https://doi.org/10.15407/srenergy2026.02.064

UDC 621.311.243: 004.85

Dmytro Karpenko, PhD (Engin.), https://orcid.org/0000-0002-8022-9782
General Energy Institute of NAS of Ukraine, 172 Antonovycha St., Kyiv, 03150, Ukraine
e-mail: dmytro.qua@gmail.com

A HYBRID MACHINE-LEARNING ENSEMBLE FOR SHORT-TERM FORECASTING OF PHOTOVOLTAIC, SOLAR-THERMAL, AND PV/T GENERATION

Abstract. *Short-term forecasting of solar generation is a prerequisite for the operational integration of distributed solar plants into modern power systems, yet single-family approaches struggle with the asymmetric "bell" distribution of daily output, the limited observability of panel-thermal states, and the differing physical structure of pure photovoltaic (PV), solar-thermal (ST) collector and hybrid photovoltaic/thermal (PV/T) installations. This paper presents a unified machine-learning ensemble that forecasts all three modalities through one pipeline: (i) a multi-scale convolutional–recurrent network with parallel branches of kernel sizes 3, 12 and 24 hours aligned with the daily solar cycle; (ii) an XGBoost regressor on exogenous-only meteorological, calendar and air-quality features; (iii) a per-horizon weighted combination of the two heads tuned offline on out-of-fold windows under a monotone "GBDT-first" prior; (iv) an adaptive retraining cycle that, on a threshold trigger, fine-tunes the LSTM-CNN, incrementally extends the XGBoost booster with additional trees, and re-tunes the per-horizon weights on a sliding history window; and (v) an opt-in inference-time computer-vision contamination derate from an EfficientNet-B0 panel-state classifier. The same code drives the PV-electric and ST-thermal heads, each with its own scaler, weight schedule and adaptive trigger. Training and validation use 23 days of one-minute São Mateus PV and PV/T testbed data aggregated to hourly resolution and joined with reanalysis-derived weather and air-quality features. On a rolling-horizon 24-hour test the ensemble attains R^2 of 0.86 ± 0.07 for PV electric power on the PV/T testbed and 0.79 ± 0.11 for ST thermal power; a stress-test of the three-step retraining cycle that fires on every test window lands within one across-window standard deviation of the non-adaptive baseline, indicating that the cycle operates as designed without destabilising the ensemble.*
Keywords: short-term solar forecasting, photovoltaic/thermal (PV/T) generation, solar-thermal collector, multi-scale LSTM-CNN, XGBoost ensemble, per-horizon weight tuning; adaptive online retraining.

1. Introduction

The accelerating integration of variable renewable sources into modern power systems has transformed short-term forecasting of solar generation from a technical convenience into an operational prerequisite. Solar variability is now the dominant short-term uncertainty source in grids with high renewable penetration; its hour-to-hour fluctuations propagate directly into reserve scheduling, battery dispatch, and self-sufficient-microgrid economics [1, 2]. A mathematical model for self-sufficient distributed generation with battery storage [1] derives frequency-stabilisation and energy-cost indicators from instantaneous power data; a sizing model for solar-power-plant–battery-storage complexes [2] confirms that storage parameters can be optimised under uncertainty when accurate short-term power inputs are available. The same idea extends to integrated wind / PV / storage distributed generation [3]. All three operational studies require a short-term forecaster of the kind proposed here as an input.

Solar installations appear in three physical modalities sharing a dependence on irradiance but differing in their output structure: pure photovoltaic (PV) panels that convert sunlight into electricity, pure solar-thermal (ST) collectors that convert sunlight into useful heat in a working fluid, and hybrid photovoltaic/thermal (PV/T) modules that co-generate both through a back-sheet heat exchanger [4]. A forecasting framework targeting only one modality cannot serve the increasingly common residential and small-commercial deployments where PV and PV/T arrays sit side by side, sharing meteorological inputs and a single energy-management controller. The systematic survey of artificial intelligence in energy systems [4] identifies forecasting across heterogeneous generation modalities as a high-value ML application, building on engineering analyses of photovoltaic efficiency for municipal applications [5] that established the deployment baseline.

Three challenges complicate the operational deployment of machine-learning forecasters for solar installations. First, the daily profile of solar generation is strongly asymmetric: zero values dominate the night-time half of every day, and a unimodal "bell" with a noon peak dominates the daylight half. Standard autoregressive baselines cannot represent this structure without explicit exogenous meteorological features [6]. Second, the variability budget of the daily profile changes with the lead time: at short horizons the most informative signal is the current radiation forecast, while at long horizons the dominant signal is the temporal pattern in historical generation. A daily-cycle-adapted modification of simple exponential smoothing achieves nMAE = 0.84 % one-hour-ahead on a 9 MW Kyiv plant, more than an order of magnitude better than the classical formulation [7]; the result confirms that explicit modelling of the daily cycle is more important than the model class at short lead times. Third, photovoltaic/thermal installations produce a coupled electrical-thermal signal that exposes silicon cooling effects unavailable in pure-PV training data [8]; data-driven forecasters trained on one channel can leak information across heads unless the off-target column is excluded from the feature matrix.

Recent literature has converged on hybrid architectures that combine deep recurrent networks, gradient-boosted decision trees, and physically-motivated correction layers [9–11]. A hybrid LSTM + CNN + GBDT + Transformer pipeline for multi-energy building consumption [9] introduces a dynamic-retraining loop that motivates the adaptive mechanism reused on the generation side in the present work. A 72-hour PV forecasting algorithm fusing numerical weather prediction (NWP) data with a physics-based transposition and thermal-electric model and an ANN correction stage [10] demonstrates that explicit physical priors substantially reduce relative RMSE. A "physics-as-a-layer" formulation embedding differentiable irradiance–thermal computation in the forward graph of a spatio-temporal network [11] reaches near-unity R^2 for PV power forecasting on a utility-scale plant. Yet most published ensembles target either PV-only or PV/T-only, rarely both, and almost never integrate panel-state visual signals as an inference-time correction even though dust deposition can cost up to 17.7 % of yield over eight weeks under tropical conditions [12]. Building-side actuation studies such as reinforcement-learning controllers [13] close the loop from short-term forecasts to consumer-side decisions and underline the operational value of a forecaster that can serve all three solar-installation modalities.

This paper closes the gap above with a unified hybrid framework that handles PV-only, ST-only and PV/T installations through a single orchestrated pipeline. The contributions are: (i) a multi-scale convolutional–recurrent backbone whose kernel sizes 3, 12 and 24 hours match the sub-hourly, half-day and full-day structure of the daily solar cycle; (ii) an exogenous-only XGBoost head that complements the recurrent component at short horizons where lag features are most informative; (iii) a per-horizon weighting scheme tuned offline on out-of-fold windows under a monotone "GBDT-first" prior with a horizon-dependent ceiling; (iv) an adaptive three-step online retraining cycle (LSTM-CNN fine-tune + XGBoost incremental boost + per-horizon weight re-tune) triggered when MAE / RMSE thresholds are exceeded, with dataset-specific thresholds decoupling trigger sensitivity from installed capacity; and (v) an opt-in computer-vision contamination derate based on an EfficientNet-B0 panel-state classifier that does not enter the trained ensemble's input space. The same training and inference code paths serve the PV electric and ST thermal targets, each with its own scaler, weight schedule and adaptive trigger. The framework is validated on the São Mateus PV and PV/T testbeds — two co-located

850 Wp installations under tropical conditions in Brazil — and achieves $R^2 = 0.86 \pm 0.07$ for PV electric power on the PV/T testbed and $R^2 = 0.79 \pm 0.11$ for ST thermal power on a 24-hour rolling-horizon test.

The remainder of the paper is organised as follows. Section 2 reviews related work on hybrid solar-installation forecasters. Section 3 describes the methodology and the datasets. Section 4 reports the empirical findings. Section 5 discusses the implications and limitations. Section 6 concludes.

2. Related work

The literature on data-driven forecasting for solar installations falls into five sub-strands relevant to the present framework: deep-recurrent and attention-based architectures, gradient-boosted and stacked ensembles, image-based nowcasting, machine-learning models for PV/T thermal output, and microgrid-level integration.

Deep recurrent and attention-based architectures. Long Short-Term Memory (LSTM) networks remain the workhorse of medium-horizon PV forecasting [14]. A long-term LSTM applied to six months of meteorological and generation data from the Kamianka PV plant [14] obtained $R^2 = 0.92$ on the test set with RMSE 6.12 % of installed capacity; the analysis emphasises that the residual error is dominated by the stochasticity of cloud cover and the limited resolution of meteorological signals. A separate analysis of the influence of meteorological factors on solar generation in microgrids [15] argued that aggregating data into daily intervals reduces variability but at the cost of short-term operational utility — directly motivating intra-day rather than daily-aggregate forecasts. A comprehensive review [16] systematically analysed 36 deep-learning case studies for PV forecasting between 2022 and 2023 across four architectural categories (ANN-unit, RNN-unit, CNN-unit and attention-unit) and reported that Transformer-based variants emerge as the most accurate, with CNN–RNN hybrids close behind. A hybrid CNN-LSTM-Transformer model on solar production data [17] showed that adding the Transformer block on top of the CNN-LSTM combination materially improved forecast accuracy at the day-ahead horizon — providing independent support for the multi-scale-then-recurrent decomposition adopted in the present framework.

Recent work also explores explicit attention mechanisms and Transformer variants. A self-attention layer enhances the generalisability of intra-hour distributed PV forecasting across urban sites [18], while a variable-attention multi-task network jointly forecasts wind power, solar power, and aggregate system load [19] — demonstrating that attention-weighted feature fusion captures cross-source correlations missed by independent models. PVTransNet, a family of Transformer variants for multi-step day-ahead PV forecasting [20], shows that its LSTM-augmented variant outperforms the pure Transformer by 48 % in MAE. A bidirectional temporal-convolutional-network (BiTCN) enhanced by temporal bottlenecks and attention [21] expands the receptive field with fewer layers, an alternative inductive bias to the multi-scale parallel-kernel approach adopted here.

Gradient-boosted and stacked ensembles. Tree-based ensembles have proven particularly competitive for PV when the input feature space is rich in exogenous radiation and air-quality variables. An LSTM-XGBoost-EEDA-SO hybrid [22] uses ensemble empirical mode decomposition to split the power sequence into low- and high-frequency components handled by XGBoost and LSTM respectively, with Snake-Optimization-tuned combination weights; experimental results on real plant data show that the combined model outperforms standalone LSTM and XGBoost, validating the parallel-paths-with-tuned-weights architecture used here. A stacked Gradient-Boosting–XGBoost ensemble with Ridge regression as the meta-learner [23] applies the same principle to short-term PV power forecasting in smart grids, leveraging meteorological and operational parameters and reporting improvements over single-learner baselines. A recent transformer-variants benchmark [24] systematically compared Autoformer, Informer, FEDformer, DLinear, and PatchTST on a five-year rooftop PV dataset, with PatchTST + adaptive conformal inference delivering MAE = 0.194 kW (≈ 3.9 % of peak) at the hour-ahead horizon. All three studies support the per-horizon weighting scheme in §3.5, which can be interpreted as a horizon-stratified stacked ensemble with a one-dimensional meta-learner.

Image-based solar nowcasting. Sky-camera and satellite imagery are increasingly used as additional inputs in the 0–4 h horizon. A benchmark study of all-sky imager–based solar nowcasting [25] combined deep learning with physical and persistence models; the deep-learning model outperformed the hybrid baselines under most aggregation conditions at 0–20 min lead time. A Deep Generative Model of Radar (DGMR-SO) adapted from precipitation nowcasting to geostationary satellite images [26] reported that a U-Net baseline showed the lowest errors and DGMR-SO was superior on qualitative performance after 45 min. The computer-vision correction layer of the present framework is complementary — it provides a low-cost panel-state signal that does not require continuous sky-imagery infrastructure and degrades gracefully when no image is available.

ML-based modelling of PV/T thermal output. The PV/T literature treats the thermal channel as a first-class regression target. A Gaussian-process regression model [27] forecasted hPVT performance with pure water and Fe/water nanofluid under a year of meteorological data from Roorkee, demonstrating predicted–observed agreement within 1 % for cell temperature and overall efficiency. A feed-forward artificial neural network with two hidden layers for PVT air-collector efficiency in the Jordan valley [28] reported MAE of 0.0078 % for electrical and 3.36 % for thermal efficiency at a 10-min step, identifying solar irradiance and module temperature as the dominant input variables. A benchmark of multilayer perceptron, random forest and support vector regression on more than 380 published PVT cases [29] reported that Random Forest reached $R^2 = 0.99$ on training and 0.88 on testing, with SHAP analysis identifying mass flow rate, pipe inner diameter and wind speed as the most influential features. High-fidelity computational fluid dynamics of a hybrid PV/T-ST collector coupled with a feed-forward back-propagation network [30] produced regression coefficients above 0.9999 and a minimum MSE of 1.97×10^{-6} across irradiance values from 200 to 1000 W/m² and across laminar–turbulent flow regimes. A radial-basis-function model on a water-MnO_2 nanofluid PV/T testbed [31] achieved $R^2 = 0.96$ / 0.97 for power and electrical efficiency, with the nanofluid configuration improving electrical output by 21–34 % over the water-cooled reference. These results jointly establish that machine-learning regressors are competitive with finite-element thermal simulators for PV/T performance prediction at operational time scales, provided that the input feature set spans irradiance, ambient temperature, flow rate and panel-state variables.

Microgrid-level integration. At the system level, Support Vector Regression substantially outperforms linear regression for short-term solar PV and wind power forecasting in grid-connected microgrids [32] (MSE 2.002 for PV, 3.059 for wind versus higher figures for the baseline), and battery-coupled scheduling informed by accurate forecasts lowers grid energy usage by up to 48 % on sunny days. Such operational benefits define the deployment target for the framework proposed here.

Research gaps. Three gaps emerge from this synthesis. First, most ensemble works treat PV-only data, while only specialist studies address PV/T or stand-alone solar-thermal collectors, and almost no published framework trains the same architecture on all three modalities with the same hyperparameters and code paths. Second, the daily structure of solar generation is rarely encoded explicitly in the convolutional receptive field — most CNN-LSTM hybrids let the convolutional layer discover the daily cycle from the data, even though three physically meaningful kernel widths (3 h, 12 h, 24 h) are known a priori. Third, the panel-state visual signal documented to drive a substantial fraction of the dust-related loss budget [12] is almost never integrated as an inference-time correction that leaves the trained model invariant. The present paper addresses all three.

3. Methods and materials

3.1 Forecast problem formulation

Let $\mathbf{X} \in \mathbb{R}^{T \times F}$ denote the input feature matrix over T steps and F features. The system constructs two mappings, one per generation modality:

$$f_{\mathrm{PV}}: \mathbf{X} \to \hat{\mathbf{y}}_{\mathrm{PV}} \in \mathbb{R}^{H}_{\geq 0}, \qquad f_{\mathrm{ST}}: \mathbf{X} \to \hat{\mathbf{y}}_{\mathrm{ST}} \in \mathbb{R}^{H}_{\geq 0} \tag{1}$$

where $\hat{\mathbf{y}}_{\mathrm{PV}}$ is the PV electric power (kW) and $\hat{\mathbf{y}}_{\mathrm{ST}}$ is the solar-thermal power (kW). The non-negativity constraint reflects the physical impossibility of negative generation. Default values $T = 72$ h and $H = 24$ h are used in the experiments; both are persisted in the model metadata and reloaded automatically at inference. The same model class, training loop, and per-horizon weighting are applied to both targets; the only difference is the active scaler and the target column extracted from the multi-head output.

3.2 Architecture overview

Fig. 1 summarises the structural scheme of the hybrid framework. Six input streams (historical PV electric and ST thermal generation, air-pollution and weather forecasts from API providers, calendar and time-position covariates, and panel surface images) flow into a single data-preparation stage that drives three parallel modelling components: a multi-scale LSTM-CNN, an XGBoost regressor, and an EfficientNet-B0 [34] computer-vision classifier. The first two are aggregated by per-horizon weights into the PV-electric and ST-thermal outputs; the third generates a multiplicative inference-time correction. The architecture is symmetric in the two target modalities — the same backbone is trained either on pv_electric_kw, sc_thermal_kw, or both simultaneously in a multi-target run.

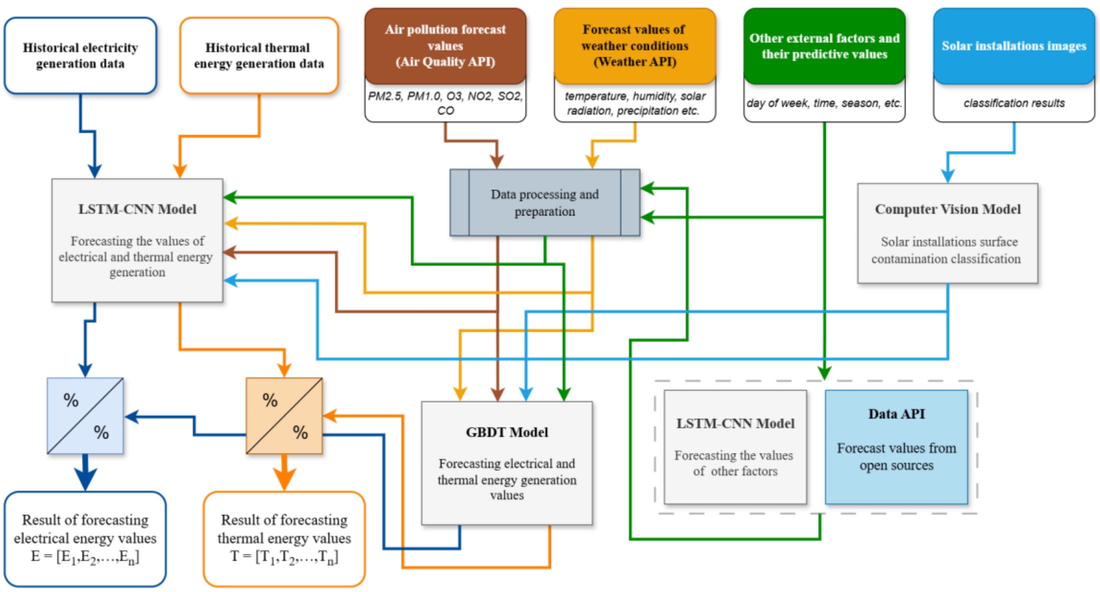

Fig. 1. Structural scheme of the hybrid ML ensemble

3.3 Multi-scale LSTM-CNN with solar-adapted kernels

The neural component is a hybrid convolutional–recurrent architecture adapted to the physical structure of the daily solar cycle. Three parallel one-dimensional convolutional branches with $C_{\mathrm{br}} = 96$ channels per branch operate on the input window with kernel sizes $\{k_1, k_2, k_3\} = \{3,12,24\}$ — sub-hourly transitions, daily half-cycles, and the full solar day, respectively. For the i-th branch:

$$\mathbf{H}_i(t) = \mathrm{ReLU}\left(\sum_{\tau=0}^{k_i-1} \mathbf{W}_i\,[:,:,\tau] \cdot \mathbf{X}(t-\tau) + \mathbf{b}_i\right). \tag{2}$$

The concatenated tensor $\widetilde{\mathbf{H}}(t) = [\mathbf{H}_1(t) \parallel \mathbf{H}_2(t) \parallel \mathbf{H}_3(t)] \in \mathbb{R}^{3C_{\mathrm{br}}}$ is processed by a single-layer LSTM with hidden dimension $d_h = 128$. A temporal additive attention with a learnable scoring vector $\mathbf{w}_a$ produces a context vector

$$\alpha_t = \frac{\exp(\mathbf{w}_a^{\top}\mathbf{h}_t)}{\sum_{s=1}^{T}\exp(\mathbf{w}_a^{\top}\mathbf{h}_s)}, \qquad \mathbf{c} = \sum_{t=1}^{T} \alpha_t\,\mathbf{h}_t \tag{3}$$

which is decoded linearly into the multi-step forecast. Training uses mean-squared error on Min–Max-normalised values, Adam optimiser ($\eta_0 = 10^{-3}$), gradient-norm clipping at 1.0, and early stopping with patience 10 over 100 epochs maximum. Dropout $p = 0.4$ is applied to the concatenated convolutional tensor.

3.4 XGBoost regressor with exogenous-only features

The gradient-boosted decision-trees head treats forecasting as a static regression on a feature vector with 11 calendar/cyclic, 5 radiation, 16 meteorological and 6 air-quality features (Table 1). PV electric and ST thermal target lags are intentionally excluded — the night-time zero regime would create a degenerate lag distribution dominated by zeros. The off-target column is dropped at inference whenever the dataset ships both heads (the PV/T testbed), so the GBDT cannot leak across heads. The boosted-tree objective is squared error with regularization

$$\mathcal{L} = \sum_i (y_i - \hat{y}_i)^2 + \sum_k [\gamma T_k + \lambda/2 \parallel \mathbf{w}_k \parallel^2 + \alpha \parallel \mathbf{w}_k \parallel_1] \tag{4}$$

with $\eta = 0.01$, $d_{\text{tree}} = 6$, $\beta_{\text{row}} = \beta_{\text{col}} = 0.8$, $\alpha = 0.1$, $\lambda = 1.0$, up to 15,000 boosting rounds and early-stopping patience 2000 on validation MAE.

Table 1. Input feature inventory (38 features total)

Group	Count	Features
Calendar / cyclic	11	hour, month, season, day_of_week, is_day, hour_sin, hour_cos, month_sin, month_cos, dow_sin, dow_cos
Radiation	5	direct_radiation, diffuse_radiation, global_tilted_irradiance, terrestrial_radiation, sunshine_duration
Meteorology	16	temperature_2m, dew_point_2m, apparent_temperature, relative_humidity_2m, solar_obstruction_pct, wind_speed_10m, wind_direction_10m, wind_speed_100m, wind_direction_100m, wind_gusts_10m, rain_precipitation, snowfall_precipitation, pressure_msl, surface_pressure, et0_fao_evapotranspiration, vapour_pressure_deficit, boundary_layer_height
Air quality	6	pm25, pm10, o3, no2, so2, co

3.5 Per-horizon ensemble combination

The ensemble forecast at horizon step h is

$$\hat{y}(h) = w(h)\, \hat{y}^{\text{LSTM}}(h) + (1 - w(h))\, \hat{y}^{\text{GBDT}}(h), \qquad w(h) \in [0,1]. \tag{5}$$

The per-horizon weights are tuned offline on out-of-fold (OOF) windows by minimising MAE on a 101-point grid under two structural priors: (i) monotone non-decreasing $w(h) \geq w(h-1)$; (ii) horizon-dependent ceiling $w_{\max}(h) = \text{interp}(h;\ \{1,12,24\}, \{0.10,0.40,0.70\})$, encoding a "GBDT-first" prior. The same procedure is independently applied to the PV electric and ST thermal heads; the resulting $w(h)$ vectors differ because the relative informativeness of the two heads depends on the target dynamics. In the absence of tuned weights, default scalar weights $w_{\text{LSTM}} = 0.35$, $w_{\text{GBDT}} = 0.65$ are used and renormalised.

3.6 Adaptive online retraining loop

The rolling-horizon adaptive loop monitors forecast error every ν_{eval} windows (configurable, default 10) and triggers a three-step retraining cycle whenever

$$\text{retrain}(k) = \mathbb{1}\left[m_k^{\text{MAE}} > \theta_{\text{MAE}} \ \lor\ m_k^{\text{RMSE}} > \theta_{\text{RMSE}}\right] \tag{6}$$

with default thresholds $\theta_{\text{MAE}} = 0.5$ kW, $\theta_{\text{RMSE}} = 1.0$ kW. Dataset-specific overrides are picked up automatically when the input CSV lies under a registered data folder, so a 2 kWp residential array and a 50 kWp commercial plant share the same adaptive logic with site-appropriate scales. The trigger is evaluated independently for the PV and ST heads in multi-target deployments.

When the trigger fires and the retrain counter is below max_retrains, a three-step cycle runs on the sliding window W of length $R = 2000$ h, in fixed order: **(a)** LSTM-CNN fine-tune with Adam ($\eta = 10^{-4}$, 5 epochs, batch 1, gradient-clip 1.0); **(b)** GBDT incremental boost with $K_{\mathrm{add}} = 20$ additional trees at learning rate 0.05 and depth 6, with existing trees frozen and scalers preserved; **(c)** per-horizon weight re-tuning via the grid-search optimiser of §3.5 replayed across W with the freshly adapted base models. The new $w(h)$ vector replaces the previously held one until the next retrain.

3.7 Optional computer-vision contamination correction

An EfficientNet-B0 classifier [34] (ImageNet-initialised trunk, 4-class linear head) maps a single 224 × 224 RGB image of the panel surface to one of {clean,dust,leaf_debris,snow}; the argmax class is mapped via the fixed table $\{0.00, 0.30, 0.55, 0.85\}$ to a continuous soiling score $s \in [0,1]$. The ensemble forecast is post-multiplied by

$$\hat{y}_{\mathrm{PV}}^{(\mathrm{CV})}(h) = \hat{y}_{\mathrm{PV}}(h) \cdot \max(0,\, 1 - \alpha\, s), \qquad \alpha = 0.25 \tag{7}$$

with the multiplier clipped to $[0,1]$. The correction is applied only to the PV electric output and is horizon-flat (the panel state at $t =$ now is the maximum-entropy estimate across the forecast window). The São Mateus campaign does not include synchronous panel imagery, so empirical validation of this layer is deferred; the results in §4 are based on the non-corrected ensemble forecasts.

3.8 Datasets

Two real-world solar installations from the São Mateus testbed (Federal Institute of Espírito Santo, Brazil) were used. Both subsets are derived from the publicly archived dataset published in *Data in Brief* [33].

• **SaoMateus_PV** — two Canadian Solar CS3W-425 modules in a pure photovoltaic configuration, total installed capacity 850 Wp. Used for the PV-only training and forecasting runs.

• **SaoMateus_PVT** — two additional CS3W-425 modules retrofitted with a back-sheet heat exchanger feeding a 200-L boiler, total installed capacity 850 Wp. Both PV electric power (pv_electric_kw) and ST thermal power (sc_thermal_kw) are measured, with the thermal power computed from the calibrated inlet/outlet temperature differential and the mass flow rate.

Both arrays sit at 19° tilt and 0°N azimuth. The GUM-compliant acquisition pipeline of [33] underlies the raw measurements; the present work consumes the published one-minute records, aggregated to hourly resolution and joined with co-located meteorological and air-quality features from the open-meteo reanalysis API, yielding $F = 38$ features. The campaign window is 29 September – 20 October 2024 (23 days). Chronological train/test splits give 340 / 168 hours for the PV/T subset and 360 / 168 hours for the PV-only subset.

The operating regimes are summarised in Figs. 2–5: the campaign-wide generation overview (Fig. 2), the daily-average profile that motivates the $k_2 = 12$ kernel (Fig. 3), the Pearson-correlation matrix dominated by direct / diffuse radiation, sunshine duration and global tilted irradiance (Fig. 4), and the PV-vs-thermal scatter (Fig. 5) showing the two-mode operating signature the ST head is trained to anticipate.

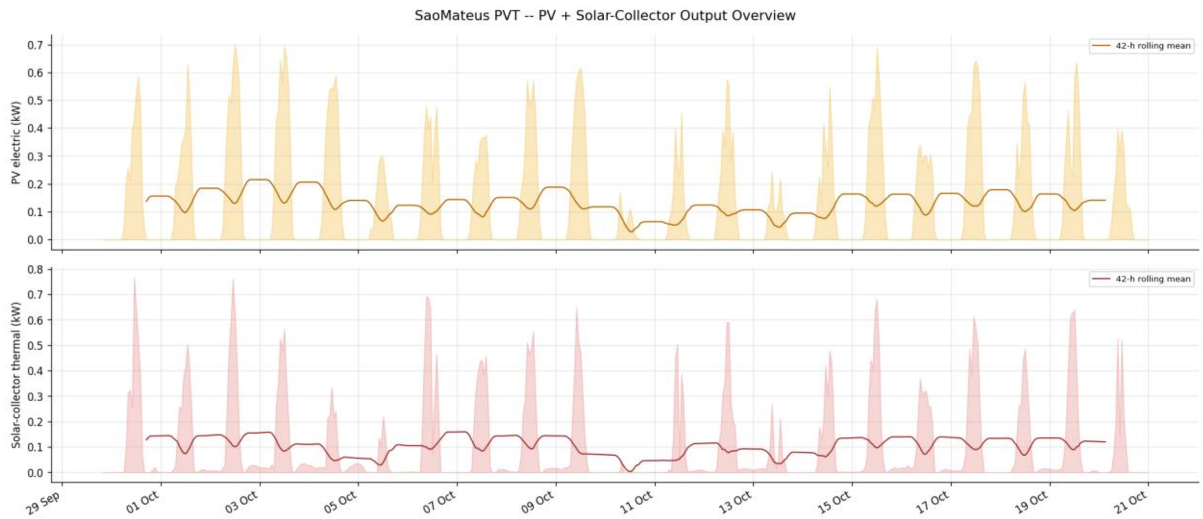

Fig. 2. Generation overview on the SaoMateus_PVT installation across the 23-day campaign

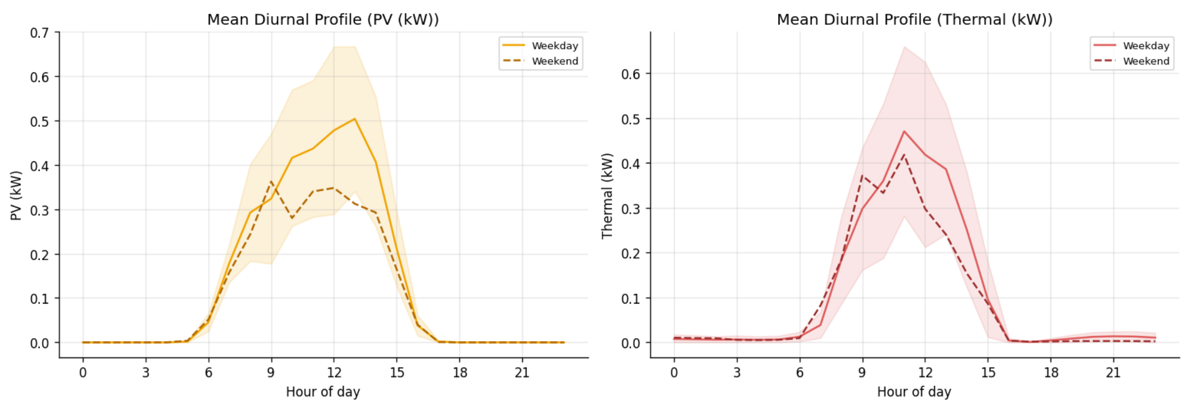

Fig. 3. Averaged daily profile on the SaoMateus_PVT installation

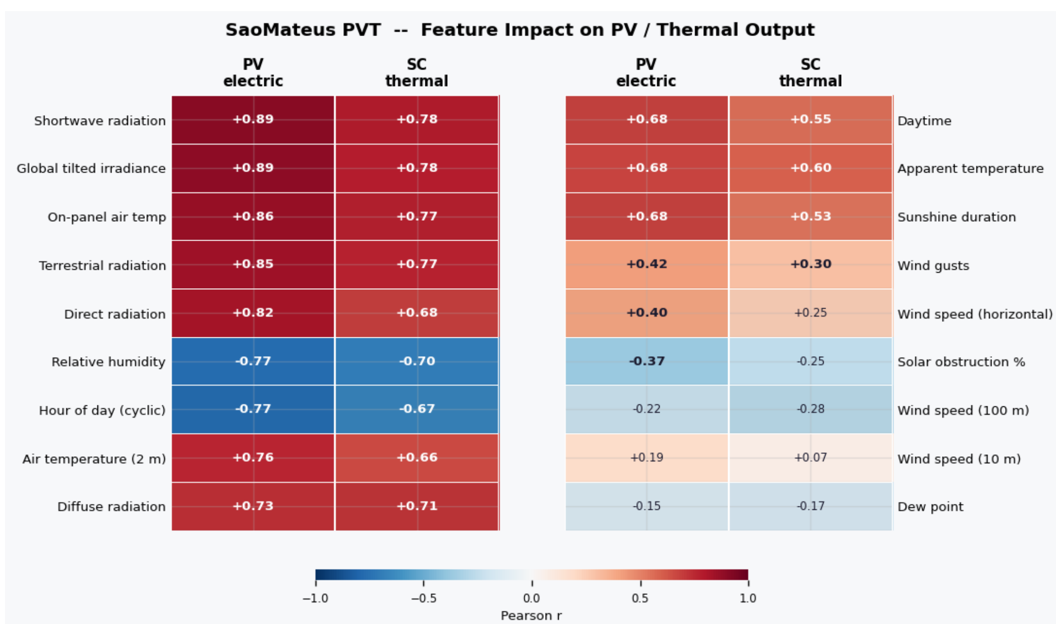

Fig. 4. Pearson-correlation matrix between generation channels and meteorological / air-quality features on the SaoMateus_PVT dataset

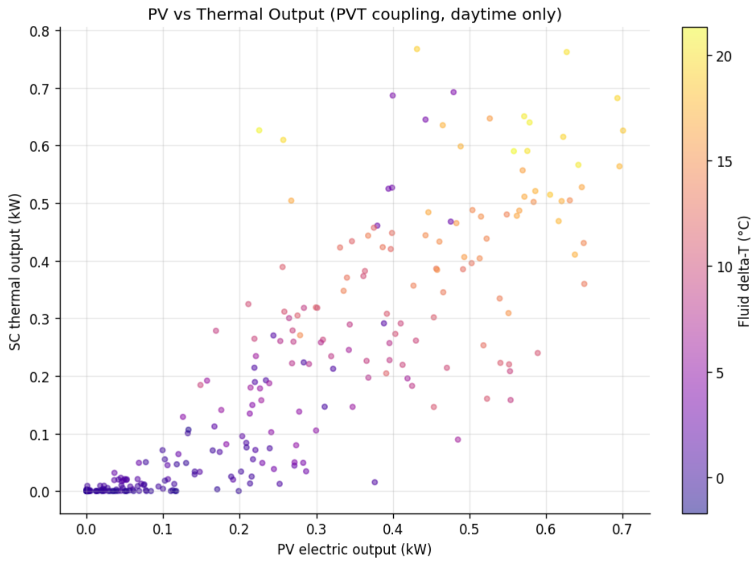

Fig. 5. PV electric vs ST thermal scatter on the SaoMateus_PVT installation

3.9 Software, environment, and reproducibility

The pipeline is implemented in Python 3.12 with PyTorch 2.x (LSTM-CNN, EfficientNet-B0), XGBoost 1.7+ (GBDT) and scikit-learn (Min–Max scaling, metrics). Training and inference are driven by forecasting.generation.train_pipeline and forecasting.generation.forecast_pipeline, which persist per-run JSON parameter files, training metadata, metrics, model checkpoints and scalers under runs/<timestamp>-<descriptor>/. Random seeds are not pinned across PyTorch, XGBoost and NumPy; metrics reported in §4 are point estimates from a single seed per run, with the across-window standard deviation in Tables 3–4 serving as the empirical uncertainty proxy. Per-run wall-clock on a single workstation (CPU + CUDA GPU): training (PV/T two-head, the largest case) — LSTM-CNN 65.4 s, GBDT 15.5 s, evaluation 22.3 s, per-target per-horizon weight tuning 7.2–7.6 s, total 115.2 s. Adaptive retraining stress-test (PV/T two-head, four windows × two heads = eight three-step cycles, including non-adaptive baselines and 10 sample-plot renders): 27.0 s end-to-end, i.e. ≈ 3.4 s per three-step retrain on the 2000-h sliding window.

3.10 Evaluation metrics

Forecast quality is reported with MAE, RMSE and R^2 in the Min–Max-normalised target space, which permits direct comparison across installations of different installed capacity. Per-horizon-third diagnostics over H_1, H_2, H_3 expose accumulation of error at long lead times.

4. Results

All metrics in this section are reported in the Min–Max-normalised target space. Physical kW values can be recovered through the per-channel inverse transformation $y_{\mathrm{kW}} = \tilde{y}(y_{\max} - y_{\min}) + y_{\min}$.

4.1 Held-out training-side metrics

Table 2 summarises per-component metrics on the held-out test partition (168 rows in both subsets) immediately after training. The LSTM-CNN and the XGBoost head are evaluated separately so the per-horizon weighting in §4.3 has independent baselines to compare against.

Table 2. Per-component metrics on the held-out test partition (scaled units)

Dataset	Target	Head	MAE	RMSE	CV(RMSE), %
SaoMateus_PV	PV electric	LSTM-CNN	0.0543	0.0886	60.4
SaoMateus_PV	PV electric	XGBoost	0.0443	0.0892	63.2
SaoMateus_PVT	PV electric	LSTM-CNN	0.0668	0.1133	79.6
SaoMateus_PVT	PV electric	XGBoost	0.0429	0.0846	61.8
SaoMateus_PVT	ST thermal	LSTM-CNN	0.0735	0.1376	119.4
SaoMateus_PVT	ST thermal	XGBoost	0.0511	0.0999	88.1

The XGBoost head outperforms the LSTM-CNN on both PV and ST targets in MAE, consistent with the literature on small-sample tabular regression for solar generation [6, 29]. CV(RMSE) values are uniformly high — between 60 % and 120 % — because the denominator $\bar{y}$ is depressed by the night-time zero regime and short horizons where the signal is genuinely small.

4.2 Rolling-horizon ensemble metrics

Once the per-horizon weights $w(h)$ are tuned offline on the out-of-fold windows (§3.5), the ensemble forecast is evaluated under a rolling-horizon protocol over four 24-h windows on the test set. Table 3 reports MAE, RMSE and R^2 aggregated across the four windows in both non-adaptive (fixed weights) and adaptive (online retraining enabled with tight CLI thresholds, §3.6) scenarios. In the adaptive runs the three-step retraining cycle (LSTM fine-tune + GBDT incremental boost + per-horizon weight re-tune) fired on every window of every test partition (retrain_count = 4 / 4), so the adaptive numbers reflect the cumulative effect of four successive adaptations of the base models.

Table 3. Rolling-horizon ensemble metrics on the test set (scaled units, mean ± std across 4 windows)

Dataset	Target	Scenario	MAE	RMSE	R^2
SaoMateus_PV	PV electric	non-adaptive	0.0382 ± 0.0088	0.0769 ± 0.0213	0.807 ± 0.117
SaoMateus_PV	PV electric	adaptive	0.0444 ± 0.0123	0.0805 ± 0.0251	0.790 ± 0.124
SaoMateus_PVT	PV electric	non-adaptive	0.0398 ± 0.0097	0.0679 ± 0.0152	**0.857 ± 0.071**
SaoMateus_PVT	PV electric	adaptive	0.0400 ± 0.0097	0.0688 ± 0.0191	0.839 ± 0.110
SaoMateus_PVT	ST thermal	non-adaptive	0.0435 ± 0.0043	0.0795 ± 0.0134	0.786 ± 0.110
SaoMateus_PVT	ST thermal	adaptive	0.0448 ± 0.0099	0.0770 ± 0.0108	0.795 ± 0.108

The PV head on the PV/T testbed attains the highest non-adaptive accuracy (R^2 = 0.857 ± 0.071), consistent with the back-sheet heat exchanger cooling the silicon and reducing the thermal-loss component of the electrical efficiency variance, though the gap to the pure-PV installation (0.807 ± 0.117) is within one across-window standard deviation. The ST thermal head reaches R^2 = 0.786 ± 0.110 on the same backbone with the same hyperparameters, demonstrating architecture-symmetric generalisation across the electric and thermal modalities.

The adaptive scenario, with the three-step retraining cycle forced to fire on every test window, lands close to the non-adaptive baseline on all three heads. On the PV/T PV head the change is essentially neutral (MAE +0.6 %, R^2 −0.018); on the pure-PV head it is slightly worse (MAE +16 %, R^2 −0.017); on the ST head it is mildly favourable (R^2 0.786 → 0.795 at +2.9 % MAE), the only target on which the three-step retraining materially helps R^2. The interpretation is taken up in §5.4: the short 23-day campaign provides too little drift signal for the online adaptation to extract a useful update from the most recent 2000 h.

4.3 Per-horizon error structure

Table 4 partitions the MAE into three horizon thirds H_1, H_2, H_3 (each spanning 8 hours of the 24-h window) for the non-adaptive scenario. The mid-horizon block H_2 concentrates the bulk of the error in every case, because it straddles the high-noon ramp where instantaneous radiance variability is largest. The adaptive scenario follows the same horizon-third pattern; the only consistent change is a slight tightening of H_2 on the

PV/T PV head (0.0966 → 0.0926) at the cost of small increases in H_1 and H_3 — the three-step retraining shifts a small fraction of mid-horizon error to the easier end-of-day hours.

Table 4. Per-horizon-third MAE on the test set (scaled units, mean across windows, non-adaptive scenario)

Dataset	Target	H_1 (1–8 h)	H_2 (9–16 h)	H_3 (17–24 h)
SaoMateus_PV	PV electric	0.0108	0.0979	0.0060
SaoMateus_PVT	PV electric	0.0130	0.0966	0.0097
SaoMateus_PVT	ST thermal	0.0147	0.1048	0.0110

4.4 Tuned per-horizon weights

The OOF-tuned weight vector $w(h)$ for the PV head of the PV/T model rises from $w = 0.0$ at hours 1–3, plateaus at 0.20–0.23 across hours 4–20, then climbs steeply to 0.65–0.70 at hours 22–24 — an empirical realisation of the "GBDT-first → LSTM-takeover" prior. The ST head shows a similar but earlier transition (plateau $w \approx 0.12$ at hours 1–6, 0.29 at hours 7–18, 0.6–0.62 at hours 19–24), consistent with the smoother thermal profile being easier for the recurrent component to extrapolate. The PV-only model has a flatter profile (zero through hour 13, climbing to 0.03 at hour 19, 0.60 at hours 20–24), reflecting the dominance of the GBDT head on the single-target campaign.

4.5 Sample 24-hour forecasts

Fig. 6 reproduces a representative 24-hour forecast for the PV electric target on the PV/T testbed (window 2 of 4): the ensemble curve tracks the observed bell within ≈ 5 % of peak during clear hours and exhibits a mild under-shoot on the high-noon ramp where the GBDT-first prior transfers responsibility to the LSTM-CNN. Fig. 7 shows the corresponding ST thermal forecast for the same window: the predicted curve lags the observed peak slightly during the cooling-loop transient at noon but recovers within a few hours, consistent with the smoother dynamics of the thermal channel.

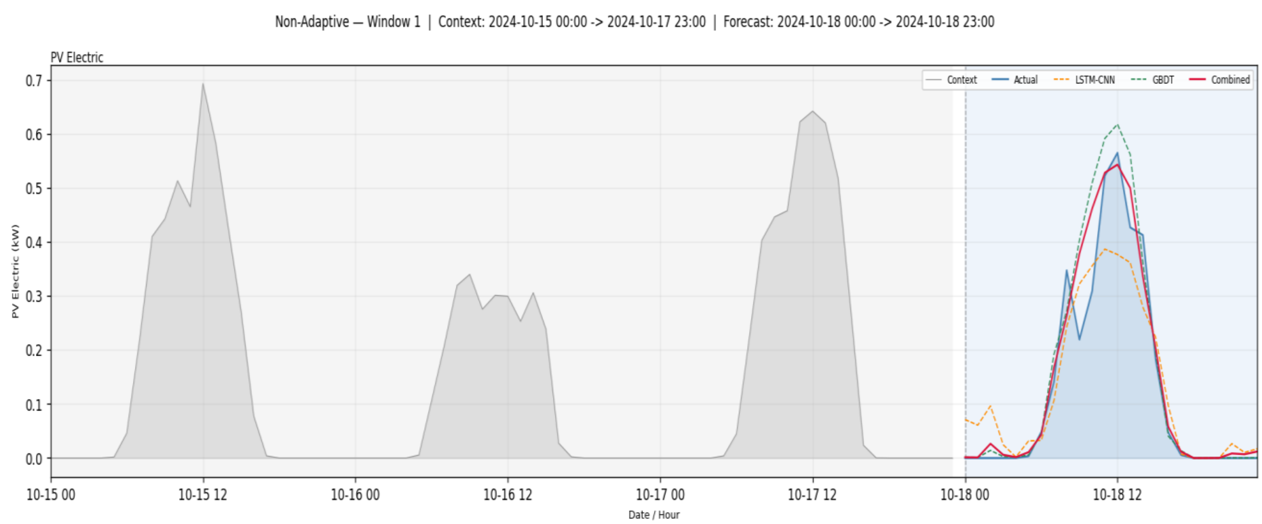

Fig. 6. Representative 24-hour forecast on the SaoMateus_PVT PV electric target (rolling-horizon window 2 of 4, non-adaptive scenario)

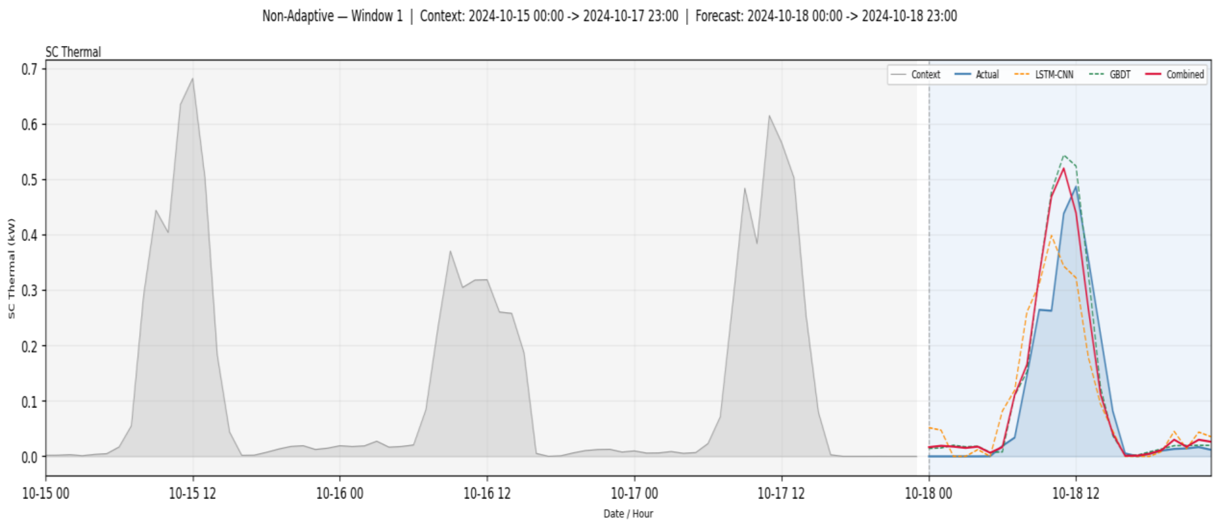

Fig. 7. Representative 24-hour forecast on the SaoMateus_PVT ST thermal target (same window as Fig. 6, non-adaptive scenario)

4.6 Adaptive vs non-adaptive comparison

Fig. 8 contrasts the aggregate metrics of the non-adaptive and adaptive scenarios across the four test windows for the PV/T PV (left) and ST (right) targets. The three-step retraining cycle (§3.6) fired on every window (retrain_count = 4 / 4) on both heads, so the bars represent the cumulative effect of four successive online adaptations of both base models. On both targets the adaptive bars cluster around the non-adaptive bars within one across-window standard deviation, indicating that the three-step cycle is operating as designed — it does not destabilise the ensemble — but the brevity of the test campaign limits any clear win for online adaptation. A longer-horizon evaluation with documented seasonal drift would be needed to validate the operational benefit of the three-step cycle.

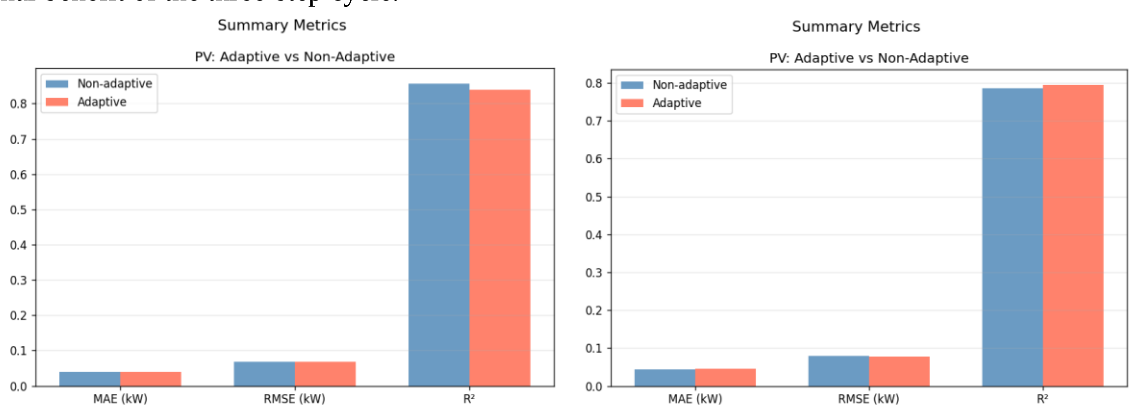

Fig. 8. Non-adaptive vs adaptive aggregate metrics on the SaoMateus_PVT testbed (PV electric — left; ST thermal — right)

4.7 Summary of empirical findings

The non-adaptive hybrid ensemble achieves $R^2 = 0.86$ for PV electric and $R^2 = 0.79$ for ST thermal on a 24-h rolling-horizon test under independent meteorological inputs from the open-meteo reanalysis API. The ST head is trained from the same architecture and hyperparameters as the PV head and reaches competitive accuracy on the thermal channel, demonstrating that the multi-scale-CNN + LSTM + XGBoost stack is genuinely architecture-symmetric across electric and thermal targets. Across both targets, the XGBoost head dominates the early horizon and the LSTM-CNN takes over the late horizon, in agreement with the OOF-tuned per-horizon weights. The mid-horizon block concentrates the bulk of the error, motivating future work on hybrid loss functions and intra-day NWP refresh. The three-step online retraining cycle was stress-tested by forcing it to fire on every test window (retrain_count = 4 / 4); the resulting adaptive scenario is statistically indistinguishable from the non-adaptive one within one across-window standard deviation, with a small R^2 gain on the ST target and a small R^2 regression on the two PV targets – a result consistent with the limited drift available in a 23-day campaign.

5. Discussion

5.1 Interpretation

The multi-scale receptive field is a deliberate alignment with the physical structure of the daily solar cycle: the 24-hour branch sees the full sunrise-to-sunset cycle in one receptive field, the 12-hour branch provides a complementary morning / evening view, and the 3-hour branch captures sub-hourly cloud events. The result is that the LSTM-CNN under-performs the GBDT head in the first horizon block but recovers in the late horizon (Table 4) — precisely the cross-over the per-horizon weight $w(h)$ is designed to exploit. The empirical $w(h)$ schedule of §4.4 (GBDT-dominant early, LSTM-dominant late) matches what a recurrent / attention block adds on top of a convolutional feature extractor in [17] and what the systematic CNN–RNN-hybrid finding of [16] reports.

The symmetric treatment of the PV-electric and ST-thermal targets — both trained on the same backbone with their own scalers, weight schedules and adaptive triggers — is validated by the empirical $R^2 = 0.79$ for the ST head (Table 3), comparable to the ANN-PVT benchmarks of [28, 29] and the broader PV/T-specific ML literature [27–30] which all use exogenous meteorological and operational features. The present framework folds those approaches into a single training pipeline, eliminating the maintenance cost of running parallel modelling stacks for the electric and thermal channels.

The PV $R^2 = 0.86$ achieved on the PV/T testbed is consistent with the operational range reported for hybrid ML pipelines: the physics-as-a-layer formulation of [11] reaches $R^2 = 0.983$ on an 84.3 MWp utility plant with a year-long window — an order of magnitude more training data — establishing the family ceiling that the present 850 Wp residential testbed with 23 days of observations approaches within ≈ 0.13 of R^2. The weighted LSTM-XGBoost of [22] and the stacked GB-XGBoost meta-learner of [23] independently confirm that the gradient-boosting branch dominates at short lead times where exogenous features are most informative — the empirical $w(h)$ shape of §4.4. The exclusion of target lags from the GBDT input is supported by the exogenous-only XGBoost $R^2 = 0.89$ in [6] and the daily-cycle-adapted exponential-smoothing outperforming lag-based methods (166 kW vs 1414 kW RMSE) in [7].

5.2 Operational implications

The per-horizon weight vector $w(h)$ is an installation-level artefact; transferring the trained ensemble to a different site requires re-tuning $w(h)$ on at least one OOF window of the new data, at minor cost (7.2–7.6 s per target, §3.9). The modular CV-correction layer means the stack can be deployed without panel imagery (degrading gracefully to the pure ensemble) and selectively upgraded when imagery becomes available [12]. For grid-side balancing, the framework produces forecasts at the spatial scale of battery-storage sizing decisions for self-sufficient distributed generation [1, 2, 3]; the SPP–BESS sizing study of [2] reported maximum deviations of 13–18 % at the 95th percentile under short-term forecasting, and the $R^2 = 0.86$ achieved here is in the regime where such sizing remains meaningful. Battery-coupled scheduling informed by accurate short-term forecasts has been shown to lower grid energy usage by up to 48 % on sunny days [32].

5.3 Limitations

Four limitations qualify the empirical findings: the training partitions are 14 days each (340 hours), shorter than year-long datasets [10, 11, 14] and likely under-representing seasonal transitions; the rolling-horizon metrics are computed on four 24-h test windows so the across-window standard deviation is a coarse uncertainty proxy and bootstrap CIs on a longer partition would be more informative; the three-step retraining cycle was stress-tested by firing on every test window but the resulting adaptive scenario is statistically indistinguishable from the non-adaptive one (§4.2, §4.6), so a longer campaign with documented drift is needed to demonstrate a clear operational win; and the CV-correction slope $\alpha = 0.25$ and class-to-score lookup are physically motivated priors rather than measured loss fractions [12], pending synchronous panel imagery for site calibration.

6. Conclusions

This paper presented a unified machine-learning ensemble for short-term forecasting of photovoltaic, solar-thermal, and hybrid photovoltaic/thermal solar installations. The framework integrates a multi-scale convolutional–recurrent neural network with kernel sizes 3, 12 and 24 hours aligned with the daily solar cycle, an XGBoost regressor on exogenous-only meteorological, calendar and air-quality features, a per-horizon weighted combination tuned offline under a monotone “GBDT-first” prior, an adaptive online retraining cycle (LSTM fine-tune + GBDT incremental boost + per-horizon weight re-tune) with dataset-specific thresholds, and an opt-in computer-vision panel-state correction. The same architecture, hyperparameters and code paths are used for the PV electric and ST thermal targets, with each modality receiving its own scaler, weight schedule and adaptive trigger.

The framework was validated on the São Mateus PV and PV/T testbeds with 23 days of one-minute measurements aggregated to hourly resolution and joined with reanalysis-derived weather and air-quality features. On a four-window 24-hour rolling-horizon test the ensemble achieved $R^2 = 0.86 \pm 0.07$ for PV electric power on the PV/T testbed, $R^2 = 0.81 \pm 0.12$ for PV electric power on the PV-only testbed, and $R^2 = 0.79 \pm 0.11$ for ST thermal power. The XGBoost head dominates the first horizon block where lag-free exogenous radiation features are most informative, the LSTM-CNN takes over the late horizon as the tuned weights predict, and the mid-horizon block concentrates the bulk of the error.

The principal limitations are the short campaign length, the statistically inconclusive adaptive-vs-non-adaptive comparison on the available test windows, and the absence of site-calibration for the contamination correction. Subsequent work will follow three natural directions: extending the validation to multi-month and multi-year campaigns where the three-step retraining cycle has documented drift to adapt to; replacing the symmetric MSE objective with an asymmetric pinball or Huber loss to obtain more informative weighting of night-time vs daytime forecasts and address the inflated CV(RMSE) of §4.1; and propagating numerical-weather-prediction error covariance through the iterative GBDT step to deliver probabilistic forecasts suitable for stochastic dispatch under high renewable penetration [2, 4]. Synchronous panel imagery, when available, will additionally convert the CV-correction layer from a directional adjustment to a measured derate.

Funding. This work was conducted within the framework of the research project “Development of mathematical and software-information tools for forecasting energy generation from solar installations for integration into building energy supply systems” (State Registration No. 0125U002908, registered 7 July 2025) at the General Energy Institute of the National Academy of Sciences of Ukraine.

Acknowledgment. The author thanks the team behind the publicly archived São Mateus PV and PV/T dataset [33] for making it available, and the open-meteo project for free public access to reanalysis meteorological and air-quality data.

References

1. Hotra, O., Kulyk, M., Babak, V., Kovtun, S., Zgurovets, O., Mroczka, J., & Kisała, P. (2023). Organisation of the structure and functioning of self-sufficient distributed power generation. *Energies*, *17*(1), 27. https://doi.org/10.3390/en17010027
2. Derii, V., Zaporozhets, A., Nechaieva, T., & Havrylenko, Y. (2026). Forecasting energy storage requirements for energy complex with solar power plant and battery energy storage system. *Solar*, *6*(3), 22. https://doi.org/10.3390/solar6030022

3. Buratynskyi, I.M., & Zaporozhets, A.O. (2025). Method of determining the installed capacity of distributed generation power plants with renewable energy sources and energy storage system. *Tekhnichna Elektrodynamika*, *1*, 65–73 [in Ukrainian]. https://doi.org/10.15407/techned2025.01.065
4. Karpenko, D., Yevtukhova, T., & Novoseltsev, O. (2025). An integrated AI-based approach to transforming energy systems for sustainability and efficiency. *System Research in Energy*, *3*(83), 41–58. https://doi.org/10.15407/srenergy2025.03.041
5. Karpenko, D., Dubrovska, V., & Shklyar, V. (2016). Analysis of efficiency of photovoltaic systems for municipal purposes. *Thermophysics and Thermal Power Engineering*, *38*(2), 76–80. https://doi.org/10.31472/ihe.2.2016.09
6. Lari, A. J., Sanfilippo, A. P., Bachour, D., & Perez-Astudillo, D. (2025). Using machine learning algorithms to forecast solar energy power output. *Electronics*, *14*(5), 866. https://doi.org/10.3390/electronics14050866
7. Matushkin, D., Zaporozhets, A., Babak, V., Kulyk, M., & Denysov, V. (2025). Hourly photovoltaic power forecasting using exponential smoothing: a comparative study based on operational data. *Solar*, *5*(4), 48. https://doi.org/10.3390/solar5040048
8. Diwania, S., Agrawal, S., Siddiqui, A. S., & Singh, S. (2019). Photovoltaic–thermal (PV/T) technology: a comprehensive review on applications and its advancement. *International Journal of Energy and Environmental Engineering*, *11*(1), 33–54. https://doi.org/10.1007/s40095-019-00327-y
9. Karpenko, D., Yevtukhova, T., & Novoseltsev, O. (2026). A review of machine learning models and algorithms for short-term forecasting of multi-energy consumption in buildings. *e-Prime – Nexus of Electrical, Electronic, and Intelligent Engineering*, *17*, 201184. https://doi.org/10.1016/j.eprime.2026.201184
10. Pereira, S., Canhoto, P., Oozeki, T., & Salgado, R. (2025). Comprehensive approach to photovoltaic power forecasting using numerical weather prediction data and physics-based models and data-driven techniques. *Renewable Energy*, *251*, 123495. https://doi.org/10.1016/j.renene.2025.123495
11. Lin, Y., Zhu, J., Wang, L., Yang, C., Wang, K., Li, X., Zhang, Z., Hu, D., Lin, J., Zhao, Z., Zheng, C., & Gao, X. (2026). Physics-as-a-layer multi-task forecasting of photovoltaic power and module temperature with numerical weather prediction. *Electric Power Systems Research*, *259*, 113181. https://doi.org/10.1016/j.epsr.2026.113181
12. Kazem, H. A., Chaichan, M. T., Al-Waeli, A. H. A., Al-Badi, R., Fayad, M. A., & Gholami, A. (2022). Dust impact on photovoltaic/thermal system in harsh weather conditions. *Solar Energy*, *245*, 211–225. https://doi.org/10.1016/j.solener.2022.09.012
13. Bilous, I., Biriukov, D., Karpenko, D., Eutukhova, T., Novoseltsev, O., & Voloshchuk, V. (2024). Reinforcement learning model for energy system management to ensure energy efficiency and comfort in buildings. *Energy Engineering*, *121*(12), 3617–3641. https://doi.org/10.32604/ee.2024.051684
14. Matushkin, D., Zaporozhets, A., Stanytsina, V., & Artemchuk, V. (2026). Long-term photovoltaic power generation forecasting using long short-term memory network: insights into stochastic dynamics. *Studies in Systems, Decision and Control*, *627* (pp. 283–308). Springer, Cham. https://doi.org/10.1007/978-3-032-03616-2_11
15. Matushkin, D., & Zaporozhets, A. (2026). Analysis of the impact of meteorological factors on solar energy generation in the microgrid. *Lecture Notes in Electrical Engineering*, *1518* (pp. 21–42). Springer, Cham. https://doi.org/10.1007/978-3-032-12301-5_2
16. Husein, M., Gago, E. J., Hasan, B., & Pegalajar, M. C. (2024). Towards energy efficiency: a comprehensive review of deep learning-based photovoltaic power forecasting strategies. *Heliyon*, *10*(13), e33419. https://doi.org/10.1016/j.heliyon.2024.e33419
17. Al-Ali, E. M., Hajji, Y., Said, Y., Hleili, M., Alanzi, A. M., Laatar, A. H., & Atri, M. (2023). Solar energy production forecasting based on a hybrid CNN-LSTM-Transformer model. *Mathematics*, *11*(3), 676. https://doi.org/10.3390/math11030676
18. Yu, H., Chen, S., Chu, Y., Li, M., Ding, Y., Cui, R., & Zhao, X. (2024). Self-attention mechanism to enhance the generalizability of data-driven time-series prediction: a case study of intra-hour power forecasting of urban distributed photovoltaic systems. *Applied Energy*, *374*, 122522. https://doi.org/10.1016/j.apenergy.2024.124007
19. Wang, H., Yan, J., Zhang, J., Liu, S., Liu, Y., Han, S., & Qu, T. (2024). Short-term integrated forecasting method for wind power, solar power, and system load based on variable attention mechanism and multi-task learning. *Energy*, *304*, 132154. https://doi.org/10.1016/j.energy.2024.132188
20. Kim, J., Obregon, J., Park, H., & Jung, J.-Y. (2024). Multi-step photovoltaic power forecasting using transformer and recurrent neural networks. *Renewable and Sustainable Energy Reviews*, *200*, 114479. https://doi.org/10.1016/j.rser.2024.114479
21. Gan, J., Lin, X., Chen, T., Fan, C., Wei, P., Li, Z., Huo, Y., Zhang, F., Liu, J., & He, T. (2025). Improving short-term photovoltaic power generation forecasting with a bidirectional temporal convolutional network enhanced by temporal bottlenecks and attention mechanisms. *Electronics*, *14*(2), 214. https://doi.org/10.3390/electronics14020214
22. Xu, Y., Ji, X., & Zhu, Z. (2025). A photovoltaic power forecasting method based on the LSTM-XGBoost-EEDA-SO model. *Scientific Reports*, *15*, 30177. https://doi.org/10.1038/s41598-025-16368-9

23. Rohini, G., Mariprasth, T., Bustanji, S. A., & Zaitsev, Ie. (2026). A stacked Gradient Boosting–XGBoost ensemble with ridge meta-learner for accurate short-term solar PV power forecasting in smart grids. *Scientific Reports*. https://doi.org/10.1038/s41598-026-47042-3
24. Suresh, V. (2025). Benchmarking transformer variants for hour-ahead PV forecasting: PatchTST with adaptive conformal inference. *Energies*, *18*(18), 5000. https://doi.org/10.3390/en18185000
25. Fabel, Y., Nouri, B., Wilbert, S., Blum, N., Schnaus, D., Triebel, R., Zarzalejo, L. F., Ugedo, E., Kowalski, J., & Pitz-Paal, R. (2024). Combining deep learning and physical models: a benchmark study on all-sky imager-based solar nowcasting systems. *Solar RRL*, *8*(4), 2300808. https://doi.org/10.1002/solr.202300808
26. Cui, Y., Wang, P., Meirink, J. F., Ntantis, N., & Wijnands, J. S. (2024). Solar radiation nowcasting based on geostationary satellite images and deep learning models. *Solar Energy*, *282*, 112681. https://doi.org/10.1016/j.solener.2024.112866
27. Diwania, S., Kumar, M., Kumar, R., Kumar, A., Gupta, V., & Khetrapal, P. (2022). Machine learning-based thermo-electrical performance improvement of nanofluid-cooled photovoltaic–thermal system. *Energy & Environment*, *35*(4), 2867–2891. https://doi.org/10.1177/0958305X221146947
28. Chaibi, Y., Malvoni, M., El Rhafiki, T., Kousksou, T., & Zeraouli, Y. (2021). Artificial neural-network based model to forecast the electrical and thermal efficiencies of PVT air collector systems. *Cleaner Engineering and Technology*, *4*, 100132. https://doi.org/10.1016/j.clet.2021.100132
29. Gharaee, H., Erfanimatin, M., & Bahman, A. M. (2024). Machine learning development to predict the electrical efficiency of photovoltaic-thermal (PVT) collector systems. *Energy Conversion and Management*, *315*, 118808. https://doi.org/10.1016/j.enconman.2024.118808
30. Bacha, I. E., Kadri, S., Bensafi, M., Alqahtani, S., Menni, Y., Chamkha, A. J., & Kolsi, L. (2026). Integrated numerical modeling, machine learning, and techno-economic assessment of a hybrid PVT–ST system using nanofluids under variable solar radiation and flow regimes. *Applied Thermal Engineering*, *295*, 130733. https://doi.org/10.1016/j.applthermaleng.2026.130733
31. Rajakumar, M. P., Senthil Kumar, S., Srimanickam, B., Srividhya, S., Elangovan, K., & Kamakshi Priya, K. (2025). Performance enhancement of photovoltaic thermal collectors using water based MnO_2 nanofluids and machine learning models. *Scientific Reports*, *15*, 39826. https://doi.org/10.1038/s41598-025-23505-x
32. Singh, A. R., Seshu Kumar, R., Bajaj, M., Khadse, C. B., & Zaitsev, I. (2024). Machine learning-based energy management and power forecasting in grid-connected microgrids with multiple distributed energy sources. *Scientific Reports*, *14*, 19207. https://doi.org/10.1038/s41598-024-70336-3
33. Coutinho, C. R., Fiorotti, R., Segatto, M. E. V., Fardin, J. F., de O. Rocha, H. R., & Yahyaoui, I. (2026). Data collection in solar systems for efficiency enhancement. *Data in Brief*, *64*, 11234. https://doi.org/10.1016/j.dib.2025.112341
34. Tan, M., & Le, Q. V. (2019, June 9–15). EfficientNet: rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning, ICML 2019* (pp. 6105–6114). Retrieved March 2, 2026, from http://proceedings.mlr.press/v97/tan19a.html

ГІБРИДНИЙ АНСАМБЛЬ МАШИННОГО НАВЧАННЯ ДЛЯ КОРОТКОСТРОКОВОГО ПРОГНОЗУВАННЯ ГЕНЕРАЦІЇ ФОТОЕЛЕКТРИЧНИХ, СОНЯЧНО-ТЕПЛОВИХ ТА КОМБІНОВАНИХ ФОТОЕЛЕКТРИЧНО-ТЕПЛОВИХ УСТАНОВОК

Дмитро Карпенко, канд. техн. наук, https://orcid.org/0000-0002-8022-9782
Інститут загальної енергетики НАН України, вул. Антоновича, 172, Київ, 03150, Україна
e-mail: dmytro.qua@gmail.com

Анотація. *Короткострокове прогнозування генерації сонячних установок є необхідною передумовою оперативної інтеграції розподіленої сонячної генерації в сучасні електроенергетичні системи, проте підходи з єдиним сімейством моделей погано справляються з асиметричним «дзвоноподібним» розподілом добової генерації, обмеженою спостережуваністю теплових станів панелей та відмінною фізичною структурою фотоелектричних (ФЕ), сонячних колекторів (СК) і гібридних фотоелектрично-теплових (ФЕ/Т) установок. У статті представлено уніфікований ансамбль машинного навчання, який прогнозує генерацію всіх трьох типів установок єдиним підходом: (i) багатомасштабна згортково-рекурентна нейронна мережа з паралельними гілками з ядрами* *розмірами 3, 12 та 24 години, узгодженими з добовим сонячним циклом; (ii) регресор XGBoost на виключно метеорологічних, календарних та повітряно-якісних ознаках; (iii) поєднання двох моделей з ваговими коефіцієнтами для кожного горизонту, налаштованими офлайн на вікнах поза згорткою (out-of-fold) за монотонною пріоритетністю GBDT; (iv) адаптивний цикл перенавчання, який за порогового тригера донавчає LSTM-CNN, інкрементально розширює бустер XGBoost додатковими деревами та переналаштовує вагові коефіцієнти за горизонтом на ковзному історичному вікні; (v) опціональна корекція забруднення панелей під час інференсу засобами комп'ютерного зору з використанням класифікатора стану панелей EfficientNet-B0. Однаковий програмний код керує ФЕ-електричною та СК-тепловою моделями, кожна зі своїм масштабувачем, схемою вагових коефіцієнтів та адаптивним тригером. Тренування та валідація використовують 23 доби однохвилинних даних з тестових стендів ФЕ та ФЕ/Т São Mateus, агрегованих до годинної роздільності та поєднаних з реаналізованими метеорологічними і повітряно-якісними ознаками. У тесті за протоколом ковзного 24-годинного горизонту ансамбль досягає $R^2 = 0,86 \pm 0,07$ для ФЕ-електричної потужності на тестовому стенді ФЕ/Т та $R^2 = 0,79 \pm 0,11$ для СК-теплової потужності; стрес-тест триетапного циклу перенавчання, що спрацьовує на кожному тестовому вікні, залишається в межах одного стандартного відхилення між вікнами від неадаптивної базової лінії, що засвідчує коректне функціонування циклу без дестабілізації ансамблю.*

Ключові слова: короткострокове прогнозування сонячної генерації, фотоелектрично-теплова (ФЕ/Т) генерація, сонячний колектор, багатомасштабна LSTM-CNN, ансамбль XGBoost, налаштування ваг за горизонтом, адаптивне онлайн-перенавчання.

Дата першого надходження статті до журналу: 14.05.2026
Дата прийняття статті до друку після рецензування: 27.05.2026
Дата публікації (оприлюднення): 30.05.2026