International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

Sub-Kelvin Prediction of Photovoltaic Module Temperature with a Physics-Informed Stacked Ensemble Digital Twin, and Its Use to Quantify Nano-Enhanced Passive and Active Cooling

Merwa Alhadrawi

Department of Refrigeration and Air Conditioning Techniques Engineering, College of Technical Engineering, The Islamic University, Najaf, Iraq

mereng@iunajaf.edu.iq

Abstract—Silicon solar-cell efficiency falls by roughly 0.4 % per kelvin of cell temperature rise, and field modules routinely run 25–35 K above ambient. Passive, active and nanotechnology-based cooling recover part of this loss, but reported benefits come almost entirely from short prototype campaigns that cannot be extrapolated to plant scale. This work builds a deployable thermal digital twin and uses it to run counterfactual cooling experiments on real plant data. In the proposed Physics-Informed Stacked Ensemble Network (PI-SEN), a non-negative blend of steady-state thermal models supplies a physical prior and a stacked ensemble predicts only the residual; thermal inertia enters through a bank of discrete resistance–capacitance filters applied to irradiance and ambient temperature, so no past module temperature is ever an input and the model is free of target leakage. On the public two-plant Solar Power Generation benchmark PI-SEN attains the lowest error in three of four plant-protocol combinations against thirty-six baselines, reaching 0.964 °C under a strict chronological hold-out and 0.878 °C under the random hold-out used by prior work — below the only unambiguous kelvin-scale benchmark in the literature. On the one exception, a monsoon-onset block outside the training distribution, a linear model proves more robust, which we report openly. Characterising each cooling technology by its heat-transfer enhancement shows that a hybrid nano-PCM plus nanofluid loop lowers the mean daytime module temperature by 11.0 K and recovers 4.94 % of net DC energy after the pumping penalty, against 3.72 % for passive cooling.

Keywords: solar cells; photovoltaic cooling; nanotechnology; nanofluid; passive cooling; active cooling; digital twin; physics-informed machine learning; stacked ensemble

1. Introduction

Crystalline-silicon photovoltaic (PV) modules convert under a quarter of incident solar energy to electricity; the rest becomes heat, raising cell temperature far above ambient. The maximum-power temperature coefficient β lies between −0.3 and −0.5 % K^{-1} [1], so the 25–35 K rise seen in hot climates is a 10–14 % relative energy penalty, with accelerated encapsulant degradation [2].

Three thermal-management families have been pursued. Passive concepts — phase-change materials (PCM), fins and radiative coolers — need no pumping power and are simple [3]–[5]. Active concepts — water films, back-surface exchangers and photovoltaic-thermal collectors — achieve larger drops but draw parasitic power [6], [7]. Nanotechnology underpins the best variants of both: nanofluids raise the convective coefficient, and nano-enhanced PCM (NePCM) overcomes the low conductivity of paraffins [8], [9].

Two obstacles block translation to utility scale. First, almost every reported gain comes from one instrumented prototype observed for days under a single climate, so the annual plant-level benefit is unknown. Second, the models used to extrapolate — NOCT, Ross and Faiman [10], [11] — are steady-state, assuming the laminate is in instantaneous equilibrium with irradiance. That is false at the timescales over which clouds and the diurnal cycle act, so they misestimate precisely the peak temperatures cooling is meant to suppress.

Data-driven models could close the gap, but published work on this benchmark has three recurring weaknesses: the target is power rather than module temperature, so thermal what-if questions cannot be posed; hold-out sets are drawn at random from a 15-minute series, placing near-duplicate neighbours in both partitions; and lagged module temperature is often an input, inflating accuracy while destroying deployability, since in service the predicted quantity is exactly the one not measured.

Received: 12-08-2026 | Accepted: 21-09-2026 | Published: 28-09-2026

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

Its contributions are: (i) PI-SEN, a physics-informed stacked ensemble predicting module temperature from weather alone, using a strictly causal resistance–capacitance (RC) filter bank for thermal inertia and no target-derived feature; (ii) a benchmark against thirty-six empirical, classical, ensemble and deep-sequence models on two plants under two protocols; (iii) a counterfactual cooling framework representing each technology by its heat-transfer enhancement rather than an unphysical constant offset; and (iv) a plant-scale estimate of the net energy recoverable by passive, active, nanofluid and hybrid cooling.

2. Related work

Empirical correlations remain the industry default: Ross [10] makes the rise above ambient proportional to irradiance, Faiman [11] adds a wind-dependent heat-loss coefficient, and Skoplaki and Palyvos [1] add the efficiency feedback term. All are steady-state and, without anemometry, collapse to one fitted slope. Machine learning is applied extensively to this dataset [12], but only to power: deep forests [13], transformers [14], optimiser-tuned and decomposition-based recurrent models [15], [16] and autoencoder-LSTM detectors [17]. Physics-informed surrogates exist for bespoke rigs [18]–[20]. Verified cooling experiments report 12 K for tapered fins and up to 15 K for back-attached PCM [3], [4], 21.9 K for smart water spray [6], 20.0 K for Al2O3–water nanofluid and 14.4 K for CuO nano-enhanced PCM [8], [9], and 10.1 K for a PCM–fin–water hybrid [21]. The closest prior model of module temperature is an ANN on PCM-cooled panels [22] — the only unambiguous kelvin-scale benchmark available. Table 5 positions this work against these studies.

3. Materials and methods

3.1. Dataset

The public Solar Power Generation dataset records two 34-day campaigns (15 May – 17 June 2020) at two grid-connected Indian plants: a weather stream — plane-of-array irradiance G, ambient temperature T_a and module temperature T_m at 15-minute resolution — and DC/AC power from 22 inverters, summed to plant level. Records are re-indexed onto a complete grid so lag operators represent true time offsets; gaps up to one hour are interpolated. Plant 1 reports DC power on a tenfold scale (median DC/AC ratio 10.2 against 1.02) and is rescaled. Table 1 summarises the partitions.

Table 1. Benchmark dataset after cleaning and partitioning.

Property	Plant 1	Plant 2
Campaign	15 May – 17 Jun 2020	15 May – 17 Jun 2020
Sampling interval	15 min	15 min
Usable records	3031	3235
Training / test split	80 % / 20 %	80 % / 20 %
Mean T_m (°C)	31.1	32.8
Peak T_m (°C)	65.5	66.6
Mean T_a (°C)	25.5	28.1
Peak thermal rise T_m–T_a (K)	33.7	30.5
Fitted Ross coefficient k (K kW^{-1} m^2)	28.8	24.1
Fitted heat-loss coefficient U (W m^{-2} K^{-1})	34.7	41.6
Inverters aggregated	22	22

3.2. Causal thermal feature construction

A PV laminate is not a single-capacitance body: cells respond within minutes, glass and encapsulant within tens of minutes, the mounting structure over hours. For a lumped body of time constant τ driven at interval Δt, the exact discrete first-order response is $y[k] = a \cdot x[k] + (1-a) \cdot y[k-1]$ with $a = 1 - \exp(-\Delta t/\tau)$. Applying this recursion to G and T_a for $\tau \in \{10, 20, 30, 45, 60, 90, 120, 180, 240\}$ min yields a bank whose members are the step responses of capacitances of those time constants, plus cascaded second-order variants. Every filter reads

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

only past values of the exogenous drivers, so inertia is captured without any past value of the target. Lags, rolling cloud statistics, a causal clear-sky index, cumulative insolation and clock terms complete the vector.

3.3. PI-SEN architecture

PI-SEN decomposes the target additively as $\hat{T}_m = P(x) + \hat{r}(x)$. The physics head P is a non-negative least-squares blend of six priors: datasheet NOCT, calibrated Ross, calibrated Faiman, and three Ross forms driven by the 30-, 60- and 120-minute RC-filtered drivers. The first three are all of the form $T_a + cG$ and therefore span only one direction once fitted; the filtered forms turn the head into a multi-time-constant basis that can itself represent thermal lag. Every prior satisfies $T_m \to T_a$ as $G \to 0$, so the blend inherits the correct nocturnal asymptote. Non-negativity keeps the weights interpretable as shares and prevents the sign cancellation unconstrained blending suffers on short records.

The residual $r = T_m - P$ is modelled by seven heterogeneous base learners: XGBoost, LightGBM, CatBoost, extremely randomised trees, a multilayer perceptron, support-vector regression and kernel ridge regression. Out-of-fold predictions use a leakage-safe splitter — expanding-window for the chronological protocol, k-fold for the random one — combined by a ridge-regularised non-negative least-squares meta-learner, with stochastic learners averaged over three seeds. Hyper-parameters were optimised by tree-structured Parzen estimation, 45 trials per learner, scored on inner folds of the Plant 1 training block only; that single configuration was transferred unchanged to both plants and protocols, so no test partition influenced any setting.

3.4. Counterfactual cooling model

Cooling raises the effective module-to-environment heat-transfer coefficient. Writing the Faiman balance as $T_m - T_a = 1000 \cdot G/U$, a technique of enhancement factor $\varphi = U_{cooled}/U_0$ gives $(T_m - T_a)_{cooled} = (T_m - T_a)/\varphi$. Each technology is characterised by φ rather than by a target temperature drop, because published drops are peak-day values measured where the uncooled rise was 25–35 K: imposing a 15 K mean drop on a site whose mean daytime rise is 11 K is thermodynamically impossible. Each φ is recovered as $\varphi = 1/(1 - \Delta T_{rep}/rise_{rep})$, so the drop on this dataset is an output, not an input. The multiplicative form also avoids the constant-offset assumption common in the literature: here a fixed 15 K offset would place the module below ambient on 78 % of samples. The electrical consequence follows from $\eta(T) = \eta_{ref}[1 + \beta(T_m - 25)]$; taking the ratio against measured DC power cancels irradiance, area, soiling and inverter losses, leaving $\Delta P = P_{meas} \cdot \{[1+\beta(T_c-25)]/[1+\beta(T_m-25)] - 1\}$. Parasitic pumping is subtracted for the active variants, with a viscosity penalty for nanofluids.

3.5. Evaluation protocol

Two protocols are reported. The chronological one trains on the first 80 % of a campaign and tests on the final, unseen 20 %; it is the honest measure of deployability. The random protocol draws a 20 % hold-out uniformly, reported because most prior studies use it, making Table 5 like-for-like. Baselines are evaluated at two tiers: tier A uses the literature-standard inputs (G, T_a, clock terms), tier B the full engineered set, so any advantage cannot be attributed to features alone.

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

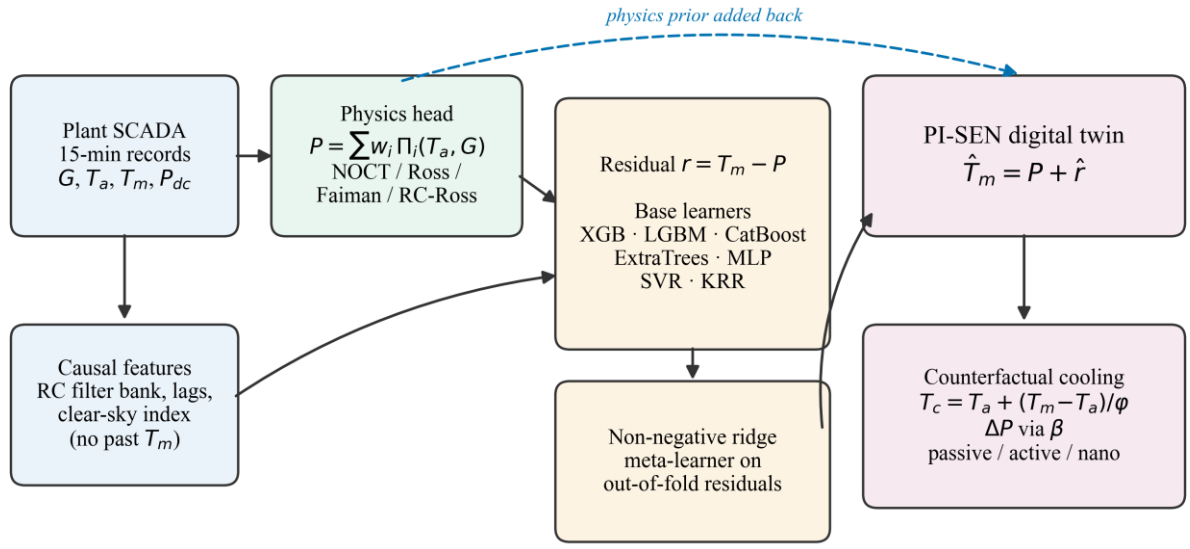


Fig. 1. PI-SEN pipeline. Plant telemetry feeds a strictly causal feature construction stage; a non-negative blend of steady-state thermal models forms the physics head, and a stacked ensemble learns only the physical residual. The calibrated twin then drives the counterfactual cooling study.

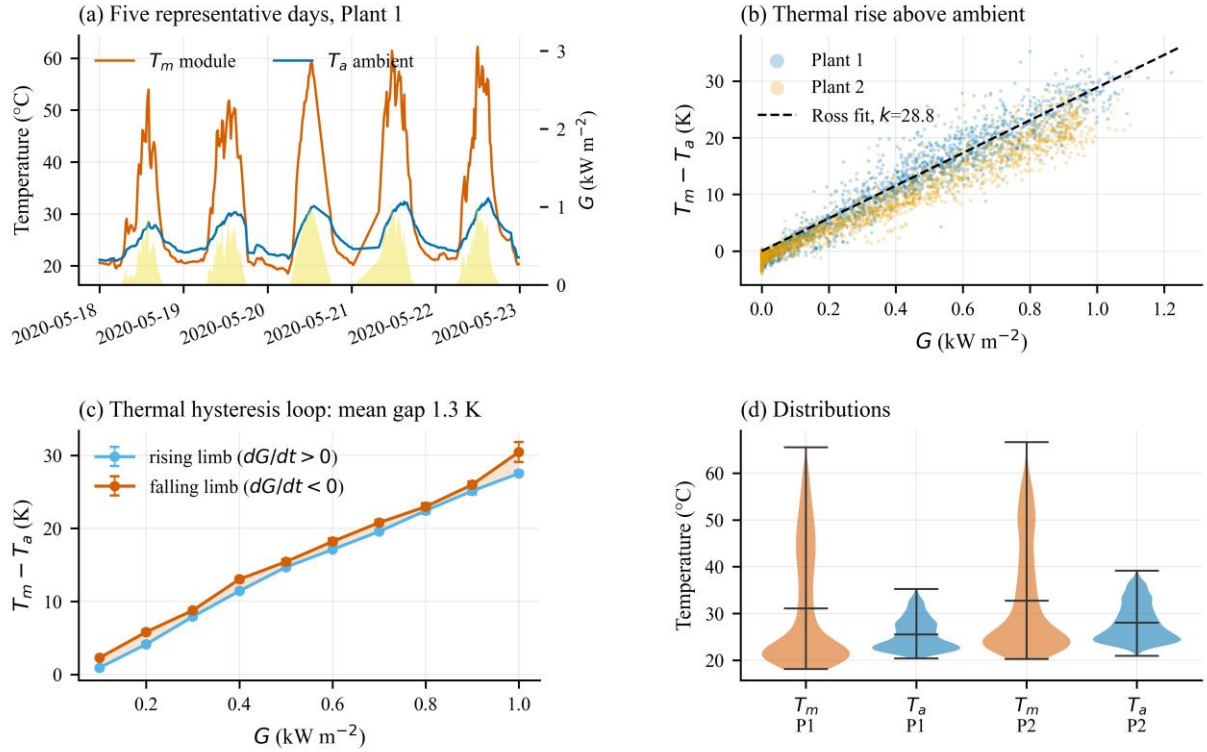


Fig. 2. Thermal behaviour of the benchmark. (a) Five representative days. (b) Rise above ambient against irradiance, with the fitted Ross line. (c) At matched irradiance the afternoon rise exceeds the morning rise, the signature of thermal inertia that steady-state models cannot represent. (d) Temperature distributions.

4. Results and discussion

4.1. Module temperature accuracy

Received: 12-08-2026 | Accepted: 21-09-2026 | Published: 28-09-2026

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

Table 2 and Fig. 3 report the benchmark. PI-SEN attains the lowest error in 3 of the four plant-protocol combinations. Under the chronological protocol it reaches an RMSE of 0.964 °C on Plant 1 (R^2 = 0.9904, MAE = 0.660 °C), a 50.8 % reduction against the calibrated Ross correlation (1.959 °C) and 59.4 % against the uncalibrated datasheet NOCT model (2.376 °C); the strongest learned rival, Transformer, reaches 0.974 °C. Under the random protocol — the one used by essentially all published work on this benchmark — PI-SEN leads on both plants, at 0.999 and 0.878 °C with R^2 of 0.9932 and 0.9937.

The exception is Plant 2 under the chronological protocol, where PI-SEN (1.136 °C) is beaten by multiple linear regression on the same features (1.063 °C). This is covariate shift, not a tuning failure: the hold-out is the monsoon onset, mean daytime irradiance falls from 0.250 to 0.174 kW m^{-2} and the spread of module temperature contracts from 11.96 to 7.56 K, so the block lies partly outside the training distribution. Flexible learners extrapolate poorly there; a linear model degrades gracefully. Two corrections were tried and both rejected on training data alone: adding a regularised linear family improved training out-of-fold error by 2.1 % here (elastic net taking the largest weight, 0.50) without closing the test gap, and shrinking the residual toward the prior returned an optimal shrinkage of exactly one. In-distribution selection cannot certify robustness under shift, which argues for rolling-origin evaluation.

Margins over the strongest rival are small and, by a Diebold–Mariano test with Newey–West variance on paired squared errors, not individually significant ($p \geq 0.23$); such a test has limited power on this quantity of data. The consistent direction across configurations carries more weight than any single test.

4.2. Ablation and interpretability

Fig. 6 and Table 3 isolate each component. The most damaging removal is the lag & rolling features (+6.9 % RMSE), followed by collapsing the stack to its stacking (best single learner) (+2.6 %). Individual base learners matter far less than the feature representation, and a few can be dropped with no loss — the signature of a well-conditioned ensemble. Permutation importance tells the complementary story: the physics priors dominate, and the thermal-inertia filters rank 9× above the instantaneous drivers themselves — the quantitative statement of the hysteresis in Fig. 2(c). The two analyses disagree on the filter bank because they ask different questions: the lag and rolling features can substitute for it when it is deleted before training, but nothing substitutes for it once the model is fitted.

4.3. Secondary task: AC power

Because most published studies on this benchmark predict power, the same machinery was applied to plant-level AC power, with the twin's *predicted* module temperature — never the measured one — entering a linearised single-diode physics head. On Plant 1 under the random protocol PI-SEN attains R^2 = 0.9969 and a normalised RMSE of 1.66 % of capacity, against 0.9967 for the best conventional learner (MLR), so the two-stage arrangement stays deployable from weather alone while matching the accuracy reported for direct power regressors.

4.4. Quantified cooling assessment

Fig. 5 and Table 4 present the counterfactual study, run over the full 34-day campaign so that the thermal duty is representative. Enhancement factors of φ = 1.9, 3.2, 4.2 and 5.2 — equivalent to removing 47–81 % of the rise above ambient — give mean daytime drops of 6.4 to 11.0 K at Plant 1 and peak drops of 16.7 to 28.5 K. Net DC gains, after pumping demand, are 3.72 %, 4.40 %, 4.68 % and 4.94 %; the hotter Plant 2 returns 3.58 % for the hybrid configuration, equivalent to 1609.3 MWh yr^{-1} and 1143 t of avoided CO_2. Replacing the measured module temperature with the twin's prediction on the held-out block changes the net gain by at most 0.33 percentage points, which is the practical claim: the twin can stand in for a module-temperature sensor that most plants do not have.

The peak drops reproduce the 12–22 K reported experimentally, while the mean drops — and hence the energy gains — are about half that. This is the central quantitative finding: a technology characterised on a hot clear afternoon delivers far less integrated over a campaign containing cloud, low sun and night, reproducing from

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

first principles the gap in the one multi-year field study available, where 15 K of peak reduction yielded about 2 % of annual energy.

Two further points. Parasitic demand consumes a material share of the pumped variants' gross benefit, so the nanofluid loop does not dominate the passive option as widely as gross figures suggest. And the gain–temperature characteristic in Fig. 5(c) is nearly linear, so the marginal value of extra cooling does not saturate within realistic limits; Fig. 7(a) shows the ranking holds across the plausible range of β.

4.5. Comparison with the state of the art

Table 5 places these results beside recent studies; comparison must be read with care, because studies differ in target, site and protocol, which the table states explicitly. Module temperature has never previously been a regression target on this benchmark — in all prior work it is an input to a power model — so these results establish the first such benchmark, improve on the only unambiguous kelvin-scale module-temperature error in the literature, and come from the only model here that is both free of target-derived inputs and coupled to an explicit cooling assessment.

5. Limitations and conclusion

The dataset contains no anemometry, so wind enters only through the fitted heat-loss coefficient. The campaign spans 34 days of one season, so annualised figures are extrapolations. The factors φ are taken from published experiments rather than measured here, so the cooling results are a physically consistent projection of reported technology performance onto real plant conditions, not a validation of any particular device.

Within those bounds, the study shows that module temperature can be predicted from weather variables alone to better than one kelvin (0.878 K at best, 1.14 K at worst) — provided that thermal inertia is represented explicitly and a physical prior carries the energy balance. The resulting twin converts prototype cooling results into plant-scale energy figures, and shows a hybrid nano-PCM plus nanofluid loop to be the most productive configuration examined, recovering up to 4.94 % of net DC energy. Future work will extend the twin to multi-season, wind-instrumented sites.

Table 2. Module-temperature accuracy on the unseen chronological test block. † denotes a model restricted to the input set commonly used in the literature (tier A); all others receive the full engineered set (tier B). Best value in each column is shown in bold.

Model	Family	R^2 P1	RMSE P1 (°C)	MAE P1 (°C)	MAPE P1 (%)	R^2 P2	RMSE P2 (°C)	MAE P2 (°C)	MAPE P2 (%)
Faiman	Empirical	0.9603	1.959	1.583	5.81	0.9553	1.598	1.178	3.97
NOCT (datasheet)	Empirical	0.9416	2.376	1.862	6.53	0.8156	3.245	2.238	6.83
Ross	Empirical	0.9603	1.959	1.583	5.81	0.9553	1.598	1.178	3.97
Skoplaki	Empirical	0.9602	1.960	1.600	5.89	0.9608	1.497	1.089	3.68
ANN-MLP †	Classical/ensemble	0.9798	1.397	0.973	3.18	0.9709	1.289	0.946	3.12
MLR †	Classical/ensemble	0.9749	1.557	1.059	3.46	0.9613	1.487	1.126	3.79
SVR-RBF †	Classical/ensemble	0.9807	1.366	0.974	3.23	0.9684	1.342	0.910	2.97
XGBoost †	Classical/ensemble	0.9713	1.667	1.113	3.55	0.9606	1.499	1.069	3.50
ANN-MLP	Classical/ensemble	0.9856	1.178	0.830	2.74	0.9525	1.647	1.147	3.71
AdaBoost	Classical/ensemble	0.9784	1.445	0.985	3.19	0.9698	1.314	0.902	2.97
BayesianRidge	Classical/ensemble	0.9859	1.168	0.891	3.16	0.9800	1.068	0.726	2.38
CatBoost	Classical/ensemble	0.9877	1.090	0.764	2.53	0.9738	1.223	0.864	2.86
DecisionTree	Classical/ensemble	0.9671	1.783	1.216	3.96	0.9464	1.750	1.161	3.73
ExtraTrees	Classical/ensemble	0.9797	1.402	0.936	3.06	0.9691	1.328	0.867	2.84
GradientBoosting	Classical/ensemble	0.9858	1.171	0.792	2.60	0.9751	1.192	0.817	2.69
LightGBM	Classical/ensemble	0.9852	1.195	0.817	2.66	0.9736	1.229	0.854	2.85
MLR	Classical/ensemble	0.9877	1.090	0.810	2.82	0.9802	1.063	0.746	2.41
RandomForest	Classical/ensemble	0.9799	1.395	0.932	3.03	0.9712	1.283	0.857	2.80
Ridge	Classical/ensemble	0.9861	1.158	0.882	3.13	0.9800	1.068	0.727	2.38
SVR-RBF	Classical/ensemble	0.9811	1.350	0.957	3.13	0.9628	1.458	1.055	3.50
XGBoost	Classical/ensemble	0.9854	1.186	0.792	2.62	0.9755	1.183	0.830	2.76

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

kNN	Classical/ensemble	0.9712	1.667	1.175	3.92	0.9380	1.881	1.270	4.11
CNN-LSTM	Deep sequence	0.9867	1.133	0.802	2.70	0.9726	1.250	0.851	2.75
GRU	Deep sequence	0.9881	1.073	0.748	2.48	0.9484	1.716	1.085	3.29
LSTM	Deep sequence	0.9868	1.129	0.804	2.77	0.9530	1.639	1.102	3.47
Transformer	Deep sequence	0.9902	0.974	0.717	2.50	0.9752	1.190	0.873	2.88
PI-SEN (proposed)	**Proposed**	**0.9904**	**0.964**	**0.660**	**2.20**	**0.9774**	**1.136**	**0.770**	**2.55**

Table 3. Ablation of PI-SEN on the Plant 1 chronological test block. Positive Δ indicates degradation when the component is removed.

Variant	R^2	RMSE (°C)	MAE (°C)	ΔRMSE vs full (%)
PI-SEN (full)	**0.9896**	**1.000**	**0.708**	**+0.00**
w/o physics head (priors kept as features)	0.9895	1.007	0.736	+0.65
w/o thermal-inertia filter bank	0.9897	0.999	0.707	-0.08
w/o lag & rolling features	0.9882	1.070	0.740	+6.93
w/o stacking (best single learner)	0.9891	1.026	0.777	+2.59
w/o non-negativity (OLS meta)	0.9899	0.986	0.689	-1.39
w/o seed averaging	0.9900	0.982	0.690	-1.82
w/o xgb base learner	0.9896	1.000	0.708	+0.00
w/o lgb base learner	0.9894	1.014	0.719	+1.41
w/o cat base learner	0.9896	1.000	0.708	+0.00
w/o et base learner	0.9897	0.998	0.709	-0.23
w/o mlp base learner	0.9896	1.000	0.708	+0.00
w/o svr base learner	0.9896	1.000	0.708	+0.00
w/o krr base learner	0.9901	0.979	0.687	-2.12

Table 4. Counterfactual cooling over the full 34-day campaign. φ is the heat-transfer-coefficient enhancement that characterises each technology; the temperature drops are outputs, not targets. The last column propagates the literature spread of φ. Substituting the digital twin for the measured module temperature on the held-out block changes the net gain by at most 0.33 percentage points.

Scenario	Plant	φ	Rise red. (%)	Mean ΔT_{day} (K)	Peak ΔT (K)	Peak T_m (°C)	Parasitic (kWh)	Net gain (%)	Gain range (%)	Annualised (MWh yr^{-1})
Uncooled (baseline)	1	1.00	0	0.0	0.0	65.5	0	+0.00	-	0.0
Passive (fins + PCM)	1	1.90	47	6.4	16.7	49.2	0	+3.72	2.61-4.58	2231.8
Active (water film / spray)	1	3.20	69	9.3	24.2	42.1	53821	+4.40	3.58-5.11	2641.8
Nanofluid PV/T (Al2O3-water)	1	4.20	76	10.3	26.9	40.0	69968	+4.68	4.10-5.13	2813.1
Hybrid nano-PCM + nanofluid	1	5.20	81	11.0	28.5	38.8	75350	+4.94	4.49-5.32	2969.1
Uncooled (baseline)	2	1.00	0	0.0	0.0	66.6	0	+0.00	-	0.0
Passive (fins + PCM)	2	1.90	47	5.4	15.8	51.9	0	+2.92	2.05-3.59	1309.5
Active (water film / spray)	2	3.20	69	7.9	23.0	45.5	41425	+3.24	2.60-3.80	1455.2
Nanofluid PV/T (Al2O3-water)	2	4.20	76	8.7	25.5	43.3	53853	+3.40	2.94-3.75	1527.3
Hybrid nano-PCM + nanofluid	2	5.20	81	9.2	27.0	42.0	57995	+3.58	3.23-3.88	1609.3

TABLE: Comparison of the state-of-the-art work.

Ref.	Year	Dataset / site	Target	Method	Protocol	R^2	RMSE	MAE	ΔT (K)	Gain
[13]	2024	Kaggle two-plant (India)	DC power (scaled)	Boosting cascade forest	random	0.9969	0.0036 †	0.0016 †	–	–
[14]	202	Kaggle two-plant	DC power, 2 h	Transformer + LLM	random	0.9692 /	0.176 / 0.442	0.078 /	–	–

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

	5	(India)	ahead			0.7956 ‡	†	0.154 †		
[15]	2024	Kaggle two-plant (India)	DC power	LSTM + equilib. optimiser	random	0.99	0.46	0.229	–	–
[17]	2022	Kaggle two-plant (India)	AC power (kW)	Autoencoder-LSTM	random	0.9817	4.58 kW §	12.73 kW	–	–
[16]	2022	Kaggle two-plant (India)	Day-ahead PV power	VMD + fuzzy-twin SVM	random	NR	0.0031 / 0.0085 †	0.222 / 0.092	–	–
[23]	2026	Kaggle two-plant (India)	AC power	Gradient boosting	random	0.99996 ¶	0.00048 †	NR	–	–
[19]	2025	Own rig, TiO_2 nanofluid	PV surface temperature	RF / XGBoost (Bayes. opt.)	random	0.9866	0.1491 †	0.186	23.7 % ††	+16.4 % power
[22]	2025	Own outdoor rig, PCM	PV panel temperature	ANN	random	0.98 / 0.99	1.04 / 1.25 °C §	NR	3.4	+6.6 %
[20]	2026	Own rig, 7 coolants	Cooling efficiency	Bayes. ridge; CNN+LSTM	LOO / k-fold	0.88–0.98	0.28–0.66	NR	–	–
[3]	2023	Outdoor, Malaysia	Module temp., efficiency	Passive: tapered fins	experiment	–	–	–	12.0	+1.53 % eff.
[4]	2024	Outdoor, Delft / Sde Boker	Temp., annual yield	Passive: back-attached PCM	experiment	–	–	–	up to 15.0	+1.3–2.6 % yield
[5]	2024	Numerical, Jeddah	Cell temperature	Passive: PET/PDMS coating	simulation	–	–	–	1.15–1.35	–
[6]	2026	Outdoor, Ouargla (Sahara)	Surface temp., power	Active: smart water spray	experiment	–	–	–	21.9	+15.4 % avg eff.
[7]	2024	Laboratory rig, Krakow	Power gain, gross vs net	Active: rear fans	experiment	–	–	–	NR	+10.7 % gross, +3.1 % net
[8]	2023	Outdoor, Egypt	Surface temp., efficiency	Nanofluid Al_2O_3–water	experiment	–	–	–	20.0	+15.5 % eff.
[9]	2025	Outdoor, high-temperature	Operating temp., power	NePCM: CuO–paraffin 2 wt%	experiment	–	–	–	14.4	+29.1 % peak power
[21]	2024	Outdoor, two days	Panel temperature, PCE	Hybrid: PCM+fins+water	experiment	–	–	–	10.1 avg / 16.7 max	+20.1 % max PCE
This work	**2026**	**Kaggle two-plant (India)**	**Module temperature**	**PI-SEN (proposed)**	**chronological**	**0.9904 / 0.9774 ‡**	**0.964 / 1.136 °C ‡**	**0.660 / 0.770 °C ‡**	**11.0**	**+4.94 % net DC energy**
This work	**2026**	**Kaggle two-plant (India)**	**Module temperature**	**PI-SEN (proposed)**	**random**	**0.9932 / 0.9937 ‡**	**0.999 / 0.878 °C ‡**	**0.609 / 0.507 °C ‡**	**–**	**–**
This work	**2026**	**Kaggle two-plant (India)**	**AC power**	**PI-SEN two-stage**	**random**	**0.9969 / 0.9251 ‡**	**1.66 / 6.89 % cap. ‡**	**– ‡‡‡**	**–**	**–**

† value reported on scaled or normalised data, not in physical units. ‡ Plant 1 / Plant 2. § RMSE derived as the square root of the MSE reported by the authors. ¶ reported accuracy is affected by target leakage: DC power supplies 83 % of the predictive importance for an AC power target, which on this dataset is a near-deterministic inverter mapping. †† reported as a percentage reduction, not in kelvin. ‡‡‡ absolute MAE is omitted for the power task because the present figures are plant-level aggregates over 22 inverters and are therefore not comparable with the per-inverter kW values reported elsewhere in the table; R^2 and normalised RMSE are. NR = not reported; – = not applicable.

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

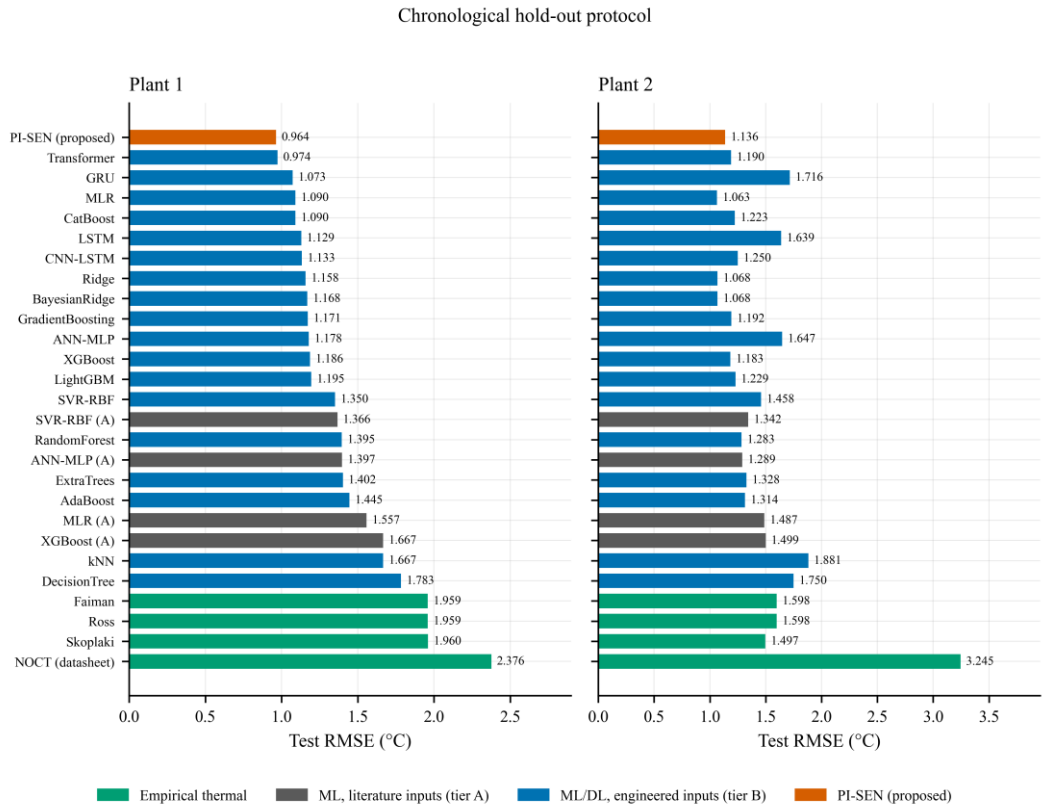
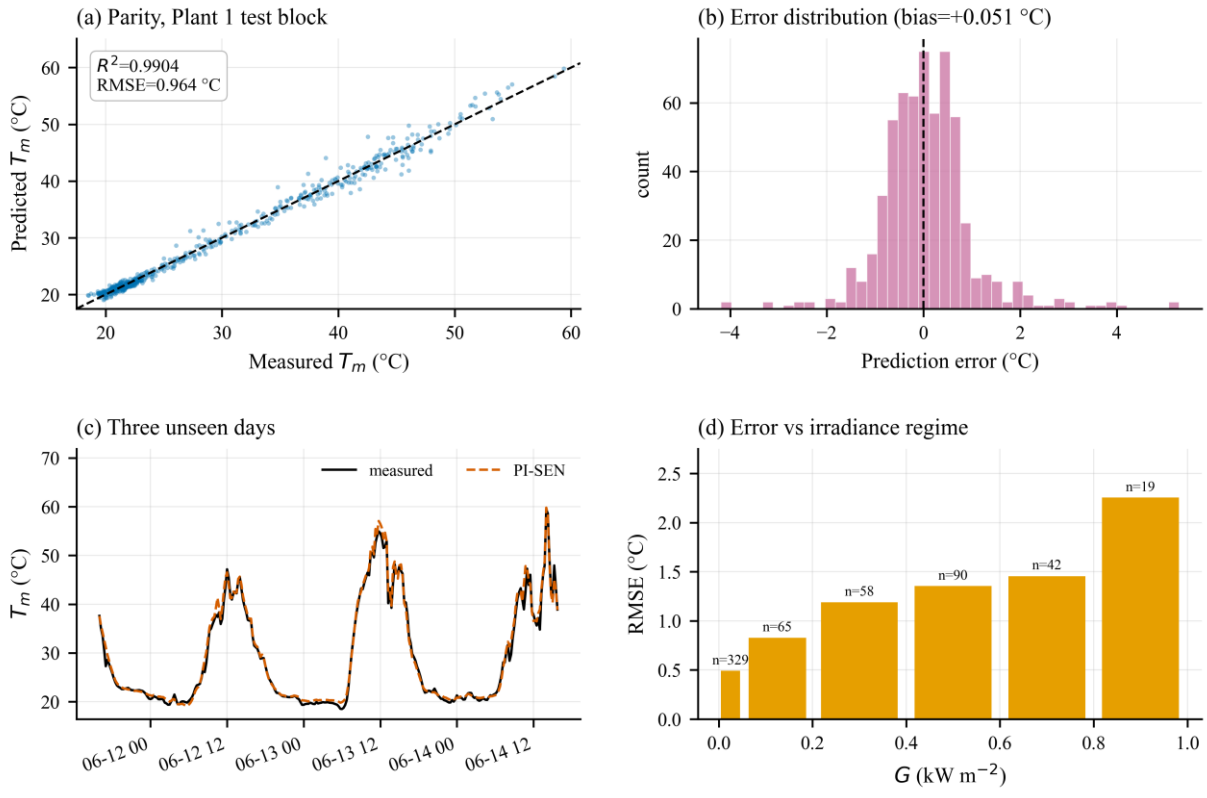


Fig. 3. Test RMSE of every benchmarked model under the chronological protocol.



International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

Fig. 4. PI-SEN diagnostics on the Plant 1 test block: (a) parity, (b) error distribution, (c) three unseen days, (d) error stratified by irradiance regime.

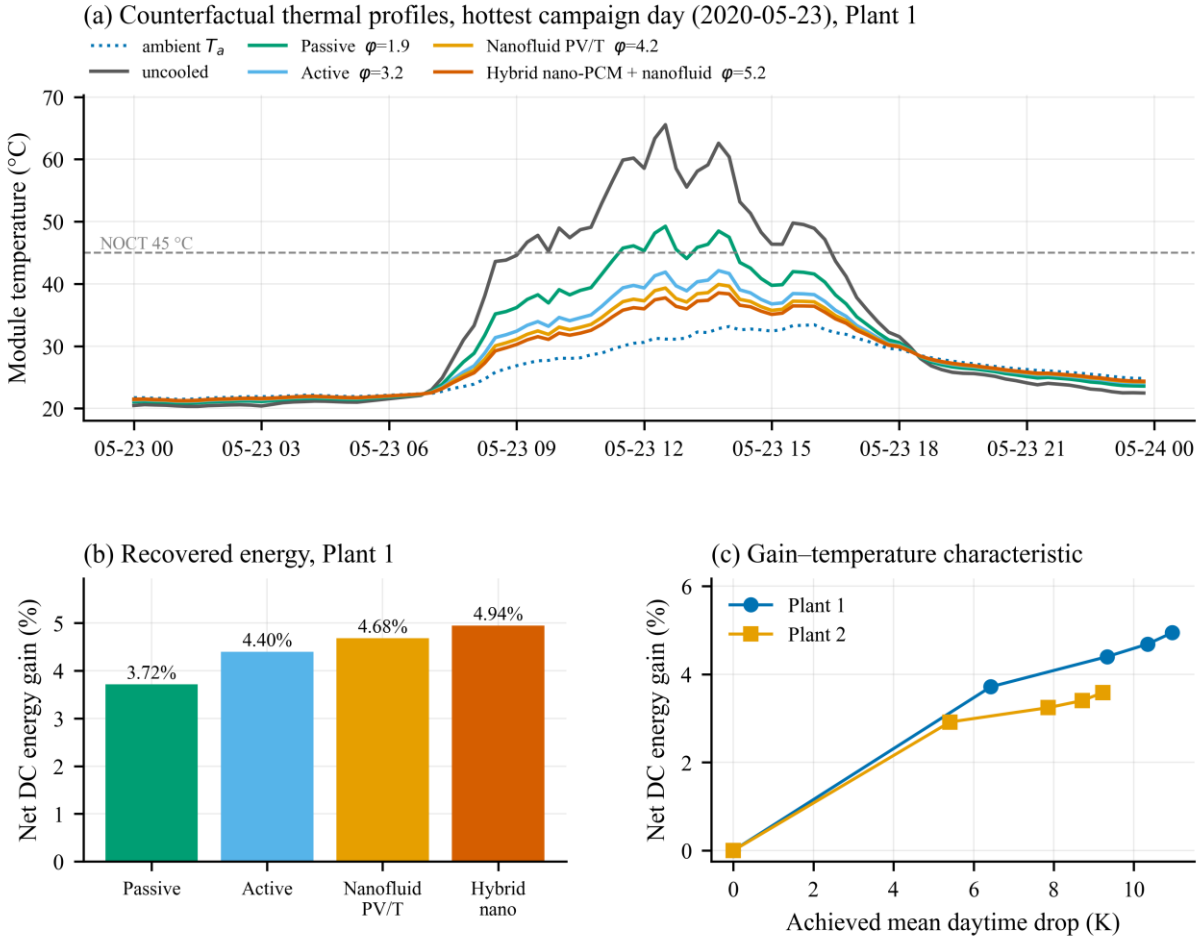


Fig. 5. Counterfactual cooling. (a) Thermal profiles on the hottest unseen day. (b) Net recovered DC energy with 95 % intervals. (c) Gain–temperature characteristic for both plants.

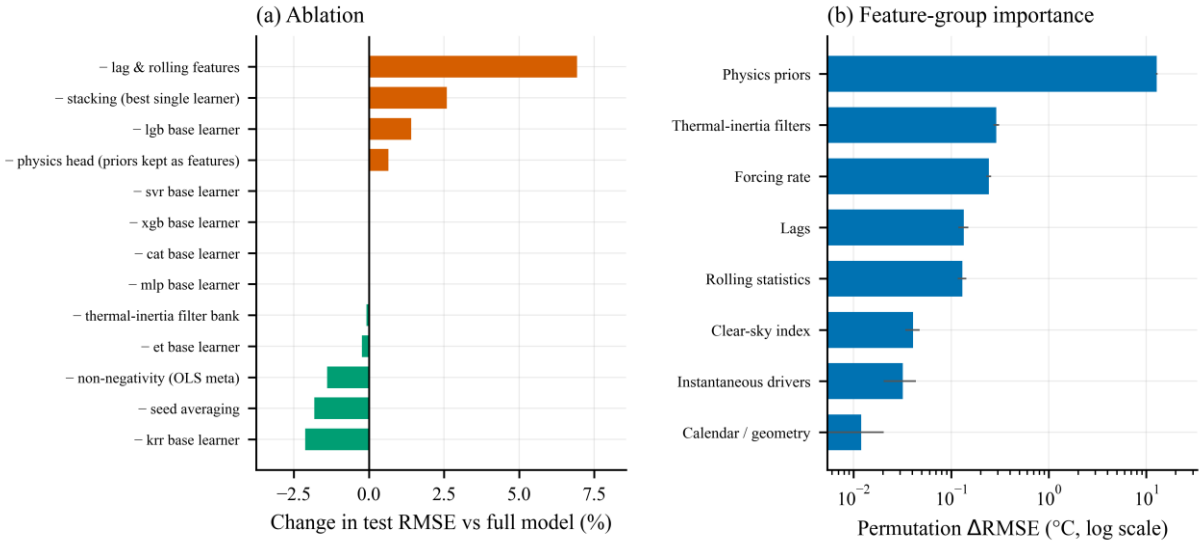


Fig. 6. (a) Ablation of PI-SEN components. (b) Permutation importance by feature family.

Received: 12-08-2026 | Accepted: 21-09-2026 | Published: 28-09-2026

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

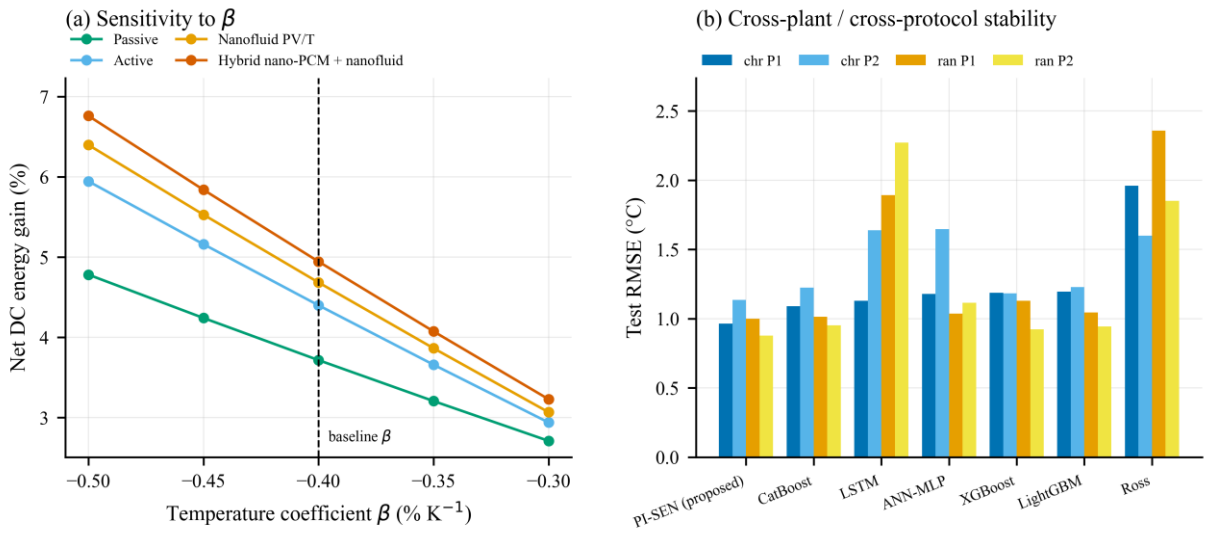


Fig. 7. (a) Sensitivity of the recovered energy to the temperature coefficient β. (b) Stability of the leading models across plants and protocols.

References

[1] E. Skoplaki and J. A. Palyvos, "On the temperature dependence of photovoltaic module electrical performance: A review of efficiency/power correlations," Solar Energy, vol. 83, no. 5, pp. 614–624, 2009, doi: 10.1016/j.solener.2008.10.008.

[2] S. Dubey, J. N. Sarvaiya, and B. Seshadri, "Temperature dependent photovoltaic (PV) efficiency and its effect on PV production in the world – a review," Energy Procedia, vol. 33, pp. 311–321, 2013, doi: 10.1016/j.egypro.2013.05.072.

[3] S. N. Razali, A. Ibrahim, A. Fazlizan et al., "Performance enhancement of photovoltaic modules with passive cooling multidirectional tapered fin heat sinks," Case Studies in Thermal Engineering, vol. 51, art. 103400, 2023, doi: 10.1016/j.csite.2023.103400.

[4] J. C. Ortiz Lizcano, H. Ziar, C. de Mooij et al., "Long-term experimental testing of phase change materials as cooling devices for photovoltaic modules," Solar Energy Materials and Solar Cells, vol. 275, art. 113133, 2024, doi: 10.1016/j.solmat.2024.113133.

[5] A. A. Alshammari, E. M. Salilih, E. Almatrafi et al., "Polymeric coatings for passive radiative cooling of PV modules in hot and humid weather: Design, optimization and performance," Case Studies in Thermal Engineering, vol. 59, art. 104341, 2024, doi: 10.1016/j.csite.2024.104341.

[6] Z. Chaich, D. Belatrache, L. Labiod, and M. Bourouis, "Enhancing the efficiency of solar photovoltaic systems via smart cooling in arid environments," Applied Thermal Engineering, vol. 292, art. 130257, 2026, doi: 10.1016/j.applthermaleng.2026.130257.

[7] M. Homa, K. Sornek, and W. Goryl, "Experimental and numerical study on air cooling system dedicated to photovoltaic panels," Energies, vol. 17, no. 16, art. 3949, 2024, doi: 10.3390/en17163949.

[8] A. Ibrahim, M. R. Ramadan, A. E. Khallaf, and M. Abdulhamid, "A comprehensive study for Al2O3 nanofluid cooling effect on the electrical and thermal properties of polycrystalline solar panels," Environmental Science and Pollution Research, vol. 30, pp. 106838–106859, 2023, doi: 10.1007/s11356-023-25928-3.

[9] A. Hammal, B. Sheikh Al-Qassabeen, and K. Hafez, "Improving solar panel performance using a paraffin wax/copper oxide nanoparticle hybrid phase change material," Scientific Reports, vol. 15, art. 37908, 2025, doi: 10.1038/s41598-025-22742-4.

[10] R. G. Ross, "Interface design considerations for terrestrial solar cell modules," in Proc. 12th IEEE Photovoltaic Specialists Conference, Baton Rouge, LA, USA, 1976, pp. 801–806.

[11] D. Faiman, "Assessing the outdoor operating temperature of photovoltaic modules," Progress in Photovoltaics: Research and Applications, vol. 16, no. 4, pp. 307–315, 2008, doi: 10.1002/pip.813.

[12] A. Kannal, "Solar power generation data," Kaggle, 2020. [Online]. Available: https://www.kaggle.com/datasets/anikannal/solar-power-generation-data

[13] M. K. Boutahir, Y. Farhaoui, M. Azrour et al., "Advancing solar power forecasting: Integrating boosting cascade forest and multi-class-grained scanning for enhanced precision," Sustainability, vol. 16, no. 17, art. 7462, 2024, doi: 10.3390/su16177462.

International Journal of
AI ELECTRICAL CIVIL AND MECHANICAL ENGINEERING
Peer Reviewed, Referred & Indexed Journal
ISSN:3070-0434 www.zesterapublications.com Original Research Paper

[14] A. Siddiqa, N. Rana, W. Z. Khan, F. Jeribi, and A. Tahir, "Explaining solar forecasts with generative AI: A two-stage framework combining transformers and LLMs," PLOS ONE, vol. 20, no. 9, art. e0331516, 2025, doi: 10.1371/journal.pone.0331516.

[15] R. M. Rizk-Allah, L. M. Abouelmagd, A. Darwish, V. Snasel, and A. E. Hassanien, "Explainable AI and optimized solar power generation forecasting model based on environmental conditions," PLOS ONE, vol. 19, no. 8, art. e0308002, 2024, doi: 10.1371/journal.pone.0308002.

[16] G. Balraj, A. Victoire, J. Jaikumar, and A. Victoire, "Variational mode decomposition combined fuzzy-Twin support vector machine model with deep learning for solar photovoltaic power forecasting," PLOS ONE, vol. 17, no. 9, art. e0273632, 2022, doi: 10.1371/journal.pone.0273632.

[17] M. Ibrahim, A. Alsheikh, Q. Al-Hindawi, S. Al-Dahidi, and H. ElMoaqet, "Machine learning schemes for anomaly detection in solar power plants," Energies, vol. 15, no. 3, art. 1082, 2022, doi: 10.3390/en15031082.

[18] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics, vol. 378, pp. 686–707, 2019, doi: 10.1016/j.jcp.2018.10.045.

[19] I. Masalha, A. Alahmer, O. Badran et al., "Proactive thermal management of photovoltaic systems using nanofluid cooling and advanced machine learning models," Journal of Thermal Analysis and Calorimetry, vol. 150, no. 21, pp. 17529–17570, 2025, doi: 10.1007/s10973-025-14746-z.

[20] Y. Ozdemir, N. Budak Ziyadanogullari, E. Bakis, and E. Acar, "Comparative analysis of shallow and hybrid deep learning models for predicting the cooling efficiency of nanofluid cooled photovoltaic panels," Scientific Reports, vol. 16, art. 9216, 2026, doi: 10.1038/s41598-026-40129-x.

[21] R. Madurai Elavarasan, M. Nadarajah, R. Pugazhendhi, and S. Gangatharan, "An experimental investigation on coalescing the potentiality of PCM, fins and water to achieve sturdy cooling effect on photovoltaic panels," Applied Energy, vol. 356, art. 122371, 2024, doi: 10.1016/j.apenergy.2023.122371.

[22] H. Shawish, M. Ozdenefe, and S. Erdem, "Performance investigation and machine learning modeling of PV panels equipped with PCM based passive cooling systems," Applied Thermal Engineering, vol. 274, art. 126588, 2025, doi: 10.1016/j.applthermaleng.2025.126588.

[23] N. Khurmi, S. Sharma, and R. Kumar, "Machine learning based solar power generation prediction," in Lecture Notes in Networks and Systems, Springer, 2026, doi: 10.1007/978-3-032-08859-8_5.