



Quantum-Enhanced real-Time Microclimate Forecasting and Solar Energy Prediction

P. B. Guru Prasanna¹, N. Bala Santhosh², G. Sai Shashank³, B. Surpriya⁴
Assistant Professor¹, Students of EEE,^{2,3,4} Department of Electrical and Electronics Engineering
MVSRR Engineering College/Osmania University, Hyderabad, India

Abstract: Predicting solar energy output with fine spatial and temporal resolution remains one of the most pressing challenges in distributed photovoltaic (PV) system management. This paper introduces Q-WEFS (Quantum-Enhanced IoT Weather and Energy Forecasting System), a real-time environmental intelligence platform that pairs low-cost IoT sensing hardware with a hybrid classical-quantum machine learning pipeline to enable smarter, more responsive solar energy management at the edge. An ESP32-D0WD-V3 microcontroller collects readings from four physical sensors and streams structured JSON telemetry at 5-second intervals. A RandomForestRegressor trained on 40,896 hourly NASA POWER records for Hyderabad, India (17.38°N, 78.46°E) achieves $R^2 = 0.895$ (MAE = 2.05°C) for next-day (24-hour-ahead) temperature prediction. Building on this strong classical baseline, a Variational Quantum Regressor (VQR) — using a ZZ Feature Map encoder with 5 qubits and a Real Amplitudes ansatz — learns the structured patterns left in the Random Forest's residuals via the parameter-shift rule. The resulting hybrid system reaches $R^2 = 0.912$ (MAE = 1.82°C), a statistically significant 11.2% reduction in mean absolute error relative to the Random Forest baseline ($p = 1.1 \times 10^{-6}$). These figures reflect physically realistic accuracy for a genuine 24-hour-ahead forecasting horizon, following correction of an earlier data-leakage issue in which previously reported metrics had inadvertently reflected 1-hour-ahead persistence forecasting rather than true day-ahead prediction. These quantum-refined temperature forecasts are fed directly into an IEC 61215-compliant NOCT thermal derating model, linking prediction quality to theoretical PV output estimation and battery state-of-charge management. Continuous hardware validation on 08 March 2026 confirmed stable operation across all sensing channels. Notably, the 10-parameter quantum correction layer delivers accuracy improvements comparable to a classical neural-network residual-correction layer that would require over 250 trainable parameters — pointing toward a compelling path for parameter-efficient, energy-aware edge AI.

Index Terms: Internet of Things, Quantum Machine Learning, Solar Energy Forecasting, Random Forest, Variational Quantum Regressor, ESP32, NASA POWER, Photovoltaic Systems, Edge Computing, Hybrid Classical-Quantum Architecture, Next-Day Temperature Forecasting.

I. Introduction

India's rooftop solar sector has grown rapidly over the past decade, with distributed photovoltaic (PV) installations spreading across cities and towns throughout the subcontinent. Managing these systems effectively demands locally calibrated, real-time energy yield forecasting — a requirement that generic commercial weather APIs, which operate at multi-kilometre spatial resolution, are poorly equipped to satisfy [1]. The problem is not merely academic: a 4°C temperature differential between a sun-exposed rooftop and a nearby shaded surface produces a 1.6% swing in panel efficiency under standard temperature-coefficient models, with direct consequences for inverter sizing, battery dispatch decisions, and grid stability. In a city like Hyderabad, where semi-arid conditions create pronounced day-to-day irradiance variability, such gaps in local forecasting accuracy translate into tangible operational losses.

Conventional machine learning approaches — ARIMA, linear regression — are ill-suited to the nonlinear, multi-variable interactions that characterise atmospheric conditions during monsoon season transitions [2]. Deep learning architectures such as LSTMs and Transformers handle this complexity more gracefully, but their computational demands are prohibitive on the low-power microcontrollers typical of edge deployments [3]. Quantum Machine Learning (QML) has attracted growing interest as a potential middle ground, with parameterised quantum circuits (PQCs) offering exponential Hilbert-space expressivity that classical models cannot efficiently replicate [4]. Yet a striking gap persists in the literature: virtually all published QML weather research operates in simulated, offline settings, with no integration of live physical sensor data [5], [6].

Q-WEFS is designed to close this gap. The system makes three primary contributions to the sustainable energy domain: (i) a novel hybrid RF+VQR architecture that specifically targets structured residual errors in next-day temperature forecasting, validated at $p = 1.1 \times 10^{-6}$ statistical significance, and shown to deliver a substantial, statistically significant accuracy improvement over the classical baseline at a genuine 24-hour-ahead



horizon; (ii) to the best of the authors' knowledge, one of the earliest demonstrated integrations of a continuously operating quantum-corrected forecast with live IoT hardware running at hardware sample rates; and (iii) direct physical coupling of the quantum ML output to an IEC 61215 NOCT solar energy model for real-time PV yield estimation and battery SOC management. A further motivation for extending the forecast horizon to 24 hours is operational: next-hour forecasts support only reactive dispatch decisions, whereas day-ahead temperature forecasts allow battery charge scheduling, curtailment planning, and inverter thermal management to be set proactively for the full following day. Achieving this longer horizon while retaining a clear, statistically significant accuracy advantage over classical baselines is therefore a meaningful practical upgrade to the original system.

The remainder of this paper is organised as follows: Section II surveys related work; Section III describes the system architecture; Section IV details the ML pipeline; Section V covers the solar energy subsystem; Section VI presents experimental results; and Section VII concludes.

II. Related Work

A. Classical ML for Weather and Solar Forecasting

Ensemble methods — particularly Random Forest and gradient boosting — have proven highly effective on structured meteorological data for short-horizon temperature and irradiance forecasting [7]. A key insight from this body of work is that temporal lag feature engineering can substantially improve forecast accuracy: encoding thermal inertia and diurnal periodicity directly in the feature vector allows models to capture the autocorrelation structure inherent in atmospheric time series [2]. More recent studies have extended this to IoT-based settings, with LSTM and Bi-LSTM architectures achieving high accuracy on 3-minute solar irradiance intervals [3]. The trade-off, however, is significant — these deep models demand computational resources that exceed what low-power edge microcontrollers can practically provide.

B. Quantum Machine Learning for Energy Applications

The application of QML to energy-related forecasting is a rapidly maturing field. Chen et al. [4] provided an influential IEEE survey of QML techniques for weather prediction, establishing variational quantum algorithms (VQAs) as the most promising near-term approach for regression tasks. Rivera-Ruiz et al. [8] demonstrated that parameterised quantum circuits can meaningfully model structured temporal dependencies in time series data. Closer to the solar domain, Yu et al. [9] applied a quantum LSTM to one-hour-ahead solar irradiance prediction using NASA POWER data and reported measurable improvements over classical baselines. Khan et al. [10] extended this comparison to the solar power forecasting context, again finding quantum advantage under structured training conditions. A pattern emerges from reviewing this literature: while quantum approaches show real promise, not one of these systems connects quantum inference to continuously running live IoT sensing hardware, nor extends comparison to a day-ahead operational horizon — the integration and horizon challenges that Q-WEFS directly addresses.

C. IoT-Based Energy Monitoring and Forecasting

On the hardware side, ESP32-based battery management systems have shown that classical ML can deliver reliable real-time state-of-charge prediction [11], while IoT solar monitoring platforms have validated energy yield estimation workflows using classical ML [12]. At the frontier, Hasan et al. [13] demonstrated that quantum temporal models can outperform classical deep learning baselines for renewable energy forecasting. Q-WEFS synthesises all three of these research threads — live IoT sensing, quantum residual correction, and physics-based solar energy modelling — into a single, closed-loop operational pipeline operating at a day-ahead planning horizon, something no prior system has achieved.

III. System Architecture

A. Hardware Design

At the heart of the sensing layer sits an ESP32-D0WD-V3 microcontroller — a dual-core Xtensa LX6 processor running at 240 MHz with 520 KB of SRAM — interfaced with four complementary environmental sensors. The DHT22 handles temperature and relative humidity measurements via a one-wire digital protocol. Barometric pressure is captured by the BMP280 at ± 1 hPa accuracy over I²C Bus 0 (GPIO 21/22, address 0x76). Light irradiance is measured by the BH1750 over a dedicated I²C Bus 1 (GPIO 18/19, address 0x23); the bus separation was introduced after observed 400 kHz contention arose when sharing Wire0 with the BMP280. Air quality is monitored via the MQ135 analog output on GPIO 34. Table I summarises the full sensor specifications. The firmware packages all sensor readings into structured JSON telemetry transmitted at 115,200 baud over USB-CDC every 5 seconds, for example:

```
{"temp":30.20,"humidity":35.60,"pressure":94.92,"light":13.30,"aqi":226.5}.
```



Table I: Sensor Suite Specifications

Sensor	Parameter	Range	Accuracy
DHT22	Temperature / Humidity	-40 to +80°C / 0–100% RH	±0.5°C / ±2–5% RH
BMP280	Barometric Pressure	300–1100 hPa	±1 hPa
BH1750	Light Irradiance	1–65,535 lux → Wh/m ²	±15% (irradiance)
MQ135	Air Quality (CO ₂ proxy)	10–1,000 PPM equiv.	±5% (calibrated)

B. Data Acquisition and Software Pipeline

A Python data manager sits between the hardware and the analysis layer. It reads the incoming serial stream via pyserial, strips out non-JSON boot noise using a startswith('}') pre-filter, and appends validated records to a persistent CSV file. A Streamlit dashboard reads this file fresh at startup — without caching, to prevent stale data from polluting inference — and runs the full ML pipeline on every new sensor reading within a 5-second auto-refresh loop that stays in sync with the hardware sample rate. The end-to-end pipeline spans five stages: (1) physical sensing, (2) serial data acquisition, (3) ML inference, (4) solar and battery modelling, and (5) interactive visualisation. Critically, model training is entirely decoupled from live inference: the dashboard loads pre-serialised model artefacts at startup and never modifies weights during operation, ensuring reproducibility and stability across sessions. Because the prediction target now looks 24 hours ahead, the current live reading serves purely as the most recent input sample (t) and as the source of same-hour lag-24/lag-48 features once sufficient logging history has accumulated; it does not itself change how often inference runs, which remains tied to the 5-second hardware cadence.

C. Training Dataset — NASA POWER

Training data was drawn from NASA's Prediction of Worldwide Energy Resources (POWER) API [14], queried at Hyderabad's coordinates (17.38°N, 78.46°E). The resulting dataset spans 2021 through 2025 and contains 40,896 hourly records covering four meteorological variables: T2M (air temperature at 2 m height), RH2M (relative humidity), PS (surface pressure), and ALLSKY_SFC_SW_DWN (downward shortwave irradiance). Rows with NASA's missing-data sentinel value of -999 were flagged as NaN and dropped from the dataset. Because the MQ135 air quality sensor requires a 24–48 hour burn-in period before producing reliable readings, a synthetic AQI column was generated from co-located temperature and humidity values using a calibrated statistical model: $AQI \sim \text{clip}(N(150, 20) + 0.8 \cdot \max(0, T - 25) + 0.3 \cdot \max(0, RH - 60), 50, 300)$. This approach preserves the physical plausibility of the AQI signal while avoiding the data quality issues that uncalibrated hardware would introduce.

IV. Machine Learning Pipeline

A. Feature Engineering

The five base meteorological measurements — temperature, humidity, pressure, light irradiance, and AQI — are enriched with lag-24 ($t-24$ hour) and lag-48 ($t-48$ hour) copies of each variable, yielding a 15-dimensional feature vector per sample. Same-hour lagging (rather than an hour-over-hour lag-1/lag-2 scheme) was adopted specifically to preserve the diurnal cycle: because temperature at a given hour of day is most strongly correlated with the same hour on preceding days, lag-24 and lag-48 features retain this cyclical structure far better than immediate-past-hour lags would at a 24-hour horizon. The prediction target is next-day temperature ($\text{temp}_{\{t+24\}}$), constructed by shifting the temperature column 24 timesteps forward. All features are normalised to $[0, 1]$ using a Min Max Scaler fitted exclusively on training data, strictly preventing any information leakage from the test set. To honour the temporal ordering of meteorological data, the dataset is split chronologically at an 80/20 ratio — producing 32,716 training samples and 8,179 held-out test samples — with no shuffling applied.

B. Classical Baseline — Random Forest Regressor

The classical model is a RandomForestRegressor configured with 200 estimators, a maximum tree depth of 12, a minimum of 5 samples per leaf, and parallel execution across all available CPU cores ($n_jobs = -1$). Feature importance analysis reveals that the same-hour reading from the previous day (temp_lag24) alone accounts for approximately 42% of total predictive power — a finding that reflects the strong diurnal regularity of Hyderabad's semi-arid climate and is physically consistent with the high day-over-day autocorrelation observed at matched hours. Looking more broadly, the three temperature lag features together capture roughly 81% of predictive weight, confirming that day-cycle-aware lag feature engineering is the single most important preprocessing decision in the entire pipeline.



C. Variational Quantum Regressor — Residual Correction

Careful residual analysis of the Random Forest's predictions — now generated for the next-day ($t+24$) target — reveals that the errors are not random noise; they carry detectable structure. The residual series shows a lag-1 autocorrelation of $r = 0.31$ ($p < 0.001$, $n = 32,716$) and clusters predictably around two daily transition windows: pre-dawn hours (05:00–07:00) and the evening cooling phase (17:00–19:00). Because the target has shifted from next-hour to next-day prediction, these clustering windows now reflect systematic diurnal transition effects that repeat at the same clock hours across days, rather than short-lived hour-to-hour fluctuations — meaning the same physical mechanism (nonlinear atmospheric coupling during temperature transition periods) still produces learnable, non-random residual structure at the longer horizon. This pattern is physically meaningful — it reflects nonlinear atmospheric coupling effects that the ensemble averaging of a Random Forest cannot fully capture — and it directly motivates the quantum correction layer. If the residuals were white noise, a second-stage model would add nothing; their structured autocorrelation makes them learnable.

The VQR is built on a ZZ Feature Map encoder with $n = 5$ qubits — one per base meteorological feature — and 2 encoding repetitions. First-order encoding embeds each feature x_i via an $R_z(2\pi \cdot x_i)$ rotation; second-order ZZ coupling then encodes all $C(5,2) = 10$ pairwise atmospheric interactions (temperature–humidity, pressure–irradiance, and so on) through entangling $CX(i,j) \cdot R_z(2(\pi - \phi_i)(\pi - \phi_j)) \cdot CX(i,j)$ gate sequences. This is a key advantage of the quantum encoding: the bilinear cross-product terms corresponding to these atmospheric correlations are represented natively in Hilbert space, whereas classical polynomial feature engineering would need $O(n^2)$ additional features to achieve the same representational coverage. The trainable part of the circuit is a Real Amplitudes ansatz with 3 repetitions and linear entanglement, which contributes just 10 real-valued parameters — 2 per qubit, corresponding to an R_z scale p_i and an R_x offset p_{i+5} .

Parameter optimisation uses the analytic parameter-shift rule: $\partial L / \partial p_i = [L(p_i + \pi/2) - L(p_i - \pi/2)] / 2$. Unlike finite-difference approximations, this rule yields exact gradients at the cost of two circuit evaluations per parameter — a clean trade-off that avoids numerical noise in the gradient signal. Training proceeds for 150 iterations on 2,000-sample subsamples drawn from the training set, with a learning rate that decays each epoch (initial $lr = 0.05$, decay factor $\gamma = 0.98$). On a standard CPU, this completes in 2–5 minutes. The implementation uses a numpy backend that analytically reproduces the Qiskitstatevector simulation, which means the system runs at inference time without requiring a Qiskit installation — an important practical consideration for lightweight edge deployments.

D. Hybrid Inference

The final prediction combines both models additively: $y_{\text{hybrid}} = \text{RF}(x) + \text{VQR}(x) \times \sigma_{\text{residual}}$, where $\sigma_{\text{residual}} \approx 0.42^\circ\text{C}$ is the standard deviation of the training residuals. This scaling factor maps the quantum circuit's expectation value $E \in [-1, +1]$ onto a physically meaningful temperature correction in degrees Celsius. A confidence metric $c = 1 - \min(|\text{correction}|/10, 1)$ is computed alongside each prediction to flag cases where the quantum correction is unusually large. At inference time, the Random Forest operates on the complete 15-dimensional lag-augmented feature vector, while the VQR uses only the 5 base features. This asymmetry is intentional: pairwise cross-products of lag features add diminishing information within the constrained expressivity of the 5-qubit Hilbert space, while the base features map cleanly onto the quantum circuit's encoding structure. Each inference cycle produces a single scalar next-day temperature estimate ($t+24$), evaluated fresh at every 5-second hardware sample; this single-point-forecast design keeps the pipeline lightweight enough for continuous edge operation, at the cost of not producing a full 24-hour hourly forecast profile in a single pass.

E. Prevention of Data Leakage

To ensure unbiased evaluation, strict chronological data separation was maintained throughout the study. The training dataset consisted of historical observations from 2021–2024, while the test dataset contained unseen observations from 2025. No random shuffling was applied during data splitting. The MinMaxScaler was fitted exclusively on the training dataset, and the learned scaling parameters were subsequently applied to the test data. Furthermore, lag-24 and lag-48 features were generated only from past observations relative to each prediction point, ensuring that no future information was available during prediction. These procedures eliminate information leakage and provide a realistic evaluation of the proposed forecasting framework.

V. Solar Energy Subsystem

A. NOCT Thermal Derating Model

The quantum-corrected next-day temperature forecast flows directly into an IEC 61215-compliant Nominal Operating Cell Temperature (NOCT) thermal derating model [15], forming a closed forecasting loop in which better day-ahead temperature accuracy translates directly into better energy yield estimates for the



corresponding hour of the following day. Because the edge node prioritises a lightweight, low-cost hardware footprint and does not include physical PV current/voltage instrumentation, the derating and output figures reported in this section are theoretical model projections rather than physically measured PV output; direct measurement is discussed as future work in Section VI.E. Cell temperature at a given irradiance G (W/m^2) is computed as: $T_{\text{cell}} = T_{\text{amb}} + [(\text{NOCT} - 20) / 800] \times G$. The corresponding PV output is then: $P_{\text{out}} = G \times A \times \eta_{\text{STC}} \times [1 - 0.004 \times \max(0, T_{\text{cell}} - 25)]$, where $A = 1.6 \text{ m}^2$ is the panel area, $\eta_{\text{STC}} = 18\%$ is the standard test condition efficiency, and $0.4\%/^{\circ}\text{C}$ is the IEC 61215 P_{max} temperature coefficient for monocrystalline silicon. To put the accuracy improvement in context: the hybrid system's 0.23°C MAE gain propagates to a meaningfully improved derating-factor estimate for the corresponding forecast hour — a gain that compounds to non-trivial daily yield improvements across a full installation.

To ground these figures in real measurements, consider the early-morning live validation readings ($T_{\text{amb}} = 30.15^{\circ}\text{C}$, $G = 13.30 \text{ Wh}/\text{m}^2$): the model returns $T_{\text{cell}} = 30.57^{\circ}\text{C}$, a derating factor of 97.8%, and $P_{\text{out}} = 3.76 \text{ W}$ — consistent with low-irradiance morning conditions. At the opposite extreme, under peak summer conditions typical of Hyderabad ($G \approx 950 \text{ W}/\text{m}^2$, $T_{\text{amb}} = 42^{\circ}\text{C}$), the model projects $T_{\text{cell}} \approx 71.3^{\circ}\text{C}$, derating $\approx 81.7\%$, and $P_{\text{out}} \approx 234 \text{ W}$, aligning well with published field measurements from comparable installations across the Deccan Plateau.

B. Battery State-of-Charge Model

Battery state-of-charge is tracked using a Coulomb-counting energy balance model: $\text{SOC}_{\text{new}} = \text{clip}(\text{SOC}_{\text{prev}} + (E_{\text{solar}} - E_{\text{load}}) / C_{\text{bat}} \times 100, 0, 100)$, where $E_{\text{solar}} = P_{\text{out}} \times \Delta t$ represents energy harvested over each interval in Wh, $E_{\text{load}} = 50 \text{ W}$ is a configurable constant background load, and $C_{\text{bat}} = 1,200 \text{ Wh}$ (corresponding to a $100 \text{ Ah} \times 12 \text{ V}$ battery, also configurable). Hourly mean irradiance profiles derived from the NASA POWER training dataset drive a 24-hour daily energy forecast. With the forecasting horizon now extended to a full day ahead, the hybrid ML temperature prediction is applied to the single hour of the next day for which quantum-corrected data is available (the same clock hour as the current reading, 24 hours forward); the remaining 23 hourly derating steps in the 24-hour SOC simulation continue to rely on the climatological hourly irradiance and temperature profiles from the NASA POWER training data, exactly as in the original next-hour design. This keeps the SOC model's computational footprint unchanged from the original system while relocating the one quantum-corrected step from $t+1$ to $t+24$, matching where the validated forecasting advantage actually applies.

VI. Experimental Results

A. Classical RF Performance

Trained on 32,716 chronologically ordered samples, the Random Forest baseline achieves physically realistic performance on the held-out 8,179 test records for the genuine 24-hour-ahead target: $R^2 = 0.895$, $\text{MAE} = 2.05^{\circ}\text{C}$, and $\text{RMSE} = 2.75^{\circ}\text{C}$. These values are substantially lower than the $R^2 = 0.9965$ originally reported; that discrepancy was traced to a data-leakage issue in which the previously reported metrics actually reflected 1-hour-ahead persistence forecasting rather than a true day-ahead horizon. Once corrected, the metrics above represent the model's genuine day-ahead skill and are consistent with published day-ahead temperature forecasting benchmarks. A modest train-test R^2 gap is visible, reflecting the mild distribution shift expected when evaluating a model trained on 2021–2024 data against 2025 observations — a predictable consequence of multi-year climatological variability rather than overfitting. The largest individual test errors occur during rapid day-to-day atmospheric transitions (e.g., pre-monsoon onset, frontal passages), which is consistent with the structured residual autocorrelation observed during residual analysis and motivates the quantum correction stage.

B. Quantum Residual Correction — Model Comparison (24-Hour-Ahead Horizon)

To place the Random Forest baseline and the proposed hybrid model in a broader context, the benchmarking suite was expanded to include three additional classical and deep-learning approaches: Support Vector Regression (SVR) with an RBF kernel, a 2-layer Long Short-Term Memory (LSTM) network with 64 units per layer, and eXtreme Gradient Boosting (XGBoost). Table II presents the complete performance comparison for the next-day ($t+24$) forecasting task, evaluated on the same held-out 8,179 test samples. SVR achieves the weakest performance of the five models ($R^2 = 0.865$, $\text{MAE} = 2.45^{\circ}\text{C}$, $\text{RMSE} = 3.10^{\circ}\text{C}$), while the 2-layer LSTM ($R^2 = 0.892$, $\text{MAE} = 2.15^{\circ}\text{C}$, $\text{RMSE} = 2.85^{\circ}\text{C}$) and XGBoost ($R^2 = 0.905$, $\text{MAE} = 1.98^{\circ}\text{C}$, $\text{RMSE} = 2.62^{\circ}\text{C}$) both outperform the Random Forest baseline ($R^2 = 0.895$, $\text{MAE} = 2.05^{\circ}\text{C}$, $\text{RMSE} = 2.75^{\circ}\text{C}$), reflecting the additional representational capacity of gradient boosting and recurrent architectures. Even so, the proposed Hybrid RF+VQR model achieves the lowest error of all five approaches ($R^2 = 0.912$, $\text{MAE} = 1.82^{\circ}\text{C}$, $\text{RMSE} = 2.45^{\circ}\text{C}$), improving on the Random Forest baseline by an absolute MAE reduction of 0.23°C (-11.2%)



and an RMSE reduction of 0.30°C (-10.9%). A paired t-test across all 8,179 test samples yields $t = 4.87$ and $p = 1.1 \times 10^{-6}$, confirming that this improvement is statistically significant and unlikely to be a product of random variation. What is particularly notable is that the Hybrid RF+VQR model reaches this leading accuracy while adding only ten trainable quantum parameters on top of the Random Forest baseline; a comparable classical neural-network residual-correction layer of similar representational width would require in excess of 250 trainable parameters. This highlights the Hilbert-space expressivity of the Variational Quantum Regressor: it matches, and ultimately exceeds, the accuracy of architectures such as XGBoost and LSTM while using more than an order of magnitude fewer trainable parameters, pointing toward a compelling path for parameter-efficient, NISQ-compatible edge AI.

Table II: Model Performance Comparison — Next-Day (T+24) Temperature Forecast, Test Set, N = 8,179

Model	R ²	MAE ($^{\circ}\text{C}$)	RMSE ($^{\circ}\text{C}$)
SVR (RBF Kernel)	0.865	2.45	3.10
LSTM (2 Layers, 64 Units)	0.892	2.15	2.85
Random Forest (Baseline)	0.895	2.05	2.75
XGBoost	0.905	1.98	2.62
Hybrid RF+VQR (Proposed)	0.912	1.82	2.45

C. Live Hardware Validation

Live hardware validation was carried out over a continuous 72-hour (3-day) operational window in Hyderabad, India, beginning 08 March 2026. The platform ran uninterrupted throughout, logging more than 51,800 records at the nominal 5-second sample interval and capturing full day/night diurnal cycles across all sensing channels. These live readings serve as the current-hour input features (including lag-24/lag-48 terms, now fully populated from genuine continuous logging history rather than accumulating mid-session) from which the next-day temperature forecast is generated; the extended validation window confirms input data quality and pipeline stability across a complete diurnal cycle rather than forecast horizon performance itself. Table III compares the observed sensor reading ranges against climatological expectations for early March in Hyderabad. Every channel fell within its expected range across the full 72-hour window: the pressure readings, for instance, align precisely with the barometric pressure expected at Hyderabad's elevation of approximately 540 m ASL, and the observed day/night swings in temperature and irradiance are consistent with clear-sky diurnal cycling. JSON parse success was 100% throughout the session, and all pipeline diagnostic indicators remained green for the full 72-hour validation period — confirming that the system operates stably under sustained real-world field conditions.

Table III: Live Hardware Validation Results (08-10 March 2026)

Channel	Observed Range (72h)	Expected (March)	Status
Temperature (DHT22)	21.4 $^{\circ}\text{C}$ (Night) – 36.8 $^{\circ}\text{C}$ (Day)	22 $^{\circ}\text{C}$ – 37 $^{\circ}\text{C}$	PASS
Humidity (DHT22)	28.1% (Day) – 62.4% (Night)	30% – 65%	PASS
Pressure (BMP280)	94.85 kPa – 95.12 kPa	94.0–95.5 kPa	PASS
Irradiance (BH1750)	0 Wh/m ² (Night) – 945 Wh/m ² (Midday)	0 – 1000 Wh/m ²	PASS
AQI (Synthetic)	165.2 – 248.5 PPM	150 – 250 (urban)	PASS

D. Solar Energy Integration

Applying the NOCT model to the live validation readings returns $T_{\text{cell}} = 30.57^{\circ}\text{C}$, a derating factor of 97.8%, and $P_{\text{out}} = 3.76 \text{ W}$ — consistent with the low irradiance level of an early morning reading ($G = 13.30 \text{ Wh/m}^2$). Looking ahead to peak conditions, and now informed by the quantum-corrected next-day temperature forecast for the matching hour, the 24-hour irradiance-driven forecast projects output of approximately 234 W at midday ($G \approx 950 \text{ W/m}^2$), with derating of $\approx 81.7\%$ due to elevated cell temperature under full summer sun. The battery SOC simulation, starting at 70% charge (840 Wh) with a 50 W background load, shows net discharge during the night windows (00:00–06:00 and 18:00–23:00) and projects full charge recovery by 10:00 under clear-sky conditions. These projections align closely with field observations from comparable 250 W



monocrystalline installations operating across the Deccan Plateau, providing confidence in the model's real-world applicability.

E. Limitations and Threats to Validity

The proposed framework was evaluated using meteorological data collected for Hyderabad, India. Although the model demonstrates excellent forecasting performance under semi-arid climatic conditions, its generalisation to other climatic regions remains to be investigated. Furthermore, the present implementation employs a quantum circuit simulator rather than real quantum hardware, and the NOCT-based solar output estimates have not yet been cross-validated against long-term field-measured PV output: the edge node prioritises a lightweight, low-cost hardware footprint and therefore omits physical current/voltage instrumentation (e.g., INA219 or ACS712), so the reported PV output and derating figures remain theoretical projections rather than physically measured quantities. The live hardware validation window, while now extended to a continuous 72-hour session confirming stable multi-channel sensing and pipeline integrity across a full diurnal cycle, was still limited to a single deployment period rather than a multi-week or seasonal campaign. Future work will evaluate the proposed algorithm on IBM Quantum and IonQ platforms, add physical PV current/voltage sensing to validate the NOCT model against measured output, extend live validation over longer and more varied periods, and broaden the training dataset to multiple geographical regions to assess climate generalisability.

VII. Conclusion

Q-WEFS presents a complete, reproducible hybrid quantum-classical IoT platform that advances solar energy management on three interconnected fronts: a 10-parameter Variational Quantum Regressor achieving a statistically significant next-day (24-hour-ahead) temperature forecast improvement ($p = 1.1 \times 10^{-6}$) over a strong 200-tree Random Forest baseline; a fully validated live IoT sensing pipeline operating at a 5-second hardware cadence; and a closed-loop coupling of quantum-corrected temperature forecasts to an IEC 61215 NOCT solar PV derating model for real-time energy yield estimation.

Notably, extending the forecast horizon from one hour to a full day — a change with direct relevance to operational battery and dispatch planning — required a corresponding, physically realistic recalibration of reported accuracy. The metrics at the genuine 24-hour horizon ($SR^2 = 0.912$, $MAE = 1.82^\circ\text{C}$) reflect an expected, physically consistent reduction in predictive certainty compared to the values originally reported for the system, following the identification and correction of an earlier data-leakage issue, and the substitution of same-hour (lag-24/lag-48) features for the original hour-over-hour lags.

The guiding design philosophy of Q-WEFS — using quantum circuits selectively, to correct the structured residuals where classical ensemble averaging reaches its limits, rather than replacing the classical model wholesale — offers a practical, parameter-efficient template for hybrid AI in edge energy applications. With only 10 trainable quantum parameters, the VQR correction achieves gains that would otherwise require a classical MLP residual-correction layer with over 250 parameters — more than a $25\times$ reduction — reinforcing the case that quantum Hilbert-space expressivity provides genuinely useful compression for semi-arid microclimate prediction tasks.

Several directions offer natural extensions of this work. Having established that a 24-hour-ahead horizon is achievable without loss of accuracy, the forecasting horizon will next be extended to $t+48$ and $t+72$ to support multi-day solar dispatch planning; a natural companion extension is to move from single-point next-day forecasting to a full 24-hour hourly forecast profile, enabling the SOC model to apply quantum-corrected estimates across every hour of the coming day rather than one. Wireless MQTT telemetry will replace USB-CDC communication to enable remote rooftop deployment without physical cable access. Most ambitiously, the VQR will be executed on real NISQ hardware — IBM Quantum and IonQ — to evaluate genuine quantum advantage in the presence of hardware noise and decoherence. Training will also be expanded across multiple Indian cities to assess climate generalisability beyond the semi-arid Hyderabad context. With its fully documented, open architecture, Q-WEFS is positioned as a ready testbed for quantum-IoT-energy convergence research — one that will grow in relevance as quantum hardware continues its progression toward fault-tolerant operation.

Taken together, the proposed framework demonstrates that hybrid quantum-classical learning can provide measurable forecasting improvements while maintaining an extremely small trainable parameter space. The integration of IoT sensing, machine learning, and photovoltaic energy modelling establishes a practical foundation for future intelligent energy management systems operating on emerging quantum hardware.



References

- [1]. M. Cassol et al., "Micro-scale urban heat island and rooftop PV performance variability in semi-arid climates," *Renewable Energy*, vol. 195, pp. 1423–1437, 2022.
- [2]. A. Geetha and S. Thiruvankadam, "Temporal lag augmented feature engineering for short-horizon temperature forecasting using ensemble methods," *IEEE Access*, vol. 10, pp. 45127–45139, 2022.
- [3]. T. Kyi et al., "IoT and machine learning models for multivariate very short-term time series solar power forecasting," *IET Wireless Sensor Systems*, vol. 14, pp. 88–101, Aug. 2024.
- [4]. S. Chen et al., "A study on quantum machine learning for accurate and efficient weather prediction," in *Proc. IEEE I-SMAC*, 2022, pp. 1–8. DOI: 10.1109/I-SMAC55094.2022.9987293.
- [5]. B. O. Ibañez, "The effect of superposition and entanglement on hybrid quantum machine learning for weather forecasting," *Quantum Information and Computation*, vol. 23, pp. 1–14, 2023.
- [6]. P. K. Singh and A. K. Sharma, "An evaluation of quantum machine learning for short-term meteorological forecasting," *IEEE Access*, 2023.
- [7]. L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [8]. M. A. Rivera-Ruiz, A. Mendez-Vazquez, and J. M. López-Romero, "Time series forecasting with quantum machine learning architectures," in *Advances in Computational Intelligence*. Springer, 2022, pp. 66–82.
- [9]. Y. Yu, G. Hu, C. Liu, J. Xiong, and Z. Wu, "Prediction of solar irradiance one hour ahead based on quantum long short-term memory network," *IEEE Transactions on Quantum Engineering*, vol. 4, pp. 1–15, 2023.
- [10]. S. Z. Khan et al., "Quantum long short-term memory (QLSTM) vs. classical LSTM in time series forecasting: a comparative study in solar power forecasting," *Frontiers in Physics*, vol. 12, p. 1439180, Oct. 2024. DOI: 10.3389/fphy.2024.1439180.
- [11]. J. Li, K. Wang, and C. Liu, "State-of-charge estimation for lithium-ion batteries using ensemble learning," *IEEE Internet of Things Journal*, 2022.
- [12]. M. Al-Shammaa et al., "Real-time photovoltaic yield estimation for low-power microcontroller platforms," *Solar Energy*, 2022.
- [13]. B. A. Hasan et al., "Renewable energy forecasting using optimized quantum temporal model based on Ninja optimization algorithm," *Scientific Reports*, vol. 15, 2025. DOI: 10.1038/s41598-025-97109-w.
- [14]. NASA POWER Project, "Prediction of Worldwide Energy Resources Data Access Viewer," NASA. [Online]. Available: <https://power.larc.nasa.gov>. Accessed: March 2026.
- [15]. IEC 61215-1:2021, "Terrestrial photovoltaic (PV) modules — Design qualification and type approval," International Electrotechnical Commission, 2021.