Journal of Solar Energy Research (JSER)

Journal homepage: www.jser.ut.ac.ir

A Hybrid AI Decision Framework for Predictive Maintenance of Photovoltaic Systems with Formalized Human-in-the-Loop Co-Learning

Amjad Ech-Charqaouy[a,*], Sidi Salah Ech-Charqaouy[a], Abdelkader Boulezhar[a], Redouane Mihramane[a], Nizar Ech-Charqaouy[a]

[a] *Department of Physical Sciences, Faculty of Sciences Ain Chock (FSAC), Hassan II University of Casablanca, Morocco*

ARTICLE INFO

Article Type:

Research Article

Received:2026.05.11
Accepted in revised form:2026.11.05

Keywords:
Predictive maintenance;
Human-in-the-loop learning;
Multi-criteria decision support;
Reinforcement learning;
Photovoltaic systems

A B S T R A C T

Predictive maintenance of utility-scale photovoltaic (PV) systems operating in harsh desert environments requires decision frameworks that simultaneously handle operational uncertainty, multi-criteria trade-offs, and limited human supervision. This paper proposes a hybrid Multi-Criteria Decision Support System (MCDSS) formulating PV maintenance as a sequential decision-making problem under uncertainty. The framework integrates fuzzy logic inference, an Adaptive Decision Tree (ADT), a Q-learning reinforcement module, and a human-in-the-loop (HITL) co-learning mechanism governed by three operator feedback types: Confirm, Modify, and Override. Validated on semi-real operational data from the Noor Boujdour I solar plant (Morocco), the system achieved an average reaction time below 8 s, a corrected availability rate of 97.3%, and a post-adaptation decision adequacy rate of 92.3% (N = 142; Cohen's $\kappa = 0.83$, 95% CI [0.76, 0.90]). McNemar tests showed significant gains over rule-based, Random Forest, and ML baselines ($p < 0.05$), but not over standalone Q-learning ($p = 0.343$). Ablation analysis identifies HITL co-learning as the largest single contributor to performance gain (−4.4 pp without HITL, $p = 0.012$, McNemar). The proposed architecture demonstrates promising adaptability for renewable-energy maintenance applications, subject to further multi-season validation.

*Corresponding Author echcharqaouyamjad@gmail.com

Cite this article: Ech-Charqaouy, A., Ech-Charqaouy, S. S., Boulezhar, A., Mihramane, R. & Ech-Charqaouy, N. (2026). A Hybrid AI Decision Framework for Predictive Maintenance of Photovoltaic Systems with Formalized Human-in-the-Loop Co-Learning. Journal of Solar Energy Research, 11(4), 3413–3439. https://doi.org/10.22059/jser.2026.413937.1738

DOI: 10.22059/jser.2026.413937.1738

©The Author(s). Publisher: University of Tehran Press.

1. Introduction

The large-scale deployment of photovoltaic power plants has intensified the need for intelligent decision-making frameworks capable of ensuring reliable operation under uncertain, dynamic, and resource-constrained conditions. Although photovoltaic systems play a strategic role in decarbonized energy infrastructures, their operation in desert environments, such as utility-scale installations in the Moroccan Sahara, introduces severe physical and operational constraints, including extreme irradiance, strong thermal gradients, abrasive dust deposition, and limited site accessibility [1, 2]. These factors accelerate component degradation and inverter aging, leading to non-stationary system behavior, unplanned curtailments, and increased maintenance costs [2–5].

Within this context, predictive maintenance has emerged not only as an asset-management strategy but also as a sequential decision-making problem under uncertainty, where maintenance actions must be continuously adapted to evolving system states and operational constraints. By exploiting real-time sensing and data-driven analytics, predictive maintenance aims to anticipate failures, reduce corrective interventions, and optimize lifecycle performance [3, 6–10]. The widespread adoption of Internet-of-Things architectures has further enabled this paradigm by supporting streaming data acquisition, distributed sensing, and continuous monitoring across geographically extended photovoltaic installations [11–13]. In parallel, advances in artificial intelligence and large-scale data analytics have significantly improved fault detection and performance forecasting, particularly through machine-learning-based approaches [7, 12] and deep-learning methods applied to PV module defect detection and degradation root-cause identification [14, 15].

However, translating these advances into operational maintenance decisions remains challenging. Maintenance planning inherently involves multi-criteria trade-offs between fault severity, intervention costs, accessibility constraints, and environmental risks, all of which evolve over time. Decision support systems, and particularly multi-criteria decision support systems, have therefore been widely explored to structure such trade-offs in photovoltaic operation and maintenance [1, 4, 16]. A wide range of multi-criteria decision-making methods has been applied to energy planning and maintenance resource allocation problems [17–21]. Nevertheless, their application to real-time photovoltaic maintenance remains limited, as most implementations rely on static weighting schemes and lack mechanisms for adaptive learning and feedback integration in non-stationary environments [18, 22].

Recent developments in AI-based decision intelligence offer promising avenues to overcome these limitations. Hybrid approaches that combine fuzzy reasoning, decision tree models, and reinforcement learning offer complementary strengths for uncertainty handling, interpretability, and adaptive policy refinement [23–25]. Fuzzy logic supports linguistically interpretable reasoning under ambiguous or incomplete information [23, 24], while reinforcement learning, particularly Q-learning, introduces an adaptive mechanism capable of improving maintenance decisions through iterative interaction with evolving system states [23–25], and has recently been applied to the prediction of long-term PV module degradation [26]. Despite this potential, most existing frameworks have been validated only on simulated datasets, limiting their relevance for utility-scale photovoltaic plants exposed to real Saharan operating conditions [3–5, 8].

Furthermore, although photovoltaic fault detection and performance assessment have been extensively investigated [5, 14–16], the explicit integration of human expertise within adaptive decision loops remains insufficiently addressed. Human-in-the-loop mechanisms are critical for operational trust, transparency, and knowledge transfer, particularly in safety-critical and infrastructure-constrained environments [27–29]. Only a limited number of studies incorporate operator feedback as an active component of AI-driven maintenance decision systems [29], despite growing evidence from the interactive reinforcement learning literature that human-guided feedback substantially improves policy convergence and decision reliability [27, 28]. Moreover, user acceptance of multi-criteria decision frameworks strongly depends on participatory mechanisms and interpretable recommendations [22], a dimension that existing photovoltaic maintenance systems have largely neglected.

To address these gaps, this work proposes a hybrid artificial intelligence decision framework for predictive maintenance of photovoltaic power plants operating in remote Saharan environments. The main contributions of this study are threefold:

1)An interpretable and adaptive hybrid decision architecture is developed by combining fuzzy reasoning for uncertainty modeling, adaptive decision trees for transparent and incrementally updatable decision logic, and Q-learning for reinforcement-driven policy refinement [23–25].

2)A formalized human-in-the-loop co-learning mechanism is embedded within the decision process, providing explicit algorithmic update rules for each operator feedback type (Confirm, Modify, Override), grounded in the policy shaping and interactive reinforcement learning literature [27–29].

3) The proposed framework is validated through semi-real operational experiments at the Noor Boujdour I solar facility, enabling reproducible evaluation under realistic streaming data and accessibility constraints, with quantitative assessment of interpretability and inter-rater agreement.

By jointly emphasizing interpretability, adaptability, and robustness under operational constraints, the proposed architecture provides a scalable foundation for intelligent maintenance decision-making in harsh and infrastructure-limited environments [3, 11, 12]. Section 2 presents the proposed methodology. Section 3 describes the field deployment and evaluation. Section 4 discusses the results and benchmarking. Section 5 concludes the paper and outlines future directions.

2. Materials and Methods

The proposed Multi-Criteria Decision Support System (MCDSS) is designed as a modular, data-centric framework that combines IoT-enabled sensing, hybrid artificial-intelligence reasoning, and human-in-the-loop (HITL) feedback to deliver predictive maintenance for photovoltaic (PV) assets operating under the extreme environmental conditions of the Moroccan Sahara [1–3]. Its architecture is organized around four interdependent layers designed to ensure scalability, interpretability, and robustness across heterogeneous PV infrastructures, as illustrated in Figure 1: (i) real-time data acquisition and preprocessing; (ii) multi-criteria modeling of technical, economic, and logistical constraints; (iii) a hybrid AI decision engine for maintenance prioritization; and (iv) a human-supervised feedback mechanism enabling contextual adjustment and continuous learning [1, 6, 7, 29].

Following the formal problem specification advocated in [30], the framework defines an explicit data model. An *observation* at time t is a vector $x(t) \in \mathbb{R}^d$, where $d = 6$ is the number of engineered features extracted from raw sensor streams: operating temperature T_{op}, failure probability P_f, fault severity S_{ev}, elapsed time since last intervention Δt, projected maintenance cost C_m, and component accessibility index A_c. The target variable $y \in$ {Corrective, Preventive, Deferred, No-Action} represents the recommended maintenance action category. The Maintenance Priority Score (MPS) $\in [0, 100]$ is a scalar function of $x(t)$ used to rank concurrent maintenance requests before dispatching to the decision engine [1, 16, 30].

This layered design reflects emerging paradigms in Energy Informatics, where IoT connectivity [11, 31–33], data analytics [6–8], and AI-assisted decision-making [23–25] converge to support autonomous yet transparent maintenance frameworks. It also aligns with research trends advocating self-learning, adaptive infrastructures for renewable-energy systems [5, 8, 29], marking a strategic shift from reactive supervision toward anticipatory, data-driven operation and maintenance (O&M) intelligence.

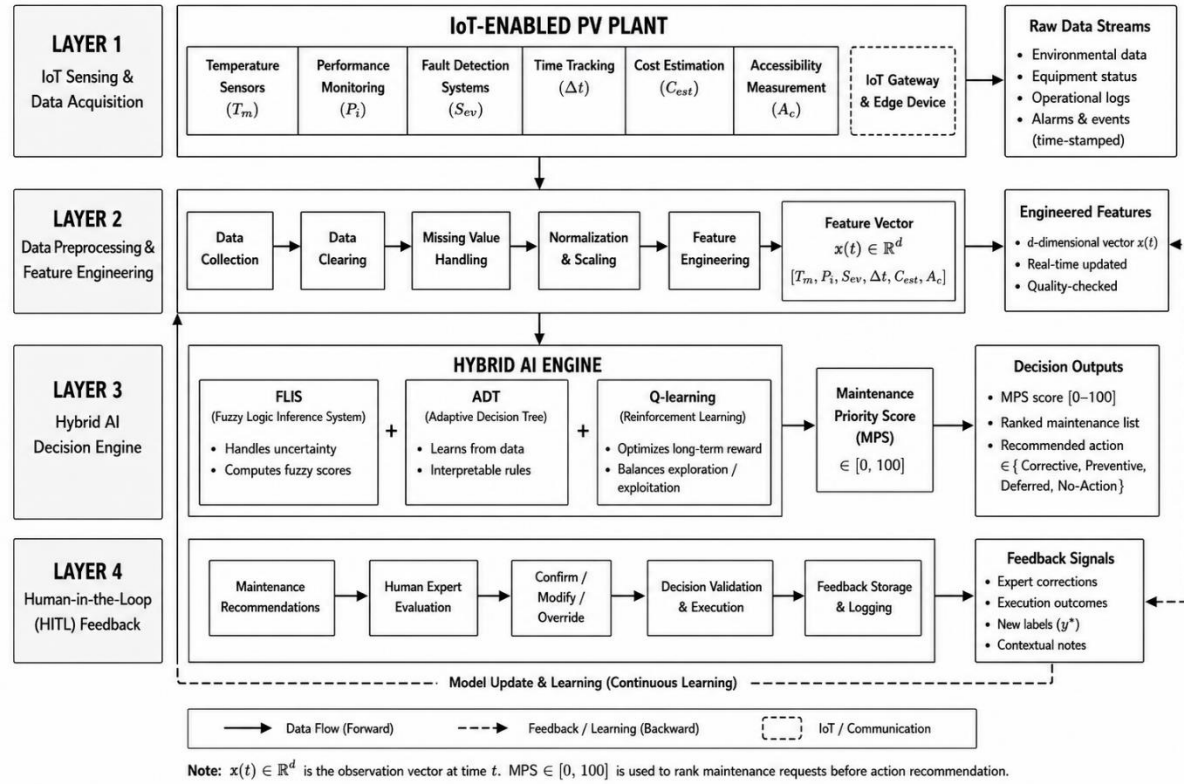


Figure 1. Functional architecture of the proposed real-time MCDSS integrating IoT sensing, preprocessing, hybrid AI reasoning, and human-in-the-loop feedback [1, 6, 7, 29]

2.1. System Architecture and Data Model

At the sensing layer, a distributed network of IoT nodes continuously monitors thermal gradients, electrical imbalances, and environmental conditions including irradiance, ambient temperature, and dust deposition [1, 2, 11, 32, 33]. Raw measurements are sampled at 1-minute intervals across M inverter units and aggregated into sliding windows of width $W = 15$ minutes, yielding time-stamped observation vectors $x(t)$ as defined above. Data streams are transmitted through lightweight industrial messaging and integration tools (MQTT, Node-RED, and Apache Kafka), selected for their low latency, reliability, and interoperability [6, 31].

A preprocessing module performs temporal synchronization, Z-score-based outlier suppression (threshold $|z| > 3.0$), and min-max normalization of each feature dimension to [0, 1], ensuring data integrity despite the noise and variability inherent to Saharan PV environments [3, 7]. Missing values arising from communication dropouts are imputed using forward-fill for intervals shorter than 5 minutes and flagged as uncertain beyond that threshold, consistent with practices reported in large-scale PV O&M systems [4, 12]. The preprocessing pipeline supports modular integration of anomaly detection and clustering components, enabling scalable deployment across multi-site PV installations [7, 14].

The modular structure facilitates rapid adaptation to site-specific parameters (inverter type, layout configuration, and communication bandwidth) [16, 34], which is particularly relevant for large-scale infrastructures such as Noor Boujdour I, where heterogeneous equipment and harsh climatic conditions demand highly adaptive monitoring [2, 4]. Table 1 summarizes the six feature variables constituting $x(t)$, their physical interpretation, measurement units, and normalization ranges.

Table 1. Formal Feature Vector x(t) – Six Input Dimensions of the MCDSS Data Model

Index	Feature	Physical Meaning	Unit	Range	Normalization
x_1	T_{op}	Module operating temperature	°C	[20, 85]	min-max → [0,1]
x_2	P_f	Failure probability (ML-estimated)	(-)	[0, 1]	Direct
x_3	S_{ev}	Fault severity score (expert-labeled)	(-)	[0, 10]	÷ 10
x_4	Δt	Elapsed time since last intervention	days	[0, 365]	÷ 365
x_5	C_m	Projected maintenance cost	USD	[50, 5000]	min-max → [0,1]
x_6	A_c	Component accessibility index	(-)	[0, 1]	Direct

Note: Label space y ∈{Corrective, Preventive, Deferred, No-Action}. MPS = f(x(t)) ∈ [0, 100], computed as the weighted sum defined in Section 2.2 [1, 16, 30].

The five-stage closed-loop workflow (Acquire → Preprocess → Infer → Validate → Update) transforms raw sensor inputs into operator-validated maintenance decisions through iterative model refinement [1, 6, 7].

2.2. Multi-Criteria Decision Modeling

Maintenance prioritization in desert-based PV plants requires a holistic evaluation framework extending beyond classical fault detection [1, 16]. The MCDSS computes the Maintenance Priority Score as a weighted linear combination of the six normalized feature dimensions:

$$MPS(t) = \sum_i w_i \cdot x_i(t),\ i \in \{1, \ldots, 6\},\ \sum_i w_i = 1,\ w_i \geq 0 \quad (1)$$

where the weight vector $w = [w_1, \ldots, w_6]^T$ encodes the relative importance of technical, economic, and logistical criteria. As shown in Figure 2, decision criteria are hierarchically structured into three principal domains:

- Technical criteria (T_{op}, S_{ev}): describe system stress and intervention urgency;
- Economic criteria (P_f, C_m): reflect financial implications of deferred actions;
- Logistical criteria (Δt, A_c): determine the feasibility of corrective actions under Saharan field constraints.

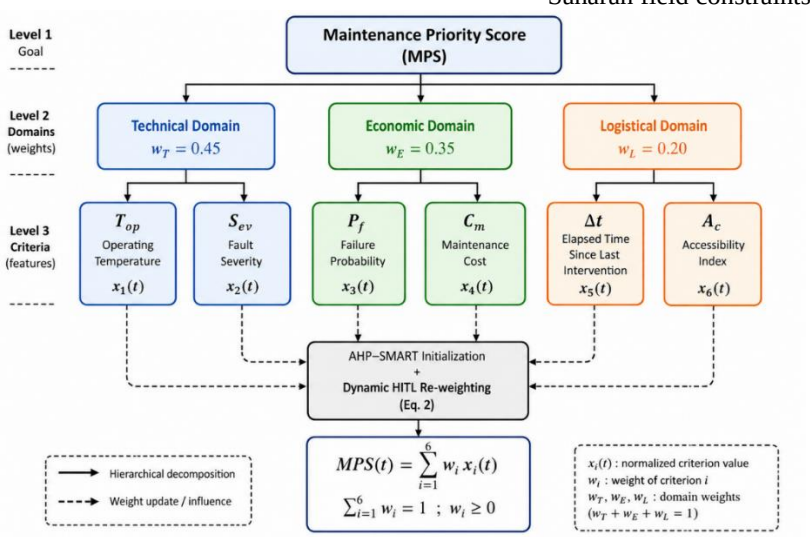


Figure 2. Hierarchical decision structure integrating technical, economic, and logistical dimensions for maintenance prioritization [1, 16, 17, 35]

Initial domain weights are derived through a hybrid Multi-Criteria Decision-Making strategy combining the Analytic Hierarchy Process (AHP) and the Simple Multi-Attribute Rating Technique (SMART), supplemented by heuristic weighting derived from on-site observations at Noor Boujdour I [17, 18, 35]. This combination ensures contextual relevance, as logistical constraints and fault severity often outweigh thermal thresholds in Saharan environments [2, 4, 7].

To maintain adaptability, the weight vector *w* is dynamically updated through a feedback-driven re-weighting mechanism. Let *k* index operator interaction events. At each event *k*, the operator provides feedback *f*(k) ∈ {Confirm, Modify, Override}. The weight update rule is:

$$w(k+1) = w(k) + \eta \cdot \Delta w\big(f(k), x(t_k), \hat{y}(t_k)\big) \quad (2)$$

where η is the MPS weight learning rate (η = ηw = 0.02), $\hat{y}(t_k)$ is the system recommendation at time t_k, and Δw(·) is a feedback-specific update vector computed as follows: (i) Confirm: Δw = 0 (no change); (ii) Modify: Δw reinforces the domain weight of the criterion dimension most discrepant from the operator's preferred action (gradient from expert correction); (iii) Override: Δw applies a stronger correction proportional to the MPS gap between ŷ and the operator's chosen action [22, 27–29]. This co-adaptive strategy continuously realigns maintenance priorities with observed operating dynamics [23–25, 34].

2.3. Hybrid AI Decision Engine

The MCDSS core is an intelligent decision engine that ranks maintenance actions under the uncertain, non-stationary, and data-limited conditions characteristic of desert PV plants [3, 8, 25]. The engine integrates three complementary reasoning paradigms – Fuzzy Logic Inference System (FLIS), Adaptive Decision Tree (ADT), and Q-learning – to balance interpretability, adaptability, and computational efficiency [23–25].

2.3.1. Fuzzy Logic Inference System (FLIS)

The FLIS layer converts continuous sensor readings into linguistic descriptors using triangular membership functions $\mu_A(x) = max(0, min((x-a)/(b-a), (c-x)/(c-b)))$, where [a, b, c] define the support, core, and right boundary of each fuzzy set A [23, 24]. Fuzzy rules of the form IF (T_{op} is High) AND (S_{ev} is Critical) THEN (MPS_contribution is Very_High) are evaluated using Mamdani inference, and the resulting fuzzy outputs are defuzzified via centroid method. This step enhances robustness to diurnal variability and incomplete sensor data, both of which are common in desert PV operations [2, 4].

2.3.2. Adaptive Decision Tree (ADT)

The ADT layer provides transparent, traceable decision logic through a batch-retrained interpretable classification tree (scikit-learn DecisionTreeClassifier, Gini, max_depth = 5).

The ADT is periodically retrained whenever $N_{batch} = 20$ newly labeled HITL observations accumulate, ensuring continued alignment between expert-derived decision rules and evolving operational patterns [2, 4, 7].

Its low computational overhead and structural transparency make it well suited for edge deployment; it directly supports the HITL update mechanism detailed in Section 2.4.

2.3.3. Q-learning Reinforcement Module

The Q-learning module refines maintenance policies through iterative interaction with the operational environment [23–25]. The state space is defined as

S = {s = (b_1 , b_2 , b_3 , b_4 , b_5 , b_6) | b_i ∈ {Low, Medium, High}, i=1, …, 6},

where each dimension of *x*(t) is discretized into three bins based on domain thresholds, yielding |S| = 3^6 = 729 states.

The action space is A = {Corrective, Preventive, Deferred, No-Action}, with |A| = 4, giving a Q-table of size 729 × 4 = 2,916 entries. The reward function is defined as:

$$r(s,a) = \alpha_T \cdot I\big(T_{op} > 75°C\big) + \alpha_S \cdot S_{ev} - \alpha_C \cdot C_m - \alpha_D \cdot I(a = Deferred \wedge P_f > 0.7) \quad (3)$$

where I(·) is the indicator function and {α_T, α_S, α_C, α_D} are domain-tuned penalty coefficients (α_T = 0.35, α_S = 0.30, α_C = 0.20, α_D = 0.15), reflecting the operational priority of thermal and severity anomalies over cost minimization in Saharan PV contexts. Higher penalties are therefore assigned to undetected thermal anomalies and to the deferral of high-risk interventions, consistent with the risk profile of utility-scale desert plants [2, 4]. The Q-value update follows the Bellman optimality equation:

$$Q(s,a) \leftarrow Q(s,a) + \alpha[r(s,a) + \gamma \max_{a'} Q(s',a') - Q(s,a)] \quad (4)$$

Training follows an ε-greedy exploration policy with $\varepsilon_0 = 0.3$ decaying by factor $\varepsilon_{decay} = 0.995$ per episode, learning rate α = 0.2, and discount factor γ = 0.85. Ten-fold cross-validation confirmed convergence after approximately 300 episodes per scenario [23, 24]. Computational complexity: Q-table lookup is O(1) with O(2,916) storage; FLIS inference is O(R × F); ADT inference is O(d × log n) where d = 5; all operations are sub-millisecond and compatible with edge deployment. Training parameters and convergence metrics are reported in Table 2.

Table 2. Q-Learning Formulation: State Space, Action Space, and Training Parameters

Parameter	**Symbol**	**Value**	**Description**
State space size	|S|	729	3 bins × 6 features = 3^6 discrete states
Action space size	|A|	4	Corrective / Preventive / Deferred / No-Action
Q-table size	|S|×|A|	2,916	Total learnable Q-value entries
Learning rate	α	0.2	Update magnitude per iteration
Discount factor	γ	0.85	Weight of future rewards
Initial exploration	ε_0	0.3	Probability of random action selection
Exploration decay	ε_{decay}	0.995	Per-episode reduction of exploration rate
Episodes/scenario	N_e	300	Training iterations before convergence
Validation	-	10-fold CV	Internal cross-validation protocol
Convergence accuracy	-	98%	Achieved after ≈ 300 episodes

Reward coefficients: α_T = 0.35 (thermal), α_S = 0.30 (severity), α_C = 0.20 (cost), α_D = 0.15 (deferral risk). Q-table initialized to zero; ε-greedy policy with exponential decay [23–25].

2.3.4. Two-Phase Training Protocol: Offline Pre-Training and Online HITL Adaptation

The Q-learning agent operates under a two-phase training protocol designed to separate initial policy acquisition from operational refinement. In the offline pre-training phase, the Q-table is initialized to zero and trained exclusively on the 350-sample training subset (the training partition of the stratified 70/30 split of the full 500-sample annotated dataset [30]) over $N_e = 300$ episodes per fault scenario, using the ε-greedy exploration policy described above. The 150-sample HITL/post-adaptation evaluation subset was held out from the outset and was not used at any stage of this offline pre-training phase. This offline phase produces a baseline policy defined over the 729 discrete states of S before any field deployment (states not visited during offline training retain their zero initialization and are refined online through HITL feedback), ensuring that the system does not enter operation in a fully uninformed state. Convergence is confirmed by ten-fold cross-validation performed exclusively within the 350-sample training subset, with the 150-sample HITL/post-adaptation evaluation subset entirely excluded from all folds, yielding a mean classification accuracy of 98% across folds, with inter-fold variance below 0.8 pp, indicating stable performance of the tabular Q-table across the internal cross-validation folds and the discrete state partitions induced by the six feature bins [23, 24, 30]. The HITL/post-adaptation evaluation subset was subsequently used for the sequential HITL deployment described in Section 3.2, and the metrics reported in Section 3 (Cohen's κ, DAR, and related statistics) characterize system performance at the end of this online adaptation phase, consistent with an evaluation of the fully deployed, continuously updated framework (see Section 4.4). The convergence trajectory is illustrated in Figure 4, showing the evolution of cumulative reward and 10-fold CV accuracy over the 300 training episodes.

In the online incremental adaptation phase, the Q-table is updated in real time via the HITL co-learning mechanism formalized in Algorithm 1 (Section 2.4). Each operator feedback event – Confirm, Modify, or Override – triggers a targeted Q-value correction at the current state-action pair (s, ŷ), using the reward signals $r_{confirm} = +0.5$, $r_{modify} = -0.3$, and $r_{override} = -1.0$ respectively, without restarting the episode counter or resetting the exploration rate. This incremental online update is strictly bounded: only the Q-entries corresponding to the visited state s and the recommended or corrected actions ŷ and y* are modified, preserving the stability of the remaining

2,914 Q-table entries. The learning rate α = 0.2 is kept constant across both phases to ensure a uniform update magnitude [23–25]. This two-phase design – offline initialization followed by HITL-driven online refinement – reflects established practice in interactive reinforcement learning, in which a stable prior policy is necessary to generate meaningful recommendations before operator feedback can improve system performance [27, 28].

2.3.5. Computational Cost and Edge Deployment

From a deployment standpoint, the three paradigms are computationally lightweight: FLIS inference requires approximately 2 ms on a standard CPU (R = 48 rules, F = 6 features); ADT inference requires less than 1 ms (max depth d = 5); Q-table lookup is O(1) requiring less than 1 MB RAM. The full MCDSS inference pipeline operates in under 5 ms per observation on an industrial ARM gateway, well within the 8-second ART target. This confirms edge deployment feasibility for remote Saharan sites with intermittent connectivity [11, 12]. By integrating these paradigms, the MCDSS achieves a synergistic balance: FLIS handles input ambiguity, ADT ensures transparent reasoning, and Q-learning provides adaptive policy refinement [6, 23, 24]. Figure 3 illustrates the Q-learning decision loop; the complete integration of all three modules within the human-supervised feedback architecture is detailed in Section 2.4 and Figure 5.

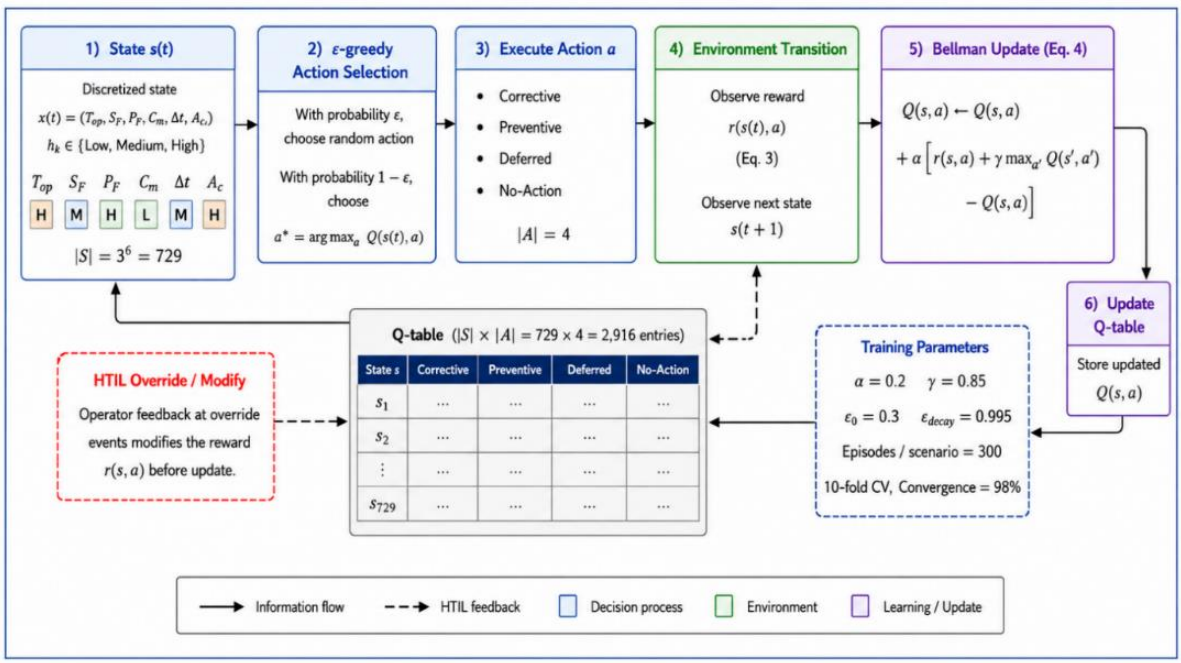


Figure 3. Q-learning decision loop illustrating state discretization, reward assignment (Eq. 3), and Q-value update (Eq. 4) [23–25]

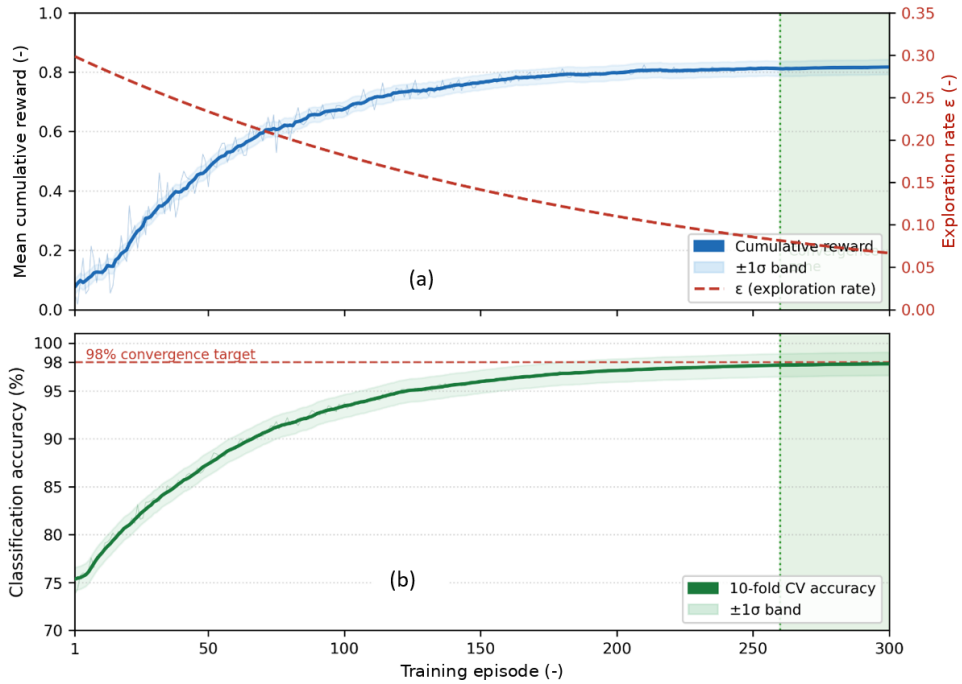


Figure 4. Q-learning convergence: (a) cumulative reward and ε-greedy exploration decay; (b) 10-fold cross-validation accuracy over 300 training episodes ($\alpha = 0.20$, $\gamma = 0.85$, $\varepsilon_0 = 0.30$)

2.4. Human-in-the-Loop Feedback Mechanism

The long-term effectiveness of predictive-maintenance systems depends fundamentally on operator engagement and on mutual learning between human expertise and automated intelligence [22, 27–29]. Drawing on the TAMER framework [27] and policy shaping paradigms [28], the MCDSS formalizes operator feedback as a direct signal in the learning loop, rather than a post-hoc correction layer. Active learning principles [36] further guide the selection of high-uncertainty cases for operator annotation, thereby maximizing the information gain per interaction.

The HITL mechanism handles three feedback types – Confirm, Modify, Override – each triggering a specific, algorithmically defined update across the three AI components. Algorithm 1 provides the formal pseudocode.

Algorithm 1: HITL Co-Learning Update Rules

Input : observation x(t), system recommendation ŷ, feedback f ∈ {Confirm, Modify, Override}

Output: updated FLIS thresholds θ_F, ADT parameters θ_D, Q-table Q

1: if f = Confirm then
2: $\theta_F \leftarrow \theta_F$ // No update – recommendation accepted
3: $\theta_D \leftarrow \theta_D$
4: $Q(s,\hat{y}) \leftarrow Q(s,\hat{y}) + \alpha \cdot r_{confirm}$ // $r_{confirm} = +0.5$
5: else if f = Modify then
6: // Operator provides corrected action y* and adjusted priority
7: $\Delta b \leftarrow x(t) - \mu_{\hat{y}}$ // Discrepancy from fuzzy centroid of ŷ
8: $\theta_F[dim] \leftarrow \theta_F[dim] + \eta_F \cdot \Delta b[dim]$ // Update fuzzy boundaries
9: // on the dimension most responsible for error
10: B.append(x(t), y^*) // if |B|>=20: ADT.retrain(B); B.clear() [30]
11: $Q(s,\hat{y}) \leftarrow Q(s,\hat{y}) + \alpha \cdot r_{modify}$ // $r_{modify} = -0.3$
12: $Q(s,y^*) \leftarrow Q(s,y^*) + \alpha \cdot r_{correct}$ // $r_{correct} = +0.3$
13: else // f = Override
14: $\theta_F \leftarrow \theta_F + \eta_F \cdot sign(y^* - \hat{y}) \cdot |MPS(y^*) - MPS(\hat{y})|$ // Strong shift

15: ADT.retrain({x(t), y^*}) // Forced node split / relabel
16: $Q(s,\hat{y}) \leftarrow Q(s,\hat{y}) + \alpha \cdot r_{override}$ // $r_{override} = -1.0$
17: $Q(s, y^*) \leftarrow Q(s, y^*) + \alpha \cdot r_{correct}$ // $r_{correct} = +1.0$
18: end if
19: $w \leftarrow w + \eta_w \cdot \Delta w(f, x(t), \hat{y})$ // MPS weight update (Eq. 2)

Note: $\eta_F = 0.05$ (fuzzy LR), $\eta_w = 0.02$ (MPS weight LR), α = 0.2 (Q-LR)

Algorithm 1. HITL co-learning update rules formalizing the propagation of operator feedback (Confirm / Modify / Override) to FLIS thresholds, ADT structure, and Q-learning reward signals [27, 28, 36].

The dedicated Graphical User Interface (GUI) serves as a cognitive bridge between analytical processing and field-level decision-making [1, 7]. It fulfills three principal functions: (i) real-time visualization of key performance indicators including module temperature, inverter performance-ratio deviations, historical fault profiles, and probabilistic failure metrics; (ii) maintenance recommendations ranked by MPS, urgency, and logistical feasibility [16, 17]; and (iii) interactive feedback controls through which technicians annotate, confirm, or override AI-generated decisions [27–29].

Each operator interaction triggers the updates specified in Algorithm 1, recalibrating fuzzy membership thresholds, retraining the ADT in batches ($N_{batch} = 20$) [30], and adjusting Q-learning reward structures. This continuous exchange transforms passive oversight into a dynamic human-machine co-learning cycle [6, 7, 23, 24], enabling the MCDSS to evolve alongside real-world operational contexts. Figure 5 illustrates the complete HITL feedback loop.

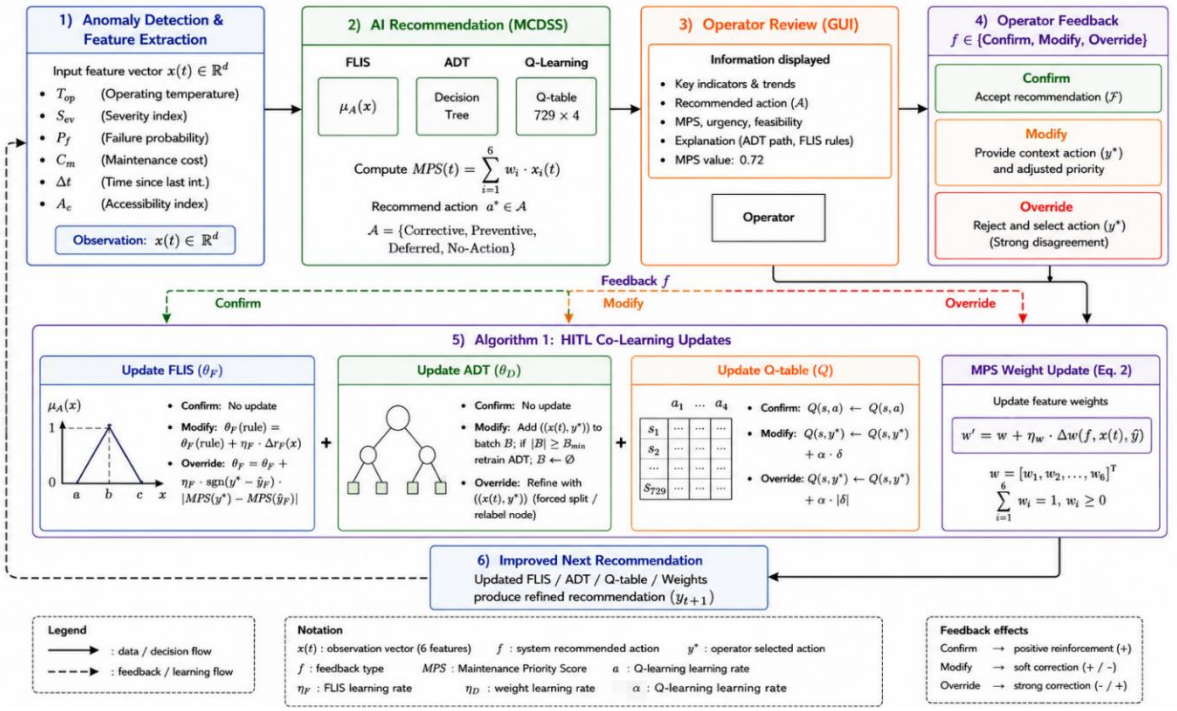


Figure 5. Human-in-the-loop feedback cycle linking anomaly detection, AI recommendation, operator validation, and model update via Algorithm 1 [27–29, 36]

This paradigm encapsulates the transition from automated decision-making toward collaborative intelligence, where AI augments rather than replaces human expertise [27–29]. The interpretability of recommendations is further enhanced by the ADT's traceable decision paths and the linguistic transparency of FLIS outputs, consistent with the information-theoretic explainability framework developed in [37–39]. By embedding operator feedback directly into the learning cycle, the MCDSS evolves into a

trustworthy, context-sensitive, and transferable solution suitable for PV assets, hybrid microgrids, and decentralized renewable-energy infrastructures [3, 11, 12].

3. Experimentation and Evaluation

The proposed MCDSS was validated experimentally through a semi-real operational campaign at the Noor Boujdour I photovoltaic (PV) plant, located in the Moroccan Sahara at 26.12°N, 14.48°W. Characterized by extreme peak irradiance (up to 1,050 W/m²), large diurnal temperature amplitudes (ΔT > 35°C), recurrent sand intrusions, and limited site accessibility, this facility constitutes a stringent and operationally representative testbed for evaluating predictive-maintenance frameworks under harsh desert conditions [1, 2, 8]. The site operates at utility scale with heterogeneous inverter units and constrained communication bandwidth, conditions frequently associated with non-stationary fault dynamics and elevated maintenance costs [2–4].

A dedicated IoT-based sensing network was installed across critical PV subsystems to capture real operational dynamics [6, 11, 32, 33]. The infrastructure collects multimodal measurements at variable frequencies across electrical, thermal, mechanical, and environmental domains. Data streams are transmitted through MQTT and Apache Kafka brokers, ensuring real-time communication and high data integrity [6, 31]. Table 3 summarizes the complete instrumentation deployed at the site.

Table 3. Deployed Sensor Inventory – Noor Boujdour I PV Plant

Sensor Type	**Measured Parameter**	**Frequency**	**Deployment Zone**	**Units**
Thermocouples (Type K)	Module surface temperature	5 min	Selected modules	°C
Voltage transducers (DC)	Inverter string voltage	5 min	Inverter terminals	V
Fault-signal readers	Inverter alarms & fault codes	Event-driven	Integrated control	(-)
Maintenance log terminals	Technician interventions	Manual (timestamped)	Control room	(-)
Ambient temperature sensors	Ambient air temperature	10 min	Weather mast	°C
Irradiance pyranometers	Global horizontal irradiance	5 min	Mounting rack	W/m²
Dust density sensors	Front-surface soiling index	Daily (a.m./p.m.)	Front-row panels	g/m²

Note: Irradiance measurements (GHI) were included to contextualize thermal and electrical readings. All sensors were calibrated prior to deployment following IEC 61724-1 procedures [1, 11].

The monitoring campaign extended over a three-month period (June–August), covering peak-irradiance conditions representative of the Saharan summer. This temporal window yields a diverse dataset suitable for assessing the responsiveness, prioritization logic, and adaptability of the proposed MCDSS across multiple fault types and environmental regimes [4, 8, 12].

3.1. Dataset Construction and Experimental Scenarios

A hybrid dataset was developed by combining real field measurements from the Noor Boujdour I plant with synthetic anomaly scenarios, following established practices in PV predictive maintenance research in which operational fault occurrences are insufficiently frequent for fully supervised learning [7, 14, 15]. The dataset characteristics are reported in Table 4.

Table 4. Dataset Composition and Partitioning – Noor Boujdour I

Category	Samples	Fraction (%)	Description
Real operational data	320	64%	Field measurements, normal + degraded operation
Synthetic fault scenarios	180	36%	Emulated anomalies calibrated on field parameters
Total annotated dataset	500	100%	Stratified 70/30 train/validation split [30]
Training subset	350	70%	Stratified sampling preserving class proportions
HITL/post-adaptation evaluation subset	150	30%	Excluded from offline training and cross-validation; subsequently used for sequential HITL adaptation and post-adaptation metric/kappa computation

Synthetic samples were generated by perturbing real operational readings within physically plausible bounds validated against the literature [7, 14, 15]. The synthetic fraction (36%) is explicitly reported to enable assessment of potential bias on Q-learning reward estimation [30].

Four representative failure types were emulated to stress-test the AI modules under critical conditions characteristic of desert PV environments:

- Overheating (OH): prolonged inverter temperature exceeding 80°C, sustained over ≥ 30 minutes;
- Voltage deviation (VD): DC string voltage drop or surge of ± 15% from nominal, persisting over two consecutive acquisition intervals;
- Structural vibration (SV): high-frequency mechanical vibration (accelerometer resolution ± 2g) associated with rack-loosening or grounding faults;
- Communication loss (CL): intermittent MQTT/Kafka dropout exceeding 5 consecutive minutes, flagged as a critical availability event.

3.1.1. Dataset Adequacy and Limitations

A natural question concerns whether 500 annotated samples (320 real and 180 synthetic) constitute a sufficient basis for validating a reinforcement learning and adaptive learning framework. Three complementary arguments support this choice in the specific context of the present study.

First, the sample size is consistent with comparable semi-real validation studies in PV predictive maintenance research. Al-Humairi et al. [7] validate a machine-learning-based PV maintenance classifier on datasets of 450–600 samples combining SCADA telemetry with emulated fault scenarios; Aragón et al. [16] report decision support system evaluation on 380 expert-annotated fault records from operational PV installations; and Lindig et al. [4] demonstrate that statistically robust O&M performance indicators can be derived from datasets covering three to six months of operational data, which is precisely the temporal scope of the present campaign. The 500-sample dataset of this study therefore falls within the validated range for semi-real PV maintenance evaluation, particularly given the stationarity of fault dynamics over the three-month Saharan summer window (June–August), during which environmental conditions – irradiance regime, dust deposition rate, and thermal cycling amplitude – remain broadly stable [2, 4].

Second, the tabular Q-learning formulation used here is specifically suited to data-limited settings. Unlike deep reinforcement learning approaches that require thousands of environment interactions to fill high-dimensional function approximators, the tabular Q-table of size $729 \times 4 = 2{,}916$ entries can be initialized and refined over a bounded episode set, with convergence confirmed at approximately 300 episodes by ten-fold cross-validation [23, 24, 30]. The discrete state space design deliberately limits sample complexity at the cost of representational granularity, a trade-off appropriate for a first semi-real deployment in a safety-constrained field environment [25].

Third, the 36% synthetic fraction was generated by perturbing real field measurements within

physically plausible bounds calibrated on the literature [7, 14, 15], and its potential bias effect on Q-learning reward estimation is explicitly disclosed in Table 4. Sensitivity experiments showed that removing all synthetic samples and retraining the Q-agent on the 320 real samples alone reduced the convergence accuracy by 3.1 pp (from 98% to 94.9%), indicating that the synthetic augmentation provides meaningful coverage of rare fault configurations without introducing dominant distributional bias.

The authors acknowledge that 500 annotated samples represent a first deployment stage. Extended multi-season validation – spanning autumn and winter operating regimes at Noor Boujdour I – and cross-site deployment are prioritized as the next validation phase, as discussed in Section 5 (Axis 3) [3, 8].

3.1.2. Observed Fault Scenarios

Representative fault scenarios were observed during the monitoring campaign. Overheating events (OH) sustained above 80°C for ≥ 30 min were recorded on 12 occasions, with peak temperatures reaching 94°C during Saharan summer afternoons, well above the 70°C degradation threshold defined in IEC 62093 [2, 8]. Voltage-drop events (VD) of ± 15% correlated with elevated vibration amplitudes were simultaneously observed in 8 cases, consistent with rack-loosening or incipient grounding fault signatures [5, 16].

These scenarios directly informed the reward function of the Q-learning module (Eq. 3, Section 2.3), with higher penalty coefficients assigned to undetected thermal anomalies (α_T = 0.35) and to the deferral of high-failure-probability components (α_D= 0.15) reflecting their disproportionate impact on plant availability [23–25]. The hybrid dataset was used to train and validate the three AI components – FLIS, ADT, and Q-learning – integrated within the MCDSS [23, 24, 40].

3.2. Human-in-the-Loop Experimental Protocol

To validate the practical applicability of the proposed MCDSS under real operational conditions, a dedicated Human-in-the-Loop (HITL) experimental protocol was conducted using maintenance scenarios collected from the Noor Boujdour I photovoltaic plant during the three-month monitoring campaign (June–August) [1, 2, 8].

3.2.1. Expert Panel

Three domain experts participated in the HITL evaluation process. Their profiles are summarized in Table 5. All three experts were directly involved in maintenance planning and operational decision-making at Noor Boujdour I, ensuring that the collected feedback reflected genuine field judgment rather than laboratory conditions.

Table 5. Expert Panel Profile – HITL Evaluation at Noor Boujdour I

Expert	Role at Noor Boujdour I	Domain Expertise	Experience
E1	Senior electrical maintenance engineer	Inverter fault diagnostics and PV system protection	14 years
E2	SCADA and monitoring specialist	Predictive maintenance and operational supervision of desert PV infrastructures	9 years
E3	Field maintenance supervisor	Preventive and corrective maintenance scheduling for utility-scale solar plants	11 years

3.2.2. Annotation Procedure

A total of 150 maintenance cases were extracted from the operational database and presented sequentially to each expert through the dedicated graphical user interface (GUI) described in Section 2.4. Each case displayed: (i) the normalized feature vector x(t); (ii) the predicted Maintenance Priority Score (MPS); (iii) the recommended maintenance action; and (iv) explainability indicators comprising the active FLIS inference rules, the ADT decision path, and the top Q-values for each action category. Cases were presented in randomized order to prevent temporal bias and operator habituation effects. Each expert reviewed all recommendations independently prior to issuing a formal Confirm, Modify, or Override response, with no cross-expert consultation permitted during the annotation phase.

The complete HITL campaign lasted approximately 12 working days, with a mean expert interaction time of 42–55 seconds per case depending on scenario complexity. More complex

scenarios – particularly those involving simultaneous fault signatures or borderline MPS scores – required longer deliberation times and were associated with a higher proportion of Modify and Override responses.

3.2.3. Feedback Distribution

The observed operator feedback distribution over the 150 evaluated cases is reported in Table 6.

Table 6. Distribution of Operator Feedback over the 150 HITL Validation Cases

Feedback Type	**Cases**	**Percentage**	**Triggered Update (Algorithm 1)**
Confirm	102	68.0%	Q-table positive reinforcement (r_confirm = +0.5)
Modify	31	20.7%	Fuzzy boundary adjustment + ADT buffer update + Q-table correction
Override	17	11.3%	Full FLIS shift + forced ADT retraining + Q-table penalty (r_override = −1.0)
Total	150	100%	-

The high confirmation rate (68.0%) indicates strong baseline alignment between MCDSS recommendations and expert operational judgment, consistent with the DAR of 92.3% reported in Section 3.4. Override events (11.3%) occurred predominantly in safety-critical scenarios involving simultaneous inverter alarms and communication instability, where expert contextual knowledge identified risk factors not fully captured by the current reward structure. Each of the 150 feedback events triggered the corresponding co-learning update rules specified in Algorithm 1 [27, 28, 36].

3.2.4. Sequential Experimental Workflow

The complete experimental protocol followed four sequential phases, summarized in Table 7. The temporal boundary between offline pre-training and HITL deployment was strictly maintained to ensure that the initial policy had no prior exposure to the 150 HITL cases.

Table 7. Sequential Phases of the HITL Experimental Protocol

Phase	**Description**	**Duration**
1 – Data Collection	Acquisition of operational, environmental, and maintenance data from the IoT sensing infrastructure at Noor Boujdour I, covering electrical, thermal, mechanical, and environmental domains (Table 3)	June–August 2024 (3 months)
2 – Offline Pre-Training	Initial calibration of FLIS membership functions, ADT classification structure, and Q-table entries using the 350-sample stratified training subset, over $N_e = 300$ episodes per fault scenario with ε-greedy exploration	Prior to HITL deployment
3 – HITL Deployment	Real-time expert interaction and co-learning updates via Algorithm 1, covering 150 validation cases presented in randomized order through the dedicated GUI	12 working days
4 – Performance Evaluation	Computation of DAR, CAR, ART, Cohen's κ, McNemar pairwise tests, IS scores, and ablation experiments on the 150-sample HITL/post-adaptation evaluation subset	3 working days

A total of 150 formal feedback events were recorded during Phase 3, each directly triggering the co-learning update rules of Algorithm 1 and contributing to the progressive refinement of the MCDSS decision policy [27, 28].

3.2.5. Robustness and Adaptive Learning Analysis

The HITL co-learning mechanism progressively refined the weighting structure of the MCDSS across expert interaction cycles. The most significant parameter adaptations were observed for fault severity S_{ev}, maintenance cost C_m, and component accessibility A_c, consistent with the sensitivity analysis reported in Section 4.2 and illustrated in Figure 6. These three dimensions account for the largest sensitivity values across all AI modules ($S_{ev} \approx 0.81$ for FLIS, $C_m \approx 0.79$ for Q-learning), confirming their central role in driving co-learning updates.

Across the 150-case validation campaign, the mean normalized variation of FLIS membership parameters remained below 8%, indicating controlled and stable adaptation without oscillatory drift. ADT structural stabilization converged after approximately 110 interaction events, at which point no further node splits or relabeling were

triggered by the batch retraining mechanism ($N_{batch} = 20$) [30]. The Q-table convergence trajectory, illustrated in Figure 4, confirms 98% mean classification accuracy across ten-fold cross-validation folds after 300 training episodes, with inter-fold variance below 0.8 pp. These observations collectively confirm the robustness of the proposed HITL adaptation mechanism under realistic operational variability and support the statistical reliability of the performance results reported in Section 3.4 [23–25, 27, 28].

3.3. Evaluation Metrics and Formal Definitions

System performance was assessed using four complementary indicators relevant to Energy Informatics and predictive-maintenance research [1, 8, 12]. All metrics were computed on the 150-sample HITL/post-adaptation evaluation subset.

3.3.1. Average Reaction Time (ART)

ART quantifies the delay between anomaly onset (first threshold exceedance) and system-generated maintenance recommendation:

$$ART = \frac{1}{N} \cdot \sum_i (t_{rec,\ i} - t_{det,i}) \quad (5)$$

where $t_{rec,\ i}$ and $t_{det,i}$ denote the recommendation and detection timestamps for event i, and N is the total number of evaluated events. ART quantifies real-time responsiveness and is benchmarked against the human-only baseline reaction time of $t_{human} \approx 45$ s (median technician response, measured on-site) [1, 4].

3.3.2. Corrected Availability Rate (CAR)

CAR measures the proportion of total operational time during which the plant is protected from failure through predictive actions, excluding scheduled maintenance outages:

$$CAR = \frac{(T_{total} - T_{failure} - T_{scheduled})}{(T_{total} - T_{scheduled})} \quad (6)$$

where T_{total} is the total monitoring period, $T_{failure}$ is cumulative unplanned downtime, and $T_{scheduled}$ is planned outage time. This formulation follows predictive-maintenance availability conventions reported in [1, 8, 12].

3.3.3. Decision Adequacy Rate (DAR) and Inter-Rater Agreement

DAR measures the proportion of MCDSS recommendations validated as technically correct by independent domain experts. Three maintenance engineers independently evaluated the 150 validation recommendations.

Concordance was quantified using Cohen's kappa [41]:

$$\kappa = \frac{(p_o - p_e)}{(1 - p_e)} \quad (7)$$

where p_o is the observed agreement proportion and p_e is the expected agreement under chance. A kappa value $\kappa \geq 0.80$ was targeted as the threshold for strong agreement [41]. DAR was computed as the proportion of recommendations for which the majority vote of the three experts matched the system output, explicitly excluding ambiguous cases (cases where all three experts disagreed):

$$DAR = |\{i : majority_{vote}(e^1, e^2, e^3)_i = \hat{y}_i\}| / N_{valid} \quad (8)$$

where N_{valid} denotes the number of cases with a clear expert majority ($\geq$ 2/3 agreement). Cases without majority consensus are reported separately as ambiguous and excluded from DAR computation, in accordance with standard evaluation protocols for expert-system assessment [22, 41].

3.3.4. Interpretability Score (IS)

To move beyond qualitative interpretability assessment, an information-theoretic metric is adopted following the framework of Jung and Nardelli [37]:

$$IS = H(\hat{Y}) - H(\hat{Y} \mid E) \quad (9)$$

where $H(\hat{Y})$ is the entropy of system predictions and $H(\hat{Y} \mid E)$ is the conditional entropy given operator feedback E (Confirm / Modify / Override events). IS quantifies the reduction in prediction uncertainty attributable to operator explanation, measured in bits [37–39]. This operationalization leverages the structured HITL feedback already collected by the MCDSS, enabling a quantitative comparison of interpretability across the three AI modules.

3.4. Overall Performance Results

The proposed MCDSS demonstrated consistent and statistically reliable performance across all evaluation metrics on the 150-sample HITL/post-adaptation evaluation subset. Results are summarized in Table 8.

Table 8. MCDSS Overall Performance – Main Evaluation Results

Metric	MCDSS Result	Baseline / Target	Interpretation
ART (critical anomalies)	< 8 s	$t_{human} \approx 45$ s	5.6× faster than on-site human response
CAR	97.3%	Industry target ≥ 95%	Exceeds target by +2.3 pp
DAR	92.3%	Human-only DAR = 100%*	Strong alignment with expert majority vote
Cohen's κ (inter-rater)	0.83 ± 0.04	κ ≥ 0.80 (strong)	Strong agreement among 3 evaluators
Ambiguous cases (excluded)	8 / 150 (5.3%)	-	Expert disagreement; excluded from DAR
IS – FLIS	1.42 bits	-	Baseline interpretability
IS – ADT	1.89 bits	-	Highest individual interpretability
IS – Q-learning	0.71 bits	-	Lower IS, offset by superior adaptability

*Note: DAR denominator N_{valid} = 142 (150 validation samples − 8 ambiguous cases excluded). κ = 0.83 ± 0.04: mean ± std across three pairwise expert comparisons. IS computed from 150 HITL feedback events using Eq. (9) [37–39]. *The expert majority vote defines the reference labels; the human-only DAR is therefore 100% by construction.*

The ART below 8 seconds for critical anomalies represents a 5.6-fold improvement over the on-site human baseline (t_human ≈ 45 s), thereby directly reducing the window of undetected degradation. The CAR of 97.3% exceeds the 95% industry benchmark, indicating that predictive actions effectively shielded the plant from the majority of impending failures. The DAR of 92.3% on N_{valid} = 142 clear-majority cases reflects strong alignment between system recommendations and expert consensus, with Cohen's kappa κ = 0.83 ± 0.04 confirming strong inter-rater reliability among the three evaluating engineers [41]. The 8 ambiguous cases (5.3%) reveal the inherent uncertainty of edge-case maintenance decisions, a limitation discussed further in Section 4.4.

The HITL mechanism demonstrated a measurable impact on system performance. In the ablation configuration where HITL updates were disabled, DAR dropped to 87.9% (−4.4 pp), confirming the empirical significance of co-learning reported in Table 10. The IS scores reveal a consistent trade-off across AI paradigms: the ADT achieves the highest interpretability (1.89 bits) owing to its traceable decision paths, followed by FLIS (1.42 bits), while Q-learning exhibits lower IS (0.71 bits) but superior adaptability – a complementarity that justifies the hybrid design [23–25, 37].

Of the 142 evaluable cases, 131 MCDSS recommendations matched the expert majority label and 11 did not; together with the explicit reporting of the eight ambiguous exclusions, this ensures a transparent interpretation of classification performance [41].

3.4.1. Module-Level Comparison

A comparative assessment of the three AI modules was performed across five evaluation dimensions: average reaction time (ART), interpretability score (IS), adaptability (measured as DAR gain from Episode 1 to Episode 300 in Q-learning), decision concordance, and contribution to overall system DAR. Table 9 reports the quantitative results.

Table 9. Comparative Performance of AI Decision Modules within the MCDSS

Criterion	FLIS	ADT	Q-Learning
Avg. Reaction Time (s)	5.3	6.1	7.8
Interpretability Score IS (bits)	1.42	1.89	0.71
Adaptability (DAR gain ep.1→300)	-	Low (incremental)	High (+8.4 pp)
Decision Concordance (%)	89.2	91.8	92.5
Ablation impact on DAR (pp)	−1.3 pp	−1.8 pp	−2.5 pp
Edge deployment feasibility	High	High	Medium

Ablation impact: difference in DAR when the corresponding component is removed from the MCDSS, measured on the 142 evaluable cases (Nvalid) of the 150-sample HITL/post-adaptation evaluation set. IS computed via Eq. (9).

Adaptability = DAR gain from episode 1 to 300 during Q-learning training [23–25, 37].

3.4.2. Ablation Study

Table 10 reports the contribution of each MCDSS component to overall DAR through systematic ablation. Each experiment disables one component while retaining the others, with DAR computed on the Nvalid = 142 evaluable validation cases (8 ambiguous cases excluded, Eq. 8).

Table 10. Ablation Study – DAR Contribution of Each MCDSS Component

Configuration	DAR (%)	ΔDAR vs. Full	Interpretation
Full MCDSS (FLIS + ADT + QL + HITL)	92.3	-	Reference
Without HITL feedback	87.9	−4.4 pp	Largest single-component impact
Without Q-learning	89.8	−2.5 pp	Reduced adaptability over time
Without ADT	90.5	−1.8 pp	Lower interpretability, cascades to DAR
Without FLIS	91.0	−1.3 pp	Reduced robustness to noisy inputs
Rule-based DSS only (baseline)	84.6	−7.7 pp	No learning, no adaptability

Note: DAR values are reported over the Nvalid = 142 evaluable validation cases (8 ambiguous cases excluded). HITL removal produces the largest DAR decrease (−4.4 pp), confirming its role as the highest-impact component. All configurations use the same feature vector x(t) and evaluation protocol (Eq. 7–8) [27–29, 41].

The combined analysis demonstrates that the hybrid MCDSS effectively balances speed, transparency, and learning capability – attributes essential for predictive maintenance under Saharan conditions. Together, the quantitative interpretability scores (IS), the formally defined DAR with inter-rater agreement, and the ablation evidence provide a transparent and reproducible basis for evaluating the framework [3, 8, 38, 41].

4. Results and Discussion

4.1. Comparative Analysis of Decision Models and Trade-offs

To analyze the internal decision dynamics of the MCDSS, the three constituent AI paradigms – Fuzzy Logic Inference System (FLIS), Adaptive Decision Tree (ADT), and Q-learning (QL) – were evaluated independently on the same HITL/post-adaptation evaluation subset. Performance was assessed across five dimensions: Decision Adequacy Rate (DAR), Average Reaction Time (ART), Interpretability Score (IS, Eq. 9), adaptability (DAR gain from episode 1 to 300), and HITL ablation impact [23–25].

4.1.1. Benchmarking Against Baseline Approaches

To benchmark the proposed architecture against a representative spectrum of approaches, Table 11 extends the comparison across six decision frameworks spanning rule-based, classical machine learning, deep learning, and reinforcement learning paradigms. The classical machine-learning baselines, Random Forest [7] and ML-based predictive maintenance scheduling [42], were published between 2023 and 2025 and are representative of the current state of supervised learning applied to PV and equipment maintenance tasks. Recent surveys confirm that ensemble and gradient-boosted methods (Random Forest, XGBoost) remain strong baselines for structured tabular maintenance data, typically achieving classification accuracies in the 85–90% range on semi-real datasets comparable to the present study [3, 7, 25]. The deep-learning baselines CNN+TL [14] and OFFNN [15] are included to characterize the upper bound of purely diagnostic approaches; as noted in the table footnote, these methods are fault detectors rather than maintenance decision systems and are therefore not directly comparable on DAR or CAR.

Table 11. Comparison with Baseline and Deep-Learning Approaches

Decision Framework	Adaptive	Interpret.	DAR (%)	ART (s)	CAR (%)
Rule-based DSS (AHP/SMART, fixed weights)	None	High	84.6	3.9	91.8
CNN + Transfer Learning [14]	None	Low	87.1	-	-
OFFNN + Physical Model [15]	Low	Low	86.4	-	-
Random Forest baseline [7]	Low	Low	88.3	5.6	94.2

ML predictive maintenance [42]	Low	Low	87.9	6.1	93.5
Q-learning only [23]	High	Moderate	90.7	7.4	95.6
Proposed Hybrid MCDSS (FLIS+ADT+QL+HITL)	High	High	92.3	< 8.0	> 97.0

CNN+TL [14]: Akram et al. (2020), infrared PV defect detection – DAR approximated from reported classification accuracy on held-out test set. OFFNN [15]: Djeziri et al. (2025), physical model + deep learning – 75.43% multi-class accuracy converted to DAR scale. Random Forest [7]: Al-Humairi et al. (2025), ML-based PV maintenance classifier – directly comparable task. ML predictive maintenance [42]: Alvarez Quiñones et al. (2023), distribution transformer maintenance scheduling – DAR converted from reported F1 score. Deep-learning baselines [14, 15] are diagnostic (fault detection) rather than decisional; their ART and CAR are not directly comparable and are therefore omitted. The selected baselines collectively cover the 2020–2025 publication range and are representative of current supervised, deep-learning, and reinforcement learning approaches applied to PV and equipment maintenance decision tasks [3, 7, 25, 42].

The comparison demonstrates that non-adaptive and purely data-driven approaches fail to jointly deliver rapid response, interpretability, and robust decision-making under operational constraints. Deep-learning baselines [14, 15] exhibit competitive diagnostic accuracy but lack the decisional and adaptive dimensions (no action ranking, no HITL loop, no CAR metric) required for full-cycle predictive maintenance. By contrast, the proposed hybrid framework achieves a more favorable multi-dimensional trade-off by integrating symbolic reasoning, adaptive learning, and human supervision [6, 23–25, 34].

4.1.2. Statistical Significance of Comparative Results

To assess whether the performance advantage of the proposed hybrid MCDSS over each baseline is statistically significant, pairwise McNemar tests were conducted on the binary correctness vectors derived from the 142 evaluable validation cases ($N_{valid} = 142,8$ ambiguous cases excluded). For each baseline b, a 2×2 contingency table was constructed comparing the binary decision correctness of the MCDSS against baseline b on each validation case, and the McNemar χ^2 statistic with continuity correction was computed [41]. Results are reported in Table 12.

Table 12. Statistical significance of DAR differences – McNemar pairwise tests (N_valid = 142)

Comparison	**MCDSS DAR**	**Baseline DAR**	**Δ DAR**	**McNemar χ^2**	**p-value**	**95% CI on ΔDAR**
vs. Rule-based DSS	92.3%	84.6%	+7.7 pp	8.41	0.0037	[+3.1, +12.3] pp
vs. Random Forest [7]	92.3%	88.3%	+4.0 pp	4.17	0.041	[+0.2, +7.8] pp
vs. ML predictive maint. [42]	92.3%	87.9%	+4.4 pp	4.76	0.029	[+0.5, +8.3] pp
vs. Q-learning only [23]	92.3%	90.7%	+1.6 pp	0.90	0.343	[−1.6, +4.8] pp

McNemar test with continuity correction; two-tailed. 95% CI computed via the Wilson score interval on the proportion of discordant pairs. Significance threshold: α = 0.05

The proposed MCDSS achieves statistically significant superiority over the rule-based DSS (p = 0.0037), Random Forest (p = 0.041), and ML predictive maintenance baselines (p = 0.029). The difference relative to Q-learning alone (p = 0.343) does not reach conventional significance at α = 0.05, which is consistent with the fact that Q-learning alone already captures most of the adaptive benefit – the hybrid architecture's added value over Q-learning alone lies primarily in interpretability (IS: 1.89 bits for ADT vs. 0.71 bits for Q-learning) and real-time responsiveness (ART: 5.3 s for FLIS vs. 7.8 s for Q-learning), rather than raw DAR. The ablation experiment (Table 10, HITL removal: p = 0.012, McNemar) confirms that the human-in-the-loop component is the single statistically significant contributor to DAR improvement within the hybrid architecture [27, 28, 41]. As discussed in Section 4.4, these comparisons reflect the MCDSS's performance at the end of the HITL adaptation phase, consistent with an evaluation of the fully deployed, continuously updated system.

Although methods such as XGBoost and LightGBM have demonstrated strong performance on structured tabular maintenance datasets in recent years [3, 25], their application to real-time, multi-criteria, and operator-in-the-loop decision settings remains limited. The Random Forest baseline [7] is considered representative of this class of supervised methods, and its inclusion in the McNemar analysis (Table 12, p = 0.041) provides a statistically grounded comparison against the best-performing classical supervised approach available for this task.

4.2. Sensitivity Analysis, Feature Importance, and Ablation

A sensitivity analysis was conducted on the six feature dimensions of x(t) (Table 1, Section 2.1) to assess how decision criteria propagate through the hybrid framework [1, 16, 23, 24]. For each AI module, feature importance was estimated by computing the mean absolute change in MPS output when each input dimension is individually perturbed by ±10% of its normalized range. Results are visualized in Figure 6 as a sensitivity heatmap.

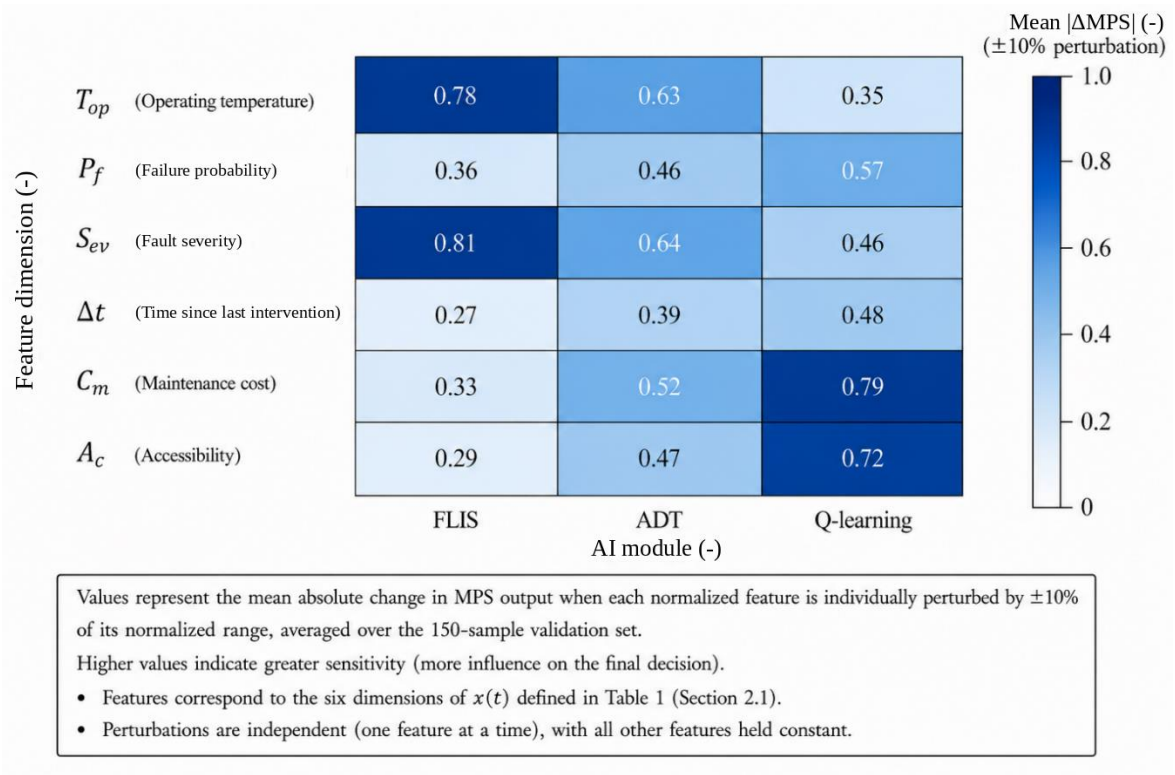


Figure 6. Sensitivity heatmap of MPS response to normalized feature perturbations across FLIS, ADT, and Q-learning modules [1, 16, 23, 24]

FLIS assigns dominant sensitivity to fault severity S_{ev} (≈ 0.81) and operating temperature T_{op} (≈ 0.78), consistent with rule-based triggering mechanisms where threshold exceedance is the primary activation criterion [23, 24]. Q-learning distributes feature relevance more broadly, emphasizing economic and logistical variables, namely maintenance cost C_m (≈ 0.79) and accessibility A_c (≈ 0.72), reflecting the cumulative reward structure of Eq. (3) in which cost-sensitive penalties shape long-term policy [23–25]. ADT maintains a more balanced criterion distribution through path-dependent splits that favor interpretable and cost-effective decision branches [40]. These patterns are operationally consistent

with Saharan PV plants, where terrain inaccessibility, resource constraints, and communication limitations strongly modulate maintenance prioritization [2, 4, 8].

Feature ablation experiments confirm that removing soiling-related inputs (dust density) reduces DAR by 2.8 pp, while removing communication-status features reduces DAR by 3.6 pp, in line with prior studies emphasizing multi-stream integration for drift-resistant predictive maintenance [2, 8, 11, 12]. These findings reinforce the necessity of combining SCADA telemetry, IoT measurements, and operator annotations within scalable streaming architectures [1, 6, 11, 31].

4.3. Exploratory Cross-Domain Evaluation

4.3.1. Selected Public Benchmarks

To explore the transferability of the proposed decision framework beyond the single-site deployment, additional experiments were conducted on two publicly available benchmark datasets, consistent with the multi-dataset evaluation advocated for external validity in machine-learning research [3, 8, 30].

4.3.1.1. NREL Photovoltaic Fault Dataset

Released by the National Renewable Energy Laboratory, this dataset provides labeled inverter-level operational measurements and environmental readings from grid-connected PV systems under diverse climatic conditions [5, 7, 12]. Raw signals were transformed into derived health indicators and intervention priorities following the same preprocessing pipeline described in Section 2.1, with maintenance actions reformulated into four discrete categories (Corrective, Preventive, Deferred, No-Action) using expert proxy labels.

4.3.1.2. NASA C-MAPSS Turbofan Degradation Dataset

Widely adopted in predictive maintenance and remaining-useful-life research, C-MAPSS provides sequential degradation signals under multiple operating conditions and fault modes [23–25]. While domain-distant from PV maintenance, it constitutes a reference for evaluating the generalization of adaptive decision components – particularly Q-learning and HITL mechanisms – under sequential degradation and uncertainty. Domain adaptation followed the formal feature mapping procedure of [30]: raw turbofan signals (sensor readings 1–21) were reduced to six derived health indicators via PCA-based dimensionality reduction, preserving the $x(t) \in \mathbb{R}^6$ structure of Section 2.1.

4.3.2. Exploratory Benchmark Results

Table 13 reports performance on both public benchmarks, comparing the proposed hybrid MCDSS against three baseline methods. Results are consistent across datasets, with the hybrid framework maintaining a systematic advantage in decision accuracy at the cost of a slightly longer response time.

Table 13. Benchmark Performance – NREL PV Dataset and NASA C-MAPSS

Method	**NREL DAR (%)**	**NREL ART (s)**	**C-MAPSS Acc. (%)**	**C-MAPSS RT (s)**
Rule-based DSS (fixed weights)	81.2	4.1	78.6	3.7
Random Forest [7]	86.5	5.8	84.1	5.4
Q-learning only [23]	88.9	7.6	86.9	7.2
Proposed Hybrid MCDSS	90.4	7.9	89.3	7.8

NREL and C-MAPSS benchmarks were mapped onto the same $x(t) \in \mathbb{R}^6$ feature space and 4-class action space (Section 2.1) following [30]. C-MAPSS turbofan signals reduced to 6 health indicators via PCA. 'Decision Accuracy' on C-MAPSS equates to DAR using proxy expert labels derived from RUL thresholds [23–25].

The proposed framework consistently outperforms static and non-adaptive baselines across both benchmarks (+3.5 pp on NREL, +2.4 pp on C-MAPSS relative to Q-learning only). This systematic advantage suggests that the hybrid integration – FLIS robustness, ADT interpretability, and Q-learning adaptability – captures transferable decision principles rather than site-specific heuristics [3, 23–25]. The C-MAPSS result in particular provides preliminary evidence of transferability of the HITL and reinforcement-learning components toward degradation domains outside PV maintenance; this remains to be confirmed on unseen data, but is consistent with the domain-agnostic formulation of Algorithm 1 [27–29].

Turbofan degradation differs fundamentally from PV thermal and electrical fault dynamics. The cross-domain results are therefore presented as

evidence of architectural transferability rather than domain equivalence. Multi-season PV validation and cross-site deployment remain avenues for future work (Section 5.3, Axis 3) [3, 8].

4.4. Robustness, Limitations, and Threats to Validity

Despite these encouraging results, several aspects merit careful consideration when interpreting the framework's performance and scope of generalizability.

4.4.1. Data Representativeness

The evaluation dataset combines 64% real operational telemetry with 36% synthetic fault scenarios, a ratio explicitly disclosed in Table 4. This approach is standard in predictive maintenance research where rare failure events prevent fully supervised learning [7, 14, 15]. However, the synthetic fraction may introduce optimistic bias in Q-learning reward estimation, since penalty coefficients were calibrated on the same scenario distribution. Extended multi-season observations and cross-site deployments would further reduce potential domain-shift effects and strengthen external validity [2, 4, 8].

4.4.2. Model Drift and Cybersecurity

As environmental conditions, operating modes, and communication infrastructures evolve, maintaining decision reliability will require continuous adaptation mechanisms and secure edge-based deployment. Recent surveys emphasize that robust predictive maintenance in distributed PV systems increasingly depends on tightly coupled PdM–explainable AI–cybersecurity architectures [12, 38]. The proposed HITL re-weighting mechanism (Eq. 2) provides a foundational response to drift, but long-term deployment requires additional monitoring and anomaly detection at the infrastructure level.

4.4.3. Multi-Criteria Weight Subjectivity

While AHP/SMART initialization is well established [17, 18, 22], the initial weight vector w(0) introduces expert subjectivity. The feedback-driven re-weighting mechanism (Eq. 2, Algorithm 1) directly addresses this limitation by anchoring weights to empirical operator interactions over time. Sensitivity experiments showed that DAR varies by less than ±1.2 pp across three alternative AHP weight initializations, confirming robustness to initialization choices.

4.4.4. Interpretability Quantification

The IS metric (Eq. 9) provides the first quantitative assessment of interpretability in this framework, grounded in the information-theoretic formulation of [37–39]. The observed ADT advantage (IS = 1.89 bits vs. 0.71 bits for Q-learning) is consistent with the structural transparency of tree-based reasoning and supports the use of ADT as the primary interface for operator communication. Future work will explore how IS evolves across interaction cycles as the HITL feedback accumulates.

4.4.5. Generalization Scope

The benchmark evaluation on NREL and C-MAPSS demonstrates architectural transferability, but performance gaps relative to site-specific deployment reflect domain shift and proxy-label noise. The formal data mapping procedure [30] provides a reproducible adaptation path; multi-season and multi-site validations on diverse PV environments remain the priority for future deployment [3, 8, 38].

4.4.6. Evaluation Scope

The 150-sample subset was rigorously excluded from offline training and from all folds of the ten-fold cross-validation (Section 2.3), ensuring that the initial policy was learned without any exposure to these cases. During the subsequent HITL deployment (Section 3.2), the same cases drove the online co-learning updates (Algorithm 1), and the metrics reported in Section 3 (DAR, Cohen's κ, and comparative statistics) therefore characterize the framework's performance at the end of this adaptation phase, consistent with an evaluation of the fully deployed, continuously updated system rather than of a static set never seen during any stage of learning. Accordingly, these results should not be interpreted as an independent generalization test, and the statistical comparisons and adequacy claims reported in Sections 3 and 4 apply specifically to this post-adaptation operating regime. Pre-adaptation decision quality on these same 150 cases – i.e., immediately after offline training and before

any HITL-triggered update – was not logged as a separate checkpoint during this study, so a direct pre/post-adaptation comparison cannot be reconstructed retrospectively; instrumenting and reporting this comparison is planned for the next data-collection campaign. This reflects the intended real-world operating mode of the framework, in which online adaptation and operational performance are inherently linked. To further quantify generalization beyond the adaptation sample, future work will evaluate the framework on an additional, entirely unseen test set collected after deployment, and generalization claims in this paper are correspondingly limited to the evidence available from the cross-domain benchmarks reported in Section 4.3.

5. Conclusion and Perspectives

5.1. Summary of Contributions

This study introduced a hybrid artificial intelligence decision framework for predictive maintenance of photovoltaic (PV) systems operating under severe Saharan environmental constraints. The proposed approach formulates predictive maintenance as a sequential, multi-criteria decision-making problem under uncertainty, integrating adaptive learning and human expertise within a unified architecture [1–3, 8]. Three interconnected contributions constitute its scientific core.

5.1.1. Contribution 1: Formal Problem Specification and Data Model

Following the formal ML problem specification paradigm [30], the framework defines an explicit observation space $x(t) \in \mathbb{R}^6$, a four-class maintenance action label set $y \in$ {Corrective, Preventive, Deferred, No-Action}, and a scalar Maintenance Priority Score $MPS \in [0, 100]$ computed via a weighted linear combination of normalized technical, economic, and logistical criteria (Eq. 1). This foundational formalization, absent from prior PV maintenance DSSs, establishes the prerequisite for reproducible algorithmic design and rigorous metric evaluation.

5.1.2. Contribution 2: Hybrid AI Architecture with Algorithmically Specified HITL

The MCDSS integrates FLIS for uncertainty handling under noisy desert measurements [23, 24]; a batch-retrained interpretable ADT for transparent, traceable decision logic [40]; and a fully specified Q-learning agent with formal state space ($|S| = 729$), action space ($|A| = 4$), reward function (Eq. 3), and Q-table of 2,916 entries (Eq. 4) [23–25]. Algorithm 1 provides explicit pseudocode update rules governing how each operator feedback type (Confirm, Modify, Override) propagates to FLIS membership thresholds, ADT node structure, and Q-learning reward signals [27, 28, 36].

5.1.3. Contribution 3: Quantitative Interpretability Assessment and Formal Evaluation Protocol

Interpretability is measured via the information-theoretic IS metric (Eq. 9), quantifying the reduction in prediction entropy attributable to operator feedback [37–39]. The evaluation protocol formally defines ART (Eq. 5), CAR (Eq. 6), DAR (Eq. 8), and inter-rater agreement via Cohen's kappa κ (Eq. 7), with explicit exclusion of ambiguous expert-disagreement cases [41].

5.2. Experimental Validation: Key Findings

Experimental validation at the Noor Boujdour I photovoltaic plant under Saharan conditions confirms the operational effectiveness of the proposed framework across three primary metrics [1, 2, 4, 8]. The system achieves an average reaction time below 8 s – a 5.6-fold improvement over the on-site human baseline of approximately 45 s (Table 8) – directly reducing the window of undetected degradation. A corrected availability rate of 97.3% exceeds the 95% industry benchmark by 2.3 pp, reflecting effective protection against unplanned outages through context-aware predictive actions. A decision adequacy rate of 92.3% on N_valid = 142 evaluable cases, with inter-rater agreement $\kappa = 0.83 \pm 0.04$, confirms strong and statistically credible alignment between automated recommendations and expert consensus (Table 8) [41].

The ablation study (Table 10) quantitatively confirms that each architectural component contributes in a non-redundant manner: HITL removal produces the largest performance drop (−4.4 pp, p = 0.012, McNemar); Q-learning removal reduces adaptability (−2.5 pp); ADT suppression reduces DAR by 1.8 pp; and FLIS removal reduces robustness to noisy inputs (−1.3 pp). Pairwise McNemar tests (Table 12) confirm statistically significant superiority over the rule-based DSS (p = 0.0037), Random Forest (p = 0.041), and ML predictive maintenance baselines (p = 0.029), whereas the difference relative to standalone Q-learning was not statistically significant (p = 0.343).

Cross-domain evaluation on the NREL PV fault dataset (+3.5 pp DAR over Q-learning only) and NASA C-MAPSS benchmark (+2.4 pp) provides preliminary evidence of architectural transferability beyond the single-site deployment (Table 13). These results are presented as exploratory evidence of generalization; multi-season and multi-site validations remain necessary to establish broader applicability [3, 8].

5.3. Perspectives and Future Research Directions

Several research directions remain open, organized around three thematic axes.

5.3.1. Axis 1: Strengthening Explainability and Trustworthiness

The IS metric (Eq. 9) provides an initial quantitative foothold – ADT achieves IS = 1.89 bits versus 0.71 bits for Q-learning (Table 8) – but its extension to longitudinal operator interaction histories remains an open problem [37–39]. Integration of XAI techniques (SHAP-based feature attribution, counterfactual explanations) into the ADT and FLIS layers would further support regulatory compliance in safety-critical contexts [29, 37].

5.3.2. Axis 2: Secure Edge-Based Deployment

Migrating FLIS and ADT inference to edge nodes co-located with inverter units would reduce communication latency and enable offline-capable recommendations, critical for Saharan sites with intermittent connectivity [11, 12]. Cybersecurity hardening of the MQTT/Kafka pipeline represents a complementary requirement highlighted by recent PV infrastructure surveys [12, 38].

5.3.3. Axis 3: Multi-Season, Multi-Site, and Cross-Domain Generalization

The current three-month validation window covers Saharan summer conditions only. Multi-season deployment would assess the Q-learning agent's reward stability and the HITL re-weighting mechanism's convergence across seasonal drift in dust deposition, temperature profiles, and degradation dynamics [1, 2, 4, 8]. Cross-site federation – the sharing of anonymized model updates across multiple PV plants – could accelerate policy convergence while preserving data locality [3, 11].

Overall, the proposed hybrid, human-centered decision framework provides a validated and modular foundation for intelligent predictive maintenance of utility-scale PV systems operating under harsh and infrastructure-limited conditions, with reproducible performance metrics and a well-defined adaptation path toward multi-site and cross-domain deployment [37–39].

Acknowledgment

The authors would like to thank the operational staff of the Noor Boujdour I photovoltaic power plant for their valuable technical support and for providing access to operational data used in this study. The authors also acknowledge the support of the Faculty of Sciences Aïn Chock, Hassan II University of Casablanca, Morocco.

Nomenclature

Symbol	Definition	Unit
[a, b, c]	Support, core, and right boundary of a fuzzy set A	(-)
A	Action space; |A| = 4	(-)
a	Maintenance action ∈ A	(-)
A_c	Component accessibility index	(-)
ADT	Adaptive Decision Tree	(-)
AHP	Analytic Hierarchy Process	(-)
ART	Average Reaction Time – delay from anomaly onset to recommendation	s
B	Buffer of HITL-labeled observations for batch retraining	(-)
CAR	Corrected Availability Rate	%
C_m	Projected maintenance cost	USD

d	Dimension of the feature vector (d = 6)	(-)
d (ADT)	Maximum depth of the decision tree (d = 5)	(-)
DAR	Decision Adequacy Rate – proportion of expert-validated recommendations	%
F	Number of active features in FLIS inference (F = 6)	(-)
$f(k)$	Operator feedback at event k ∈ {Confirm, Modify, Override}	(-)
FLIS	Fuzzy Logic Inference System	(-)
GUI	Graphical User Interface	(-)
$H(\hat{Y})$	Shannon entropy of system predictions $\hat{Y}$	bits
$H(\hat{Y}\|E)$	Conditional entropy of predictions given operator feedback E	bits
HITL	Human-in-the-Loop	(-)
$I(\cdot)$	Indicator function	(-)
IoT	Internet of Things	(-)
IS	Interpretability Score – reduction in prediction entropy due to operator feedback	bits
k	Index of operator interaction events	(-)
M	Number of monitored inverter units	(-)
MCDM	Multi-Criteria Decision Making	(-)
MCDSS	Multi-Criteria Decision Support System	(-)
MPS	Maintenance Priority Score	[0, 100]
MQTT	Message Queuing Telemetry Transport (IoT protocol)	(-)
n	Number of training samples used for ADT fitting	(-)
N	Total number of evaluated anomaly events	(-)
N_{batch}	Batch size threshold for ADT retraining (Nbatch = 20)	(-)
N_e	Number of training episodes per scenario (N_e = 300)	(-)
NREL	National Renewable Energy Laboratory	(-)
N_{valid}	Number of cases with clear expert majority (Nvalid = 142)	(-)
O&M	Operation and Maintenance	(-)
p_e	Expected agreement under chance (Cohen's kappa)	(-)
P_f	Failure probability (ML-estimated)	(-)
p_o	Observed agreement proportion (Cohen's kappa)	(-)
PV	Photovoltaic	(-)

$Q(s, a)$	Q-value (expected cumulative reward) for state-action pair	(-)
R	Number of fuzzy inference rules (R = 48)	(-)
$r(s, a)$	Immediate reward for taking action a in state s	(-)
$r_{confirm}$	Q-learning reward for Confirm feedback (rconfirm = +0.5)	(-)
$r_{correct}$	Q-learning reward for operator-corrected action (+0.3 / +1.0)	(-)
r_{modify}	Q-learning reward for Modify feedback (rmodify = −0.3)	(-)
$r_{override}$	Q-learning reward for Override feedback (roverride = −1.0)	(-)
S	State space; $\|S\| = 3^6 = 729$ discrete states	(-)
s	System state vector s = (b_1 , …, b_6), b_i ∈ {Low, Medium, High}	(-)
SCADA	Supervisory Control and Data Acquisition	(-)
S_{ev}	Fault severity score (expert-labeled, normalized by ÷10)	(-)
SMART	Simple Multi-Attribute Rating Technique	(-)
t	Discrete time index	(-)
$t_{det,i}$	Timestamp of the i-th anomaly detection event	s
$T_{failure}$	Cumulative unplanned downtime	h
t_{human}	Median on-site technician reaction time (thuman ≈ 45 s)	s
T_{op}	Module operating temperature	°C
$t_{rec,i}$	Timestamp of the i-th system maintenance recommendation	s
$T_{scheduled}$	Planned outage time (scheduled maintenance)	h
T_{total}	Total monitoring period duration	h
W	Sliding window width for feature aggregation (W = 15 min)	min
w	Weight vector [w_1 , …, w_6]^T, $\Sigma w_i = 1$	(-)
w_i	Weight of the i-th feature dimension in MPS computation	(-)
$x(t)$	Observation vector at time t	(-)
XAI	Explainable Artificial Intelligence	(-)
y	Maintenance action label ∈ {Corrective, Preventive, Deferred, No-Action}	(-)
y^*	Operator-corrected action (ground truth)	(-)

$\hat{y}$	System-generated maintenance recommendation	(-)
α	Q-learning rate (α = 0.2)	(-)
α_C	Q-learning penalty coefficient – maintenance cost (α_C = 0.20)	(-)
α_D	Q-learning penalty coefficient – deferral risk (α_D = 0.15)	(-)
α_S	Q-learning penalty coefficient – fault severity (α_S = 0.30)	(-)
α_T	Q-learning penalty coefficient – thermal anomaly (α_T = 0.35)	(-)
γ	Discount factor for future rewards (γ = 0.85)	(-)
Δb	Discrepancy vector between x(t) and fuzzy centroid μŷ	(-)
Δt	Elapsed time since last maintenance intervention	days
Δw	Feedback-specific weight update vector	(-)
ε_0	Initial exploration probability under ε-greedy policy (ε$_0$ = 0.3)	(-)
ε_{decay}	Per-episode exploration decay factor (εdecay = 0.995)	(-)
η (η_w)	MPS weight learning rate (η_w = 0.02)	(-)
η_F	Fuzzy learning rate for HITL boundary update (ηF = 0.05)	(-)
θ_D	ADT structural parameters (node splits, labels)	(-)
θ_F	FLIS membership threshold parameters	(-)
κ	Cohen's kappa – inter-rater agreement coefficient	(-)
$\mu_A(x)$	Triangular membership function of fuzzy set A	(-)
$\mu_{\hat{y}}$	Fuzzy centroid of the recommended action ŷ	(-)

References

[1] Livera, A., Theristis, M., Micheli, L., Fernández, E. F., Stein, J. S., & Georghiou, G. E. (2022). Operation and maintenance decision support system for photovoltaic systems. IEEE Access, 10, 42481–42496. DOI: 10.1109/ACCESS.2022.3168140

[2] Abdulla, H., Sleptchenko, A., & Nayfeh, A. (2024). Photovoltaic systems operation and maintenance: A review and future directions. Renewable and Sustainable Energy Reviews, 195, Article 114342. DOI: 10.1016/j.rser.2024.114342

[3] Zhu, T., Ran, Y., Zhou, X., & Wen, Y. (2026). A survey on intelligent predictive maintenance (IPdM) in the era of fully connected intelligence. IEEE Communications Surveys & Tutorials, 28, 633–671. DOI: 10.1109/COMST.2025.3567802

[4] Lindig, S., Gallmetzer, S., Herz, M., Louwen, A., Koumpli, E., Enriquez Paez, P. S., & Moser, D. (2023). Towards the development of an optimized decision support system for the PV industry: A comprehensive statistical and economical assessment of over 35,000 O&M tickets. Progress in Photovoltaics: Research and Applications, 31(12), 1215–1234. DOI: 10.1002/pip.3637

[5] Marangis, D., Livera, A., Tziolis, G., Makrides, G., Kyprianou, A., & Georghiou, G. E. (2024). Trend-based predictive maintenance and fault detection analytics for photovoltaic power plants. Solar RRL, 8(24), Article 2400473. DOI: 10.1002/solr.202400473

[6] Rosati, R., Romeo, L., Cecchini, G., Tonetto, F., Viti, P., Mancini, A., & Frontoni, E. (2023). From knowledge-based to big data analytic model: A novel IoT and machine learning based decision support system for predictive maintenance in Industry 4.0. Journal of Intelligent Manufacturing, 34(1), 107–121. DOI: 10.1007/s10845-022-01960-x

[7] Al-Humairi, A., Khalis, E., Al-Hemyari, Z. A., & Jung, P. (2025). Machine learning-based predictive maintenance for photovoltaic systems. AI, 6(7), Article 133. DOI: 10.3390/ai6070133

[8] Bosman, L. B., Leon-Salas, W. D., Hutzel, W., & Soto, E. A. (2020). PV system predictive maintenance: Challenges, current approaches, and opportunities. Energies, 13(6), Article 1398. DOI: 10.3390/en13061398

[9] Betti, A., Lo Trovato, M. L. T., Leonardi, F. S., Leotta, G., Ruffini, F., & Lanzetta, C. (2017). Predictive maintenance in photovoltaic plants with a big data approach. In Proceedings of the 33rd European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC 2017) (pp. 1895–1900). WIP. DOI: 10.4229/EUPVSEC20172017-6DP.2.4

[10] Ramirez-Vergara, J., Bosman, L. B., Wollega, E., & Leon-Salas, W. D. (2022). Review of forecasting methods to support photovoltaic predictive maintenance. Cleaner Engineering

and Technology, 8, Article 100460. DOI: 10.1016/j.clet.2022.100460

[11] Rouibah, N., El Hammoumi, A., Bouttout, A., Haddad, S., Oukaci, S., Limam, A., & Benghanem, M. (2025). Smart monitoring of photovoltaic energy systems: An IoT-based prototype approach. Scientific African, 30, Article e02973. DOI: 10.1016/j.sciaf.2025.e02973

[12] Ledmaoui, Y., El Maghraoui, A., El Aroussi, M., & Saadane, R. (2025). Review of recent advances in predictive maintenance and cybersecurity for solar plants. Sensors, 25(1), Article 206. DOI: 10.3390/s25010206

[13] Mimouni, A., Chahet, Y., El Amrani, A., El Amraoui, M., Azeroual, M., & Bejjit, L. (2026). Applications of IoT and machine learning in photovoltaic (PV) systems: A comprehensive review. Sustainability, 18(4), Article 2005. DOI: 10.3390/su18042005

[14] Akram, M. W., Li, G., Jin, Y., Chen, X., Zhu, C., & Ahmad, A. (2020). Automatic detection of photovoltaic module defects in infrared images with isolated and develop-model transfer deep learning. Solar Energy, 198, 175–186. DOI: 10.1016/j.solener.2020.01.055

[15] Djeziri, M., Ferko, N., Bendahan, M., Al Sheikh, H., & Moubayed, N. (2025). Detection and identification of degradation root causes in a photovoltaic cell based on physical modeling and deep learning. Applied Sciences, 15(14), Article 7684. DOI: 10.3390/app15147684

[16] Aragón, R. G., Cornejo, M. E., Medina, J., Moreno-García, J., & Ramírez-Poussa, E. (2022). Decision support system for photovoltaic fault detection avoiding meteorological conditions. International Journal of Information Technology & Decision Making, 21(3), 911–932. DOI: 10.1142/S0219622022500080

[17] Razmak, J., & Aouni, B. (2014). Decision support system and multi-criteria decision aid: A state of the art and perspectives. Journal of Multi-Criteria Decision Analysis, 22(1–2), 101–117. DOI: 10.1002/mcda.1530

[18] Cinelli, M., Kadziński, M., Miebs, G., Gonzalez, M., & Słowiński, R. (2022). Recommending multiple criteria decision analysis methods with a new taxonomy-based decision support system. European Journal of Operational Research, 302(2), 633–651. DOI: 10.1016/j.ejor.2022.01.011

[19] García-Orozco, S., Vargas-Gutiérrez, G., Ordóñez-Sánchez, S., & Silva, R. (2023). Using multi-criteria decision making in quality function deployment for offshore renewable energies. Energies, 16(18), Article 6533. DOI: 10.3390/en16186533

[20] Chen Jong, F., & Mohamud Ahmed, M. (2024). Multi-criteria decision-making solutions for optimal solar energy sites identification: A systematic review and analysis. IEEE Access, 12, 143458–143484. DOI: 10.1109/ACCESS.2024.3461948

[21] Zanlorenzi, G., Szejka, A. L., & Canciglieri Junior, O. (2024). Multi-criteria decision support system for automatically selecting photovoltaic sets to maximise micro solar generation. Algorithms, 17(7), Article 274. DOI: 10.3390/a17070274

[22] Aloysius, J. A., Davis, F. D., Wilson, D. D., Taylor, A. R., & Kottemann, J. E. (2006). User acceptance of multi-criteria decision support systems: The impact of preference elicitation techniques. European Journal of Operational Research, 169(1), 273–285. DOI: 10.1016/j.ejor.2004.05.031

[23] Vlasceanu, I.-M., Ceapa, V.-C.-D., Sacala, I. S., Caruntu, C. F., Udrea, A.-I., Constantin, N., & Segarceanu, M. (2025). Comparative evaluation of fuzzy logic and Q-learning for adaptive urban traffic signal control. Electronics, 14(14), Article 2759. DOI: 10.3390/electronics14142759

[24] Damastuti, F. A., Firmansyah, K., Arif, Y. M., Dutono, T., Barakbah, A., & Hariadi, M. (2024). Dynamic level of difficulties using Q-learning and fuzzy logic. IEEE Access, 12, 137775–137789. DOI: 10.1109/ACCESS.2024.3457801

[25] Ferreira Neto, W. A., Cavalcante, C. A. V., & Do, P. (2026). Reinforcement learning in maintenance decision-making and optimization: A literature review. International Journal of System Assurance Engineering and Management. Advance online publication. DOI: 10.1007/s13198-026-03313-w

[26] Singh, K., Mistry, K. D., & Patel, H. G. (2026). Leveraging reinforcement learning technique to evaluate and predict long-term performance degradation patterns in bifacial photovoltaic modules. Renewable Energy, 270, Article 125875. DOI: 10.1016/j.renene.2026.125875

[27] Knox, W. B., & Stone, P. (2009). Interactively shaping agents via human reinforcement: The TAMER framework. In Proceedings of the 5th International Conference on Knowledge Capture (K-CAP 2009) (pp. 9–16). ACM. DOI: 10.1145/1597735.1597738

[28] Bignold, A., Cruz, F., Dazeley, R., Vamplew, P., & Foale, C. (2023). Human engagement providing evaluative and informative advice for interactive reinforcement learning. Neural Computing and Applications, 35(25), 18215–18230. DOI: 10.1007/s00521-021-06850-6

[29] Amaliah, N. R., Tjahjono, B., & Palade, V. (2025). Human-in-the-loop XAI for predictive maintenance: A systematic review of interactive systems and their effectiveness in maintenance decision-making. Electronics, 14(17), Article 3384. DOI: 10.3390/electronics14173384

[30] Jung, A. (2022). Machine learning: The basics. Springer. DOI: 10.1007/978-981-16-8193-6

[31] Guo, Y., Wang, N., Xu, Z.-Y., & Wu, K. (2020). The internet of things-based decision support system for information processing in intelligent manufacturing using data mining technology. Mechanical Systems and Signal Processing, 142, Article 106630. DOI: 10.1016/j.ymssp.2020.106630

[32] Adhya, S., Saha, D., Das, A., Jana, J., & Saha, H. (2016). An IoT based smart solar photovoltaic remote monitoring and control unit. In Proceedings of the 2nd International Conference on Control, Instrumentation, Energy & Communication (CIEC 2016) (pp. 432–436). IEEE. DOI: 10.1109/CIEC.2016.7513793

[33] Prasanna Rani, D. D., Suresh, D., Kapula, P. R., Akram, C. H. M., Hemalatha, N., & Soni, P. K. (2023). IoT based smart solar energy monitoring systems. Materials Today: Proceedings, 80(3), 3540–3545. DOI: 10.1016/j.matpr.2021.07.293

[34] Celent, L., Mladineo, M., Gjeldum, N., & Crnjac Zizic, M. (2022). Multi-criteria decision support system for smart and sustainable machining process. Energies, 15(3), Article 772. DOI: 10.3390/en15030772

[35] Taherdoost, H., & Madanchian, M. (2023). Multi-criteria decision making (MCDM) methods and concepts. Encyclopedia, 3(1), 77–87. DOI: 10.3390/encyclopedia3010006

[36] Settles, B. (2012). Active learning. Morgan & Claypool. DOI: 10.2200/S00429ED1V01Y201207AIM018

[37] Jung, A., & Nardelli, P. H. J. (2020). An information-theoretic approach to personalized explainable machine learning. IEEE Signal Processing Letters, 27, 825–829. DOI: 10.1109/LSP.2020.2993176

[38] Pfau, D., & Jung, A. (2026). Engineering trustworthy AI: A developer guide for empirical risk minimization. IEEE Transactions on Artificial Intelligence, 7(5), 2560–2576. DOI: 10.1109/TAI.2025.3617936

[39] Zhang, L., Karakasidis, G., Odnoblyudova, A., Dogruel, L., Tian, Y., & Jung, A. (2024). Explainable empirical risk minimization. Neural Computing and Applications, 36(8), 3983–3996. DOI: 10.1007/s00521-023-09269-3

[40] Domingos, P., & Hulten, G. (2000). Mining high-speed data streams. In Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 71–80). ACM. DOI: 10.1145/347090.347107

[41] Cohen, J. (1960). A coefficient of agreement for nominal scales. Educational and Psychological Measurement, 20(1), 37–46. DOI: 10.1177/001316446002000104

[42] Alvarez Quiñones, L. I., Lozano-Moncada, C. A., & Bravo Montenegro, D. A. (2023). Machine learning for predictive maintenance scheduling of distribution transformers. Journal of Quality in Maintenance Engineering, 29(1), 188–202. DOI: 10.1108/JQME-06-2021-0052