{\tau=t}^{t+H_c} \pi_\tau E_{i,\tau}^{rec} \tag{23}$$

where π_τ represents the electricity tariff or market price at time τ. The total cost associated with a cleaning operation includes water consumption, labor requirements, equipment operation, and additional maintenance expenses. The cleaning cost is formulated as:

$$C_i^{equip} + C_i^{water} W_i + C_i^{labor} L_i \tag{24}$$

where C_i^{equip} represents equipment-related expenses, C_i^{water} denotes the unit water cost, W_i is the required water volume, C_i^{labor} is the labor cost coefficient, and L_i represents the required labor effort.

Before solving the scheduling problem, all quantities derived from the hybrid estimation and simulation layers are evaluated for the current optimization horizon and treated as fixed parameters. These quantities include the estimated soiling state $\hat{s}_{\mathrm{i,t}}$ predictive uncertainty $\sigma_{\mathrm{i,t}}$ expected cleaning benefit $B_{i,t}^{clean}$ and effective rainfall probability $\rho_{i,t}^{rain}$. Consequently, Eq. (25) is evaluated as a preprocessing step to obtain the rain-adjusted cleaning benefit $B_{i,t}^{clean}$; neither $B_{i,t}^{clean}$ nor $\rho_{i,t}^{rain}$ are optimization variables.

Accordingly, the direct economic benefit obtained from cleaning array group (i) at time (t) is expressed as:

$$\tilde{B}_{i,t}^{clean} - B_{i,t}^{clean}\left(1 - \rho_{i,t}^{rain}\right) \tag{25}$$

where a positive value indicates that the expected revenue recovered from improved generation exceeds the associated maintenance cost.

Eqs. (19)–(25) relate the estimated soiling condition to recoverable energy and economic benefit while accounting for time-varying electricity prices and cleaning costs. These quantities are combined with predictive uncertainty and operational constraints in the revenue-aware scheduling model presented next [31].

2.5. Revenue-aware cleaning scheduling optimization

The revenue-based scheduler determines the cleaning schedule by combining projected soiling accumulation, expected energy recovery, electricity prices, precipitation, cleaning costs, and operational constraints. Unlike fixed-interval or threshold-based policies, the formulation evaluates each potential cleaning action based on its risk-adjusted expected economic benefit. The binary variable $u_{i,t}$ in Eq. (5) determines whether cleaning is scheduled for panel group (i) in interval (t), resulting in the following mixed-integer optimization problem [30]. The optimization objective is defined as follows:

$$J = \max_{u} \sum_{i \in A} \sum_{t \in T} u_{i,t} \left(\tilde{B}_{i,t}^{clean} - \lambda \sigma_{i,t}\right) \tag{26}$$

The uncertainty term does not modify the estimated soiling state $\hat{s}_{\mathrm{i,t}}$. Instead, $\lambda\sigma_{i,t}$ acts as a risk penalty on the rain-adjusted expected cleaning benefit, thereby transferring predictive uncertainty from the estimation layer to the economic scheduling layer through the objective function.

In the above formulation, ($\tilde{B}_{i,t}^{clean}$) denotes the expected economic benefit adjusted for rainfall, defined in Eq. (25); ($B_{i,t}^{clean}$) represents the expected economic benefit before accounting for natural rainfall cleaning; ($\rho_{i,t}^{rain}$) denotes the probability of effective rainfall during the planning horizon; ($\sigma_{i,t}$) represents the model-derived predictive uncertainty associated with the estimated soiling state; and (λ) is a non-negative uncertainty aversion coefficient that controls the contribution of said uncertainty to the scheduling objective.

Eq. (26) maximizes the risk-adjusted expected economic benefit, taking into account the probability of natural rainfall cleaning and the uncertainty regarding the projected soiling level. Thus, cleaning is favored only when its expected economic benefit remains sufficient after considering these risk factors.

$$\sum_{i \in \mathcal{A}} u_{i,t} W_i \le W_t^{max}, \qquad \forall t \in \mathcal{T} \tag{27}$$

where W_i represents the water requirement of array group (i), and W_t^{max} denotes the maximum available water volume during interval (t).

Similarly, the available workforce imposes a limitation on the number of simultaneous cleaning operations. The labor constraint is expressed as:

$$\sum_{i \in \mathcal{A}} u_{i,t} L_i \le L_t^{max}, \qquad \forall t \in \mathcal{T} \tag{28}$$

where L_i represents the labor requirement for cleaning array group (i), and L_t^{max} is the available labor capacity.

In order to minimize repeated maintenance activities, a pace interval is introduced. An array group that has already been cleaned cannot be subjected to the cleaning process again until the required maintenance interval has passed.

$$\sum_{\tau=t}^{t+T_{min}-1} u_{i,\tau} \leq 1, \qquad \forall i \in \mathcal{A}, \forall t \in \mathcal{T} \tag{29}$$

where $T_{\min}$ denotes the minimum allowable time interval between two consecutive cleaning events.

The cleaning process is also restricted by operational availability conditions, such as daylight operation periods, crew working hours, or safety limitations. These restrictions are represented using an availability coefficient $a_{i,t}$:

$$u_{i,t} \leq a_{i,t}, \qquad a_{i,t} \in 0, 1 \tag{30}$$

where $a_{i,t} = 1$ indicates that cleaning is feasible at time (t), while $a_{i,t} = 0$ prohibits maintenance activities. For practical deployment, the optimization problem is solved using a rolling-horizon strategy. At each decision point, the hybrid estimator and the associated environmental and economic input data are first updated for the current planning horizon; subsequently, the resulting fixed parameters are supplied to the MILP scheduler to determine the cleaning decisions. The rolling optimization procedure can be represented as:

$$t, t+1, \ldots, t+H_o-1 \tag{31}$$

where H_o represents the optimization horizon length.

The resulting mixed-integer optimization problem can therefore be summarized as:

$$\begin{aligned} &\max_{u_{i,t}} J \\ &\text{subject to } \begin{cases} \text{Constraints Eqs.}(27)-(30), \\ u_{i,t} \in 0,1. \end{cases} \end{aligned} \tag{32}$$

The mixed-integer formulation presented in Eq. (32) establishes a cleaning schedule that maximizes the risk-adjusted economic return while adhering to the resource, spacing, and operational availability constraints specified in Eqs. (27) to (30) [32]. This schedule reflects the optimal decision at the current stage of the rolling horizon. As new measurements and forecasts are obtained, the optimization problem is updated and resolved. The practical implications and operational significance of these constraints will be discussed in the following subsection.

2.6. Operational constraints and practical feasibility

The limitations indicated in Eqs. (27) to (30) reflect the basic constraints involved in operation for cleaning solar panels. The water limitation restricts the use of water in times of drought whereas labor limitation takes into consideration the limited number of technicians available for cleaning. The minimum distance limitation does not allow cleaning of the same panels at certain intervals so that resource usage is not excessive and unnecessary idleness is avoided. The availability limitation restricts cleaning to appropriate intervals, which must take into consideration sunlight, weather, availability of employees, etc.

2.7. AI–Optimization coupling and closed-loop decision mechanism

A hybrid AI estimator is being combined with a revenue-based optimizer in a closed-loop process architecture that makes decisions. Multimodal measurements inform how bad the fouling is and what the uncertainties of this estimation are; after this, the optimization layer combines all these estimates with economic and operational data to come to a cleaning decision.

The connection between the estimation and optimization blocks is important because an appropriate maintenance decision depends not only on the severity of the soiling but also on the reliability of the prediction [33]. For example, a high value of contamination does not mean profit, as it might be expected with a high degree of uncertainty. Conversely, even a moderate level of soiling, predicted with a high degree of confidence, can justify cleaning when electricity prices are favorable. Thus, the means suggested in the paper account for the soiling status, as well as the uncertainty of predictions, electricity tariffs, and the potential impact of rain.

$$(\hat{s}_{i,t}, \sigma_{i,t}, \pi_t, R_t) \rightarrow u^*_{i,t} \tag{33}$$

Where $\hat{s}_{i,t}$ represents the estimated soiling ratio obtained from the hybrid AI model, $\sigma_{i,t}$ denotes the corresponding prediction uncertainty, π_t represents the electricity tariff or economic value signal at time t, and R_t represents the expected rainfall-related recovery information. The output $u^*_{i,t}$ is the optimal cleaning decision generated by the revenue-aware mixed-integer optimization model.

The complete decision-making process is therefore formulated as a closed-loop framework in which raw operational data are transformed into economically optimized maintenance actions:

$$\mathbf{x}_{i,t} \rightarrow \hat{s}_{i,t} \rightarrow \tilde{B}^{clean}_{i,t} \rightarrow J \rightarrow u^*_{i,t} \tag{34}$$

Where $\mathbf{x}_{i,t}$ denotes the multimodal observation vector introduced in Eq. (3), $\hat{s}_{i,t}$ is the estimated soiling state, $\tilde{B}^{clean}_{i,t}$ is the rain-adjusted expected economic benefit defined in Eq. (25), J represents the risk-adjusted optimization objective defined in Eq. (26), and $u^*_{i,t}$ is the final cleaning decision obtained from the scheduler.

The closed-loop mechanism updates the estimated state and the cleaning schedule as new information becomes available from SCADA systems, weather conditions, inspections, and the market. This continuous interaction enables adaptation to changing environmental and economic conditions. Collectively, Eqs. (3) through (34) establish an integrated approach ranging from multimodal state estimation and uncertainty quantification to dynamic fouling evolution and cost-effective maintenance optimization.

3. METHODOLOGY

The second section has developed the joint concept of dirty state estimation and optimization of the cleaning process. This section explains the application of this concept, which comprises the processes of multimodal data processing, hybrid TCN-XGBoost estimation, uncertainty management, and training and validation of models. The overall process is illustrated in Fig. 1.

3.1. Overall framework architecture

The suggested framework includes four phases: (i) collecting and pre-processing multimodal data, (ii) hybrid assessment of soiling state using TCN and XGBoost and uncertainty evaluation, (iii) MILP scheduling considering revenue from the equipped photovoltaic plant, and (iv) explainability. Information derived from SCADA, weather data, tariffs, and periodic visual inspections is collected at a decision layer and sent to the two branches of AI. The TCN captures temporal dependencies, while XGBoost models non-linear relationships between contemporaneous and aggregated features. Their outputs are combined to obtain the estimated soiling state and associated uncertainty; these values are subsequently incorporated into the revenue-maximization scheduling model. The resulting cleaning decision is implemented at the current decision-making moment; subsequently, new measurements and forecasts update the next optimization window.

During the development phase, simulated cases are separated by site for model validation, followed by evaluation on an independent, held-out test set. In the operational phase, the estimator receives updated plant and environmental data and

supplies the state estimate and uncertainty to the receding-horizon MILP scheduler. Fig. 1 summarizes the workflow, encompassing training, site-specific validation, independent testing, and operation.

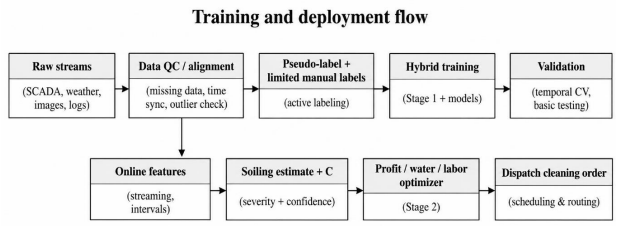

Fig. 1. Training, site-wise validation, independent testing, and rolling-horizon operational deployment workflow of the proposed hybrid AI-based PV cleaning management framework.

3.2. Multimodal data processing and feature engineering

Data regarding SCADA, meteorology, tariffs, operational status, and periodic visual inspections are synchronized with the temporal resolution of the model's decisions. SCADA variables describe electrical performance, while meteorological variables characterize environmental conditions associated with soiling accumulation and rain-induced cleaning. Visual observations are represented by compact descriptors based on color, texture, and edges, rather than raw images. The visual modality is used as a sparse supplementary source rather than as a continuous image-based sensing channel. The extracted descriptors are incorporated only at the defined inspection times and associated with the corresponding PV array group, while between inspections the estimator relies on the available electrical, environmental, and temporal information. Thus, the framework does not directly process complete images or assume continuous image acquisition; camera characteristics, illumination normalization, and detailed image-registration conditions are not explicitly modeled in the present simulation. The visual branch therefore serves as complementary information to SCADA and weather measurements rather than as a standalone sensing modality. Temporal features include time-lagged power loss measurements, rolling statistics, time elapsed since the last effective rainfall, and operational history.

Continuous variables are normalized using statistics derived exclusively from the corresponding training data, and the same parameters derived during training are subsequently applied to the validation and test samples. Missing observations are handled via variable-specific temporal interpolation or by using the most recent valid measurement, while anomalous operating periods unrelated to fouling accumulation are excluded from model development. All time-lagged and moving-average features are constructed causally using observations available no later than the corresponding decision time; thus, future observations are not used to construct these features. Similarly, each 72-hour input sequence for the TCN contains only observations that precede or coincide with its prediction time. These procedures prevent future information from being incorporated into the feature representation or the normalization process. The resulting feature groups are summarized in Table 2 and fed into the temporal and feature branches of the estimator.

The resulting feature representation provides synchronized and computationally manageable input to the two branches of the hybrid estimation architecture. The TCN branch subsequently processes the temporal information, while aggregated and non-temporal features are provided to the XGBoost branch. Their outputs are subsequently fused to obtain the soiling prediction and its associated uncertainty.

3.3. Hybrid soiling estimation model

The processed multimodal features are fed into a hybrid TCN-XGBoost estimator. The TCN receives the sequential feature history and captures temporal dependencies associated with soiling accumulation, rain-induced recovery, and changing operating conditions. XGBoost receives contemporaneous and aggregated features, modeling non-linear interactions among electrical, environmental, temporal, and visual data. The outputs from both branches are fused to obtain the estimated soiling state and its prediction uncertainty.

Table 2. Multimodal input features and their roles in the proposed framework.

Data source	Representative features	Main role
SCADA	DC power, DC voltage, DC current, module temperature, inverter status, power residual	Characterization of electrical performance and soiling-related power losses
Weather	Irradiance, ambient temperature, relative humidity, wind speed, rainfall, cumulative precipitation	Modeling environmental conditions affecting soiling accumulation and natural cleaning
Temporal features	Lagged power residuals, rolling statistics, elapsed time since effective rainfall, operating history	Representation of accumulation trends and temporal dependencies
Visual inspection	Color statistics, texture descriptors, edge density, image-derived surface indices	Auxiliary characterization of visible surface contamination
Economic data	Electricity tariff or market price	Estimation of the economic value of recoverable energy
Operational status	Availability indicators, abnormal-operation flags, maintenance status	Prevention of erroneous soiling inference during non-representative operating conditions

In order to combine the two branches, the outputs of TCN and XGBoost are concatenated. The process of making the final estimation can be described as:

$$XGBoost\left(\mathbf{x}_{i,t}^{tab}\right) \rightarrow \text{Concat}\left(\mathbf{z}_{i,t}^{TCN}, \mathbf{z}_{i,t}^{XGB}\right) \rightarrow \left(\hat{s}_{i,t}, \sigma_{i,t}\right) \quad (35)$$

where $\mathbf{x}_{i,t}^{Seq}$ and $\mathbf{x}_{i,t}^{tab}$ denote the sequential and contemporaneous/aggregated feature representations supplied to the TCN and XGBoost branches, respectively. The vectors $z_{i,t}^{TCN}$ and $z_{i,t}^{XGB}$ are their branch outputs, while $\hat{s}_{i,t}$ denotes the fused normalized soiling estimate and $\sigma_{i,t}$ its associated predictive uncertainty.

The fusion layer generates the estimated fouling state $\hat{s}_{i,t}$ along with the associated predictive uncertainty $\sigma_{i,t}$, as expressed in Eq. (33). Subsequently, this uncertainty magnitude is passed to the scheduling layer as a model-derived risk indicator. In this study, $\sigma_{i,t}$ was used to represent predictive uncertainty within the adopted simulation framework; however, no specific quantitative calibration analysis comparing $\sigma_{i,t}$ against observed prediction errors was conducted. The uncertainty estimate is retained alongside the prediction and passed to the revenue-oriented optimization layer, where it contributes to the risk-adjusted objective defined in Eq. (26). Thus, the AI module provides both the estimated state and an indication of prediction confidence, rather than a simple point estimate.

Fig. 2 summarizes the two-branch architecture and the information flow from multimodal inputs to the estimated soiling state, uncertainty, and downstream economic optimization [34]. The outputs of the two branches are subsequently combined through an uncertainty-weighted fusion mechanism to obtain a continuous soiling-state estimate and its associated predictive uncertainty. The resulting state estimate and uncertainty are then used to derive the expected soiling-related power or energy loss and are passed to the revenue-aware scheduling layer for maintenance optimization.

Overall, the hybrid architecture combines temporal modeling with nonlinear feature learning while maintaining a computationally

practical structure [34]. The resulting state and uncertainty estimates provide the information required by the revenue-aware optimizer without relying on continuous high-resolution image processing or a computationally intensive end-to-end deep-learning architecture.

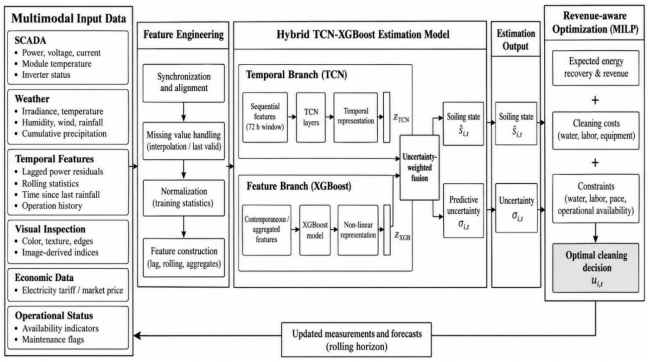

Fig. 2. Internal architecture of the proposed hybrid TCN–XGBoost soiling-estimation model with uncertainty-weighted fusion.

3.4. Model training strategy

For each Monte Carlo seed, simulation cases are evaluated using a four-fold "leave-one-site-out" (LOSO) cross-validation. In each partition, data from three solar farm sites are used for model development, while the fourth is reserved exclusively for validation. This site-based separation prevents data from the same location from being used simultaneously for training and validation, allowing for a more rigorous assessment of the model's ability to generalize across different site conditions. The validation partitions are used for model selection, hyperparameter tuning, and early stopping, while a separate, dedicated test set is reserved for the final performance evaluation presented in Section 4. Within each fold, feature construction and sequence formation were performed separately for the training and validation data, with temporal inputs restricted to information available at or before each prediction time; the independent test set was processed only after model selection was completed.

Given the lack of continuous field measurements of the actual soiling rate for the simulation study, supervised training utilized simulation-based reference labels generated from the baseline performance model (clean condition) and the prescribed soiling accumulation and precipitation processes described in Section 3. These labels served as reference targets for model development and evaluation. Consequently, the reported predictive metrics reflect performance relative to the adopted simulation-based reference labels, rather than a direct validation against field-measured soiling rates. These labels served as ground-truth targets for model development and evaluation, including the task of classifying the need for cleaning. The ground-truth cleaning decision labels were derived from simulation-based cleaning requirements and were not generated from the output of the proposed revenue-aware MILP scheduler. Consequently, the classification evaluation was conducted independently of the subsequent economic scheduling optimization.

The estimator is trained using a composite loss that combines soiling-regression accuracy, cleaning-necessity classification, and uncertainty consistency:

$$\mathcal{L} = \alpha\mathcal{L}_{MAE} + \beta\mathcal{L}_{BCE} + \gamma\mathcal{L}_{KL} \tag{36}$$

where $\mathcal{L}_{MAE}$ measures the absolute error between the predicted and reference soiling ratios, $\mathcal{L}_{BCE}$ represents the binary cross-entropy loss for the cleaning-necessity output, and $\mathcal{L}_{KL}$ provides an uncertainty-consistency regularization term. The coefficients α, β, and γ control the relative contribution of the three components. This multi-objective formulation prevents training from focusing exclusively on point-estimation accuracy and encourages the model to provide uncertainty information suitable for the subsequent risk-aware scheduling process.

The neural-network parameters of the TCN and fusion stages are optimized using the Adam optimizer with mini-batch training. XGBoost is trained separately using its gradient-boosting procedure and the specified tree parameters. The principal hyperparameters are selected using the validation folds from predefined candidate configurations. The final configuration uses a 72-h temporal window, dilation factors of 1, 2, 4, and 8, 32 TCN channels, a dropout rate of 0.15, and 300 XGBoost trees with a maximum depth of 5 and a learning rate of 0.05. The corresponding optimization and model settings are summarized in Table 4.

Early stopping is applied during neural-network training to reduce overfitting. Training is terminated when the validation loss fails to improve for a predefined number of consecutive epochs, and the parameter set associated with the lowest validation loss is retained. The validation data are not used to update the model parameters. After model selection and training, the final configuration is evaluated only on the independent test cases, which remain unseen during training and hyperparameter selection.

The resulting model provides both the estimated soiling state $\hat{s}_{i,t}$ and its predictive uncertainty $\sigma_{i,t}$. These outputs are subsequently supplied to the revenue-aware optimization layer, where the uncertainty information contributes directly to the risk-adjusted objective in Eq. (26). Thus, the training procedure establishes a consistent link between model estimation, uncertainty assessment, and the downstream cleaning decision.

3.5. Revenue-aware decision framework

The layer of decision-making combines revenue data and utilizes the results from the hybrid estimation to resolve which cleaning activities to perform. The mathematical optimization model is introduced in Section 2; therefore, the emphasis of this section lies on its practical application and information exchange between the AI estimator and the scheduling layer.

At each 6-h decision epoch, the latest SCADA, weather, tariff, and available inspection information is processed by the trained hybrid model to obtain the estimated soiling state $\hat{s}_{i,t}$ and its associated uncertainty $\sigma_{i,t}$. These results update the economic coefficients and uncertainty terms of the MILP model defined in Section 2. The model is solved using a 7-day rolling horizon, subject to constraints regarding water, labor, availability, and minimum spacing.

Only the decisions applicable to the current decision stage are implemented. In the next cycle, new measurements and updated forecasts are incorporated, and the optimization is solved again. The operational flow is therefore:

Multimodal data → Soiling estimation → Objective update → MILP scheduling → Cleaning decision → State update.

The selected decisions $u^*_{i,t}$ are transferred to the maintenance layer, while non-selected array groups remain under monitoring. After cleaning or natural rainfall has modified the soiling state, the updated measurements are incorporated into the next estimation cycle. Thus, rainfall does not introduce a separate optimization trigger; its effect is reflected through the updated state and forecast information at the subsequent decision epoch.

This rolling closed-loop structure preserves the uncertainty generated by the AI layer when the economic decision is made. Consequently, the scheduler can distinguish between a potentially valuable cleaning action supported by a reliable state estimate and one associated with greater predictive uncertainty.

Algorithm 1 summarizes the operational cycle of the proposed revenue-aware hybrid framework.

Algorithm 1. Revenue-aware hybrid scheduling loop

1. Acquire the latest SCADA, weather, tariff, and available image-derived features for each array group.

Table 3. Computational complexity and deployment characteristics of the considered approaches.

Model	Parameters / Size	Inference latency (ms)	Approx. memory	Decision-support capability
Threshold rule	—	0.2	<1 MB	No uncertainty or explanation
Random Forest + greedy	∼2.1 MB	6.2	11 MB	Limited economic optimization
LSTM + rules	0.82 M	18.4	52 MB	Limited transparency
Proposed hybrid	0.43 M + trees	9.6	39 MB	Uncertainty + SHAP + MILP

2. Update the rolling feature histories and obtain predictions from the TCN and XGBoost branches.

3. Fuse the branch outputs to obtain the estimated soiling state $\hat{s}_{i,t}$ and predictive uncertainty $\sigma_{i,t}$.

4. Update the MILP objective inputs using the estimated recoverable energy value, electricity price, cleaning cost, and rainfall-related forecast information.

5. Solve the MILP over the 7-day rolling horizon subject to water, labor, availability, and minimum-cleaning-spacing constraints.

6. Implement the selected cleaning decisions for the current 6-h decision epoch and record the decision and associated explainability information.

7. Incorporate newly available measurements and environmental information and repeat the procedure at the next decision epoch.

The resulting implementation maintains a clear separation between state estimation and economic optimization while allowing information to flow continuously between them. This modular structure also enables the AI estimator and scheduling module to be updated independently without altering the underlying mathematical decision framework.

3.6. Explainability module

SHAP (SHapley Additive exPlanations) values are applied to the hybrid AI estimator to quantify the contribution of input variables to the predicted fouling state. Global SHAP analysis characterizes the overall influence of the variables, while local SHAP analysis explains individual predictions for selected array groups and decision intervals. SHAP is used to interpret the AI-based estimation stage and is not considered a direct explanation of the MILP optimization decision.

Prediction uncertainty is presented alongside the SHAP explanation as an indicator of confidence in the estimate. This does not imply prediction accuracy and is incorporated independently into the risk-aware scheduling objective.

3.7. Computational complexity and deployment

The proposed estimator combines a TCN with approximately 0.43 million trainable parameters and an XGBoost branch composed of 300 trees with a maximum depth of 5. The MILP model operates at the array-group level rather than the individual-module level, which limits the number of decision variables compared to module-level scheduling. Scheduling is performed using the sliding horizon described in Section 3.

Table 3 below contains the key features of computational characteristics and decision-making abilities of the approaches described above.

The deployment architecture divides plant-level inference and scheduling from model development and retraining. The plant data can be processed using a local method or through edge-assisted processing; however, complicated model updates and historical analysis can occur on a central server.

3.8. Practical deployment considerations

The framework was developed and evaluated in a simulation environment; consequently, the deployment considerations outlined here represent implementation requirements rather than proven field performance. Practical deployment would require site-specific calibration, reliable integration with SCADA and inspection systems, secure data communication, and validation under the plant's specific operational, soiling, rainfall, and market conditions.

The system should operate as a decision support system instead of as an independent control system. To ensure traceability in the implementation, there should be processes implemented for documenting input data, expected stage, uncertainty, optimization, version of the model, and operator actions. The system must undergo field validation before going into operation.

4. RESULTS AND DISCUSSION

The proposed framework is evaluated through predictive, operational, ablation, robustness, explainability, and error analyses. The results allow for the assessment of both the soiling state estimation and the resulting cleaning decisions under the simulated environmental, economic, and operational conditions described in Section 3.

4.1. Experimental setup and evaluation protocol

The suggested framework was assessed in a controlled simulation setting intended to replicate the key temporal, environmental, operational, and economic variables affecting washing decisions linked to soiling in utility-scale photovoltaic (PV) plants. Four representative types of PV farms were analyzed: the desert-flat, semi-arid coastal, inland agricultural, and dusty industrial-fringe types. Each of the farms contained 12 array groups that each operated independently and was assessed for the entire 8760-hour year. The simulations included synthetically generated profiles of solar irradiance, ambient temperature, relative humidity, wind, precipitation, electricity prices, the soiling process, and recovery due to rainfall. Sparse information from visual inspections is fed in the model at the defined periods only in the form of descriptors created from images, thus modeling the intermittent nature of UAV- or ground-based inspections instead of assuming continuous image retrieval. As a result, the output dataset enables assessment of the entire decision pipeline from the estimation of the soiling state to planning cleaning activities, taking into consideration economic factors.

A simulated dataset has been generated to represent controlled variations in the key environmental, electrical, operational, and economic factors considered by the proposed framework. Consequently, the simulation reflects prescribed process relationships rather than a field-calibrated statistical reconstruction of a specific PV plant. Specifically, soiling accumulation, rainfall-induced recovery, measurement variability, and intermittent visual observations were generated according to the simulation procedures defined in Section 3. The resulting data are intended to support a controlled comparative evaluation of the proposed method across various site archetypes, rather than to reproduce the full statistical distribution of a specific plant. Accordingly, the reported results should be interpreted within the context of the adopted simulation framework; additional field-calibrated distributions and measurements would be required for site-specific replicability.

The assessment took into account both predictive ability and functional usefulness. The experimental work tested soil condition estimation, clean requirement identification, and yearly business effectiveness when there is outside interference like cloud passage, shadowing, inverter faults, and sensor malfunction. Twenty Monte Carlo seeds were used to represent the stochastic variability of the simulated inputs. For each seed, a four-fold cross-validation strategy—specifically, a "leave-one-site-out" (LOSO) approach—was employed for model development: three sites

Table 4. Experimental and evaluation settings.

Category	Setting
Simulation horizon	8760 h per site
PV-farm archetypes	4 sites × 12 array groups/site
Monte Carlo evaluation	20 independent seeds
Decision interval	6 h rolling decision
Optimization horizon	7 days (28 decision steps)
Input modalities	SCADA, weather, temporal features, sparse image descriptors, tariff
Baselines	14-day fixed, SR threshold, RF + greedy, LSTM + rules
Predictive metrics	MAE, RMSE, ROC-AUC, F1
Operational metrics	Annual net value, water consumption, cleanings/year
Validation	4-fold site-wise cross-validation (leave-one-site-out) within each Monte Carlo seed
Final evaluation	Independent held-out test set
Simulation environment	Python 3.11
Edge evaluation	ARM-equivalent latency profiling

were used for training, while the remaining site was reserved for validation. An independent, held-out test set—excluded from model fitting and hyperparameter optimization—was used exclusively for the final performance evaluation. Operational decisions were made every 6 hours using a 7-day rolling optimization horizon comprising 28 decision steps. At each decision point, the TCN-XGBoost estimator processed the most recent multimodal observations and fed the resulting estimates into the revenue-aware MILP scheduler described in Section 3.

Four reference models were selected to provide meaningful reference points: a fixed 14-day cleaning model, a threshold-based soiling ratio (SR) model, an RF-based model combined with an aggressive cleaning strategy, and an LSTM-based model combined with a rule-based cleaning strategy. These benchmarks represent fixed-period maintenance, condition-based maintenance, feature-based machine learning with greedy cleaning, and temporal deep learning with rule-based scheduling, respectively. Thus, the comparison demonstrates the efficiency of better state estimation as opposed to economically documented scheduling. The comparison considers both predictive and operational performance metrics. The estimation of soiling was analyzed using MAE and RMSE methods, whereas the performance of cleaning decisions was analyzed using ROC and F1 score measures. Three main operational indicators were analyzed: annual net economic value, water consumption, and cleaning frequency. Together, these indicators help assess whether improvements in maintenance performance result from more accurate soiling estimation and economically informed scheduling rather than from more frequent cleaning. The main experimental and evaluation conditions are summarized in Table 4, whereas Table 5 presents the main model hyper-parameters and those of optimization. Unless specified otherwise, uncertainty intervals are given in the form of bootstrap confidence intervals computed for Monte Carlo tests. Formal significance testing of the annual economic differences was not performed because the seed-level annual outcomes required for paired statistical comparison were not retained. Accordingly, the reported uncertainty intervals are used to describe the variability of the simulated annual results rather than to establish statistical significance between competing strategies.

Fig. 3 illustrates a representative section of the simulated environment over 60 days, which shows the temporal interaction between the accumulation of soiling, the power-loss proxy, as well as the recovery from rainfall. The example shows that visual inspection is intermittent and that there is a difference between

Table 5. The core hyperparameters and optimizer settings.

Hyperparameter	Value
TCN window length	72 h
TCN dilation stack	1, 2, 4, 8
TCN channels	32
Dropout	0.15
XGBoost trees	300
Tree depth	5
Learning rate	0.05
MILP horizon	28 decision steps (7 days)
Uncertainty penalty λ	0.12
Minimum cleaning spacing	3 days per array group

gradual soiling development and sudden events caused by rain.

Fig. 3 illustrates a representative 60-day segment of the simulated environment for one of the photovoltaic plant archetypes described in Section 4.1. This segment shows the temporal interaction between soiling accumulation, the associated indirect power loss indicator, and the recovery triggered by rainfall. During dry periods, the simulated soiling level and the corresponding power loss indicator rise progressively, whereas rainfall events cause relatively rapid reductions in the accumulated soiling level and associated losses. The figure also illustrates the intermittent nature of the visual modality, given that image-derived descriptors are available only at predefined inspection times. Consequently, this example demonstrates the temporal behavior the proposed hybrid estimator is designed to capture, including gradual soiling accumulation, abrupt rain-induced recovery, and the intervals between visual observations during which the model relies on electrical, environmental, and temporal information. The figure provides an illustrative representation of the simulated conditions that serve as the basis for the prediction and planning analyses presented in subsequent sections.

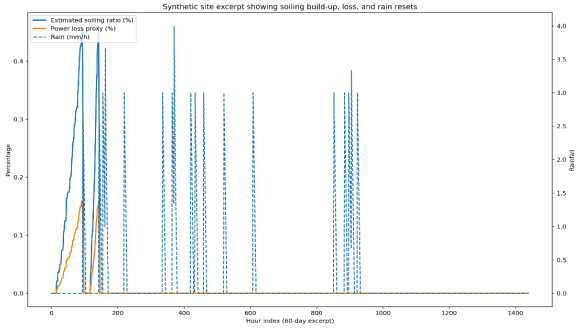

Fig. 3. Representative 60-day excerpt of the simulated environment showing soiling accumulation, the associated power-loss proxy, and rainfall-induced recovery.

4.2. Soiling estimation and cleaning-decision performance

The first stage of the evaluation examines whether the proposed hybrid estimator can reliably characterize the evolution of the fouling state and translate this information into precise cleaning decisions. These two aspects are evaluated separately because predictive and operational performance represent distinct evaluation dimensions of the proposed framework. Accordingly, the predictive analysis considers continuous soiling-ratio (SR) estimation and cleaning-decision classification for the methods that explicitly generate predictive scores, whereas the operational analysis considers the annual consequences of all evaluated maintenance policies. This separation prevents conventional predictive metrics from being applied to the fixed-interval policy, which does not perform model-based state estimation.

Table 6-(a). Predictive and cleaning-decision performance of the evaluated model-based approaches.

Method	SR MAE (%)	Cleaning-decision AUC	Cleaning-decision F1
SR threshold	1.74	0.921	0.884
RF + greedy	1.21	0.956	0.914
LSTM + rules	1.08	0.941	0.903
Proposed hybrid	0.89	0.983	0.949

Note: Value intervals for annual net value represent 95% bootstrap confidence ranges.

Table 6-(b). Annual operational performance of the evaluated cleaning strategies.

Method	Net annual value (M USD)	Water (10^3 m^3)	Cleanings/year
14-day fixed	2.81 ± 0.04	18.4	66
SR threshold	2.93 ± 0.03	14.2	49
RF + greedy	2.98 ± 0.03	12.7	41
LSTM + rules	3.01 ± 0.03	12.2	39
Proposed hybrid	3.09 ± 0.02	10.8	34

Note: Classification metrics are calculated with respect to the common simulation-based reference labels defined in Section 3.

A) Soiling-state estimation

The proposed hybrid model achieved an SR MAE of 0.89%, representing the lowest estimation error among the evaluated predictive approaches. As reported in Table 6-(a), the corresponding MAE values were 1.74% for the SR-threshold approach, 1.21% for RF + greedy, and 1.08% for LSTM + rules. The proposed hybrid model therefore reduced the SR MAE by approximately 17.6% relative to LSTM + rules, 26.4% relative to RF + greedy, and 48.9% relative to the SR-threshold approach.

The improvement is consistent with the complementary roles of the TCN and XGBoost branches described in Section 3.3. Soiling evolution exhibits a temporal structure because deposition generally accumulates during successive dry periods, whereas rainfall can rapidly reduce surface contamination. The TCN branch captures temporal dependencies using causal and dilated temporal convolutions. In contrast, the XGBoost branch captures non-linear interactions among electrical measurements, meteorological variables, temporal statistics, rainfall history, and simple image descriptors. Their combined use provides a more effective means of distinguishing continuous soiling-induced degradation from short-term fluctuations caused by environmental or operational conditions. This is crucial for PV plants, as the determination of degradation-related decrease in efficiency can hardly be based on instantaneous measurements alone. Variations in irradiance, temperature, clouds, wind speed, and inverter operation can lead to output variations similar to soiling losses. The representation in time conveys information about the persistence and evolution of these variations, while the feature branch provides a context for environmental and operational conditions. The comparison with the SR-threshold approach provides an additional distinction. A threshold-based approach determines whether the observed or estimated soiling level exceeds a predefined threshold. In contrast, the proposed estimator provides a continuous representation of the soiling state, which can be further integrated with rainfall forecasts, energy prices, forecasting uncertainty, and operational factors.

It should be emphasized that the predictive metrics in Table 6-(a) are reported only for approaches that generate an explicit predictive representation of the soiling or cleaning state. The fixed 14-day policy is therefore excluded from this panel because it follows a predetermined maintenance schedule rather than estimating soiling severity or generating a continuous cleaning-decision score.

B) Cleaning-action classification

The second predictive evaluation concerns the ability of the model-based approaches to identify decision epochs at which cleaning is warranted. The proposed hybrid model achieved a cleaning-decision AUC of 0.983 and an F1-score of 0.949, providing the strongest classification performance among the evaluated model-based approaches. The corresponding AUC and F1 values were 0.921 and 0.884 for the SR-threshold approach, 0.956 and 0.914 for RF + greedy, and 0.941 and 0.903 for LSTM + rules, respectively (Table 5).

Relative to LSTM + rules, the proposed framework improved the AUC by approximately 4.5% and the F1-score by approximately 5.1%. Relative to RF + greedy, the corresponding improvements were approximately 2.8% and 3.8%, respectively. The proposed model therefore provides the strongest discrimination between cleaning-required and non-cleaning states while also achieving the highest balance between correctly identifying beneficial cleaning opportunities and limiting incorrect classifications.

The ROC and precision–recall curves in Fig. 4 provide a graphical assessment of the classifier's ability to discriminate between cleaning decisions. The high AUC obtained by the proposed model indicates that its decision score provides strong class separability across a range of operating thresholds. The corresponding F1-score demonstrates that the classification performance remains effective at the operational decision threshold used by the framework. Together, these metrics indicate that the proposed estimator provides a reliable basis for identifying candidate cleaning events before the economic and operational constraints are considered by the scheduling layer.

The confusion matrix in Fig. 5 serves as an additional explanation for the classification performance of the suggested model at the given clean-now decision threshold. False-positive predictions indicate when a cleaning process is thought to have a potentially good impact, while in reality the operational situation says otherwise, whereas false-negative predictions stand for the possible cleaning options that are unavailable for the classifier. The consequences of these mistakes differ: false positives may lead to high costs of water and maintenance, while false negatives will mean that dirtiness-related losses occur during the time when cleaning would bring significant benefits.

Importantly, classification results are not treated as direct maintenance requirements. Instead, the state prediction and its uncertainty are provided to the revenue-aware optimization block. The block expects energy-recovery revenue, electricity prices, rainfall forecasts, maintenance costs, resource constraints, and minimum cleaning frequency requirements. Thus, the predictive classifier identifies candidate cleaning states, while the MILP block determines when the corresponding cleaning action should be implemented.

It is important to draw a line between predictive classification and operational scheduling for interpreting the results. A high classification score means that the estimator can reliably find cleaning-related states, not that the positive prediction should automatically be followed with a cleaning activity. For instance, a moderately high soiling estimate may correspond to a high chance of rain occurring soon, and the cleaning activity may be unprofitable. Conversely, a similarly estimated state may call for cleaning if a dry period is expected and the price of electricity would mean that the energy gained after the cleaning process would be worthwhile.

The final maintenance decision therefore results from the interaction between state estimation and forward-looking economic optimization.

The operational results in Table 6-(b) are considered separately from the predictive metrics because they represent the cumulative consequences of the complete maintenance policies over the annual simulation horizon. The proposed hybrid framework achieves the highest mean annual net value of 3.09 M USD, compared with 3.01 M USD for LSTM + rules, 2.98 M USD for RF + greedy, 2.93 M USD for the SR-threshold strategy, and 2.81 M USD for the fixed 14-day schedule. These results show that the predictive advantage of the proposed model is accompanied by a corresponding improvement at the plant-operation level. The resource requirements exhibit an equally important improvement. Annual water consumption decreases from 12.2 $\times 10^3 m^3$ under

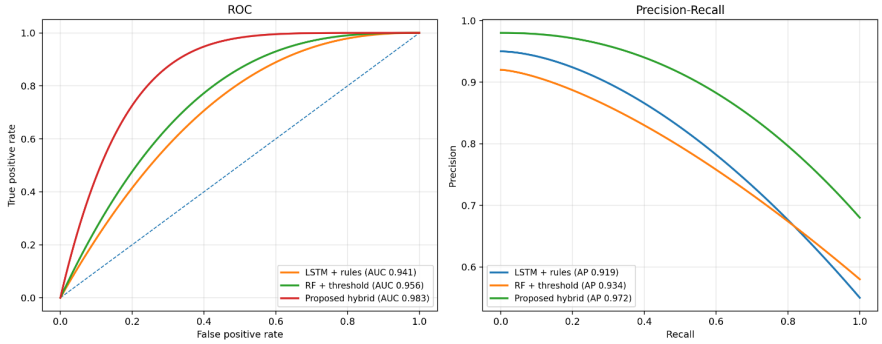

Fig. 4. ROC and precision–recall curves for cleaning-decision classification across the evaluated model-based approaches.

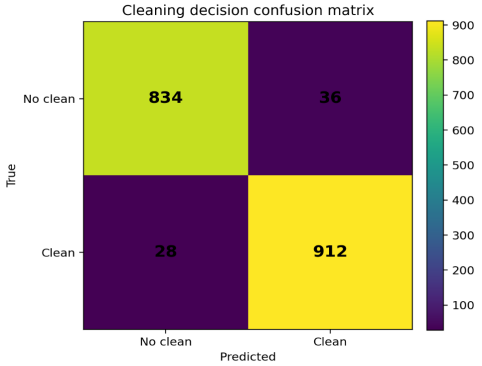

Fig. 5. Confusion matrix of the proposed hybrid model for the clean-now decision.

the LSTM + rules strategy to 10.8 $\times 10^3 m^3$ with the proposed framework, corresponding to a reduction of approximately 11.5%. Relative to the fixed cleaning strategy, for which annual water consumption reaches 18.4 $\times 10^3 m^3$, the reduction is approximately 41.3%. At the same time, the proposed strategy requires only 34 cleaning operations per year, compared with 39 for LSTM + rules and 66 for the fixed 14-day policy. The summarized results in Table 6-(b) show that the proposed framework achieves higher annual revenues while reducing water consumption and the number of cleaning interventions.

The results highlight the distinction between cleaning necessity and cleaning timing. A maintenance policy may identify a degraded state; however, it does not imply that the economically optimal period for intervention is clearly defined. The proposed framework addresses this problem by integrating the estimated soiling state with rainfall forecasts, electricity prices, predictive uncertainty, cleaning costs, and resource constraints within the rolling-horizon MILP. Thus, the obtained results allow one to determine the value of cleaning at some point in time rather than the state of soiling solely.

The findings obtained in this section confirm synergy between the stages of estimation and decision-making processes. The suggested architecture TCN–XGBoost has the minimum value of SR estimation error and the best cleaning-decision discrimination among the models being tested, while the proposed hybrid solution has the highest annual net value and the lowest level of water consumption and cleaning activity.

4.3. Revenue-aware scheduling and operational performance

Operational results indicate that the benefit of the proposed framework stems from jointly considering projected soiling accumulation, uncertainty, precipitation, electricity prices, cleaning costs, and resource constraints when determining cleaning schedules. The following analysis examines the resulting trade-offs in terms of economics and resources.

A) Economic and resource performance

The annual operational results reported in Table 6-(b) demonstrate a consistent advantage of the proposed revenue-aware strategy over the benchmark policies. The proposed hybrid framework achieves a mean annual net value of 3.09 M USD, compared with 3.01 M USD for LSTM + rules, 2.98 M USD for RF + greedy, 2.93 M USD for the soiling-ratio threshold policy, and 2.81 M USD for the 14-day fixed schedule. Relative to the strongest non-hybrid benchmark, LSTM + rules, this corresponds to an increase of approximately 2.7%, while the improvement over the fixed 14-day policy is approximately 10.0%. These results indicate that the economic benefit is not simply a consequence of replacing a conventional threshold with a more accurate predictor; rather, it emerges from coupling improved state estimation with an optimization layer that explicitly evaluates the economic consequences of alternative cleaning times.

The improvement experienced when comparing it to the benchmark known as LSTM + Rules represents a minor one, not a significant change in revenue. The 2.7% increase in revenue received is accompanied by water and cleaning savings, which indicates that the advantage obtained through this new solution is a combination of revenue-related benefits and benefits related to resource usage and efficiency. The annual net value results are presented in Table 6-(b); however, they were not collected at the level of each separate seed so that separate statistical analyses of the methods could be performed. Therefore, the study cannot make any claims regarding the statistical significance of the economic differences observed and must be interpreted taking into consideration the simulation method.

The resource requirements exhibit an equally important improvement. Annual water consumption decreases from 12.2 $\times 10^3 m^3$ under the LSTM + rules strategy to 10.8 $\times 10^3 m^3$ with the proposed framework, corresponding to a reduction of approximately 11.5%. The reduction is even more pronounced relative to the fixed cleaning strategy, for which annual water consumption reaches 18.4 $\times 10^3 m^3$. At the same time, the proposed strategy requires only 34 cleaning operations per year, compared with 39 for LSTM + rules and 66 for the fixed 14-day policy. Thus, the framework achieves higher economic value with fewer interventions and lower water consumption.

These comparative gains are summarized in Table 7, which reports the improvement of the proposed hybrid framework relative to each benchmark based on the corresponding results in Table 6-(a) and 6-(b). The comparison highlights the simultaneous reductions in estimation and classification errors, together with improvement in economic and resource-related outcomes. Because the fixed 14-day policy does not generate a predictive soiling estimate, a direct SR-MAE comparison is not applicable for this strategy.

These results are particularly relevant to utility-scale PV plants

Table 7. Relative improvement of the proposed hybrid framework over benchmark methods.

Comparison	SR MAE reduction (%)	ROC-AUC improvement (%)	F1 improvement (%)	Net annual value improvement (%)	Water reduction (%)	Cleaning reduction (%)
Proposed hybrid vs. 14-day fixed	—	22.7	23.1	10.0	41.3	48.5
Proposed hybrid vs. SR threshold	48.9	6.7	7.4	5.5	23.9	30.6
Proposed hybrid vs. RF + greedy	26.4	2.8	3.8	3.7	15.0	17.1
Proposed hybrid vs. LSTM + rules	17.6	4.5	5.1	2.7	11.5	12.8

Note: Relative improvements are calculated with respect to the corresponding benchmark value as $100\,(\mathrm{M_{hybrid}} - \mathrm{M_{baseline}})\,/\mathrm{M_{baseline}}$ for value-oriented metrics and $100\,(\mathrm{M_{baseline}} - \mathrm{M_{hybrid}})\,/\mathrm{M_{baseline}}$ for error and resource metrics. For SR MAE, "—" denotes that a direct predictive comparison with the fixed-interval policy is not conceptually applicable.

because cleaning is itself a resource-constrained maintenance activity. Increasing the number of cleaning events does not necessarily increase net revenue: each intervention consumes water and labor and may require operational downtime, while the recovered energy may have limited economic value during periods of low electricity prices or when rainfall is imminent. The proposed MILP explicitly accounts for these competing effects. Consequently, a cleaning action is selected when its expected contribution to the risk-adjusted objective is favorable under the operational constraints defined in Section 2.

An important implication is that reduced cleaning frequency should not be interpreted as lower maintenance responsiveness. The proposed method distinguishes between technically indicated cleaning and economically justified cleaning. In the case when one is predicting a large dirt accumulation and low energy recovery value, the optimizer can postpone the cleaning process. In contrast, while experiencing a similar amount of dirt and a positive fluctuation of electricity price, along with the loss of the opportunity for self-cleaning, it might be preferable to clean earlier. Such changes give the scheduler the possibility to manage the resulting operational loss at the plant and the electricity price changes.

B) Revenue–water trade-off and scheduling behavior

Fig. 6 compares the annual net value and water consumption resulting from the evaluated cleaning policies. The proposed framework achieves the highest annual net value—approximately USD 3.09 million/year—while utilizing the lowest volume of water, around 10.8 $\times 10^3 m^3$/year. In comparison, the "LSTM + rules" baseline model yields approximately USD 3.01 million/year with a consumption of 12.2 $\times 10^3 m^3$/year; this represents an economic improvement of about 2.7% and a water consumption reduction of approximately 11.5% for the proposed framework relative to that baseline. The "RF + greedy," "SR-threshold," and "fixed 14-day" strategies yield progressively lower annual net values—approximately 2.98, 2.93, and 2.81 million USD/year, respectively—while requiring about 12.7, 14.2, and 18.4 $\times 10^3 m^3$/year of water.

According to the findings, the advantages of the proposed technique arise mainly due to its timing and selectivity instead of merely increasing the frequency of cleaning actions. The moving-based scheduler modifies its cleaning schedule every six hours by using the most recent data concerning the predicted level of dirtiness, uncertainty, electricity rates, precipitation, cleaning expenses, and resources restrictions, which allows shifting cleaning actions to periods in which the obtained benefit outweighs its costs. The situation is different when dealing with fixed-interval and simple threshold strategies, as their decisions are not as much affected by the combination of temporal and financial factors.

The reduction in cleaning events—from 66 with the traditional inspection strategy to 34 with the proposed framework—confirms that maximizing annual value is not achieved by increasing cleaning frequency. Conversely, the results suggest that the proposed scheduler identifies fewer economically favorable intervention periods while simultaneously reducing water consumption. Thus, Fig. 6 supports the concept of revenue-oriented scheduling by illustrating the combined economic and resource implications of the evaluated policies.

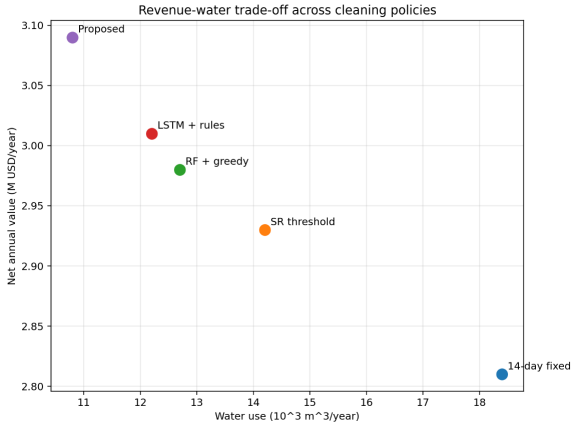

Fig. 6. Revenue–water trade-off across the evaluated cleaning policies.

4.4. Ablation, robustness, explainability, and error analysis

The previous results have shown how well the proposed framework works; however, the aggregate performance is not enough to determine if the gains were obtained from the interactions among the framework components. Therefore, other experiments were performed in order to address this challenge from three different angles: importance of the individual framework components and their parameters, robustness to certain adversities, and error analysis based on the explanation of methods used for obtaining results. Through these analyses, it was assessed whether the economic and predictive gains are due to the multimodal estimator, revenue-aware data, uncertainty-aware fusion, and constrained scheduling processes explained in Sections 2 and 3.

A) Ablation and sensitivity analysis

The ablation study demonstrates that the performance advantage of the proposed framework cannot be attributed to a single input modality or optimization component. Removing the image-derived features reduced annual net revenue by approximately 1.9% relative to the complete proposed model. Although visual data are purposefully scant in the specific design process, this finds its place in the experiments, as the results display that even irregular visual features partially contribute to different information when only electrical and physical measurements fail to distinguish between similar operational conditions. This finding is relevant to practical deployment, where continuous high-resolution images are neither assumed nor required. The second ablation test was implemented through the removal of tariff-aware economic parameters from the decision process. The resulting model showed about a 2.6% decrease in the annual net value, which proves that the timing of cleaning is economically significant and cannot be derived from condition severity. In particular, two operating states with similar estimated soiling levels may lead to different optimal actions when electricity-price conditions and the associated value of recovered energy differ. The result therefore supports the explicit incorporation of economic information into the optimization formulation rather than treating the cleaning decision as a purely condition-based classification problem.

The scheduling ablation further distinguishes the complete framework from a simplified greedy implementation. Replacing the MILP scheduler with a greedy trigger degraded the action-level performance because the latter cannot jointly account for future price conditions, anticipated rainfall, resource constraints,

and minimum cleaning-spacing requirements over the rolling horizon. The deterioration observed in this ablation analysis is consistent with the formulation in Section 2. The advantage of the proposed decision system stems not only from an accurate estimation of the plant state but also from the evaluation of alternative maintenance actions within a forward-looking, constrained optimization framework.

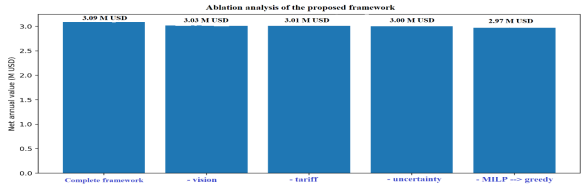

Fig. 7. Ablation analysis of the proposed framework in terms of annual net value.

Fig. 8 examines the sensitivity of the annual net value to the cleaning cost and the water scarcity penalty. An increase in either of these costs reduces the absolute economic value of cleaning, although the optimization continues to distinguish between higher- and lower-value interventions. The results indicate that the observed economic pattern is not limited to a single cost configuration within the evaluated parameter ranges.

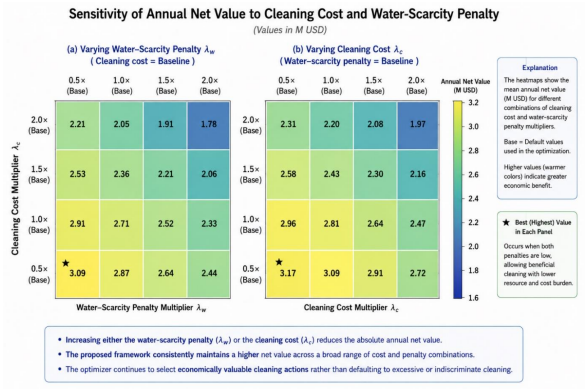

Fig. 8. Sensitivity of annual net value to cleaning cost and water-scarcity penalty.

B) Robustness to nuisance conditions

Robustness was evaluated under disruptive conditions that included cloud cover variability, partial shading, sensor drift or bias, and other disturbances unrelated to soiling. The ROC-AUC value of the proposed model decreased by 0.011 under these conditions, compared to reductions of 0.027 for RF and 0.031 for LSTM. This lower degradation indicates greater robustness in the evaluated disturbance scenarios, which is consistent with the use of temporal, multimodal, and uncertainty-aware information. Thus, the reported ROC-AUC and F1-score values evaluate the estimator's ability to identify the predefined condition of a need for cleaning, independently of the subsequent economic optimization of the scheduling.

These results should be interpreted within the context of the simulated disruptive conditions and should not be generalized to all plant failures or all field operating conditions.

C) Explainability and error analysis

The analysis of explainability was done to check if the behavior of the introduction estimator was consistent with the drivers of soiling and cleaning decisions that were significant. Global SHAP analysis led to the conclusion that estimated soiling severity was the main explanatory factor, followed by days since the last rain, tariff index, image texture, and wind speed. The obtained ranking confirms the approach of the framework proposed: accumulated contamination points to the deterioration of performance, rainfall history points out the possibilities of cleaning with the help of nature, economic factors indicate how important it is to recover the lost energy, while visual and weather-related factors reflect additional information about the state of the surface and the process of accumulation.

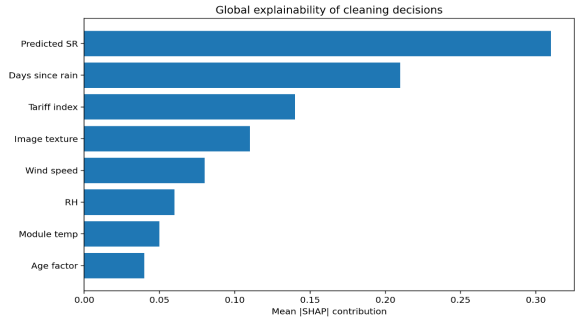

Fig. 9. Global SHAP feature attribution for the proposed hybrid soiling-estimation model.

The regional explanations give extra clarity into the workings of such factors under different operating scenarios. The system hesitated to perform cleaning in the case of medium levels of soiling combined with high probability of precipitation, as the potential gain resulting from the immediate operation was not significant enough. On the other hand, when an instance of high soiling occurred alongside hot, dry weather conditions and favorable electricity prices, cleaning took priority because the potential energy losses could be recovered. These examples illustrate a vital difference between the severity of the condition and the urgency of maintenance; the first one is a question of estimation, while the second one relies on the interplay of estimated condition, forecast, economic value, uncertainty, and operational constraints.

Error cases were concentrated in situations where electrical responses could be confused with soiling effects, such as transient cloud cover, mild diffuse soiling, and anomalous inverter operation. This supports the use of temporal features and the filtering of anomalous operations during preprocessing. Given that visual descriptors are intermittent, characterizing rapid changes occurring between inspections can also prove difficult.

The false-positive cases are limited by the economic decision-making layer. A positive cleaning suggestion from the estimation layer is not enough for cleaning to happen; within the MILP framework, it is necessary to prove that the action is economically justified. Hence, some uncertain cases on the positive side do not reach the cleaning operation level. The MI speculates that false-negative cases can mainly be found and detected during rapid environmental changes or strong winds, especially when the condition of the surface changes between the inspections. This is because visual descriptors are only taken at inspection time. These cases indicate that adaptive inspection frequency could be investigated in future field studies.

In total, the study confirms the need for the two layers – estimation and decision. SHAP makes it possible to receive an understanding of the reasons behind the estimated state, while uncertainty and money-oriented MILP are responsible for making decisions based on that estimated state.

4.5. Overall discussion, practical implications, and limitations

A) Main findings and practical implications

The results indicate that the proposed framework links the multimodal estimation of soiling accumulation with the economic

scheduling of cleaning operations under the simulated conditions. Combining state estimation, uncertainty information, precipitation forecasts, electricity prices, and resource constraints yields higher annual net value, as well as lower water consumption and cleaning frequency, compared to the benchmark policies evaluated. The performance observed across the four simulated photovoltaic plant archetypes should be interpreted as a test of robustness regarding the simulated operational scenarios considered, rather than as a proven generalization applicable to real-world photovoltaic plants. The simulations carried out did employ site-specific validation and stochastic environmental variability, as well as an independent test set, but could not cover everything that has happened in terms of module technology, tilt angle, dust, soiling accumulation, precipitation, and electricity markets. So, there has to be field measurement validation and site-specific calibration done before being able to rely on current results. Also, it should be mentioned that sparse visual descriptors may be used alongside electrical and environmental data without the need for continuous images of the objects. However, it should be noted that these results are only valid under the study framework conditions and should not be taken as performance results under real-world activities.

B) Limitations

A few limitations must be mentioned. First, the uncertainty measure is not clearly validated against actual prediction errors; as a result, the $\sigma_{i,t}$ has to be seen as a model-based indicator of uncertainty that is useful for risk-averse scheduling in the simulation context rather than as a calibrated measure of probability.

Second, the research completely relies on simulations; while four types of photovoltaic plants, the stochastic nature of the environment, site validation, and independent testing were used, the results cannot replace actual field tests with various module technologies, configurations, tilts, soil characteristics, rains, and electricity markets.

Third, the simulated methods of soiling and rainfall are abstract and do not consider local deposition and wash-off of dirt amounts. Hence, the authors do not claim that the attributed statistical parameters, ranges of variables, as well as noise inputs and rain reactions truly mirror reality. The real efficiency might be different due to the variations in the characteristics and conditions of particles and surfaces, humidity, winds, tilt angle of panels, and precipitation. Therefore, plants should receive production models that are based on calibrated testing models.

Fourth, there is no sensitivity analysis of MAE, ROC-AUC, and F1 concerning the assumptions of ground truth generation in this paper. As a result, the figures shown cannot be interpreted as firm proof of any invariance with respect to other assumptions for the generation of ground truth.

Fifth, operational planning is formulated at the panel-group level and does not explicitly model detailed cleaning logistics—such as route planning for vehicles or robots, access constraints, cleaning duration, or technology-specific efficiency. Furthermore, the visual modality employs sparse descriptors derived from images rather than continuous image streams.

Sixth, the study includes ablation tests on selected input and scheduling components—such as removing image features, omitting tariff rates, and replacing the MILP scheduler with a greedy strategy; however, separate experiments were not conducted for models based solely on TCN or XGBoost, multimodal fusion without uncertainty, or MILP without the uncertainty penalty. Consequently, it is not possible to fully isolate the individual quantitative contribution of these components based on the current results. More comprehensive ablation studies focusing on each component remain for future work.

Lastly, cybersecurity and interaction between humans and machines were taken into account as implementation requirements but were not tested experimentally. Thus, it is necessary to conduct a wide range of field measurements, site-specific calibration, detailed logistics analysis, and hardware and operational testing prior to the implementation of the actual project.

5. CONCLUSION

This study developed an integrated hybrid AI-based framework for PV soiling management that combines multimodal soiling-state estimation with uncertainty-aware, revenue-oriented cleaning scheduling. The proposed TCN–XGBoost estimator achieved the lowest soiling-ratio estimation error (MAE = 0.89%) and the highest cleaning-decision performance (ROC-AUC = 0.983; F1 = 0.949) among the evaluated model-based approaches. These predictive improvements also translated into measurable operational benefits. The proposed framework achieved a mean annual net value of 3.09 M USD while reducing annual water consumption to 10.8 $\times 10^3 m^3$ and limiting cleaning interventions to 34 events, compared with 3.01 M USD, 12.2 $\times 10^3 m^3$, and 39 events, respectively, for the LSTM + rules benchmark. The findings show that better soiling estimation alone does not suffice because of its dependence on how efficiently the estimated state and the relevant uncertainty are translated into suitable maintenance actions in economic and operational terms. Also, the analyses dealing with ablation, robustness, and explainability of results confirmed the integrated transferability of the developed framework. Furthermore, the experiment proved the positive impact of soiling estimation on decision-making quality via the use of multi-modal data, tariff-aware optimization, and constrained rolling horizon MILP. The framework also maintained stable cleaning-decision performance under the demanding conditions evaluated and revealed interpretable relationships between soiling severity, precipitation history, electricity rates, and visual indicators. Nonetheless, the results were obtained in a controlled simulation environment, and the physical, economic, rainfall, and cleaning assumptions used may not represent real PV plants in a comprehensive manner.

Thus, the subsequent research should be devoted to conducting field tests in different kinds of PV plants, using real-time sensors, images, and adaptive inspection practices for the changing soiling environment, including robotic cleaning constraints, routing and accessibility conditions, additional markets, and human practitioner verification, continuing the deployment of the suggested approach.

ACKNOWLEDGEMENT

The authors acknowledge the use of artificial intelligence tools (e.g., ChatGPT by OpenAI) for language editing and clarity improvement during the preparation of this manuscript. The authors are fully responsible for the scientific content, analysis, and conclusions.

REFERENCES

[1] B. Aïssa, R. J. Isaifan, B. W. Figgis, A. A. Abdallah, D. Bachour, D. Perez-Astudillo, A. Sanfilippo, J. Lopez-Garcia, and V. Bermudez Benito, "A comprehensive review of a decade of field pv soiling assessment in qeeri's outdoor test facility in qatar: Learned lessons and recommendations," *Energies*, vol. 16, no. 13, p. 5224, 2023.

[2] K. K. Ilse, B. W. Figgis, V. Naumann, C. Hagendorf, and J. Bagdahn, "Fundamentals of soiling processes on photovoltaic modules," *Renewable Sustainable Energy Rev.*, vol. 98, pp. 239–254, 2018.

[3] A. A. Yaghoubi, M. Gandomzadeh, A. Gholami, R. G. Ghoachani, M. Zandi, and H. A. Kazem, "Optimize photovoltaic panels cleaning scheduling framework based on variations of hourly-based active electricity pricing in the market," *Solar Energy*, vol. 275, p. 112633, 2024.

[4] C. A. Mosbah, M. Kermadi, A. Gama, and F. Yettou, "Smart and cost-effective optimisation of photovoltaic cleaning schedules," *Energy Build.*, p. 116184, 2025.

[5] H. Abdulla, A. Sleptchenko, and A. Nayfeh, "Optimizing cleaning schedules for spatially distributed photovoltaic installations with site-specific variations," *Renew. Energy*, vol. 248, p. 122971, 2025.

[6] S. S. H. Dehshiri and B. Firoozabadi, "Photovoltaic plant site selection considering dust soiling effects: A novel hybrid framework based on uncertainty and reliability with optimum cleaning schedule," *Appl. Energy*, vol. 382, p. 125252, 2025.
[7] Y. Ledmaoui, A. El Maghraoui, M. El Aroussi, and R. Saadane, "Review of recent advances in predictive maintenance and cybersecurity for solar plants," *Sensors*, vol. 25, no. 1, p. 206, 2025.
[8] A. Al-Humairi, E. Khalis, Z. A. Al Hemyari, and P. Jung, "Machine learning-based predictive maintenance for photovoltaic systems," *AI*, vol. 6, no. 7, p. 133, 2025.
[9] S. Z. Said, S. Z. Islam, N. H. Radzi, A. bin Ponniran, M. R. M. Altimania, K. S. S. Alatawi, and J. Uddin, "Model for predicting photovoltaic soiling loss and degradation of pv electrical performance using optimized machine learning," *Results Eng.*, p. 110209, 2026.
[10] U. Naeem, K. Chadda, S. Vahaji, J. Ahmad, X. Li, and E. Asadi, "Aerial imaging-based soiling detection system for solar photovoltaic panel cleanliness inspection," *Sensors*, vol. 25, no. 3, p. 738, 2025.
[11] S. Ahmed, H. Rashid, Z. Qadir, Q. Tayyab, T. Senjyu, and M. Elkholy, "Deep learning-based recognition and classification of soiled photovoltaic modules using halcon software for solar cleaning robots," *Sensors*, vol. 25, no. 5, p. 1295, 2025.
[12] F. Brunel, R. López, F. García, E. Peters, and G. Funes, "Estimation of the power loss of a soiled photovoltaic panel using image analysis techniques," *Energies*, vol. 18, no. 18, p. 4889, 2025.
[13] A. Shukla, R. K. Pachauri, R. Walia, and V. Gupta, "Machine learning-based prediction of soiling losses in photovoltaic modules under different cleaning frequencies: an experimental investigation," *Sci. Rep.*, 2026.
[14] S. Sattuf, S. A. Celtek, and F. Shahnia, "Visual deep learning-based soiling detection on photovoltaic panels with inverter-level energy validation and sustainability-aware cleaning decision support," *Sustainability*, vol. 18, no. 8, p. 4123, 2026.
[15] F. El-Batta, A. Aboghrara, R. El-Khozondar, H. El-Khozondar, and Y. Nassar, "Ai-based monitoring of solar panels in desert environments: Distinguishing dust accumulation for fault detection," *Wadi Alshatti Univ. J. Pure Appl. Sci.*, pp. 340–353, 2026.
[16] S. De, B. Ghosh, N. Shiradkar, and A. Kottantharayil, "A data-driven approach to automate cleaning event detection in pv systems for accurate quantification of soiling loss across varying installation capacities," *Solar Energy*, vol. 308, p. 114384, 2026.
[17] A. Al-Humairi, E. Khalis, Z. A. A. Hemyari, and P. Jung, "A two-stage comparative framework for predicting photovoltaic cleaning schedules: Modeling and comparisons based on real and simulated data," *Appl. Sci.*, vol. 16, no. 6, p. 2976, 2026.
[18] B. M. Khan, A. Wadood, H. Albalawi, S. Khan, A. M. Alatwi, and O. H. Albalawi, "A hybrid ensemble-based intelligent decision framework for risk-aware photovoltaic panel soiling detection and cleaning," *Electron.*, vol. 15, no. 6, p. 1192, 2026.
[19] H. An, "Reinforcement learning-based dynamic cleaning scheduling framework for solar energy system," *arXiv preprint arXiv:2603.07518*, 2026.
[20] R. Aman, A. Ali, K. A. Memon, F. Baig, and M. A. Khan, "An enhanced approach to cost-effective cleaning optimisation for enhanced solar pv power production in arid areas," *Eng. Res. Express*, vol. 8, no. 13, p. 135305, 2026.
[21] L. Guanghua, S. H. H. Shah, R. Arshad, and N. Jiahua, "Machine learning-based optimization of photovoltaic panel cleaning schedules in arid environments: A case study of ryadh, saudi arabia," *Mehran Univ. Res. J. Eng. Technol.*, vol. 45, no. 2, pp. 08–18, 2026.
[22] J. I. A. Tsanakas, J. P. Rakotoniaina, J. Montoya, D. O. Soza, E. Pilat, A. Marzo, V. D. Campo, E. V. Fuentealba, M. Gaete, J. Aimé, *et al.*, "Assessment of soiling dynamics and cleaning efficiency for photovoltaic modules under different dust environments," *Solar RRL*, vol. 10, no. 6, p. e202500792, 2026.
[23] S. Sharma, A. Singh, and S. Sinha, "Evaluating the energy and economic effects of soiling and cleaning dynamics in urban solar pv systems," *IEEE Trans. Ind. Appl.*, 2026.
[24] A. Azamian, B. Rezaeealam, T. Ghanbarih, and E. Rokrok, "Low voltage ride-through improvement of a two-stage grid-connected photovoltaic system by using an optimized spwm technique," *J. Oper. Autom. Power Eng.*, vol. 14, no. 1, pp. 60–69, 2026.
[25] A. A. AREMU, A. O. FASINA, and O. T. AFOROLAGBA-BALOGUN, "Design and performance evaluation of an intelligent condition-based automated solar photovoltaic panel cleaning system," *Harvard Int. J. Eng. Res. Technol.*, 2026.
[26] I. ALRubeei, S. Idi, I. Alsammak, H. AlRikabi, H. Mutar, and A. Alaidi, "Using artificial intelligence for enhancement of solar cell efficiency in the south of iraq," 2025.
[27] L. Micheli, M. Theristis, A. Livera, J. S. Stein, G. E. Georghiou, M. Muller, F. Almonacid, and E. F. Fernández, "Improved pv soiling extraction through the detection of cleanings and change points," *IEEE J. Photovolt.*, vol. 11, no. 2, pp. 519–526, 2021.
[28] M. Coello and L. Boyle, "Simple model for predicting time series soiling of photovoltaic panels," *IEEE J. Photovolt.*, vol. 9, no. 5, pp. 1382–1387, 2019.
[29] J. G. Bessa, L. Micheli, J. Montes-Romero, F. Almonacid, and E. F. Fernández, "Estimation of photovoltaic soiling using environmental parameters: A comparative analysis of existing models," *Adv. Sustain. Syst.*, vol. 6, no. 5, p. 2100335, 2022.
[30] Z. Wang, Z. Xu, B. Liu, Y. Zhang, and Q. Yang, "A hybrid cleaning scheduling framework for operations and maintenance of photovoltaic systems," *IEEE Trans. Syst., Man, Cybern.: Syst.*, vol. 52, no. 9, pp. 5925–5936, 2021.
[31] M. M. H. Shihab, R. N. Sajjad, and M. R. Khan, "Effects of local economy and seasonal cleaning cycle on yield and profit of soiled solar farms," *IEEE Trans. Syst., Man, Cybern.: Syst.*, vol. 14, no. 4, pp. 669–678, 2024.
[32] M. González-Castillo, P. Navarrete, T. Tapia, Á. Lorca, D. Olivares, and M. Negrete-Pincetic, "Cleaning scheduling in photovoltaic solar farms with deterministic and stochastic optimization," *Sustain. Energy, Grids Netw.*, vol. 36, p. 101147, 2023.
[33] C. Zhang, Y. Ma, G. Yang, and T. Chen, "An integrated industrial pv panel cleaning recommendation system for optimal dust removal," *Appl. Energy*, vol. 377, p. 124692, 2025.
[34] M. Fang, W. Qian, T. Qian, Q. Bao, H. Zhang, and X. Qiu, "Dgimnet: A deep learning model for photovoltaic soiling loss estimation," *Appl. Energy*, vol. 376, p. 124335, 2024.