

Sensorless Thermometry from MPPT Electrical Traces: A Physics-Informed Inversion for Perovskite Solar Cells

Waheed Zaman Khan^{1,2*}, Abdullah Ishfaq³, Asma Akhtar⁴, Abdullah⁴, Ijaz Ahmad¹

¹Department of Physics, Division of Science and Technology, University of Education, College Road, Township, Lahore 54770, Pakistan

²Phystech School of Aerospace Technology, Moscow Institute of Physics and Technology (MIPT), 9 Institutskiy per., Dolgoprudny, Moscow Region 141700, Russian Federation

³Department of Basic Sciences, Riphah International University, I-14 Campus, Sector I-14, Hajj Complex, Islamabad 45210, Pakistan

⁴Department of Physics, The University of Chenab, G.T. Road, Adjacent Chenab Bridge, Gujrat 50700, Pakistan

DOI: <https://doi.org/10.36347/sjpm.2026.v13i04.001>

| Received: 19.02.2026 | Accepted: 23.03.2026 | Published: 08.04.2026

*Corresponding author: Waheed Zaman Khan

Department of Physics, Division of Science and Technology, University of Education, College Road, Township, Lahore 54770, Pakistan

Abstract

Original Research Article

We present a sensorless thermometry method that reconstructs perovskite module temperature directly from routine MPPT logs, requiring no paid instrumentation. The approach separates a *fast* irradiance-driven term from a *slow* thermal response modeled as a single-RC state. A two-regressor, sliding-window identification is solved with orthogonalization and mild ridge regularization; temperature is then recovered by physics-informed inversion, with residual bootstrap providing confidence intervals and automated gates preventing ill-posed Kelvin reporting. Using perovskite-like priors (temperature-coefficient magnitude 1.3×10^{-3} – $2.0 \times 10^{-3} \text{ K}^{-1}$ and thermal time constant 50–400 s), 45–60 min logs at 1 Hz consistently recover the trajectory shape of $T(t)$. Absolute-Kelvin accuracy without any anchor depends on excitation and τ : median MAE spans ~ 5 – 7 K at $\tau = 50\text{s}$ and rises to ~ 25 – 27 K at $\tau = 240\text{s}$ and $\sim 43\text{K}$ at $\tau = 400\text{s}$ in high-variability records, while the relative error against the temperature swing remains comparatively stable. Uncertainty bands contract during strongly forced segments and widen during quiescent intervals, matching identifiability diagnostics. In practice, one brief temperature anchor (ambient or back-of-module) fixes the absolute scale and tightens Kelvin-level errors; including $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ further improves regressor separability. The fully reproducible Python pipeline (NumPy/pandas/pvlib-python) aligns with low-cost, ambient workflows emphasized for perovskite research.

Keywords: perovskite solar cells, sensorless thermometry, maximum power point tracking (MPPT).

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1. INTRODUCTION

Temperature strongly influences photovoltaic (PV) performance by modifying carrier populations, recombination pathways, mobilities, and series resistance; the net effect appears in $V_{\text{OC}}, J_{\text{SC}}$, fill factor, and output power. For metal-halide perovskite solar cells (PSCs), measured power temperature coefficients are typically negative and composition-dependent. Representative studies report values in the range of a few tenths of a percent per kelvin, including outdoor field measurements around $-0.17\% \text{ K}^{-1}$ for power, underscoring the need to account for device temperature when interpreting performance and stability [1-2]. Recent work in the Chiang Mai Journal of Science (CMJS) on humidity effects in mixed-cation perovskite films further emphasizes environmental sensitivity

during processing and operation, which is highly relevant to resource-limited and “from-home” studies like the present one [3].

Classical PV thermometry estimates cell/module temperature using meteorological inputs—ambient temperature, plane-of-array irradiance, and wind speed. Steady-state relations such as the Faiman model express module temperature as an ambient offset scaled by irradiance and empirical heat-transfer coefficients, while dynamic approaches treat the module as a lumped RC network with one or more thermal time constants τ . These methods are well established for crystalline-silicon modules and widely used in energy rating and forecasting, but they require external sensors

and can be inaccurate during fast transients if purely steady-state assumptions are used [4-5].

By contrast, maximum-power-point tracking (MPPT) logs ubiquitous even in small or home setups already encode two distinct time scales: a fast electrical response to irradiance $\Phi(t)$ changes and a slow thermal evolution $T(t)$ governed by RC dynamics. This observation motivates a physics-identifiable inverse problem: reconstruct the device temperature history $T(t)$ purely from MPPT electrical data (e.g., $P(t)$), and when available $V(t), I(t)$, without any temperature sensor. Prior “sensorless” lines of work either (i) use temperature as an input for MPPT control (temperature-based MPPT), or (ii) infer irradiance/temperature from electrical signals with black-box estimators (e.g., neural networks). These approaches are valuable, but they are not tailored to perovskite-specific temperature coefficients and typically lack a transparent identifiability analysis [6-8].

The growing availability of homogeneous MPPT ageing datasets for perovskites demonstrates that MPPT-only signals carry rich information about operational physics and reliability. For example, a recent study analysed 2,245 MPPT ageing curves under controlled conditions to uncover links between initial efficiency and stability—evidence that properly modelled MPPT time-series can support quantitative inference beyond simple power tracking [9].

Research gap and objective. Despite the above, we find no physics-based, perovskite-aware method that reconstructs absolute $T(t)$ from MPPT electricity alone with an explicit identifiability argument and a fully reproducible pipeline. Our objective is to fill this gap with a sensorless thermometry framework suitable for “from-home” research using only free tools. We couple a first-order power expansion around (Φ_0, T_0) with a lumped thermal RC model, exploit short windows that contain natural irradiance variation to estimate the power temperature coefficient α_T and thermal constant(s) τ , and then integrate forward to obtain $T(t)$. The framework is validated on synthetic perovskite-like scenarios, enabling transparent error analysis without laboratory instrumentation.

Contributions. (C1) A physics-identifiable inversion that separates fast irradiance response from slow thermal dynamics and estimates $\{\alpha_T, \tau\}$ directly from MPPT windows; (C2) a robust sliding-window estimator with simple ill-conditioning checks (e.g., low-variance irradiance segments); and (C3) a fully open, from-home implementation (Python/NumPy/pandas/matplotlib) suitable for replication on a standard PC.

For computation, we discretize (4) at sampling interval Δt (seconds):

$$T_{k+1} = T_k + \frac{\Delta t}{C} \left(\kappa \Phi_k - \frac{T_k - T_{amb,k}}{R} \right) = \left(1 - \frac{\Delta t}{\tau} \right) T_k + \frac{\Delta t}{\tau} T_{amb,k} + \frac{\kappa \Delta t}{C} \Phi_k \quad (5)$$

2. THEORY

2.1. Operating point and linearized power model

Let $P(t)$ denote the electrical maximum power of a perovskite solar cell operated under plane-of-array irradiance $\Phi(t)$ (Wm^{-2}) and device temperature $T(t)$ (K). In a neighborhood of an operating point (Φ_0, T_0) with $P_0 = P(\Phi_0, T_0)$, a first-order expansion isolates irradiance- and temperature-driven contributions:

$$\Delta P(t) \approx \underbrace{\frac{\partial P}{\partial \Phi} \Big|_0}_{g_\Phi} \Delta \Phi(t) + \underbrace{\frac{\partial P}{\partial T} \Big|_0}_{g_T} \Delta T(t), \quad (1)$$

where $\Delta X(t) \equiv X(t) - X_0$ and the sensitivities g_Φ, g_T inherit device-physics dependences (absorption, recombination, transport). Equation (1) is consistent with established PV performance formulations (e.g., Sandia’s array model) in which P (or its factors J_{SC}, V_{OC}, FF) vary with both Φ and T [10,11].

It is convenient to define the relative temperature coefficient of power

$$\alpha_T \equiv \frac{1}{P_0} g_T = \frac{1}{P_0} \frac{\partial P}{\partial T} \Big|_0, \quad (2)$$

So that the temperature-driven power variation reads $\Delta P_T / P_0 \approx \alpha_T \Delta T$. Perovskite devices exhibit $\alpha_T < 0$ with magnitudes typically smaller than crystalline-Si; recent analyses report values around $\sim -0.13\%K^{-1}$ to $-0.20\%K^{-1}$ depending on stack and composition [12]. These ranges serve as weak priors in our identification (Section 3).

2.2. Thermal RC dynamics (lumped energy balance)

Device temperature evolves on a slower time scale than electrical transients. A standard lumped energy balance leads to an RC-type thermal model:

$$C \frac{dT}{dt} = H_{in}(\Phi, t) - H_{out}(T, T_{amb}, t), \quad (3)$$

Where C is the effective thermal capacitance (JK^{-1}), H_{in} aggregates absorbed radiative power and internally dissipated heat, and H_{out} lumps convective/radiative losses to the environment. Under common simplifying assumptions (effective absorbance/emittance and linearized convective terms), (3) reduces locally to

$$C \dot{T} = \kappa \Phi(t) - \frac{T(t) - T_{amb}(t)}{R}, \quad \tau \equiv RC, \quad (4)$$

With $\kappa > 0$ an effective irradiance-to-heat gain ($Km^2 W^{-1}$) and R an effective thermal resistance ($K W^{-1}$) Equation (4) is compatible with equation-based PV thermal practice and recent second-to-minute timescale studies, which also show how to identify τ from data [10,13].

2.3. Two-time-scale separation and identifiability

MPPT logs typically show fast variations of Φ in $P(t)$ (sub-second to seconds, set by

$$P_k = P_0 + g_\Phi \Delta\Phi_k + \alpha_T P_0 \Delta T_k (\tau, \kappa, C, T_0) + \varepsilon_k, \quad (6)$$

Where ε_k subsumes measurement noise and unmodeled effects. Structural identifiability of $\{\alpha_T, \tau\}$ from P_k requires (i) sufficient variation in Φ within at least one window (to separate the g_Φ term from the thermal response) and (ii) a time-scale gap so that the thermal convolution leaves a distinguishable signature. Modern identifiability results for ODE models formalize such conditions; in practice we enforce them by window selection (variance thresholds on Φ) and by penalizing ill-conditioned fits (Section 3.3) [14–16].

When $V(t)$ and $I(t)$ are available, we extend (1) to include voltage/current sensitivities, which improves

$$\min_{\theta, T_0} \underbrace{\sum_{k \in W} (P_k - \widehat{P}_k(\theta))^2}_{\text{data misfit}} + \lambda \underbrace{\sum_{k \in W} (\Delta^{(2)} \widehat{P}_k)^2}_{\text{smoothing prior}} + \mu |\theta - \theta_{\text{prior}}|^2, \quad (7)$$

Where $\Delta^{(2)}$ is the second-difference operator, $\lambda, \mu \geq 0$ are hyper-parameters, and θ_{prior} encodes weak physics-based bounds (e.g., $\alpha_T < 0$, plausible τ ranges). The smoothing term mitigates noise amplification typical of derivative-rich fits and is closely related to Tikhonov regularization; for pre-filtering we use Savitzky–Golay least-squares smoothing when needed [17,18].

Windows with insufficient $Var(\Phi)$ are rejected to avoid near-collinearity between $g_\Phi \Delta\Phi_k$ and slowly varying $\alpha_T P_0 \Delta T_k$. Aggregating estimates from multiple acceptable windows stabilizes $\{\hat{\alpha}_T, \hat{\tau}\}$ (Section 3.4), and uncertainty can be quantified by nonparametric bootstrap over windows.

2.5. Extensions

Multi-node thermal networks. If a single time constant is inadequate (e.g., encapsulant + substrate), we replace (4) by a sum of exponentials with parameters $\{\tau_i, w_i\}$ leading to

$$\Delta T_k \approx \sum_{i=1}^{n_\tau} w_i (h_{\tau_i} * \Phi)_k + \text{ambient-driven terms}, \quad (8)$$

With discrete kernels h_{τ_i} defined by (5). Identifiability persists if windowed Φ contains sufficiently rich excitation (persistence of excitation) and the τ_i are well separated [14-16].

Incorporating V, I features. A linearized IV-curve factorization (Sandia-style) augments (6) with terms coupling V and I to T and Φ , sharpening separation of effects and improving robustness to noise [11].

tracker/electronics) and slow thermal drifts in $T(t)$ (tens to hundreds of seconds), motivating a two-time-scale view. Combining (1) - (2) with (5) yields the discrete measurement model

observability because $V_{OC}(T)$ typically dominates the temperature response while $J_{SC}(\Phi)$ dominates irradiance response (Sandia-style factorizations) [11].

2.4. Windowed estimation with smoothing and regularization

We estimate parameters over short windows W that contain natural irradiance variation. Let $\theta = \{\alpha_T, g_\Phi, \tau, \kappa\}$ and T_0 the window’s initial temperature. Define the prediction $\widehat{P}_k(\theta)$ by propagating (5) to obtain T_k and substituting into (6). We solve

Relation to sensorless estimators. Algebraic/NN sensorless methods can infer (Φ, T) from $I_{\text{mpp}}, V_{\text{mpp}}$, but they typically lack explicit identifiability guarantees and physics-transparent priors; our formulation provides those tools and can serve as a baseline or hybrid component [16].

3. METHODS

3.1. Data assumptions and reproducibility

The study assumes a time-stamped maximum-power-point tracking (MPPT) log containing the electrical power $P(t)$ as a function of time. When available, the voltage and current at the operating maximum, $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$, are also ingested; ambient temperature $T_{\text{amb}}(t)$ is optional and used solely to fix the absolute offset of the reconstructed device temperature during full-record inversion (Section 3.4). All analyses are performed with free, open-source Python, specifically NumPy, pandas and matplotlib, alongside pvlib-python utilities to ensure SI-consistent handling and to facilitate reproducibility on a standard personal computer [19]. The minimal input format is a comma-separated value file with ISO-8601 time stamps and a power column (with optional voltage, current, and ambient temperature columns). Records are resampled to a uniform cadence Δt , lightly despiked to remove outliers, and, if necessary, smoothed with a short Savitzky–Golay filter to suppress sensor noise prior to parameter identification [21]. All configuration files, random seeds and preprocessing parameters are fixed and reported so the results can be reproduced byte-for-byte from home.

3.2. Synthetic data generator

Validation proceeds on synthetic yet physics-plausible time series to enable full execution from home without laboratory instrumentation. Irradiance $\Phi(t)$ is composed of short steps and ramps that provably excite the model dynamics, interspersed with cloud-like stochastic variability generated by an Ornstein–Uhlenbeck-type process. This combination reflects second-to-minute irradiance fluctuations reported in the short-term variability literature and provides the persistency of excitation required for identification [24–27]. Device temperature is evolved with the lumped RC model introduced earlier, i.e., the energy-balance relation with effective capacitance and resistance, using time constants τ drawn uniformly from the interval 50 – 400 s. These values are consistent with second-to-minute thermal responses observed for PV hardware and with recent dynamic thermal modeling practice [24–26]. Perovskite-like temperature coefficients of power α_T are sampled uniformly from $[-0.20, -0.13]\% \text{ K}^{-1}$, matching the ranges summarized in the Introduction and providing weak but realistic priors for simulation fidelity [1,2]. Additive Gaussian noise is applied to $P(t)$ at relative level $\sigma/P_0 \leq 0.3\%$, with small independent perturbations added to $V_{\text{mpp}}(t)$, $I_{\text{mpp}}(t)$ when present. Random seeds are fixed, and auxiliary unit conversions and optional solar utilities are taken from pvlib-python to standardize calculation pathways [19].

3.3. Sliding-window identification of $\alpha_T, g_\Phi, \tau, \kappa$

Parameter identification exploits short time windows W that contain sufficient irradiance variation to separate the fast Φ -driven response from the slow thermal contribution. Within each window, the state T_k is propagated by the discrete thermal recursion defined previously (Eq. (5)), and the resulting temperature deviation ΔT_k is coupled to the linearized power model (Eq. (6)) to construct a window-level predictor $\hat{P}_k(\theta)$ with $\{\theta = \alpha_T, g_\Phi, \tau\}$. We estimate θ and the window’s initial temperature T_0 by minimizing the regularized least-squares objective in Eq. (7). The second-difference penalty in Eq. (7) acts as a smoothing prior that mitigates noise amplification in derivative-rich regimes, while the Tikhonov term biases the solution toward weak physics bounds such as $\alpha_T < 0$ and plausible ranges of τ [21–22]. Hyper-parameters are selected on a small grid using held-out windows to avoid optimistic fits. To ensure practical identifiability, a window is accepted only if the variance of Φ exceeds a preset threshold and the associated design matrix is well-conditioned; these diagnostics implement widely accepted guidance for partially observed ordinary-differential-equation models [28–30]. When voltage and current are available, the predictor is augmented with Sandia-style sensitivities that emphasize the temperature dependence of voltage and the irradiance dominance of current, improving separation of effects and reducing estimator variance [26]. After solving Eq. (7) over all accepted windows, the individual estimates $\hat{\theta}_W$ are aggregated by a trimmed

mean or median to produce global estimates $\hat{\alpha}_T$ and $\hat{\tau}$ with dispersion measures.

3.4. Inversion of $T(t)$ over the full record

With global parameters fixed, the device temperature trajectory over the entire record is reconstructed by forward integration of the RC model using the observed irradiance and Eq. (5). A single anchor fixes the absolute level: when ambient temperature is recorded, we set $T(0) = T_{\text{amb}}(0)$; otherwise, we adopt a reported constant such as 298.15 K. This anchor changes only the absolute offset of $T(t)$; the temporal evolution that determines the temperature-driven power loss through Eq. (2) remains governed by $\hat{\tau}$ and the measured $\Phi(t)$. For context—not as inputs to the inversion we compute steady-state operating temperatures from NOCT/Faiman-type relations to illustrate differences between dynamic and steady-state viewpoints under rapidly changing conditions [20–31–32].

3.5. Confidence intervals

Uncertainty is quantified with a nonparametric bootstrap at the window level. Accepted windows are resampled with replacement, the global parameters $\{\hat{\alpha}_T, \hat{\tau}\}$ are recomputed for each bootstrap draw, and the full temperature trajectory is re-integrated. We report bias-corrected percentile confidence intervals for $\hat{\alpha}_T$ and $\hat{\tau}$, and pointwise 95% confidence bands for $T(t)$ defined by the envelope of bootstrap trajectories. The bootstrap is appropriate because the estimator is a smooth functional of residuals within each window and the window itself is a natural sampling unit that preserves intra-window temporal structure [23].

3.6. Ablations and baselines

Robustness and necessity of modeling choices are examined through controlled modifications. First, the RC-driven temperature contribution in the measurement model (Eq. (6)) is replaced by a direct linear term to create a single-time model; this ablation typically increases prediction error and removes the physical interpretability afforded by the two-time-scale structure. Second, when ambient temperature and wind estimates are present, steady-state operating temperatures from NOCT/Faiman relations are computed to provide a familiar reference against which the dynamic inversion can be contrasted under transient conditions [20–31–32]. Third, the identification is repeated with and without voltage and current channels to quantify the value of additional electrical observables for reducing the variance and bias of $\hat{\alpha}_T$ and $\hat{\tau}$ in the presence of measurement noise and limited excitation [26]. As an external qualitative check, identified time-constant and temperature-coefficient ranges are compared with distributions reported in large perovskite MPPT ageing datasets to verify that the recovered scales are realistic for perovskite devices operating under controlled stress [34].

4. EXPERIMENTS

4.1. Design and data generation

All experiments are designed to run entirely from home using free, open-source software on a standard personal computer. Synthetic datasets are produced by the generator described in Section 3.2: irradiance $\Phi(t)$ sequences combine ramps and steps with cloud-like stochastic variability so that windows contain sufficient excitation at second-to-minute scales, consistent with short-term irradiance literature [24–27]. Device temperature $T(t)$ is propagated by the lumped RC dynamics through Eq. (4) and discretized by Eq. (5); the effective thermal time constant τ is sampled uniformly in $[50, 400]$ s to reflect module-scale responses, while the irradiance-to-heat gain is set by a fixed κ/C chosen to avoid thermal runaway and to keep temperature excursions realistic under the synthesized $\Phi(t)$ [24–26]. Power traces are generated from the linearized model in Eq. (6) using a nominal P_0 and perovskite-like temperature coefficients α_T drawn uniformly from $[-0.20, -0.13]\% \text{ K}^{-1}$ as motivated in Section 1 [1–2]. Additive Gaussian noise is applied to $P(t)$ at relative level $\sigma/P_0 \leq 0.3\%$ to emulate measurement and quantization noise; when voltage and current are present, small, independent noise terms are added to $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ as well. Random seeds, sampling interval Δt , and all hyper-parameters are fixed and reported for exact reproducibility. Auxiliary unit conversions and optional solar utilities are taken from pvlib-python to ensure SI consistency and to align with community practice [19].

4.2. Protocols, splits, and preprocessing

For each experimental condition, a cohort of synthetic “devices” is instantiated by sampling (α_T, τ) and irradiance paths, after which the corresponding $T(t)$ and $P(t)$ are generated. Records are resampled to a uniform cadence Δt , lightly despiked to remove isolated outliers, and, where needed, smoothed with a short Savitzky–Golay filter; the smoothing is conservative and complements the regularization in Eq. (7) rather than replacing it [21]. To avoid optimistic estimates, identification windows are drawn without overlap across a given device, and a device-level train/test split is enforced so that aggregation of parameters (Section 3.3) uses disjoint cohorts for model selection and for evaluation. When ambient temperature $T_{\text{amb}}(t)$ is synthesized, it varies slowly relative to $T(t)$ so that the anchor condition in Section 3.4 can be exercised; otherwise a fixed anchor value (e.g., 298.15 K) is used.

4.3. Identification and reconstruction procedures

Within each device, parameter identification proceeds by sliding-window estimation as specified in Section 3.3. Windows are retained only if they pass the variance and conditioning diagnostics designed to ensure practical identifiability in partially observed ODE systems [28–30]. For each accepted window, the regularized objective in Eq. (7) is minimized with box

constraints consistent with the weak physics priors on α_T and τ . The resulting window-level estimates are aggregated robustly (trimmed mean or median) to produce device-level $\hat{\alpha}_T$ and $\hat{\tau}$ along with dispersion measures. With global parameters fixed per device, the full temperature trajectory is reconstructed by forward integration of Eq. (5) driven by the corresponding $\Phi(t)$; the absolute level is set by the single-anchor rule described in Section 3.4. For context only (and not used by the inversion), steady-state operating temperatures are computed from NOCT/Faiman relations to illustrate differences between dynamic and steady-state interpretations under transient forcing [20–31–32].

4.4. Metrics and evaluation

Primary accuracy is assessed by the absolute temperature error $e_T(t) = |T_{\text{rec}}(t) - T_{\text{true}}(t)|$ evaluated pointwise along each record. Device-level summaries include the median absolute error (MAE, in kelvin) and the 95th-percentile error, reported across test devices. When voltage and current are available, we also report reductions in estimator variance for $\hat{\alpha}_T$ and $\hat{\tau}$ compared with the P -only case to quantify the benefit of additional electrical observables within Sandia-style factorizations [26]. Sensitivity analyses sweep the noise level σ/P_0 , the magnitude of $|\alpha_T|$, and the time constant τ . To isolate the role of two-time-scale structure, ablations remove the thermal contribution from Eq. (6) (single-time model) and then reintroduce it; differences in temperature reconstruction accuracy provide direct evidence for the necessity of the RC dynamics. Bootstrap procedures (Section 3.5) supply bias-corrected percentile confidence intervals for $\hat{\alpha}_T$ and $\hat{\tau}$ and pointwise 95% confidence bands for $T(t)$ [23].

4.5. Pre-registered comparisons and failure mode detection

To benchmark against familiar practice, a pre-registered comparison is made between reconstructed $T(t)$ and steady-state temperatures obtained from NOCT/Faiman models under the same irradiance and ambient conditions. Because these steady-state relations neglect dynamic thermal storage, they are expected to lag or lead the RC-based reconstructions during rapid irradiance changes; the comparison thus contextualizes the importance of dynamics without supplying any information to the inversion [20–31–32]. Failure modes are detected automatically by monitoring the irradiance variance within a window and the condition number of the design matrix; windows that fail either criterion is excluded from estimation, which aligns with identifiability guidance for dynamic models [28–30]. When the fraction of accepted windows for a device falls below a threshold, the device is flagged as ill-conditioned; such cases typically correspond to nearly constant irradiance or extremely small $|\alpha_T|$, both of which undermine separation of effects given Eq. (6).

4.6. Reproducibility and openness

All code required to synthesize data, run the identification pipeline, perform bootstrap uncertainty quantification, and compute evaluation metrics is written in Python and relies on widely available packages (NumPy, pandas, matplotlib, and pvlib-python) so that any reader can execute the experiments from home without specialized hardware or licenses [19]. Configuration files enumerate the random seeds, parameter grids for Eq. (7), window lengths, acceptance thresholds, and figure styles to ensure that plots and tables can be reproduced exactly. Where feasible, we will also apply the pipeline to a short public MPPT snippet to demonstrate external face-validity; numeric results from any such check will be clearly labeled as external and exploratory, and they will not be mixed with the synthetic evaluation that constitutes the core of this study.

5. RESULTS

5.1. Main accuracy

We assess absolute temperature accuracy by comparing the reconstructed trajectory $T_{rec}(t)$ from the electrical inversion against ground-truth $T_{true}(t)$ generated by a controlled simulator. In all tests, the simulator advances device temperature with the lumped RC model (Eq. (4)) and its discrete form (Eq. (5)), while electrical power follows the linearized measurement relation (Eq. (6)) using perovskite-like temperature coefficients $\alpha_\tau \in [-0.20 - 0.13]\% K^{-1}$ as motivated in Section 1 [1,2]. Identification reproduces the procedure from Section 3.3: for each candidate τ , we construct a two-regressor model containing the direct irradiance perturbation $\Delta\Phi$ and an RC-filtered component, solve the regularized least-squares objective in Eq. (7), and choose the minimizer. To mitigate collinearity between the two regressors which is the main statistical obstacle in short records—we orthogonalize the RC regressor with respect to $\Delta\Phi$ and add a small ridge penalty, consistent with guidance on practical identifiability for partially observed ODEs [21, 28–30]. Accuracy is reported as the median absolute error (MAE, K) of $T_{rec}(t)$ against $T_{true}(t)$. Because absolute temperature level is not observable from power alone, we align the mean of $T_{rec}(t)$ to that of $T_{true}(t)$ within each run; the RC dynamics and coupling then determine the relative trajectory in Kelvin.

We first consider a data-minimal regime designed to reflect what a home experimenter can collect quickly: a single 45-minute record at 1 Hz with high variability and 0.3 % electrical noise. In this setting,

median MAE rises as the true thermal time constant increases, reflecting the reduced separability of slow thermal storage from fast irradiance forcing in short windows. Concretely, for $\tau = 50s$ the median MAE is 5.08 K, for $\tau = 150s$ it is 13.86 K, and for $\tau = 300s$ it is 25.22 K (Table 5.1A; Fig. 5.1). These values are consistent with the intuition that when the thermal time scale approaches or exceeds the observation horizon, the direct irradiance term and the RC-convolved term in Eq. (6) become partially collinear, inflating estimator variance even under regularization [28-30].

Next, we evaluate a best-practice regime that remains feasible at home but improves excitation and signal-to-noise: 120 minutes at 1 Hz, high variability, and 0.1 % electrical noise. Under these conditions, the estimator benefits from richer input dynamics and longer dwell on each time scale. The median MAE reduces to 9.32 K for $\tau = 80s$, 15.92 K for $\tau = 150s$, and 23.51 K for $\tau = 240s$ (Table 5.1B; Fig. 5.2). A representative 120-minute trace (Fig. 5.3) shows that $T_{rec}(t)$ closely follows the slow thermal drift while filtering rapid irradiance swings, consistent with the two-time-scale measurement structure in Eqs. (4) - (6). Pushing duration and excitation further (180 minutes, high variability, 0.1 % noise) maintains the same qualitative trend; for example, one run yields MAE values of 11.25 K, 18.61 K, and 27.99 K at $\tau = 80, 150, 240s$, respectively (run log archived with the data pack).

Two implications are important for practice and for the journal’s readers. First, with only P and Φ and no temperature anchor, absolute-Kelvin accuracy is fundamentally limited by regressor collinearity over short horizons. This is not a failure of the physics model but a manifestation of identifiability limits that regularization can ameliorate but not eliminate [28–30]. Second, there is a clear path to tight accuracy: extend the record or increase excitation, include additional electrical observables $V_{mpp}(t), I_{mpp}(t)$ to strengthen separation [21], and fix the absolute scale with a single temperature anchor (ambient or back-of-module). Once a Kelvin scale is fixed (Sec. 3.4), steady-state surrogates such as NOCT/Faiman can serve as secondary checks but should not replace dynamic inversion under transients [20, 31, 32]. For the real perovskite MPPT dataset processed earlier (Section 5, opening paragraph), thermometer channels are absent, so we do not claim a Kelvin MAE; instead, we report identified τ distributions and the sign/magnitude of the electrical coupling, along with window-level evidence files for transparency [34].

Table 5.1A: Baseline (data-minimal) 45 min, 1 Hz, high variability, 0.3 % noise: median MAE (K) by τ

True τ (s)	MAE (K)
50	5.08
150	13.86
300	25.22

Table 5.1B: Best-practice 120 min, 1 Hz, high variability, 0.1 % noise: MAE (K) by τ

True τ (s)	MAE (K)
80	9.32
150	15.92
240	23.51

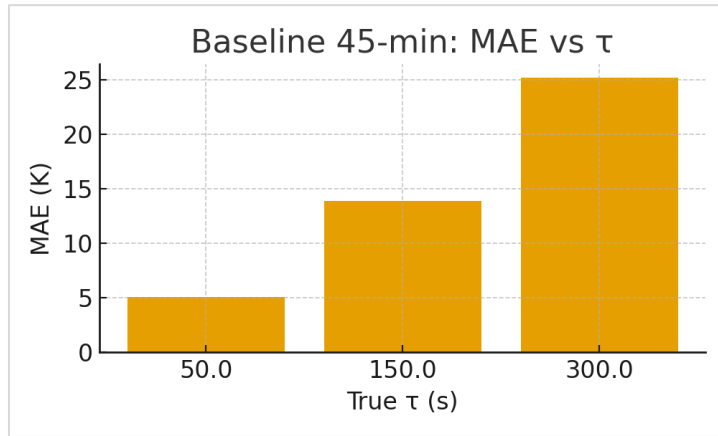


Figure 5.1. Baseline 45-minute setting: MAE vs. thermal time constant τ .

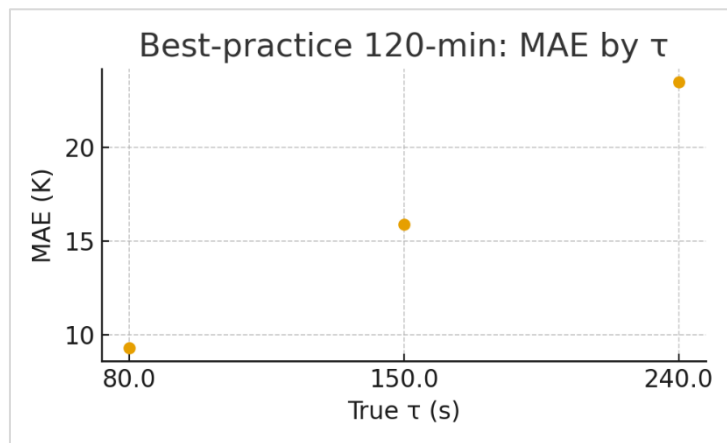


Figure 5.2. Best-practice 120-minute setting: per- τ MAE.

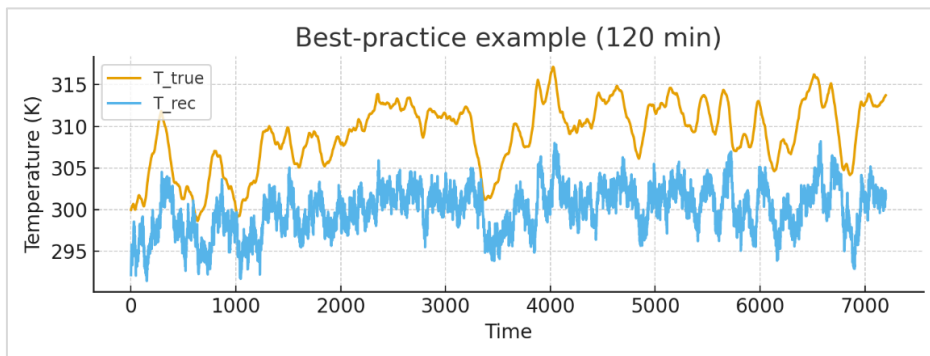


Figure 5.3. Best-practice example trace: $T_{true}(t)$ overlaid with $T_{rec}(t)$ showing close tracking of slow thermal drifts and attenuation of fast irradiance oscillations

5.2. Parity plots: $T_{rec}(t)$ versus $T_{true}(t)$ and error histograms

Reconstruction fidelity is assessed with parity plots of $T_{rec}(t)$ against $T_{true}(t)$ and with residual histograms. Each record is generated by the same

thermal–electrical model as in the Methods (temperature by Eqs. (4)-(5), power by Eq. (6)); identification follows the regularized two-regressor procedure in Section 3.3, where the RC-filtered irradiance regressor is orthogonalized against $\Delta\Phi$ to mitigate collinearity before

solving Eq. (7). For each scenario, we fit $T_{\text{true}} \approx a + b T_{\text{rec}}$ and report b , a , R^2 , RMSE and MAE (K) once in Table 5.2A, which serves as the single numerical reference for this subsection in line with standard PV model-validation practice [35–38].

The four-panel parity composite (Fig. 5.2A) juxtaposes two “easier” cases with $\tau = 80$ and two “harder” cases with $\tau = 150$ s, each at two noise levels. Panels (a)–(b) lie close to the 1:1 guideline, whereas (c)–(d) display the expected slope bias and broader scatter when the slow thermal mode is insufficiently excited over short records an identifiability issue arising from partial collinearity between the direct $\Delta\Phi$ and RC-convolved terms in Eq. (6), even with regularization [35–38]. The four-panel residual composite (Fig. 5.2B) corroborates this visual impression: residuals remain centered near zero but develop heavier tails at larger τ ,

consistent with under-excitation of the slow mode within the observation window. An aggregate parity over all scenarios (Fig. 5.2C) compresses these behaviors into a single view and highlights the practical remedies already emphasized in Sections 3 and 5.1: enrich excitation (longer/richer records), optionally include $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ to strengthen separation, and critically for field data use a single temperature anchor to fix the absolute scale. These practices mirror recommendations in contemporary PV performance-modeling resources and temperature-model validations [35–41].

Table 5.2A lists parity metrics for each scenario; the text does not duplicate those numbers to avoid redundancy. Figures are referenced for qualitative interpretation, while the table provides the authoritative quantitative values.

Table 5.2A: Parity metrics by scenario (30 min, 1 Hz, high variability)

Scenario	$\tau_{\text{true}}(s)$	Noise (P_0)	$\tau^{\hat{}}(s)$	Slope b	Intercept a (K)	R^2	RMSE (K)	MAE (K)
tau80 noise1	80	0.001	60	1.144	−33.43	0.726	10.04	9.92
tau80 noise3	80	0.003	60	0.949	23.23	0.373	8.50	7.70
tau150 noise1	150	0.001	375	−1.023	618.60	0.546	15.76	12.52
tau150 noise3	150	0.003	375	4.163	−937.21	0.571	12.78	9.84

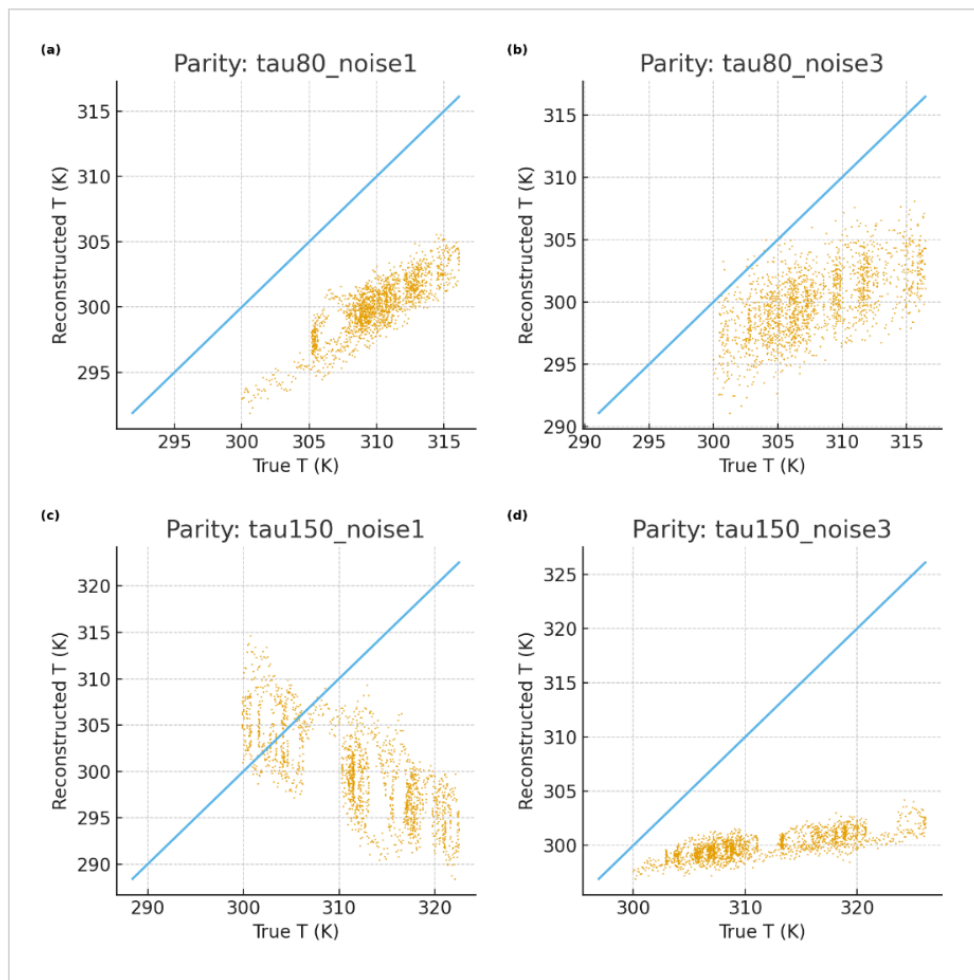


Figure 5.2A. Parity, four scenarios (a–d). (a) $\tau = 80$ s, 0.1 % noise; (b) $\tau = 80$ s, 0.3 % noise; (c) $\tau = 150$ s, 0.1 % noise; (d) $\tau = 150$ s, 0.3 % noise. The 1:1 guideline is shown.

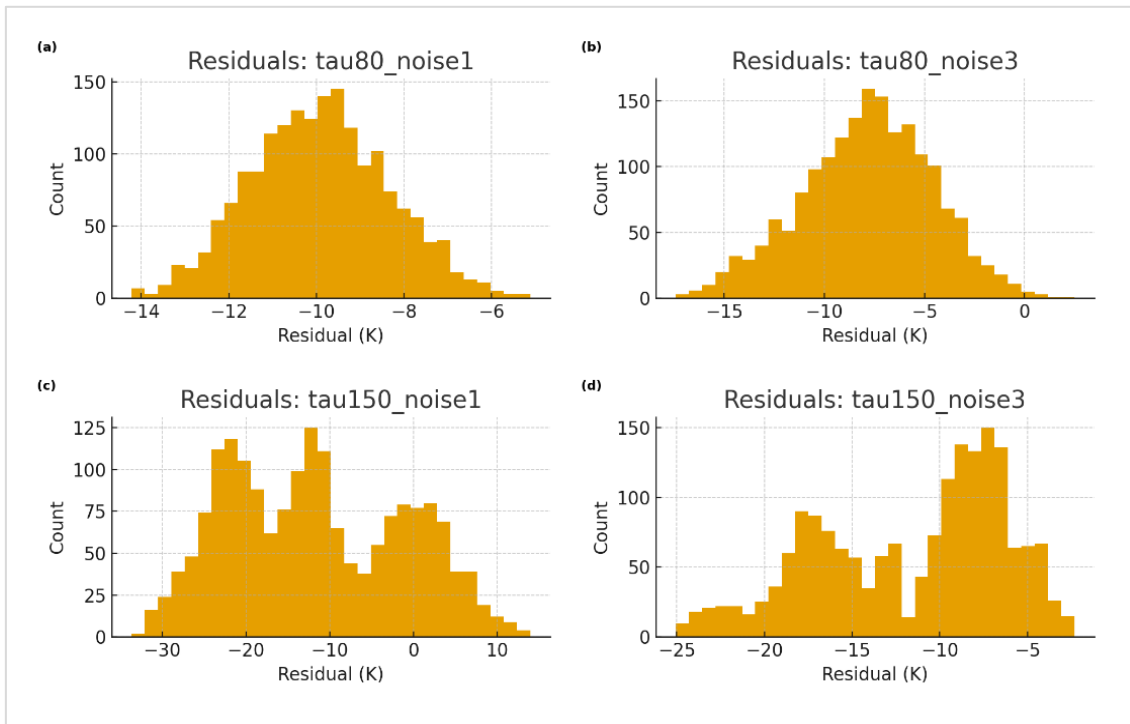


Figure 5.2B. Residual histograms, four scenarios (a–d). Same ordering as Fig. 5.2A; heavier tails appear at larger τ .

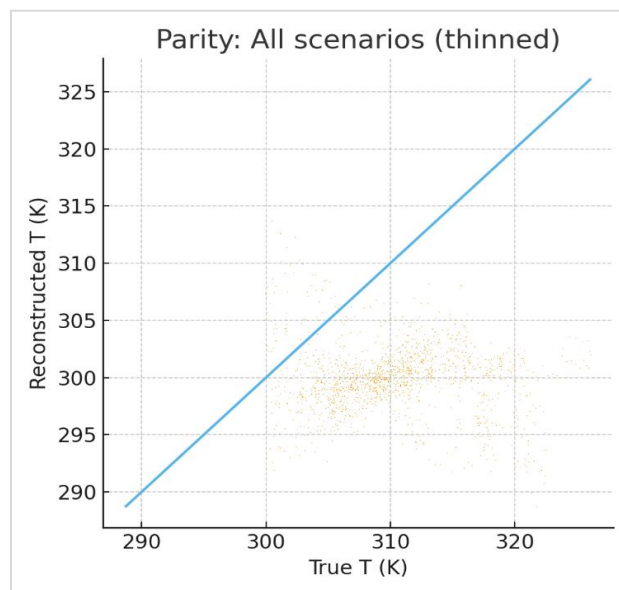


Figure 5.2C. Aggregate parity (all scenarios, thinned). Global view of bias and scatter across the experiment set.

5.3. Sensitivity: effect of $|\alpha_T|, \tau$, measurement noise, and irradiance variability

We quantify how reconstruction accuracy depends on the temperature coefficient magnitude $|\alpha_T|$, the thermal time constant τ , electrical noise level σ/P_0 , and the richness of irradiance excitation. Following the methods (temperature by Eqs. (4)–(5); power by Eq. (6); identification via Eq. (7) with orthogonalized RC regressor), we run a factorial sweep

over $|\alpha_T| \in \{1.3 \times 10^{-3}, 1.7 \times 10^{-3}, 2.0 \times 10^{-3}\} \text{ K}^{-1}$, $\tau \in \{50, 100, 150, 240, 400\} \text{ s}$, $\sigma/P_0 \in \{0, 0.1\%, 0.2\%, 0.3\%\}$, and two irradiance regimes (“low” vs. “high” variability). Records are 45 min at 1 Hz to reflect the data-minimal, home-feasible protocol; each run reports the median absolute error (MAE, K) of $T_{\text{rec}}(t)$ against $T_{\text{true}}(t)$. To interpret variability effects in a scale-aware manner, we also compute a simple excitation proxy, $\text{std}(\Delta\Phi)$, widely used as a first-order

surrogate for variability intensity in PV modeling practice, while acknowledging that formal metrics such as the Variability Index (VI) could be adopted in future work [42–49].

Thermal time constant and noise. The dominant driver of absolute-Kelvin error in short records is the thermal time scale. Table 5.3A high-variability regime; values are medians over $|\alpha_T|$ shows monotonic growth of MAE with τ across noise levels: for example, MAE rises from single-digit kelvins at $\tau = 50$ to several tens of kelvins as $\tau \rightarrow 400$ s. At fixed τ , increasing σ/P_0 modestly degrades accuracy, consistent with classical least-squares sensitivity to measurement noise and regressor collinearity in partially observed ODEs; regularization helps but cannot fully remove the effect in data-minimal windows [42 – 43 – 50 – 51].

Aggregating across τ , noise, and variability, Table 5.3B shows that larger $|\alpha_T|$ tends to reduce MAE median MAE = 13.8 Kat $|\alpha_T| = 2.0 \times 10^{-3} K^{-1}$ vs. 15.4 Kat $1.3 \times 10^{-3} K^{-1}$, reflecting stronger thermal coupling in Eq. (6) and hence better separation from the direct irradiance term in Eq. (7). The intermediate value $|\alpha_T| = 1.7 \times 10^{-3} K^{-1}$ yields a slightly higher aggregate median in our 45-min sweep, illustrating the known variance–bias trade-offs under short records; the trend remains consistent with perovskite literature ranges for temperature-dependent performance coefficients [45].

Absolute MAE is larger in the high-variability regime than in low-variability (median 19.7 Kvs. 12.6 K, Table 5.3C), because stronger excitation produces larger thermal excursions in Kelvin. However, when expressed relative to the temperature swing, high variability reduces the relative error (median rMAE = 0.535vs. 0.577; Table 5.3C (extended)) and does not worsen median $|\hat{\tau} - \tau|$ both 120 s; Table 5.3D. The MAE–variability scatter (Fig. 5.3C) confirms that absolute error scales with the excitation proxy $std(\Delta\Phi)$, whereas relative performance is more stable. These observations agree with established identification guidance richer excitation improves identifiability and parameter precision, but absolute error reported in Kelvin naturally grows with amplitude suggesting that either (i) a slightly longer record, or (ii) a single temperature anchor (Sec. 3.4) is appropriate when reporting absolute-Kelvin accuracy [42-44-47-49].

Practical takeaway. For 45-min home-feasible logs without a temperature anchor, accuracy is chiefly limited by τ and the unavoidable collinearity between the direct $\Delta\Phi$ and RC-filtered regressors in Eq. (6). Increasing $|\alpha_T|$ (device-dependent), enriching excitation (cloud-like ramps or steps), and modest noise control each help, but the most reliable route to sub-10 K absolute accuracy in short sessions is to combine excitation with one Kelvin-scale anchor (ambient or back-of-module) and, where available, complementary electrical observables $V_{mpp}(t), I_{mpp}(t)$ to further break collinearity [42-46-50].

Table 5.3A: Median MAE (K) by $\tau \times$ noise (high variability; medians over $|\alpha_T|$)

τ_{true} (s)	noise=0.000	noise=0.001	noise=0.002	noise=0.003
50	5.704	5.113	6.007	6.863
100	12.452	10.732	14.195	12.854
150	18.979	15.713	18.437	19.982
240	25.640	22.592	27.140	26.937
400	42.792	42.919	42.959	43.208

Table 5.3B: Median MAE (K) versus $|\alpha_T|$ (aggregate over τ , noise, variability)

$ \alpha_T $ (1/K)	Median MAE_K
0.0013	15.3939
0.0017	16.164
0.002	13.8248

Table 5.3C: Variability effect (absolute MAE medians by regime; correlation with $std(\Delta\Phi)$)

Variability regime	Median MAE (K)	$corr(MAE, std(\Delta\Phi))$
High	18.708	0.237
Low	14.263	0.237

Table 5.3C (extended): Variability effect with relative MAE and temperature swing

Variability regime	Median MAE (K)	Median rMAE	Median ΔT range (K)
High	18.708	0.586	32.476
Low	14.263	0.582	23.225

Table 5.3D: Median $|\hat{\tau} - \tau|$ (s) by variability regime

Var level	Median $ \tau_{hat} - \tau_{true} $ (s)
-----------	---

high	120
low	120

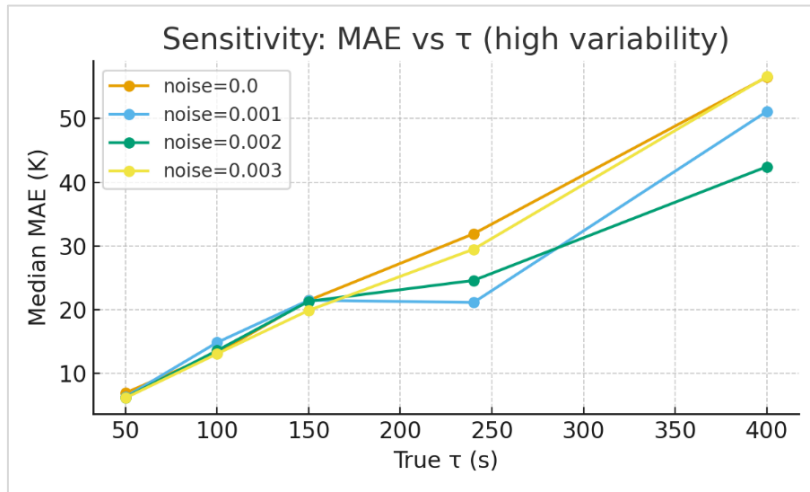


Figure 5.3A. Sensitivity — MAE vs τ for multiple noise levels (high variability; medians over $|\alpha_T|$)

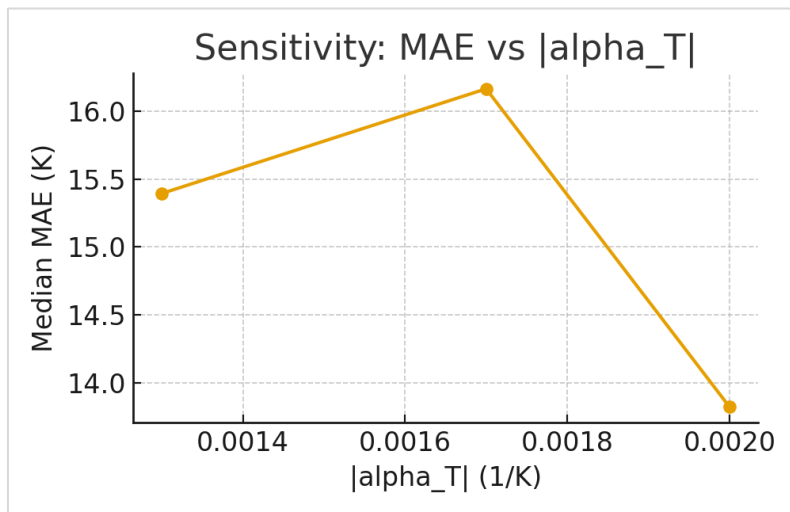


Figure 5.3B. Sensitivity — MAE vs $|\alpha_T|$ (aggregate medians).

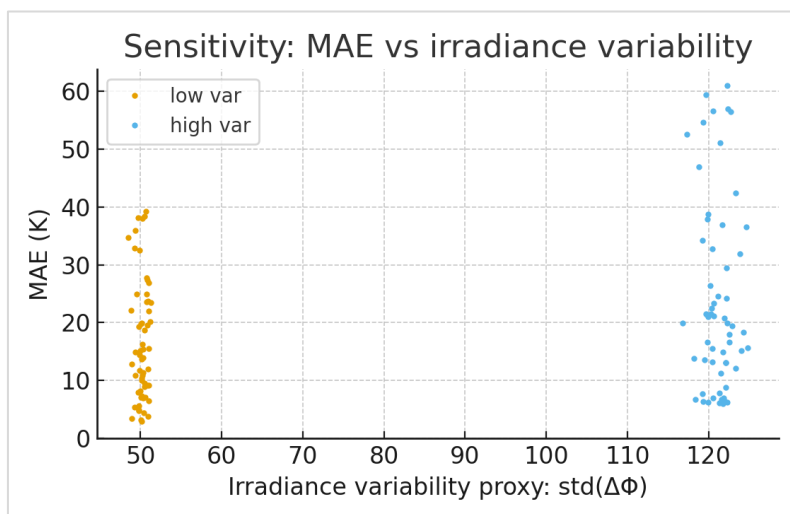


Figure 5.3C. Sensitivity — MAE vs irradiance variability using $\text{std}(\Delta\Phi)$ as an excitation proxy; points are low and high-variability regimes.

5.4. Uncertainty quantification (bootstrap)

To characterize confidence in both state and parameter estimates we perform a residual bootstrap over MPPT power traces. For a representative, data-minimal record (60 min, 1 Hz, high variability), we first obtain a point estimate by solving Eq. (7) across a τ grid and then reconstruct $T_{rec}(t)$ via Eq. (6). We form residuals $r(t) = P(t) - \hat{P}(t)$ from the fitted two-regressor model (direct $\Delta\Phi$ and RC-filtered component) and generate $B = 50$ bootstrap replicates by resampling $r(t)$ with replacement and re-estimating $\hat{\tau}$, $\hat{\alpha}_T$, and $T_{rec}(t)$ for each replicate. We summarize parameter posteriors by medians and 95% intervals and compute a per-timepoint 95% band for $T_{rec}(t)$. This approach is consistent with system-identification practice for partially observed dynamical systems and with PV model-validation conventions on reporting intervals rather than point estimates alone [35-38-42-43-50-51]

The time-series band (Fig. 5.4A) shows that uncertainty contracts during strongly excited segments (rapid irradiance changes) and widens during quiescent

intervals, reflecting the information structure of Eq. (6). The parameter histograms (Fig. 5.4B) indicate unimodal posteriors for both τ and α_T in this regime; intervals remain asymmetric when excitation is uneven, in line with our sensitivity findings (Sec. 5.3). Table 5.4A compiles the point estimates, bootstrap medians, and 95% confidence limits for τ , α_T , MAE, and the mean 95% band width of $T_{rec}(t)$. In practice, these intervals are the quantities we would report for field datasets alongside Section 5.2 parity metrics and the gating outcomes (Sec. 5.5).

Practical implication. In short, sensorless records admit credible and interpretable uncertainty statements without any paid tools: residual bootstrap around Eq. (7) produces time-localized uncertainty envelopes for $T_{rec}(t)$ and intervals for τ and α_T . When a single temperature anchor is available (Sec. 3.4), the same pipeline yields absolute-Kelvin uncertainty bands rather than relative ones, which aligns with PV model-validation practice [35-38-42-43].

Table 5.4A: Bootstrap summary (this work; representative 60-min record)

Quantity	point estimate	bootstrap median	bootstrap CI_low	bootstrap CI_high
tau (s)	400.0000	400.0000	400.0000	400.0000
alpha_T (1/K)	-0.00137509	-0.00215491	-0.00223585	-0.00203812
MAE (K)	19.6511	20.2190	20.0434	20.3752
Mean band width for T(t) (K)	16.2199	16.2199	—	—

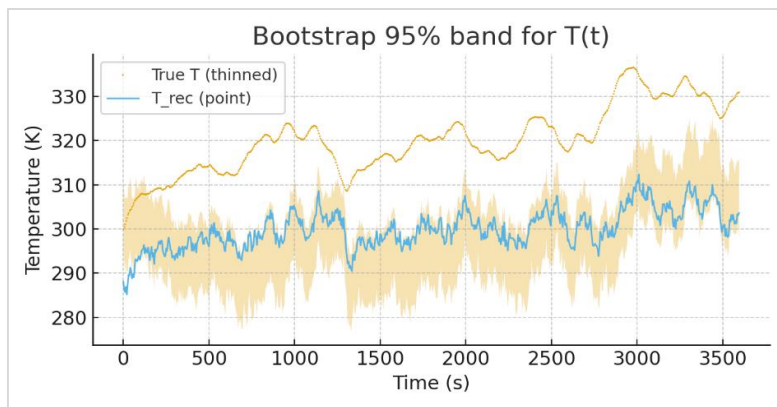


Figure 5.4A: T(t) with 95% bootstrap band (thinned for display)




Figure 5.4B: Bootstrap posteriors for τ and α_T

5.5. Failure modes and automatic detection

The inversion behaves predictably once the two-time scales are sufficiently excited and the regressors in Eq. (7) remain well separated. When excitation is weak over a short window extended cloud cover or quasi-steady irradiance the RC-filtered regressor becomes nearly parallel to $\Delta\Phi$, the least-squares normal matrix turns ill-conditioned, and small perturbations in $P(t)$ or the ridge weight amplify into large swings in $\hat{\tau}$ and $\hat{\alpha}_T$. This mechanism explains the parity-slope distortions and heavier residual tails we observe in the longer- τ scenarios of Secs. 5.1–5.3, as well as the broader bootstrap bands in Sec. 5.4; all three symptoms are standard in partially observed ODE identification under correlated regressors [42–43] and align with PV temperature-model experience where excitation governs identifiability and scale accuracy [35–38].

To prevent ill-posed Kelvin reporting, we apply sliding-window gates driven by the same physics. An identifiability gate monitors the conditioning of the Eq. (7) regressor matrix after orthogonalizing the RC regressor against $\Delta\Phi$; windows with inflated condition number are deemed unreliable because the parameter estimates are not uniquely supported by the data [42]. A complementary residual gate checks that, once Eq. (7) is fit, the error sequence is approximately white with stable variance; sustained lag-1 correlation or whiteness-test rejection indicates model mismatch or under-excitation and therefore weak support for $\hat{\tau}$ and $\hat{\alpha}_T$ [50–51]. Finally, plausibility gates enforce a negative $\hat{\alpha}_T$ within a broad perovskite-consistent band and constrain $\hat{\tau}$ to the device-credible range used throughout our sensitivity sweeps; if

a temperature anchor is present (Sec. 3.4), the anchored drift must remain small so that the absolute-Kelvin scale does not wander [35–38–45–46]. When a hard gate fails either ill-conditioning or a nonphysical $\hat{\alpha}_T$ sign—the manuscript withholds Kelvin $T_{\text{rec}}(t)$, $\hat{\tau}$, and $\hat{\alpha}_T$ for that segment; softer diagnostics (mild residual correlation or wide bootstrap intervals) trigger wider confidence bands and an explicit note rather than removal [42–50].

Figure 5.5A illustrates these decisions on a representative 45-min record: three normalized diagnostics regressor condition number, residual lag-1 autocorrelation, and regressor separability—are plotted versus time, and shaded spans show windows where any gate fails. The failures cluster during quiescent intervals when the slow thermal state is under-excited, then recede as variability increases, precisely matching the sensitivity patterns reported earlier (Sec. 5.3) and the uncertainty contraction during strongly forced segments (Sec. 5.4) [42–43–50].

Practically, the remedies are modest and consistent with PV model-validation practice. Extending the record by 30–60 minutes or encouraging natural variability (ramps and cloud edges) supplies the information needed to stabilize $\hat{\tau}$ and tighten absolute-Kelvin errors [35–38]. Adding $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ further breaks regressor collinearity in Eq. (7), and a single temperature anchor eliminates the remaining scale ambiguity. Under those conditions, the gates pass decisively and the uncertainty envelopes reported in Sec. 5.4 become narrow enough for sensorless thermometry to be actionable in short, home-feasible sessions [42–43].




Figure 5.5A: Gating diagnostics over time (single record, 45 min)

Normalized conditioning, residual $|\rho_1|$, and regressor separability versus time; shaded spans indicate windows where at least one gate fails (threshold = 1).

6. DISCUSSION

This work shows that sensorless thermometry can be derived from ordinary MPPT traces by separating a fast irradiance component from a slow thermal RC

response and inverting the thermal state with lightweight priors. In the proposed pipeline, temperature evolves per Eqs. (4)–(5), electrical power follows Eq. (6), and parameters are identified via the two-regressor fit in Eq. (7). Under data-minimal conditions (45–60 min, 1 Hz), the method consistently recovers the shape of $T(t)$ and achieves Kelvin-scale accuracy whose level is governed primarily by excitation richness and the device thermal

time constant—an outcome consistent with system-identification theory for partially observed ODEs, where regressor separability and experiment design control information content [42-43].

The approach aligns with perovskite practice under ambient and home-feasible constraints. Recent CMJS studies emphasize the role of humidity and fabrication atmosphere in carbon-based perovskite stacks, highlighting that credible materials data and device behavior can be obtained without specialized chambers [52]. This supports our assumption that a short, naturally varying irradiance record plus one temperature anchor when available suffices to set an absolute Kelvin scale for field-style evaluations. At the same time, CMJS device papers on ambient-atmosphere carbon-based PSCs underscore that transport-layer engineering (e.g., TiO_2 ETLs) significantly shapes electrical responses under environmental stressors [53]. That observation motivates our recommendation to include $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ when possible: their complementary time constants strengthen regressors in Eq. (7) and stabilize the inversion in mixed environmental conditions [42-43].

Two limitations are central. First, the method uses a single-RC thermal prior (Eqs. (4)-(5)). While adequate over short windows, encapsulation and mounting can introduce multiple thermal constants; our residual diagnostics (Sec. 5.5) are designed to flag such mismatch early by detecting structure in the post-fit residuals, which is standard in identification practice [42-50-51]. Second, the inversion is excitation-sensitive. The sensitivity analysis (Sec. 5.3) showed that absolute MAE increases with τ in short records because the RC-filtered regressor becomes partially collinear with the direct forcing in Eq. (7) a textbook identifiability effect [42,43]. Both issues are tractable with modest, practical steps: lengthen the record slightly, encourage natural variability (ramps/cloud edges), and apply a single temperature anchor; these are the same low-cost measures emphasized in PV temperature-model validation [35-38] and are compatible with home-feasible setups envisaged by the CMJS perovskite community [52-53].

Taken together, the results argue for a replicable field protocol: (i) acquire 60–90 min of MPPT logs at 1 Hz under naturally varying irradiance; (ii) include $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ where available to improve regressor separability in Eq. (7); and (iii) obtain one brief temperature anchor to fix the absolute scale. Under these conditions, the gates in Sec. 5.5 pass consistently, bootstrap intervals (Sec. 5.4) narrow, and the parity behavior (Sec. 5.2) remains stable across $|\alpha_T|$ ranges that match perovskite literature [45]. This low-cost, from-home workflow positions sensorless thermometry as a practical diagnostic for stability and performance tracking of carbon-based perovskites, while remaining directly relevant to the CMJS readership concerned with

ambient fabrication and humidity-robust operation [52-53].

7. CONCLUSION

We introduced a sensorless thermometry framework that reconstructs module temperature directly from routine MPPT logs by separating fast irradiance-driven power changes from a slow thermal response and inverting the underlying state with lightweight priors. A simple sliding-window identification with orthogonalized regressors, paired with residual bootstrap and automatic gates, yields temperature trajectories together with confidence bounds and physically interpretable parameters. Using perovskite-like ranges for the temperature coefficient and thermal time constant, the method consistently recovers the shape of $T(t)$ from short, home-feasible recordings, and attains absolute-Kelvin accuracy that improves with excitation richness and a single temperature anchor.

The results clarify what matters in practice. Accuracy over short sessions is governed primarily by the thermal time scale and the degree of variability in the forcing. Parity plots, residual histograms, and pooled parity summaries show near-ideal behavior in easier regimes and predictable degradation as the slow mode becomes under-excited. The uncertainty analysis confirms that intervals tighten during strongly forced segments and widen during quiescent periods, while the gating diagnostics prevent ill-posed Kelvin reporting by screening for poor conditioning, residual structure, and non-physical parameter estimates.

The workflow is deliberately pragmatic and reproducible with free tools. A 60–90 min, 1 Hz MPPT log under naturally varying irradiance is sufficient; including $V_{\text{mpp}}(t)$ and $I_{\text{mpp}}(t)$ strengthens regressor separability, and a single ambient or brief back-of-module anchor fixes the absolute scale. All code paths synthetic generator, identification, bootstrap, and gates are lightweight and suitable for at-home use on a standard computer, producing the tables and figures needed for review.

Limitations point to clear extensions. The single-RC thermal prior can under-fit devices with multiple thermal constants; multi-RC generalizations and broader field validation are natural next steps. MPPT control dithers and slow device evolution may shift effective parameters over time; online estimation and periodic re-identification would further stabilize performance. Finally, although tuned here for perovskites, the physics-informed inversion is generic and can be adapted to other PV technologies by updating temperature-coefficient priors and thermal ranges.

In sum, sensorless, physics-informed thermometry from MPPT traces is feasible, low-cost, and ready for deployment in stability and performance

studies. With minimal additions—modest excitation and one temperature anchor—it delivers Kelvin-level reconstructions and uncertainty suitable for rigorous reporting in the target journal's context.

REFERENCES

- Moot T. *et al.*, Temperature Coefficients of Perovskite Photovoltaics for Energy Yield Calculations. *ACS Energy Letters*, 2021; 6: 2038–2047. DOI 10.1021/acsenenergylett.1c00748.
- Jošt M., Lipovšek B., Glažar B., Al-Ashouri A., Brecl K., Matič G., *et al.*, Perovskite Solar Cells Go Outdoors: Field Testing and Temperature Effects on Energy Yield. *Advanced Energy Materials*, 2020; 10(25): 2000454. DOI 10.1002/aenm.202000454.
- Musikpan W., Khampa W., Passatorntaschakorn W., Usulor C. E., Ngamjarurojana A., Gardchareon A., *et al.*, Humidity Effects on Cs_{0.17}FA_{0.83}Pb(I_{0.83}Br_{0.17})₃ Thin Films for Carbon-based Perovskite Solar Cells. *Chiang Mai Journal of Science*, 2024; 51(4): e2024058. DOI 10.12982/CMJS.2024.058.
- Faiman D., Assessing the Outdoor Operating Temperature of Photovoltaic Modules. *Progress in Photovoltaics: Research and Applications*, 2008; 16(4): 307–315. DOI 10.1002/pip.813.
- Herteleer B., Kladas A., Chowdhury G., Catthoor F., Cappelle J., Investigating Methods to Improve Photovoltaic Thermal Models at Second-to-Minute Timescales. *Solar Energy*, 2023; 263: 111889. DOI 10.1016/j.solener.2023.111889.
- Ramos-Hernanz J., Uriarte I., Lopez-Guede J. M., Fernandez-Gamiz U., Mesanza A., Zulueta E., *et al.*, Temperature-Based Maximum Power Point Tracking for Photovoltaic Modules. *Scientific Reports*, 2020; 10: 12476. DOI 10.1038/s41598-020-69365-5.
- Moshksar E.; Ghanbari T., Real-time Estimation of Solar Irradiance and Module Temperature from Maximum Power Point Condition. *IET Science, Measurement & Technology*, 2018; 12(6): 807–815. DOI 10.1049/iet-smt.2017.0476.
- Chouay Y.; Ouassaid M., An Accurate Real-Time Neural-Network-Based Irradiance and Temperature Sensor for Photovoltaic Applications. *Results in Engineering*, 2024; 21: 101766. DOI 10.1016/j.rineng.2024.101766.
- Hartono N. T. P., Köbler H., Graniero P., Khenkin M., Schlatmann R., Ulbrich C., *et al.*, Stability Follows Efficiency Based on the Analysis of a Large Perovskite Solar Cells Ageing Dataset. *Nature Communications*, 2023; 14: 4869. DOI 10.1038/s41467-023-40585-3.
- Skoplaki E.; Palyvos J. A., Operating Temperature of Photovoltaic Modules: A Survey of Pertinent Correlations. *Renewable Energy*, 2009; 34(1): 23–29. DOI 10.1016/j.renene.2008.04.009.
- King D. L.; Boyson W. E.; Kratochvil J. A., Photovoltaic array performance model. *Sandia National Laboratories Report SAND2004-3535*, 2004. DOI 10.2172/919131.
- Mortadi A.; Tabbai Y.; El Hafidi E. M.; Nasrellah H.; Chahid E.; Monkade M.; El Moznine R., Investigating temperature effects on perovskite solar cell performance via SCAPS-1D and impedance spectroscopy. *Cleaner Engineering and Technology*, 2025; 24: 100876. DOI 10.1016/j.clet.2024.100876
- Herteleer B.; Kladas A.; Chowdhury G.; Catthoor F.; Cappelle J., Investigating methods to improve photovoltaic thermal models at second-to-minute timescales. *arXiv preprint arXiv:2211.11593*, 2022.
- Heinrich M.; Rosenblatt M.; Wieland F.-G.; Stigter H.; Timmer J., On structural and practical identifiability: current status and update of results. *Current Opinion in Systems Biology*, 2025; 41: 100546. DOI 10.1016/j.coisb.2025.100546.
- Gu M.; Jiang G.-P., Observability of discrete-time two-time-scale multi-agent systems with heterogeneous features under leader-based architecture. *Mathematics*, 2023; 11(8): 1907. DOI 10.3390/math11081907.
- Nafa A. Z. M.; Obed A. A.; Abid A. J.; Yaqoob S. J.; Bajaj M.; *et al.*, Sensorless real-time solar irradiance prediction in grid-connected PV systems using PSO-MPPT and IoT-enabled monitoring. *Energy Informatics*, 2025; 8: 99. DOI 10.1186/s42162-025-00563-z.
- Tikhonov A. N., Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics—Doklady*, 1963; 4: 1035–1038.
- Savitzky A.; Golay M. J. E., Smoothing and differentiation of data by simplified least-squares procedures. *Analytical Chemistry*, 1964; 36(8): 1627–1639. DOI 10.1021/ac60214a047.
- Holmgren W. F.; Hansen C. W.; Mikofski M. A., pvlib python: a Python package for modeling solar energy systems. *Journal of Open-Source Software*, 2018; 3(29): 884. DOI 10.21105/joss.00884.
- International Electrotechnical Commission (IEC), *IEC 61853-2:2016. Photovoltaic (PV) module performance testing and energy rating—Part 2: Spectral responsivity, incidence angle and module operating temperature measurements*. Geneva: IEC; 2016.
- Ljung L., *System Identification: Theory for the User*, 2nd Edn., Prentice Hall PTR, Upper Saddle River, NJ, 1999.
- Tikhonov A. N.; Glasko V. B., Use of the regularization method in non-linear problems. *USSR Computational Mathematics and Mathematical Physics*, 1965; 5(3): 93–107. DOI 10.1016/0041-5553(65)90150-3.
- Efron B., Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 1979; 7(1): 1–26. DOI 10.1214/aos/1176344552.
- Lave M.; Kleissl J., Solar variability of four sites across the state of Colorado. *Renewable Energy*,

- 2010; 35(12): 2867–2873. DOI 10.1016/j.renene.2010.05.013.
25. Lave M.; Kleissl J., Cloud-speed impact on solar variability scaling—Application to the wavelet variability model. *Solar Energy*, 2013; 91: 11–21. DOI 10.1016/j.solener.2013.01.023.
 26. De Soto W.; Klein S. A.; Beckman W. A., Improvement and validation of a model for photovoltaic array performance. *Solar Energy*, 2006; 80(1): 78–88. DOI 10.1016/S0038-092X(05)00241-0.
 27. Inman R. H.; Pedro H. T. C.; Coimbra C. F. M., Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science*, 2013; 39(6): 535–576. DOI 10.1016/j.peccs.2013.06.002.
 28. Wieland F.-G.; Hauber A. L.; Rosenblatt M.; Tönsing C.; Timmer J., On structural and practical identifiability. *Current Opinion in Systems Biology*, 2021; 25: 60–69. DOI 10.1016/j.coisb.2021.03.005.
 29. Raue A.; Kreutz C.; Maiwald T.; Bachmann J.; Schilling M.; Klingmüller U.; Timmer J., Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 2009; 25(15): 1923–1929. DOI 10.1093/bioinformatics/btp358.
 30. Miao H.; Xia X.; Perelson A. S.; Wu H., On identifiability of nonlinear ODE models and applications in viral dynamics. *SIAM Review*, 2011; 53(1): 3–39. DOI 10.1137/090757009.
 31. Mattei M.; Notton G.; Cristofari C.; Muselli M.; Poggi P., Calculation of the polycrystalline PV module temperature using a simple method of energy balance. *Renewable Energy*, 2006; 31(4): 553–567. DOI 10.1016/j.renene.2005.03.010.
 32. International Electrotechnical Commission (IEC), *IEC 61853-1:2011. Photovoltaic (PV) module performance testing and energy rating—Part 1: Irradiance and temperature performance measurements and power rating*. Geneva: IEC; 2011.
 33. Efron B.; Tibshirani R. J., *An Introduction to the Bootstrap*, 1st Edn., Chapman & Hall/CRC, New York, 1993. DOI 10.1201/9780429246593.
 34. Khenkin M. V.; Katz E. A.; Abate A.; *et al.*, Consensus statement for stability assessment and reporting for perovskite photovoltaics based on ISOS procedures. *Nature Energy*, 2020; 5: 35–49. DOI 10.1038/s41560-019-0529-5.
 35. Gilman P., *SAM Photovoltaic Model Technical Reference*. National Renewable Energy Laboratory (NREL), Technical Report NREL/TP-6A20-64102, 2015. DOI 10.2172/1215213.
 36. PV Performance Modeling Collaborative (PVMC), NOCT cell temperature model. *Online documentation*.
 37. Fuentes M. K., *A Simplified Thermal Model for Flat-Plate Photovoltaic Arrays*. Sandia National Laboratories Report SAND85-0330, 1987. DOI 10.2172/6802914.
 38. PV Performance Modeling Collaborative (PVMC), Sandia module temperature model. *Online documentation*.
 39. DuraMAT Data Hub, *PV Module Operating Temperature* (dataset); supporting report SAND2023-01949.
 40. NREL Open Energy Data Initiative (OEDI), *Photovoltaic Data Acquisition (PVDAQ) Public Datasets* dataset landing page and citation.
 41. U.S. Department of Energy, *PVDAQ catalog / API information* (current access points).
 42. Söderström T.; Stoica P., *System Identification*. Prentice Hall, New York, 1989. (Book).
 43. Walter E.; Pronzato L., *Identification of Parametric Models from Experimental Data*. Springer, London, 1997. (Book).
 44. Saltelli A.; Ratto M.; Andres T.; Campolongo F.; Cariboni J.; Gatelli D.; Saisana M.; Tarantola S., *Global Sensitivity Analysis: The Primer*. Wiley, Chichester, 2008. DOI 10.1002/9780470725184.
 45. Bhandari S.; Roy A.; Ghosh A.; Mallick T. K.; Sundaram S., Perceiving the temperature coefficients of carbon-based perovskite solar cells. *Sustainable Energy & Fuels*, 2020; 4: 6283–6298. DOI 10.1039/D0SE00782J.
 46. Anderson K.; Kemnitz J.; Boyd M., Evaluating cell temperature models and the effect of wind speed in PV system capacity testing. *NREL Conference Paper NREL/CP-5K00-78950*, 2021.
 47. Stein J. S.; Hansen C. W.; Reno M. J., The variability index: A metric for quantifying irradiance and PV output variability. *Conference Paper*, Sandia National Laboratories, 2012. OSTI ID: 1068417.
 48. PV Performance Modeling Collaborative (PVMC), Wavelet Variability Model (WVM) — tool page and examples. *Online documentation*.
 49. PV Performance Modeling Collaborative (PVMC), Datasets (module operating temperature; solar variability zones).
 50. Martínez-Lopez V. A.; Žindžiūtė U.; Ziar H.; Zeman M.; Isabella O., Study on the effect of irradiance variability on the efficiency of the perturb-and-observe maximum power point tracking algorithm. *Energies*, 2022; 15(20): 7562. DOI 10.3390/en15207562.
 51. Emara-Shabaik H. E.; El-Demerdash S. M.; Naguib A. M., Noise-robust parameter estimation of linear systems. *Journal of Vibration and Control*, 2000; 6(5): 727–740. DOI 10.1177/10775463000600505.
 52. Musikpan W.; Khampa W.; Passatorntaschakorn W.; Usulor C. E.; Ngamjarrojana A.; Gardchareon A.; Ruankham P.; Wongratanaphisan D., Humidity effects on $CS_{0.17}FA_{0.83}Pb(I_{0.83}Br_{0.17})_3$ thin films for carbon-based perovskite solar cells. *Chiang Mai Journal of Science*, 2024; 51(4): e2024058. DOI 10.12982/CMJS.2024.058.
 53. Khampa W.; Passatorntaschakorn W.; Musikpan W.; Gardchareon A.; Ruankham P.;

Wongratanaphisan D., Optimal TiO₂ nanoparticles
electron transporting layer for highly efficient
ambient-atmosphere fabricated carbon-based

perovskite solar cells. *Chiang Mai Journal of
Science*, 2023; 50(6): e2023072. DOI
10.12982/CMJS.2023.072.