Energy & Buildings 369 (2026) 117997

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Preventing parasitic learning in online identification of building RC thermal models: A stochastic-volatility Rao–Blackwellized particle filter

Mohamed Ramadane [a,b,*], Doris Bohnet [a]

[a] Hochschule Konstanz University of Applied Sciences, Department of Computer Sciences, Alfred-Wachtel-Str. 8, Konstanz, 78462, Germany
[b] Karlsruhe Institute of Technology, Institute AIFB, Kaiserstr. 89, Karlsruhe, 76133, Germany

A R T I C L E I N F O

Keywords:
Rao–Blackwellized particle filter
Online parameter estimation
Thermal model of buildings
Calibrated prediction
Probabilistic filtering

A B S T R A C T

Accurate online identification of building thermal models is central to energy-efficient operation; however, intermittent unmeasured heat gains can bias online estimators by being absorbed into resistance–capacitance (RC) parameters ("parasitic learning"), degrading both dynamics and uncertainty calibration. We propose a Rao–Blackwellized Particle Filter (RBPF) that introduces a stochastic-volatility state to scale the process covariance and suppress parameter adaptation during disturbance-dominated periods. We validate the approach on a semi-synthetic testbed with known ground truth and on the Ecolar Home real-building testbed. In the semi-synthetic testbed, the RBPF reduces transfer-function distortion relative to Extended and Unscented Kalman filter baselines. Its 95% prediction intervals stay close to nominal coverage in low-disturbance periods and remain far better calibrated than these baselines when the model is structurally mismatched. In two real-building measurement periods, the one-step prediction intervals of the RBPF stay close to their nominal coverage. During low-disturbance night-time windows, its open-loop simulations are more accurate than those of an offline least-squares baseline. These results support the proposed use of stochastic volatility as an inverse model-trust indicator for online RC parameter learning, not as a method for identifying individual heat-gain sources.

1. Introduction

The decarbonization of buildings is a central element of climate and energy policy. In the European Union, buildings account for roughly 40% of primary energy use [1]. Within households, space heating represented about 63% and water heating about 15% of final energy consumption in 2023 [2].

Model Predictive Control (MPC) offers a way to reduce heating consumption while maintaining thermal comfort. MPC uses an explicit model of the building's thermal dynamics to optimize future control actions and can outperform reactive strategies when the model is sufficiently accurate [3]. Drgoňa et al. [4] report energy savings between 15–50% compared with conventional rule-based building control, and Freund et al. [5] report a 30% reduction in heating energy in a large low-energy office building (and up to 75% during transition periods) with unchanged comfort conditions. However, these benefits depend critically on model quality. Prívara et al. [6] emphasize that building modelling is a crucial and often challenging part of the MPC control chain, and Maasoumy et al. [7] show that beyond a threshold of model uncertainty, simple rule-based strategies can outperform MPC. Despite the large body of scientific work on MPC, its use in real-life building energy management systems is still rare, seemingly due to the difficulty of developing and updating reliable models [8].

One approach to addressing the modelling challenge is to use purely data-driven models, but their performance relies on the availability of a large amount of high-quality training data and requires lengthy supervised training [9]. Even within the narrower task of building-level energy prediction, comparisons between simple statistical baselines and machine-learning models report mixed results depending on the prediction horizon and feature set [10]. Recent approaches claim to reduce the necessary training data and training time by transfer learning from foundational models or prior models from different buildings or simulations, as discussed in [11] and [12].

Another approach is the use of grey-box resistance-capacitance (RC) models, which are widely used because they balance physical interpretability with data-driven flexibility. These models represent buildings as low-order thermal networks, preserving an interpretable input–output structure while remaining simple enough to be calibrated from typical building measurements [13]. Fitting these models has been extensively studied [14]. While offline calibration remains the default in

* Corresponding author.
E-mail address: Mohamed.Ramadane@htwg-konstanz.de (M. Ramadane).

https://doi.org/10.1016/j.enbuild.2026.117997
Received 30 December 2025; Received in revised form 14 July 2026; Accepted 22 July 2026
Available online 23 July 2026

practice, online parameter estimation is increasingly explored to exploit streaming sensor data and to link parameter identifiability directly to excitation patterns during operation [15].

Real buildings, however, are subject to unmeasured, time-varying heat inputs: occupancy-driven internal gains, intermittent solar gains, window-opening and ventilation events, and appliance use. Deployment studies underline the practical importance of occupancy, reporting measurable HVAC performance differences when occupancy information is used in control [16]. We refer to these intermittent inputs as *unmeasured heat gains*: hidden heat inputs that are not part of the measured input vector. They act on minute-to-hour time scales and are poorly represented by stationary noise of constant variance. When an estimator absorbs them into a fixed, zero-mean Gaussian process-noise model, it cannot separate transient disturbances from structural model mismatch. It then attributes the unexplained temperature deviations partly to the RC parameters, driving them away from physically plausible values [17]. We refer to this disturbance-induced parameter drift as *parasitic learning*: the estimator absorbs unmeasured heat gains into the thermal parameters and is left with a biased model once the disturbance fades. For model-based control this is a damaging failure mode, because it corrupts both the dynamics used for prediction and the uncertainty used for robust constraint handling.

Despite a substantial body of work on online RC identification (Section 2), two questions remain open. First, the extent to which unmeasured heat gains in the identification data distort the resulting model has not been systematically quantified. Operating datasets typically contain segments in which the building's response is not well described by any low-order RC model, and the degree to which such segments distort the identified dynamics remains unclear. This distortion is not revealed by one-step prediction accuracy alone: because an online estimator continually corrects its state with incoming measurements, its one-step fit can stay accurate even when the identified model has drifted from the building's underlying dynamics, the property relevant to model-based control [7]. Second, the EKF, UKF, and particle filters all produce a predictive distribution rather than a single point estimate, yet building studies that use them (e.g. [18–21]) report predominantly point-error metrics and seldom assess interval calibration. Whether these intervals remain reliable during disturbance-dominated episodes is therefore not established. Rouchier et al. [22] show that neglecting process noise leads to overconfident estimates, but coverage and interval scores across low-disturbance and disturbance-dominated regimes have not been systematically characterised.

We address both questions with a Rao–Blackwellized Particle Filter (RBPF), which marginalises the linear–Gaussian state analytically with a per-particle Kalman filter (Section 3.4.2) and is augmented here with a stochastic-volatility state. The volatility acts as a scalar inverse model-trust indicator that scales the process covariance over time: larger volatility means lower trust in the deterministic RC model. When the measured dynamics become inconsistent with that model and its measured inputs, the filter raises the volatility, widens its prediction intervals, and reduces parameter rejuvenation (the controlled jitter that maintains particle diversity; Section 3.4.4). This protects the RC parameters from absorbing unmeasured heat gains, while the stochastic process disturbance can also absorb residual structural model mismatch. We validate the estimator on a semi-synthetic testbed with known ground-truth parameters and unmeasured heat-gain profiles [23] and on Ecolar Home, benchmarking it against EKF, UKF, and offline least-squares baselines. The evaluation covers preservation of the identifiable input–output dynamics (through transfer-function distance), suppression of disturbance-driven parameter updating (through a parasitic-learning index, Section 3.5.2), one-step prediction accuracy, probabilistic calibration through coverage and interval scores, and identification of disturbance-dominated regimes through the inferred volatility state.

2. Related work

The choice of RC topology is itself a modelling decision. Increasing the number of thermal states can improve fit, but it also introduces additional hidden states and parameters that are unstable from limited measurements. For MPC-oriented grey-box modelling, the objective is therefore not necessarily to choose the most detailed thermal network, but to select a low-order structure that captures the dominant input–output dynamics, yields weakly structured residuals, and remains physically interpretable. This trade-off has motivated both classical grey-box model-selection procedures for building heat dynamics [13], systematic studies of RC model quality and identifiability [24], and more recent automated stochastic identification frameworks for large sets of RC model candidates [25]. In this work, RC topology selection is not the main contribution; nevertheless, the model structures used in the experiments are chosen to be dynamically adequate for online identification.

In the context of online learning, Kalman filter variants are widely used for state-parameter estimation. Fux et al. [18] apply an Extended Kalman Filter (EKF) to a passive house, estimating both states and selected RC parameters (including a disturbance heat-gain term), and show that the EKF adapts effective heat-transfer coefficients to changing conditions while maintaining accurate indoor temperature predictions. Radecki and Hencey [26], Martincevic et al. [27], Maasoumy et al. [7], and Baldi et al. [28] develop EKF/Unscented Kalman Filter (UKF) schemes for online identification and report satisfactory performance. Li et al. [19] use a UKF to jointly estimate states and parameters in a residential building and demonstrate that online updates can reduce model mismatch compared to static models. Zamani et al. [29] integrate a UKF with nonlinear least squares to jointly estimate RC parameters and unmeasured inputs, such as solar and heating gains.

Particle-based methods are less commonly reported in building applications than EKF/UKF approaches, but they have also been studied for online identification. Cheng et al. [30] apply a particle filter to identify parameters of a first-order space thermal model, and later propose a window-varying particle filter that adapts the effective prediction window to better track time-varying parameters [20]. Ham and Karava [21] apply a particle filter to a model in a multifamily residential building, updating RC parameters and state temperatures online so that the model can evaluate heating and cooling energy under alternative thermostat setpoint behaviours. Rouchier et al. [15] use Sequential Monte Carlo (SMC) for online parameter estimation of a lumped RC model and explore how real-time excitation affects identifiability.

Prior work accommodates unmeasured heat gains in two main ways: by augmenting the model with an explicit disturbance state or input, or, more implicitly, by inflating a fixed process-noise covariance. Coffman [31] reconstructs occupant-induced loads using a piecewise-constant noise model, and Fux et al. [18] include a generic disturbance heat-gain term alongside the RC parameters. These formulations are effective when the disturbance matches the chosen structure, for example when it is stationary or well approximated as piecewise constant. They become insufficient when the disturbance magnitude varies rapidly and irregularly.

In summary, EKF/UKF and SMC methods fit grey-box RC models online, but largely under the assumption of stationary disturbances captured by fixed noise tuning or simple disturbance models. How intermittent, non-stationary unmeasured heat gains affect the identified dynamics and the calibration of the prediction intervals has not been systematically studied. This is the gap addressed in the present work.

3. Background: model and estimation methods

This section introduces the RC state-space model, the Bayesian filtering formulation, the baseline Kalman filters used for comparison, and the proposed Rao–Blackwellized particle filter (RBPF) with stochastic volatility. We conclude with the performance metrics used in the experiments.

3.1. RC thermal state–space model

We represent the thermal behaviour of a building zone by a lumped resistance-capacitance (RC) network. Nodes correspond to isothermal masses (e.g. indoor air, internal mass, and building-envelope layers) and edges to conductive couplings. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the undirected graph describing the network, with vertex set $\mathcal{V}$ and edge set $\mathcal{E}$. The vertices are partitioned into *state nodes* $\mathcal{V}_x$ and *boundary nodes* $\mathcal{V}_b$:

- $\mathcal{V}_x$ (e.g. indoor air, internal mass, and wall nodes) have finite thermal capacitance C_i and temperature $T_i(t)$, $i \in \mathcal{V}_x$;
- $\mathcal{V}_b$ (e.g. outdoor air, ground, and adjacent zones) are treated as prescribed boundary temperatures $T_i(t)$, $i \in \mathcal{V}_b$.

Let $n_x = |\mathcal{V}_x|$ and $m_b = |\mathcal{V}_b|$ denote the number of state and boundary nodes, respectively.

Each undirected edge $\{i, j\} \in \mathcal{E}$ is assigned a symmetric thermal conductance $G_{ij} = G_{ji} > 0$. The heat flow from node j to node i is

$$q_{ji}(t) = G_{ij}\big(T_j(t) - T_i(t)\big). \tag{1}$$

Known heat-flow inputs like heating are collected in $\mathbf{u}^{(q)}(t) \in \mathbb{R}^{m_q}$. For a state node $i \in \mathcal{V}_x$, the nodal energy balance is:

$$C_i \dot{T}_i(t) = \sum_{j \in \mathcal{N}(i)} G_{ij}\big(T_j(t) - T_i(t)\big) + \sum_{r=1}^{m_q} B_{ir}^{(q)} u_r^{(q)}(t), \tag{2}$$

where $\mathcal{N}(i) \subset \mathcal{V}$ is the neighbour set of i and $\mathbf{B}^{(q)} = \big(B_{ir}^{(q)}\big) \in \mathbb{R}^{n_x \times m_q}$ distributes the heat-flow inputs to the state nodes. Define the diagonal capacitance matrix $\mathbf{C} = \mathrm{diag}(C_i) \in \mathbb{R}^{n_x \times n_x}$, and assemble the conductance contributions in the standard way to obtain internal and boundary conductance matrices $\mathbf{G}_{\mathrm{int}} \in \mathbb{R}^{n_x \times n_x}$ and $\mathbf{G}_{\mathrm{ext}} \in \mathbb{R}^{n_x \times m_b}$ [13,14]. Collect the state-node temperatures in $\mathbf{x}(t) = [T_i(t)]_{i \in \mathcal{V}_x} \in \mathbb{R}^{n_x}$. We augment the vector of heat inputs $\mathbf{u}^{(q)}(t)$ by the boundary temperatures and define

$$\mathbf{u}(t) = \begin{bmatrix} [T_i(t)]_{i \in \mathcal{V}_b} \\ \mathbf{u}^{(q)}(t) \end{bmatrix} \in \mathbb{R}^m, \qquad m = m_b + m_q,$$

and set $\mathbf{B}_u = \begin{bmatrix} \mathbf{G}_{\mathrm{ext}} & \mathbf{B}^{(q)} \end{bmatrix} \in \mathbb{R}^{n_x \times m}$. Stacking (2) over $i \in \mathcal{V}_x$ yields the continuous-time linear time-invariant (LTI) model

$$\mathbf{C}\dot{\mathbf{x}}(t) = -\mathbf{G}_{\mathrm{int}}\mathbf{x}(t) + \mathbf{B}_u \mathbf{u}(t). \tag{3}$$

In standard state-space form,

$$\dot{\mathbf{x}}(t) = \mathbf{A}(\boldsymbol{\theta})\,\mathbf{x}(t) + \mathbf{B}(\boldsymbol{\theta})\,\mathbf{u}(t), \tag{4}$$

with $\mathbf{A}(\boldsymbol{\theta}) = -\mathbf{C}^{-1}\mathbf{G}_{\mathrm{int}}$ and $\mathbf{B}(\boldsymbol{\theta}) = \mathbf{C}^{-1}\mathbf{B}_u$. The parameter vector $\boldsymbol{\theta}$ collects the unknown capacitances and conductances (entries of $\mathbf{C}$, $\mathbf{G}_{\mathrm{int}}$, $\mathbf{G}_{\mathrm{ext}}$).

We focus on low-order RC networks (typically order ≤ 5 for single-zone prediction), which provide a practical compromise between dynamic adequacy and physical interpretability for MPC-relevant prediction horizons [13,14].

For estimation and control, (4) is discretised with sampling period Δt under a zero-order-hold (ZOH) assumption on the inputs:

$$\mathbf{A}_d(\boldsymbol{\theta}) = \exp\!\big(\mathbf{A}(\boldsymbol{\theta})\Delta t\big), \tag{5}$$

$$\mathbf{B}_d(\boldsymbol{\theta}) = \int_0^{\Delta t} \exp\!\big(\mathbf{A}(\boldsymbol{\theta})\tau\big)\mathbf{B}(\boldsymbol{\theta})\,\mathrm{d}\tau. \tag{6}$$

For completeness, we also note the explicit Euler discretisation of (4), which approximates the matrix exponential and input integral for small sampling periods Δt as

$$\mathbf{A}_d(\boldsymbol{\theta}) \approx \mathbf{I} + \mathbf{A}(\boldsymbol{\theta})\Delta t, \tag{7}$$

$$\mathbf{B}_d(\boldsymbol{\theta}) \approx \mathbf{B}(\boldsymbol{\theta})\Delta t. \tag{8}$$

This first-order scheme is widely used in practice when Δt is sufficiently small. Unless stated otherwise, we assume ZOH discretisation in the algorithmic description below. Numerical implementation choices are stated in the experimental section.

For filtering, we represent the discrete-time dynamics as the stochastic state-space model

$$\mathbf{x}_{k+1} = \mathbf{A}_d(\boldsymbol{\theta})\mathbf{x}_k + \mathbf{B}_d(\boldsymbol{\theta})\mathbf{u}_k + \mathbf{w}_k, \tag{9}$$

where $\mathbf{x}_k \approx \mathbf{x}(k\Delta t)$, $\mathbf{u}_k \approx \mathbf{u}(k\Delta t)$, and $\mathbf{w}_k \in \mathbb{R}^{n_x}$ captures unmeasured heat gains and residual model mismatch as an additive process disturbance.

Measurements are modelled as

$$\mathbf{y}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \tag{10}$$

where $\mathbf{H}$ selects the measured temperatures (typically the indoor-air temperature) and $\mathbf{v}_k$ is measurement noise. In the standard stochastic RC formulation, these disturbances are modelled as Gaussian,

$$\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \qquad \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \tag{11}$$

with constant covariances $(\mathbf{Q}, \mathbf{R})$. In Section 3.4, we relax this assumption by allowing the effective process covariance to vary in time.

3.2. Bayesian filtering framework

Given the stochastic state-space model (9)-(10), Bayesian filtering aims at computing the posterior distribution of the hidden state given noisy measurements $\mathbf{y}_{1:k} = \{\mathbf{y}_1, \dots, \mathbf{y}_k\}$ and known inputs $\mathbf{u}_{0:k-1} = \{\mathbf{u}_0, \dots, \mathbf{u}_{k-1}\}$. The one-step prediction and update take the standard form

$$p(\mathbf{x}_k \mid \mathbf{y}_{1:k-1}) = \int p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_{k-1})\, p(\mathbf{x}_{k-1} \mid \mathbf{y}_{1:k-1})\, \mathrm{d}\mathbf{x}_{k-1}, \tag{12}$$

$$p(\mathbf{x}_k \mid \mathbf{y}_{1:k}) \propto p(\mathbf{y}_k \mid \mathbf{x}_k)\, p(\mathbf{x}_k \mid \mathbf{y}_{1:k-1}). \tag{13}$$

In our setting, both the dynamic state $\mathbf{x}_k$ and the RC parameters $\boldsymbol{\theta}$ are unknown. We therefore target the joint posterior $p(\mathbf{x}_k, \boldsymbol{\theta} \mid \mathbf{y}_{1:k})$. A common practical approach for online estimation of static parameters is *state augmentation*, in which the parameters are embedded into an augmented state and assigned an artificial (typically Gaussian) random-walk evolution to enable sequential learning.

In the special case where $\boldsymbol{\theta}$ is known and (9)-(10) are linear Gaussian with fixed $(\mathbf{Q}, \mathbf{R})$, the Kalman filter provides the exact Gaussian solution for $p(\mathbf{x}_k \mid \mathbf{y}_{1:k})$.

3.3. Baseline estimators

This section summarises the baseline estimators used for comparison: (i) augmented-state EKF and UKF estimators for joint state and parameter estimation, and (ii) an offline nonlinear least-squares fit of the RC parameters.

3.3.1. Augmented-state EKF for joint state and parameter estimation

A standard online approach augments the state vector with the free parameters and applies an extended Kalman filter (EKF) [7,26]. Since RC parameters (capacitances and conductances) are strictly positive, we represent them in log-space: $\boldsymbol{\phi}_k = \log \boldsymbol{\theta}_k$, applied elementwise to the free components of $\boldsymbol{\theta}$. Define the augmented state

$$\mathbf{s}_k = \begin{bmatrix} \mathbf{x}_k \\ \boldsymbol{\phi}_k \end{bmatrix}. \tag{14}$$

We postulate a Gaussian random walk in parameter space.

$$\boldsymbol{\phi}_{k+1} = \boldsymbol{\phi}_k + \boldsymbol{\zeta}_k, \qquad \boldsymbol{\zeta}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_\phi), \tag{15}$$

and obtain the augmented dynamics

$$\mathbf{s}_{k+1} = \begin{bmatrix} \mathbf{A}_d(\exp(\boldsymbol{\phi}_k))\mathbf{x}_k + \mathbf{B}_d(\exp(\boldsymbol{\phi}_k))\mathbf{u}_k + \mathbf{w}_k \\ \boldsymbol{\phi}_k + \boldsymbol{\zeta}_k \end{bmatrix}, \qquad \mathbf{y}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \tag{16}$$

with $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_x)$ and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$. The corresponding augmented process covariance is block-diagonal,

$$\mathbf{Q}_s = \begin{bmatrix} \mathbf{Q}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_\phi \end{bmatrix},$$

and is held constant in time. The EKF linearises the augmented dynamics and measurement model around the current estimate and applies the Kalman recursion to the linearised system [32,33]. In this formulation, the process covariance is fixed and persistent mismatch is typically absorbed into the parameters.

3.3.2. Augmented-state UKF for joint state and parameter estimation

We also consider an augmented-state UKF for joint state and parameter estimation, using the augmented state of (14) and the random walk of (15) [34,35]. The UKF replaces linearisation with a sigma-point approximation of the prediction and update steps [34]. The hyperparameters $(\mathbf{Q}_x, \mathbf{Q}_\phi, \mathbf{R})$ are fixed.

3.3.3. Offline nonlinear least-squares (LS) baseline

As an offline baseline, we fit a single parameter vector θ_{LS} by minimising the open-loop squared output error under the deterministic RC model (i.e. with $\mathbf{w}_k \equiv \mathbf{0}$). Given an initial condition $\hat{\mathbf{x}}_0$ and an input sequence $\{\mathbf{u}_k\}_{k=0}^{N-1}$, the simulated trajectory is

$$\hat{\mathbf{x}}_{k+1}(\theta) = \mathbf{A}_d(\theta)\,\hat{\mathbf{x}}_k(\theta) + \mathbf{B}_d(\theta)\,\mathbf{u}_k, \qquad \hat{y}_k(\theta) = \mathbf{H}\hat{\mathbf{x}}_k(\theta). \tag{17}$$

We estimate

$$\theta_{\mathrm{LS}} \in \arg\min_{\theta\in\Theta} \frac{1}{N-1}\sum_{k=0}^{N-1}\Big(y_{k+1} - \hat{y}_{k+1}(\theta)\Big)^2, \tag{18}$$

where Θ encodes physically plausible box constraints on the capacitances and conductances. In our implementation, $\hat{\mathbf{x}}_0$ is set by a simple physical heuristic based on the first measurement and available boundary temperatures, and (18) is solved with the bound-constrained L-BFGS-B quasi-Newton method [36].

3.4. Proposed RBPF with stochastic volatility

We now introduce a Rao–Blackwellized particle filter (RBPF) for joint state estimation and online identification of RC parameters under time-varying process uncertainty. The key modelling choice is to represent unmeasured heat gains by a scalar stochastic-volatility state α_k that scales a baseline process covariance, and to use the inferred volatility to modulate the amount of parameter rejuvenation in a Liu–West kernel [37]. This yields a filter that adapts its predictive uncertainty online while limiting parameter updates during high-volatility periods.

3.4.1. Unmeasured heat gains as time-varying process uncertainty

Conceptually, the true dynamics can be written as

$$\mathbf{x}_{k+1} = \mathbf{A}_d(\theta)\mathbf{x}_k + \mathbf{B}_d(\theta)\mathbf{u}_k + \mathbf{F}\mathbf{f}_k, \tag{19}$$

where $\mathbf{f}_k$ denotes the unmeasured heat gains and $\mathbf{F}$ distributes them to the thermal states. We do not attempt to estimate $\mathbf{f}_k$ explicitly. Instead, the additive process disturbance $\mathbf{w}_k$ in Eq. (9) represents the net stochastic effect of $\mathbf{F}\mathbf{f}_k$ and can additionally absorb residual structural model mismatch. We therefore allow its covariance to vary over time:

$$\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k), \qquad \mathbf{Q}_k = \alpha_k^2 \mathbf{Q}_0, \tag{20}$$

where $\mathbf{Q}_0$ is a baseline covariance and $\alpha_k > 0$ is the scalar *stochastic-volatility state*. Larger α_k indicates greater mismatch with the deterministic RC model and therefore lower model trust. Conditioned on α_k, the state equation remains linear–Gaussian.

The volatility affects filtering and uncertainty in the following way: larger α_k increases the predicted covariance $\mathbf{P}_{k|k-1}$, which increases the Kalman gain and shifts weight from the RC model toward the measurement. As a result, one-step prediction intervals widen automatically during disturbance-dominated periods and tighten again when the system returns to behaviour consistent with the RC model.

We model α_k as a log-random walk (throughout, log denotes the natural logarithm),

$$\log\alpha_k = \log\alpha_{k-1} + \eta_{k-1}, \qquad \eta_{k-1} \sim \mathcal{N}(0, \sigma_\eta^2), \tag{21}$$

with α_k clipped to $[\alpha_{\min}, \alpha_{\max}]$ to avoid degenerate or numerically unstable values.

3.4.2. Rao–Blackwellized structure

Let $\mathbf{s}_k = [\mathbf{x}_k^\top, \theta^\top, \alpha_k]^\top$. Conditioned on a particular parameter vector θ and a stochastic-volatility state α_k, the model is linear–Gaussian in $\mathbf{x}_k$ and can be propagated analytically by a Kalman filter. The joint posterior factorises as

$$p(\mathbf{x}_k, \theta, \alpha_k \mid \mathbf{y}_{1:k}) = p(\mathbf{x}_k \mid \theta, \alpha_k, \mathbf{y}_{1:k})\; p(\theta, \alpha_k \mid \mathbf{y}_{1:k}), \tag{22}$$

where the first term is Gaussian and the second is approximated with particles.

We maintain N particles

$$\Big\{\theta^{(i)}, \alpha_k^{(i)}, \hat{\mathbf{x}}_{k|k}^{(i)}, \mathbf{P}_{k|k}^{(i)}, w_k^{(i)}\Big\}_{i=1}^N,$$

where $(\hat{\mathbf{x}}_{k|k}^{(i)}, \mathbf{P}_{k|k}^{(i)})$ is the conditional Kalman posterior and $w_k^{(i)}$ are normalised importance weights. Intuitively, particles whose $(\theta^{(i)}, \alpha_k^{(i)})$ better explain the current observation obtain higher likelihood and thus higher weight; this shifts the posterior over α_k toward larger values during disturbance episodes.

3.4.3. Parameter representation and priors

RC parameters are strictly positive. For the free components of θ we use the log-parameterisation $\boldsymbol{\phi} = \log\theta$ (elementwise) and independent log-uniform priors over physically plausible bounds:

$$\phi_{0,j}^{(i)} \sim \mathcal{U}\big(\log\theta_{j,\min},\, \log\theta_{j,\max}\big), \tag{23}$$

followed by $\theta_0^{(i)} = \exp(\boldsymbol{\phi}_0^{(i)})$. In the RBPF, the parameters are treated as static; time indices on $\theta_k^{(i)}$ denote the particle values after rejuvenation (Section 3.4.4), not a physical parameter evolution model.

3.4.4. Volatility-based learning gate and Liu–West rejuvenation

A standard particle filter with static parameters often suffers from particle impoverishment (the particle set collapsing onto a few distinct parameter values). We maintain parameter diversity using the Liu–West kernel [37], which approximately preserves the first two moments of the parameter posterior while introducing controlled jitter. In our method, we *gate* the amount of rejuvenation using the estimated volatility α_k, so that parameter adaptation is reduced during disturbance-dominated periods.

Let $\boldsymbol{\phi}_k^{(i)} = \log\theta_k^{(i)}$. We define the weighted mean and covariance of the log-parameters at time step k:

$$\bar{\boldsymbol{\phi}}_k = \sum_{i=1}^N w_k^{(i)}\boldsymbol{\phi}_k^{(i)}, \tag{24}$$

$$\mathbf{V}_k = \sum_{i=1}^N w_k^{(i)}\big(\boldsymbol{\phi}_k^{(i)} - \bar{\boldsymbol{\phi}}_k\big)\big(\boldsymbol{\phi}_k^{(i)} - \bar{\boldsymbol{\phi}}_k\big)^\top. \tag{25}$$

Following Liu and West [37], we introduce a discount factor $\delta_0 \in (0,1]$ (set to $\delta_0 = 0.95$ in our experiments) and define the baseline kernel coefficients

$$a_0 = \frac{3\delta_0 - 1}{2\delta_0}, \qquad h_0^2 = 1 - a_0^2.$$

The baseline rejuvenation step samples each particle's new log-parameter from a Gaussian centred at a shrinkage mixture of its current value and the population mean, with covariance scaled by h_0^2:

$$\boldsymbol{\phi}_{k+1}^{(i)} \sim \mathcal{N}\Big(a_0\boldsymbol{\phi}_k^{(i)} + (1-a_0)\bar{\boldsymbol{\phi}}_k,\; h_0^2\mathbf{V}_k\Big). \tag{26}$$

To suppress this adaptation when the volatility is high, we introduce a scalar *volatility-based learning gate* $g_k \in [g_{\min}, 1]$ that modulates the Liu–West coefficients:

$$a_k = 1 - g_k(1-a_0), \qquad h_k^2 = g_k h_0^2, \tag{27}$$

And thus, our parameter rejuvenation updates become

$$\boldsymbol{\phi}_{k+1}^{(i)} \sim \mathcal{N}\Big(a_k\boldsymbol{\phi}_k^{(i)} + (1-a_k)\bar{\boldsymbol{\phi}}_k,\; h_k^2\mathbf{V}_k\Big), \qquad \theta_{k+1}^{(i)} = \exp\big(\boldsymbol{\phi}_{k+1}^{(i)}\big). \tag{28}$$

The coefficient a_k controls where the new sample is centred: it is a convex combination of the particle's own value $\boldsymbol{\phi}_k^{(i)}$ and the weighted mean $\bar{\boldsymbol{\phi}}_k$. When $g_k \to 1$, $a_k \to a_0$ and the centre is shrunk toward the mean by the standard Liu–West amount, with full jitter $h_0^2 \mathbf{V}_k$ - this is the normal learning regime. When $g_k \to g_{\min}$ (with $g_{\min}$ small), $a_k \to 1$ and the centre collapses onto the particle's current value while the jitter covariance shrinks toward zero, effectively freezing the parameters. Intermediate values of g_k interpolate smoothly between the two.

It remains to specify how g_k is computed from the current volatility. To avoid relying on absolute thresholds, we map the volatility through a running empirical cumulative distribution function (CDF) derived from an online histogram, so that the gate is self-calibrating to the volatility range observed during operation.

At each step we compute the weighted median $\tilde{\alpha}_k$ of $\{\alpha_k^{(i)}\}_{i=1}^N$ under the current particle weights. We define n_{bins} bin edges $b_0 < b_1 < \cdots < b_{n_{\text{bins}}}$ uniformly covering the clipped range $[\log \alpha_{\min}, \log \alpha_{\max}]$, and initialise the histogram counts c_j (for $j = 1, \ldots, n_{\text{bins}}$) with a uniform pseudo-count $c_j \leftarrow c_{\text{pseudo}}$ for numerical stability at start-up. At each time step we identify the bin index $j_k = \text{bin}(\log \tilde{\alpha}_k)$, increment $c_{j_k} \leftarrow c_{j_k} + 1$, and compute the empirical CDF value

$$F_k = \frac{\sum_{j=1}^{j_k} c_j}{\sum_{j=1}^{n_{\text{bins}}} c_j} \in [0, 1].$$

Two thresholds $F_{\text{low}} < F_{\text{high}}$ separate low-volatility, transitional, and high-volatility regimes, and g_k is obtained as

$$g_k = \begin{cases} 1, & F_k \le F_{\text{low}}, \\ g_{\min} + (1 - g_{\min}) \dfrac{F_{\text{high}} - F_k}{F_{\text{high}} - F_{\text{low}}}, & F_{\text{low}} < F_k < F_{\text{high}}, \\ g_{\min}, & F_k \ge F_{\text{high}}. \end{cases}$$

A low quantile of the historical volatility ($F_k \le F_{\text{low}}$) yields $g_k = 1$, permitting normal learning; a high quantile ($F_k \ge F_{\text{high}}$) yields $g_k = g_{\min}$, suppressing adaptation. In between, g_k ramps linearly. This is the mechanism that protects the RC parameters from absorbing unmeasured heat gains.

3.4.5. RBPF recursion

Algorithm 1 summarises the proposed RBPF recursion. The detailed definitions of the volatility process, the volatility-based learning gate g_k, and the resulting gated Liu–West rejuvenation are given in Sections 3.4.1–3.4.4. Each particle carries an RC parameter vector, a volatility state, and a conditional Gaussian state posterior. The Kalman recursion is performed analytically inside each particle, while the particle weights approximate the posterior over parameters and volatility.

Algorithm 1 RBPF with stochastic volatility and a volatility-based learning gate.

Require: Measurements $\{\mathbf{y}_k\}_{k=1}^K$, inputs $\{\mathbf{u}_k\}_{k=0}^{K-1}$, number of particles N
Require: $\mathbf{Q}_0$, $\mathbf{R}$, volatility parameters, Liu–West parameters, resampling threshold $N_{\text{eff,min}}$
Ensure: Particle approximation of $p(\mathbf{x}_k, \boldsymbol{\theta}, \alpha_k \mid \mathbf{y}_{1:k})$

Particle initialisation
for $i = 1, \ldots, N$ **do**
  Draw $\boldsymbol{\phi}_0^{(i)}$ from (23); set $\boldsymbol{\theta}_0^{(i)} = \exp(\boldsymbol{\phi}_0^{(i)})$
  Draw $\alpha_0^{(i)}$ from the initial volatility prior and clip to $[\alpha_{\min}, \alpha_{\max}]$
  Initialise $(\hat{\mathbf{x}}_{0|0}^{(i)}, \mathbf{P}_{0|0}^{(i)})$ and set $w_0^{(i)} = 1/N$
end for
Initialise the running histogram for $\log \tilde{\alpha}_k$

Online filtering recursion
for $k = 1, \ldots, K$ **do**
  Volatility prediction: propagate the stochastic-volatility states $\alpha_k^{(i)}$ for all particles using (21), clip to $[\alpha_{\min}, \alpha_{\max}]$, and set $\mathbf{Q}_k^{(i)} = (\alpha_k^{(i)})^2 \mathbf{Q}_0$ as in (20)
  Kalman prediction: for each particle, compute
  $$(\hat{\mathbf{x}}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)}, \hat{\mathbf{y}}_{k|k-1}^{(i)}, \mathbf{S}_k^{(i)})$$
  using $\boldsymbol{\theta}_{k-1}^{(i)}$, $\alpha_k^{(i)}$, and $\mathbf{u}_{k-1}$
  One-step prediction interval: before assimilating $\mathbf{y}_k$, form
  $$p(\mathbf{y}_k \mid \mathbf{y}_{1:k-1}) \approx \sum_{i=1}^{N} w_{k-1}^{(i)} \mathcal{N}\left(\hat{\mathbf{y}}_{k|k-1}^{(i)}, \mathbf{S}_k^{(i)}\right)$$
  and compute the required mixture quantiles
  Weight update: update and normalise
  $$w_k^{(i)} \propto w_{k-1}^{(i)} \mathcal{N}\left(\mathbf{y}_k; \hat{\mathbf{y}}_{k|k-1}^{(i)}, \mathbf{S}_k^{(i)}\right)$$
  Kalman measurement update: update each conditional posterior
  $$(\hat{\mathbf{x}}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)}) \to (\hat{\mathbf{x}}_{k|k}^{(i)}, \mathbf{P}_{k|k}^{(i)})$$
  Learning gate: compute $\tilde{\alpha}_k$, update the histogram of $\log \tilde{\alpha}_k$, compute F_k, and obtain g_k using the gate mapping defined above
  Resampling: compute $N_{\text{eff}} = 1/\sum_i (w_k^{(i)})^2$; if $N_{\text{eff}} < N_{\text{eff,min}}$, resample $(\boldsymbol{\theta}, \alpha, \hat{\mathbf{x}}_{k|k}, \mathbf{P}_{k|k})$ and reset $w_k^{(i)} = 1/N$
  Parameter rejuvenation: apply the gated Liu–West rejuvenation (27)-(28) with gate g_k to obtain $\theta_k^{(i)}$ for the next filtering step
end for

3.5. Performance metrics

We report metrics for (i) preservation of input–output dynamics, (ii) disturbance-driven parameter updating, (iii) one-step prediction accuracy, and (iv) uncertainty calibration. We do not score parameter recovery directly: under single-output observation the RC parameters are not uniquely identifiable, so physical plausibility is enforced *by construction*—through the log-parameterisation, which keeps capacitances and conductances positive, and the bounded priors and box constraints that confine them to physically plausible ranges (Sections 3.3.3 and 3.4.3)—rather than reported as an outcome metric. What we quantify instead is the identifiable, control-relevant behaviour: the transfer-function distance below measures preservation of the input–output dynamics, and the parasitic-learning index measures disturbance-driven parameter updating.

3.5.1. Transfer-function distance

The discrete RC model (9)-(10), with input $\mathbf{u}_k \in \mathbb{R}^m$ and output $\mathbf{y}_k \in \mathbb{R}^p$, is linear and time-invariant, so its forced input–output behaviour is characterised completely by its transfer function

$$\mathbf{G}(z; \boldsymbol{\theta}) = \mathbf{H}\left(z\mathbf{I} - \mathbf{A}_d(\boldsymbol{\theta})\right)^{-1} \mathbf{B}_d(\boldsymbol{\theta}).$$

Two parameter vectors are input–output equivalent if and only if $\mathbf{G}(\cdot; \boldsymbol{\theta}_1) = \mathbf{G}(\cdot; \boldsymbol{\theta}_2)$, which makes the transfer function the appropriate, control-relevant object to compare, for two reasons. First, when only the indoor air temperature is observed and the excitation available during operation is limited, the parameters are not uniquely identifiable: the map $\boldsymbol{\theta} \mapsto \mathbf{G}(\cdot; \boldsymbol{\theta})$ is not injective, so distinct parameter vectors can reproduce the same—or, in practice, an indistinguishable—input–output response [15,24]. A parameter-space distance to a reference $\boldsymbol{\theta}^\star$ is therefore not well posed. Second, the transfer-function distance is invariant within an equivalence class: an estimate that converges to a different parameter vector realizing the correct dynamics is credited correctly, whereas a parameter-space error would penalize it spuriously.

We compare the transfer functions channel by channel. Let $u_k^{(\ell)}$ denote the ℓth input channel and $y_k^{(r)}$ the rth output component; the scalar entry $G_{r\ell}(z;\theta)$ of $\mathbf{G}(z;\theta)$ relates them through the frequency-domain relation

$$Y^{(r)}(e^{i\Omega}) = G_{r\ell}(e^{i\Omega};\theta)\, U^{(\ell)}(e^{i\Omega}), \qquad \Omega \in [0,\pi],$$

where $U^{(\ell)}$ and $Y^{(r)}$ are the discrete-time Fourier transforms of the respective channels.

Given a reference parameter vector $\theta^\star$ (the ground truth), define the frequency-response deviation

$$\Delta G_{r\ell}(e^{i\omega\Delta t};\theta) = G_{r\ell}(e^{i\omega\Delta t};\theta) - G_{r\ell}(e^{i\omega\Delta t};\theta^\star), \tag{29}$$

where ω denotes the physical angular frequency in rad/s and $\Omega = \omega\Delta t \in [0,\pi]$ is the corresponding discrete-time frequency.

We evaluate frequency responses on a grid $\{\omega_j\}_{j=1}^{N_\omega}$ up to the Nyquist angular frequency $\omega_{\mathrm{Nyq}} = \pi/\Delta t$ (rad/s), i.e. $\omega_j \in [0,\omega_{\mathrm{Nyq}}]$. We use an H_2-like (RMS) deviation and an H_∞-like (worst-case) deviation:

$$E_{r\ell,2}(\theta) = \sqrt{\frac{1}{N_\omega}\sum_{j=1}^{N_\omega}\left|\Delta G_{r\ell}(e^{i\omega_j\Delta t};\theta)\right|^2}, \tag{30}$$

$$E_{r\ell,\infty}(\theta) = \max_{1\le j\le N_\omega}\left|\Delta G_{r\ell}(e^{i\omega_j\Delta t};\theta)\right|. \tag{31}$$

To obtain dimensionless measures, we normalise by the corresponding norms of the reference transfer function:

$$\tilde{E}_{r\ell,2}(\theta) = \frac{E_{r\ell,2}(\theta)}{\sqrt{\frac{1}{N_\omega}\sum_{j=1}^{N_\omega}\left|G_{r\ell}(e^{i\omega_j\Delta t};\theta^\star)\right|^2}},$$
$$\tilde{E}_{r\ell,\infty}(\theta) = \frac{E_{r\ell,\infty}(\theta)}{\max_{1\le j\le N_\omega}\left|G_{r\ell}(e^{i\omega_j\Delta t};\theta^\star)\right|}. \tag{32}$$

3.5.2. Parasitic-learning index

We quantify parasitic learning by measuring whether an estimator updates the RC parameters disproportionately during disturbance-dominated periods. Since RC parameters are estimated in log-space, let

$$\hat{\boldsymbol{\phi}}_k = \log \hat{\boldsymbol{\theta}}_k$$

denote the point estimate of the log-parameter vector at time step k. For the EKF and UKF, $\hat{\boldsymbol{\phi}}_k$ is the augmented-state mean. For the RBPF, $\hat{\boldsymbol{\phi}}_k$ is the weighted posterior mean over the parameter particles.

We define the per-step parameter-update magnitude as

$$d_k = \left\|\hat{\boldsymbol{\phi}}_k - \hat{\boldsymbol{\phi}}_{k-1}\right\|_2. \tag{33}$$

In the semi-synthetic testbed, the unmeasured-heat-gain signal is known, so time indices can be partitioned into low-disturbance periods $\mathcal{L}$ and disturbance-dominated periods $\mathcal{D}$. We define the parasitic-learning index as

$$\mathrm{PLI} = \frac{\mathrm{median}_{k\in\mathcal{D}}\, d_k}{\mathrm{median}_{k\in\mathcal{L}}\, d_k + \varepsilon}, \tag{34}$$

where $\varepsilon > 0$ is a small numerical constant. A value close to 1 means that the estimator does not update the parameters more strongly during disturbance-dominated periods than during low-disturbance periods. Values substantially larger than 1 indicate disturbance-amplified parameter adaptation, which is the sequential signature of parasitic learning. The offline LS baseline is static and therefore is not assigned a PLI.

3.5.3. One-step prediction accuracy

We assess one-step prediction accuracy via root-mean-square error (RMSE). In the semi-synthetic testbed, RMSE is calculated against the noise-free ground-truth temperature state. In the real-building experiment, where ground truth is unavailable, RMSE is calculated against the measured temperature. For a sequence of one-step predictions $\hat{y}_{k|k-1}$ and reference values y_k^{ref} (either ground truth or measurement), $k = 1,\dots,N$, the RMSE is

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}\left(\hat{y}_{k|k-1} - y_k^{\mathrm{ref}}\right)^2}. \tag{35}$$

3.5.4. Uncertainty calibration and Winkler score

Uncertainty calibration is assessed by empirical coverage of nominal prediction intervals and by the Winkler score [38], a proper scoring rule for interval forecasts that rewards narrow (sharp) intervals and adds a penalty when the observation falls outside the interval, proportional to the distance from the nearest bound. A prediction interval is calibrated if its empirical coverage matches its nominal level.

For nominal γ-level one-step prediction intervals $[\hat{y}^{\mathrm{L}}_{k|k-1}, \hat{y}^{\mathrm{U}}_{k|k-1}]$ (with $\gamma = 0.95$ in our experiments), the empirical coverage is

$$\hat{c}_\gamma = \frac{1}{N}\sum_{k=1}^{N}\mathbb{I}\left\{y_k \in [\hat{y}^{\mathrm{L}}_{k|k-1}, \hat{y}^{\mathrm{U}}_{k|k-1}]\right\}, \tag{36}$$

where $\mathbb{I}\{\cdot\}$ equals 1 if the condition is true and 0 otherwise. We also report the mean band width

$$\bar{w}_\gamma = \frac{1}{N}\sum_{k=1}^{N}\left(\hat{y}^{\mathrm{U}}_{k|k-1} - \hat{y}^{\mathrm{L}}_{k|k-1}\right). \tag{37}$$

The Winkler score W_γ for a single observation y_k and interval $[L_k, U_k]$ (with $L_k < U_k$) is defined as

$$W_\gamma(y_k, L_k, U_k) = \begin{cases} U_k - L_k, & \text{if } y_k \in [L_k, U_k], \\ (U_k - L_k) + \dfrac{2}{1-\gamma}(L_k - y_k), & \text{if } y_k < L_k, \\ (U_k - L_k) + \dfrac{2}{1-\gamma}(y_k - U_k), & \text{if } y_k > U_k. \end{cases} \tag{38}$$

The overall Winkler score is the average over all time indices,

$$\overline{W}_\gamma = \frac{1}{N}\sum_{k=1}^{N} W_\gamma\left(y_k, \hat{y}^{\mathrm{L}}_{k|k-1}, \hat{y}^{\mathrm{U}}_{k|k-1}\right). \tag{39}$$

Smaller $\overline{W}_\gamma$ indicates better calibrated and sharper intervals: intervals that are too wide are penalised through their width, and intervals that systematically miss the observations incur an additional penalty proportional to the distance between y_k and the nearest bound. The symbol γ is used for nominal interval coverage so that α_k remains reserved exclusively for the stochastic-volatility state.

4. Experimental setup

4.1. Overview and study design

We evaluate the proposed estimation framework in two settings. First, we use a semi-synthetic testbed based on the public dataset of [23] (building ID 67), for which the ground-truth physical parameters are known. In this setting, we assess preservation of input–output dynamics via transfer-function metrics, disturbance-driven parameter updating via the parasitic-learning index, and the accuracy of volatility-based regime detection. Second, we apply the method to Ecolar Home to assess probabilistic calibration and open-loop simulation performance. Table 1 summarises the configuration of both testbeds.

4.2. Semi-synthetic testbed

4.2.1. Data description

We use the RC model provided by [23] (ID:67). Weather data (outdoor temperature and global solar irradiance) are sourced from the German Weather Service (DWD) [39] for the period 10 April 2024 to 10 September 2024. The data are sampled to a grid of $\Delta t = 600$ s. The system is excited by a known heating sequence u_{heat} generated as a square-wave signal (amplitude 800 W) with random on/off durations to ensure persistent excitation, i.e. input variation rich enough to identify the parameters.

Table 1
Summary of experimental configurations for the semi-synthetic and real-building testbeds.

Feature	Semi-synthetic	Real-building testbed
Data source	DWD weather (2024) + simulation	HTWG Konstanz sensors (Ecolar) 03 May – 01 Jun 2025
Period	10 Apr – 10 Sep 2024	03 Oct – 20 No. 2025
Sampling (Δt)	600 s (10 min)	300 s (5 min)
Model structure	5-state RC (5 states, 10 parameters)	4-state RC (4 states, 10 parameters)
Inputs (u)	u_{heat}, T_{out}, I_{sol}	u_{heat}, T_{out}, T_{adj}
Measured output (y)	Indoor-air temperature (T_i)	Indoor-air temperature (T_{air})
Ground truth	Parameters $\theta^\star$, unmeasured heat gains $\mathbf{f}_k$	Unavailable
Key metrics	TF error (H_2/H_∞), PLI, calibration, RMSE	One-step calibration, RMSE, night-time open-loop RMSE

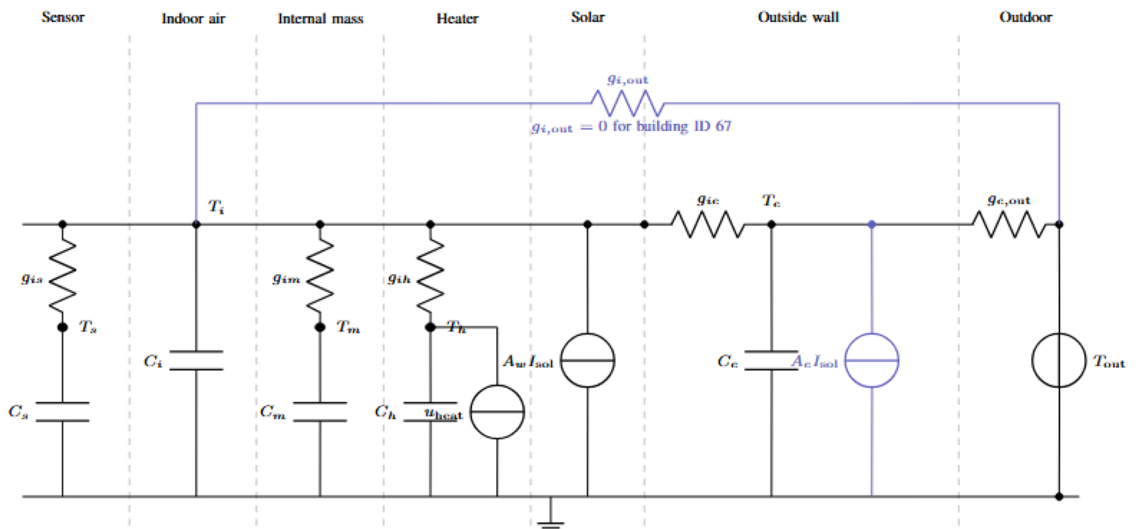

Fig. 1. Equivalent circuit representation of the fifth-order (5R5C) thermal network used in the semi-synthetic testbed, adapted from [23]. Thermal connections are labelled by conductance g; the corresponding resistance is $R = 1/g$. The blue paths are inactive for building ID 67 ($g_{i,\text{out}} = 0$ and $A_e = 0$). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.2.2. RC model structure and ground-truth parameters

The model is a fifth-order (5R5C) network, as shown in Fig. 1. The state vector is $\mathbf{x} = [T_i, T_m, T_e, T_h, T_s]^\top$, representing the temperatures of the indoor air, internal mass, outside wall, heater, and unmeasured sensor thermal mass, respectively. The parameter vector $\theta \in \mathbb{R}^{10}$ consists of five capacitances (C_i, C_m, C_e, C_h, C_s) and five active conductances ($g_{im}, g_{ie}, g_{ih}, g_{is}, g_{e,\text{out}}$). The adapted circuit also shows the direct indoor-air–outdoor conductance $g_{i,\text{out}}$ and the outside-wall solar-gain path $A_e I_{\text{sol}}$; both are inactive for building ID 67 ($g_{i,\text{out}} = 0$ and $A_e = 0$) and are highlighted in blue. The measured inputs are the heater power u_{heat}, outdoor temperature T_{out}, and solar irradiance I_{sol}. The measurement model observes the indoor-air state directly, $y_k = T_{i,k} + v_k$, with Gaussian noise $\sigma_y = 0.15\,\text{K}$; T_s is not the measured output. The notation T_i is retained for this source-model topology, whereas T_{air} denotes the corresponding indoor-air state in the Ecolar topology. The discrete-time system matrices $\mathbf{A}_d(\theta)$ and $\mathbf{B}_d(\theta)$ are obtained via Zero-Order Hold (ZOH) discretisation of the continuous dynamics.

4.2.3. Unmeasured heat-gain generation

To simulate the unmeasured heat gains described in Section 2, the scalar heat-gain signal $q_{\text{dist}}(t)$ is injected directly into the indoor-air node. This signal is unmodelled and is not part of the known input vector; it is the semi-synthetic realisation of the general heat-gain term $\mathbf{F}\mathbf{f}_k$ in Eq. (19). We generate $q_{\text{dist}}(t)$ as a stochastic square-wave process (amplitude 7000 W) with an event-start probability of $p = .0055$ per step and event durations uniformly distributed between 20 min and 8 h. Periods where $q_{\text{dist}}(t) \approx 0$ are labelled the low-disturbance regime, and periods where $q_{\text{dist}}(t)$ is active are labelled the disturbance-dominated regime; these labels serve as the ground truth for evaluating volatility-based detection.

4.2.4. Estimator configurations and hyperparameters

We benchmark the proposed RBPF against three baselines: an Extended Kalman Filter (EKF), an Unscented Kalman Filter (UKF), and an offline nonlinear least-squares (LS) fit.

- **Common settings:** All estimators infer the full 10-dimensional parameter vector in log-space. Initial parameters are drawn from uniform priors.
- **EKF/UKF (Simulator-matched baseline):** To isolate structural limitations from tuning errors, the EKF and UKF are provided with the exact process noise covariance $\mathbf{Q}_x$ used to generate the semi-synthetic data (excluding the unmeasured heat gains). The parameter random-walk covariance is set to $\mathbf{Q}_\phi = \text{diag}(10^{-4})$ for slow adaptation.
- **RBPF (Proposed):** The filter uses $N = 9{,}000$ particles. The process noise is adaptive via the stochastic-volatility state $\alpha_k \in [0.01, 25.0]$, which evolves with $\sigma_\eta = 0.1$. Parameter rejuvenation uses the volatility-based learning gate g_k with thresholds $F_{\text{low}} = 0.40$ and $F_{\text{high}} = 0.95$.

4.2.5. Evaluation protocol and reported metrics

We evaluate performance over the full 5-month horizon. We assess the identified input–output dynamics by comparing the transfer functions induced by the final estimated parameters $\hat{\theta}_N$ with those induced by the ground-truth parameters $\theta^\star$. We assess disturbance-driven parameter updating with the parasitic-learning index. Probabilistic calibration is evaluated via the coverage of 95% prediction intervals and the Winkler score. See Section 3.5.

4.3. Real-building testbed: ecolar home

4.3.1. Dataset and preprocessing

We use two measurement periods from the Ecolar Home testbed at HTWG Konstanz: 03 May–01 June 2025 and 03 October–20 November 2025. Both periods are used for the real-building online-estimation experiments with the RBPF and the offline least-squares (LS) baseline. The RC topology pre-selection reported in Section 4.3.3 is based on the May–June period; applying the same pre-selection procedure to the October–November period led to the same selected topology, and the complete outputs are provided in the public repository. Both periods include indoor-air temperature (T_{air}), outdoor temperature (T_{out}), adjacent-room temperature (T_{adj}), and infrared-heater power (u_{heat}). Temperatures are recorded in degrees Celsius and heater power in watts.

The same preprocessing is applied to both periods. Records are ordered by timestamp and aggregated on a uniform $\Delta t = 300\,\text{s}$ (5 min) grid. Temperature samples within each interval are averaged, whereas heater power retains the first logged value in the interval to preserve its piecewise-constant switching pattern. Intervals missing any required channel are discarded, and gaps longer than one sampling interval are treated as boundaries between continuous data segments; values are not interpolated across such gaps. The processing code and the data-column mapping are provided in the public repository [40].

4.3.2. Measurement setup and building description

Ecolar Home is an experimental "smart home" at HTWG Konstanz. It is not inhabited and is used only for experimental purposes. It comprises two main thermal zones: a living zone of approximately 73 m^2—an open kitchen, living room, and bathroom—which is the zone modelled in this study, and an adjacent room of approximately 16 m^2 whose temperature provides the boundary input T_{adj}. The building is closed between 18:00 and 07:00 and is used irregularly during the day by academic and maintenance staff, whose activities can include door openings, occasional occupancy, and equipment use.

The adjacent room is kept closed for long periods, is unventilated, and is used to store equipment and material. Under solar exposure, this combination lets its air temperature rise well above the outdoor air, reaching about 35 °C on warm days. Its sensor is shielded against solar radiation, so these readings represent the adjacent-room air temperature rather than a measurement artefact. The indoor-air sensor is likewise shielded against solar radiation. The outdoor-temperature sensor is mounted on the north side of the building in a shaded, air-exposed location; it is not radiation-shielded, so a slight solar influence in the early morning cannot be entirely excluded, consistent with the small morning feature visible in T_{out}. Ecolar Home is not directly attached to another structure, but nearby buildings cast time-varying shadows on its façades; the modelling implications of this local shading are discussed in Section 4.3.3. An annotated floor plan with sensor locations is provided in the public repository [40].

4.3.3. Offline RC topology pre-selection

Before applying the online estimators, we selected the RC topology used for the Ecolar zone. This pre-selection used measurements from 03 May to 01 June 2025. For each candidate topology, we fitted the unknown conductances and capacitances by nonlinear least squares over the complete measurement trajectory. The objective was the deterministic open-loop simulation error: after an initial state was assigned, the model was propagated forward using only the measured inputs, without re-correcting the state with the measured indoor temperature. This procedure selects the topology whose deterministic RC dynamics best reproduce the measured trajectory.

Table 2 summarises the candidate model families. Across these candidates, we varied the number of thermal states, the heater-input placement, the direct air–outdoor conductance, the outside-wall representation, and the inclusion of regional GHI. Outdoor temperature and adjacent-room temperature were measured by sensors. For the solar variants, regional GHI was obtained from the German Weather Service (DWD) Climate Data Center for Konstanz [39].

Here, A denotes indoor air, OW the outside-wall node, IW the adjacent-wall node (the internal wall adjoining the adjacent room), H the heater node, and OWi/OWo the inner and outer layers of a split outside wall. For each model family, we also tested variants with and without the regional GHI input.

We evaluated each fitted candidate in two ways. First, we computed the full-trajectory open-loop mean squared error (MSE) on the same trajectory used for fitting. Because the models were fitted by minimising this MSE, Table 3 reports its square root, the root-mean-square error (RMSE), so that the result is expressed on the temperature scale. Second, we evaluated shorter open-loop simulations on repeated daytime and night-time windows. For each daytime window, we re-initialised the model at the first sample of the window and simulated it open-loop from 07:00 to 17:00. For each night-time window, we re-initialised the model at the first sample of the window and simulated it open-loop from 20:00 to 06:00. We then computed the mean absolute error (MAE) for each window and averaged it over all available days or nights. These segment-based MAEs test whether the fitted topology remains predictive when restarted from realistic operating points and provide a summary that is less dominated by an atypical window. Thus, RMSE and MAE are retained here for distinct purposes rather than interpreted interchangeably.

We also inspected residual structure. Let e_k denote the open-loop residual obtained from the fitted deterministic RC model. We define

$$J_{\text{acf}} = \sum_{\ell=\ell_{\text{acf}}}^{L} \rho_e(\ell)^2, \tag{40}$$

where $\rho_e(\ell)$ is the lag-ℓ residual autocorrelation. For an input channel u, we define

$$J_{\text{xcorr},u} = \sum_{\ell=\ell_{\text{xcorr}}}^{L} \rho_{e,u}(\ell)^2, \tag{41}$$

where $\rho_{e,u}(\ell)$ is the lagged cross-correlation between the residual and the input. We evaluated these correlations over a six-hour lag horizon. Lower values indicate weaker residual structure and weaker unexplained dependence on the corresponding input.

Table 3 reports the representative candidates that drove the topology choice. The two-state model is too simple: it has large open-loop errors and strong residual structure. Adding the heater state gives a large improvement. Adding the adjacent-wall state gives the best overall trade-off. The selected four-state model has the lowest training RMSE, full-period MAE, and night-time MAE among the tested candidates. It also reduces the residual cross-correlations with outdoor temperature, adjacent-room temperature, and regional GHI to near zero. Five-state variants give comparable results, but they do not improve the main open-loop metrics enough to justify the additional hidden state and parameters.

The selected four-state model without solar also gives very low residual/input cross-correlation scores: $J_{\text{xcorr},T_{\text{out}}} = 7.31 \times 10^{-4}$, $J_{\text{xcorr},T_{\text{adj}}} = 9.16 \times 10^{-5}$, and $J_{\text{xcorr,solar}} = 1.36 \times 10^{-4}$. For comparison, the best three-state model gives substantially larger values: $J_{\text{xcorr},T_{\text{out}}} = 6.70 \times 10^{-2}$, $J_{\text{xcorr},T_{\text{adj}}} = 1.63 \times 10^{-2}$, and $J_{\text{xcorr,solar}} = 7.74 \times 10^{-2}$. Thus, moving from three to four states improves both open-loop accuracy and residual decorrelation.

The GHI variant of the selected four-state model slightly reduces daytime MAE, from 0.549 °C to 0.543 °C. However, it increases training RMSE, full-period MAE, night-time MAE, and residual autocorrelation. We do not retain this simple regional GHI input in the final deterministic input vector. Solar gains remain physically important, and the limited improvement obtained by adding the DWD GHI reflects the on-site measurement constraints rather than the irrelevance of solar input. The Ecolar Home instrumentation does not include an on-site pyranometer or façade-level irradiance sensors. The only available

Table 2
Candidate RC model families explored for the Ecolar zone.

ID	Model family	States	Main modelling choice
M1	$RC\_2C\_A\_OW$	A, OW	Air and outside-wall model
M2	$RC\_3C\_A\_OW\_H$	A, OW, H	Adds heater thermal mass
M3	$RC\_3C\_A\_OW\_IW$	A, OW, IW	Adds adjacent-wall node
M4	$RC\_4C\_A\_OW\_IW\_H$	A, OW, IW, H	Indoor air, outside wall, adjacent wall, and heater
M5	$RC\_4C\_A\_OW\_IW\_H$ heat split	A, OW, IW, H	Splits heating input between air and heater
M6	$RC\_4C\_A\_OW\_IW\_H$ no infiltration	A, OW, IW, H	Removes direct air–outdoor conductance
M7	$RC\_5C\_A\_OWi\_OWo\_IW\_H$	A, OWi, OWo, IW, H	Splits the outside wall into inner and outer nodes
M8	$RC\_5C\_A\_OWi\_OWo\_IW\_H$ no infiltration	A, OWi, OWo, IW, H	Split outside wall without direct infiltration path

Table 3
Representative results from the Ecolar RC topology pre-selection, using the model IDs defined in Table 2. The two M4 rows share the four-state topology and differ only in the GHI input. RMSE denotes root-mean-square error and MAE denotes mean absolute error; all error values are in °C. Lower values are better; the selected model is highlighted in bold.

ID	Solar	States	Train RMSE (°C)	Full MAE (°C)	Night MAE (°C)	Day MAE (°C)	J_{acf}
M1	No	2	0.637	0.508	0.314	0.934	0.136
M2	No	3	0.369	0.284	0.284	0.560	0.047
M4	**No**	**4**	**0.359**	**0.274**	**0.256**	0.549	0.021
M4	Yes	4	0.363	0.277	0.258	0.543	0.051
M8	No	5	0.365	0.277	0.270	0.533	0.016

solar signal is therefore the regional GHI provided by DWD [39]. Because nearby buildings cast time-varying shadows on the Ecolar Home façades, this regional irradiance does not faithfully represent the heat actually entering the zone. A more accurate deterministic solar input would require additional façade-level instrumentation together with shading and gain-distribution modelling, which are outside the scope of the present online-identification study. The real-building experiment therefore treats the remaining solar-related mismatch as an unmeasured heat gain.

Daytime open-loop errors remain higher than night-time errors even when the simple GHI input is included. This result supports the night-time validation protocol described in Section 4.3.7.

4.3.4. RC model structure

Based on the topology pre-selection in Section 4.3.3, the zone is represented by a four-state RC network, as illustrated in Fig. 2. The state vector is $\mathbf{x} = [T_{\mathrm{air}}, T_{\mathrm{wo}}, T_{\mathrm{wi}}, T_{\mathrm{h}}]^\top$, corresponding to the temperatures of the indoor air, outside-wall node, adjacent-wall node, and heater mass, respectively.

The parameter vector $\theta \in \mathbb{R}^{10}$ comprises the four corresponding thermal capacitances ($C_{\mathrm{air}}, C_{\mathrm{wo}}, C_{\mathrm{wi}}, C_{\mathrm{h}}$) and the six conductances $g_{\mathrm{air,out}}$, $g_{\mathrm{wo,out}}$, $g_{\mathrm{air,wo}}$, $g_{\mathrm{air,wi}}$, $g_{\mathrm{wi,adj}}$, and $g_{\mathrm{air,h}}$. The system is driven by the measured inputs $\mathbf{u} = [u_{\mathrm{heat}}, T_{\mathrm{out}}, T_{\mathrm{adj}}]^\top$. The measurement model is $y_k = T_{\mathrm{air},k} + v_k$, where v_k represents sensor noise with $\sigma_y = 0.1\,\mathrm{K}$ (taken from the sensor datasheet).

4.3.5. Representative measured trajectories

Fig. 3 shows the October–November period as a representative example, with the full time series in panel (a) and a zoomed multi-day window in panel (b) that spans both disturbance-dominated daytime behaviour and low-disturbance nights. The indoor trace is the raw measured signal, with no setpoint or estimate; the May–June period and the full-length series are available in the public repository [40].

4.3.6. Estimator configurations and hyperparameters

- **Offline LS:** For each Ecolar measurement period, a static parameter vector θ_{LS} is identified by minimising the open-loop prediction MSE over that period using the L-BFGS-B algorithm [36], serving as a static baseline (Section 3.3.3).
- **RBPF (proposed):** The filter uses $N = 20{,}000$ particles. The prediction step utilises Euler sub-stepping ($n_{\mathrm{sub}} = 5$) for computational purposes. The stochastic-volatility state is constrained to $\alpha_k \in [0.01, 15.0]$ with evolution noise $\sigma_\eta = 0.05$. Parameter rejuvenation uses the volatility-based learning gate g_k with thresholds $F_{\mathrm{low}} = 0.40$ and $F_{\mathrm{high}} = 0.95$.

4.3.7. Evaluation protocol and reported metrics

In the absence of ground-truth parameters, we assess performance via:

1. **Probabilistic calibration:** We report the empirical coverage and average width of the 95% prediction intervals over the full horizon.
2. **Open-loop simulation:** We test the learned parameters in night-time windows (20:00–06:00), without state updates. The RBPF learns from the full data stream, including daytime data. Because the available measurements do not resolve local solar gains or all daytime internal gains (Sections 4.3.2 and 4.3.3), daytime open-loop errors would confound parameter quality with unobserved inputs. The night-time windows provide comparatively low-disturbance conditions and therefore test whether the learned parameters remain valid when the building dynamics are close to the selected RC structure.

5. Results

This section evaluates the algorithms with respect to (i) preservation of input–output dynamics under unmeasured heat gains, (ii) suppression of disturbance-driven parameter updating, (iii) one-step prediction performance and interval calibration, and (iv) open-loop simulation on low-disturbance segments. We report results for the semi-synthetic testbed and for the Ecolar Home.

5.1. Semi-synthetic testbed

5.1.1. Preservation of input–output dynamics under unmeasured heat gains

We first assess whether the identified RC models preserve the dynamic input–output behaviour. As discussed in Section 3.5, this is the control-relevant quantity under partial observation. Using the dimensionless relative transfer-function errors $\tilde{E}_{r\ell,2}$ and $\tilde{E}_{r\ell,\infty}$ from Eq. (32), Table 4 reports the final deviation of the identified models (θ_N) from the ground truth ($\theta^\star$) for the three channels (solar, heater, outdoor temperature) into the indoor-air output.

Across the solar and heater channels, the constant-covariance baselines (LS, EKF, UKF) exhibit large transfer-function distortion. This indicates that the unmeasured heat gains have changed the identified input–output dynamics. In contrast, the proposed RBPF retains low transfer-function error, consistent with the volatility-based learning gate suppressing disturbance-driven parameter drift.

Fig. 4 complements Table 4 by showing the temporal evolution of $\tilde{E}_{r\ell,2}$, the relative H_2-like transfer-function error. The RBPF converges and remains stable over the full horizon, whereas the EKF and UKF maintain a higher error and larger variance throughout the entire horizon.

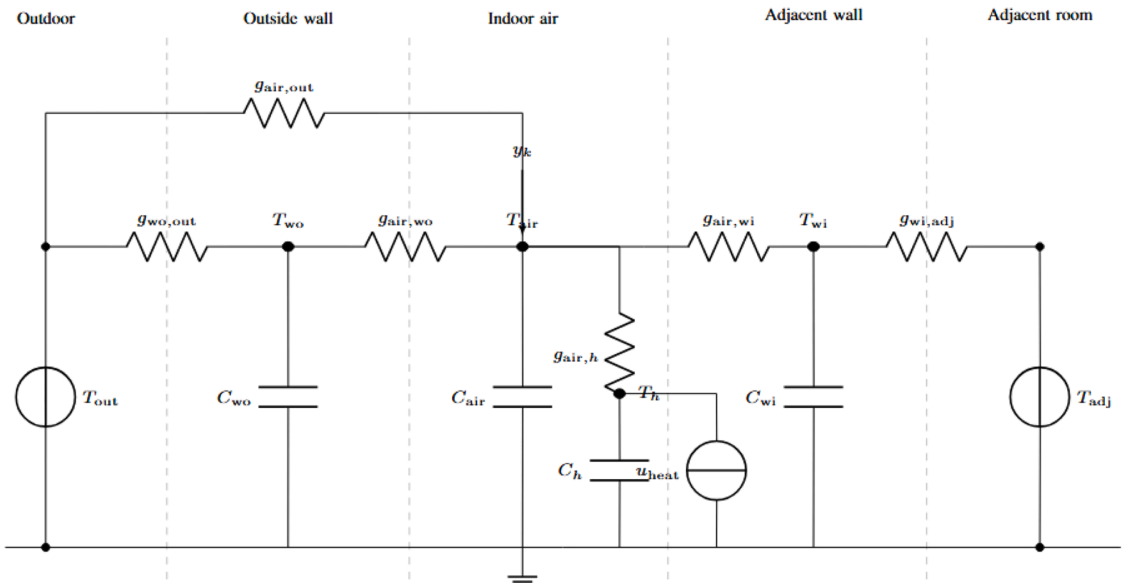

Fig. 2. Equivalent circuit representation of the fourth-order (6R4C; four capacitances and six conductances) thermal network used for Ecolar Home. Thermal connections are labelled by conductance g, with $R = 1/g$. Unlike the semi-synthetic model in Fig. 1, the Ecolar model has no separate sensor node; sensor uncertainty is represented by v_k with $\sigma_y = 0.1\,\mathrm{K}$.

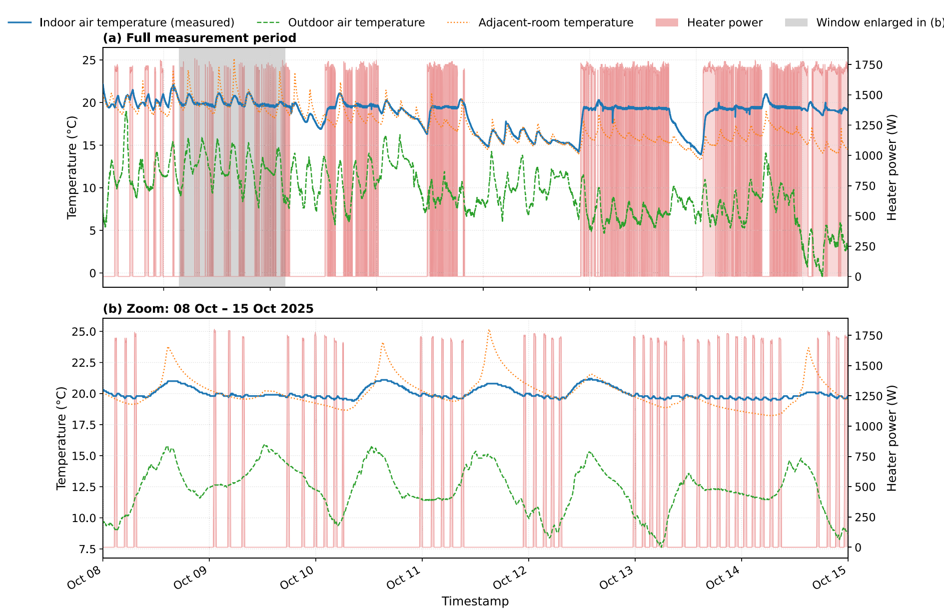

Fig. 3. Ecolar Home (October–November 2025): measured indoor-air, outdoor-air, and adjacent-room temperatures together with heater power. (a) Full measurement period; the shaded band marks the window enlarged in panel (b). (b) Detailed view of 08–15 October 2025.

Table 4
Semi-synthetic testbed: final relative transfer-function errors (θ_N vs. $\theta^\star$). Lower is better.

Channel	Solar		Heater		Outside	
Method	$\tilde{E}_{r\ell,2}$	$\tilde{E}_{r\ell,\infty}$	$\tilde{E}_{r\ell,2}$	$\tilde{E}_{r\ell,\infty}$	$\tilde{E}_{r\ell,2}$	$\tilde{E}_{r\ell,\infty}$
LS (offline)	0.844	0.866	0.845	0.866	0.085	0.046
EKF	0.846	0.868	0.847	0.868	0.107	0.030
UKF	0.753	0.771	0.753	0.771	0.109	0.031
RBPF	**0.049**	**0.048**	**0.071**	**0.048**	**0.035**	**0.008**

5.1.2. *Volatility regimes and suppression of parasitic learning*

To verify that the RBPF learns primarily from low-disturbance periods, where the dynamics are consistent with the RC model, and suppresses adaptation during disturbance-dominated periods, we analyse (i) the inferred volatility state and (ii) the magnitude of parameter updates.

Fig. 5 shows the distribution of $\log(\alpha_k)$ under both the low-disturbance and disturbance-dominated periods. Using the regime sets $\mathcal{L}$ and $\mathcal{D}$ from Eq. (34), the low-disturbance median is $\mathrm{median}_{k\in\mathcal{L}} \log(\alpha_k) = 1.26$ ($\alpha \approx 3.54$), while the disturbance-dominated median is $\mathrm{median}_{k\in\mathcal{D}} \log(\alpha_k) = 2.91$ ($\alpha \approx 18.40$), indicating that the filter assigns substantially lower trust to the RC model during disturbance-dominated periods. The baseline (disturbance-free) process covariance is $\mathbf{Q}_{\mathrm{true}} = \mathrm{diag}(0.03^2)$, whereas the RBPF uses $\mathbf{Q}_0 = \mathrm{diag}(0.01^2)$; matching $\alpha^2\mathbf{Q}_0$ to $\mathbf{Q}_{\mathrm{true}}$ implies a theoretical low-disturbance volatility level $\alpha_{\mathrm{low}} = \sqrt{(0.03^2)/(0.01^2)} = 0.03/0.01 = 3.0$, or $\log(\alpha_{\mathrm{low}}) \approx 1.10$, which is close to the observed low-disturbance median on the logarithmic scale.

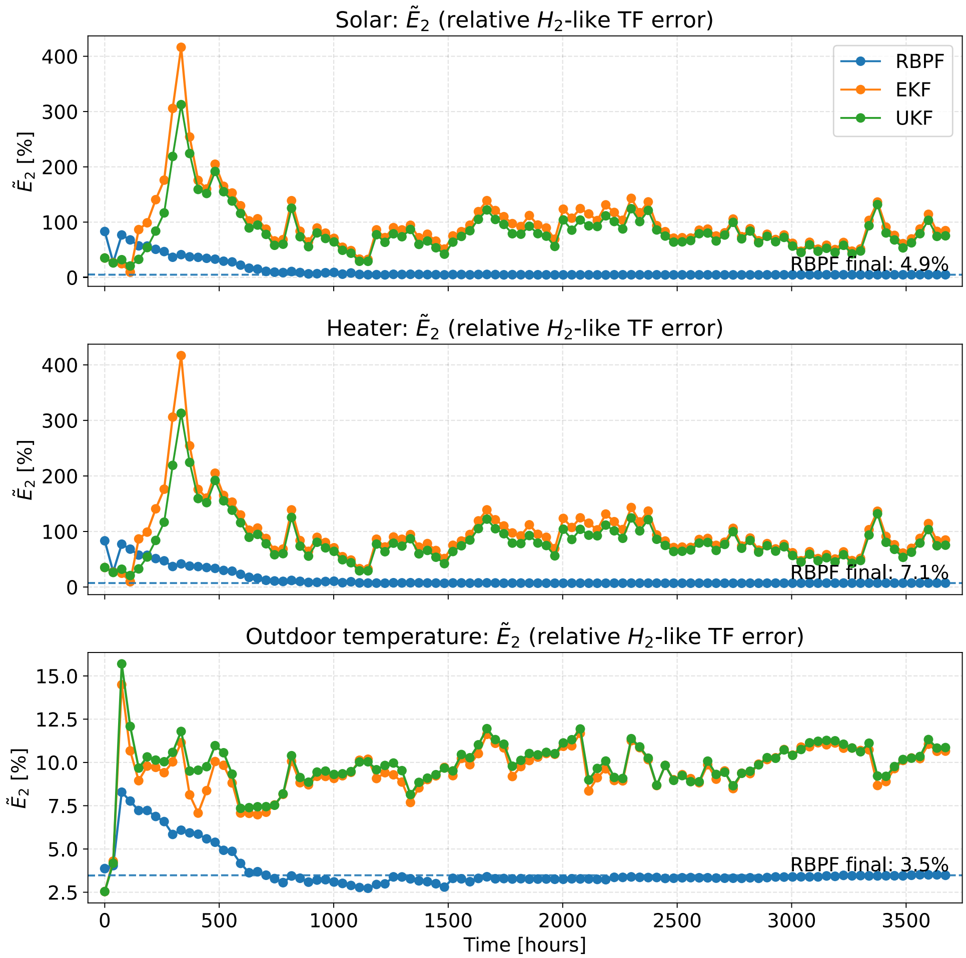

Fig. 4. Semi-synthetic testbed: time evolution of $\tilde{E}_{r\ell,2}$, the dimensionless relative H_2-like transfer-function error, for the solar, heater, and outdoor-temperature channels. The figure expresses the metric as a percentage; for example, 4.9% corresponds to $\tilde{E}_{r\ell,2} = 0.049$ in Table 4.

We quantify *parasitic learning* using the parasitic-learning index (PLI) defined in Eq. (34). The PLI compares the median log-parameter update magnitude during disturbance-dominated periods with the corresponding median during low-disturbance periods. Fig. 6 compares this statistic between low-disturbance and disturbance-dominated periods for each estimator. The RBPF exhibits only a mild increase in median parameter-update magnitude between regimes (2.31×10^{-4} in low-disturbance periods vs. 2.68×10^{-4} in disturbance-dominated periods), giving PLI = 1.16. In contrast, the EKF and UKF show strong disturbance-amplified adaptation, with PLI = 4.39 and PLI = 4.31, respectively.

5.1.3. One-step prediction accuracy and probabilistic calibration

We evaluate one-step prediction performance and calibration of 95% prediction intervals. Table 5 summarises RMSE, Winkler score, and empirical coverage, both overall and conditioned on the regime type (low-disturbance vs. disturbance-dominated periods). The RBPF achieves low RMSE and substantially improved calibration relative to EKF/UKF, with coverage close to nominal during low-disturbance periods and only moderate degradation during disturbance-dominated periods.

Fig. 7 visualises the volatility adaptation at the start of a disturbance-dominated period. Unlike the EKF and UKF, which maintain narrow intervals and track the true state temperature slowly, the RBPF reacts rapidly to the disturbance by expanding its prediction interval. This increase in volatility maintains calibration by shifting trust from the model to the measurements.

Table 5
Semi-synthetic testbed: one-step prediction accuracy and calibration. Coverage is for nominal 95% intervals. Lower Winkler is better.

Method	RMSE (°C)	Winkler	Coverage (95% nominal) Overall	Low-dist.	Dist.-dom.
RBPF	**0.088**	**0.42**	**94.7%**	**95.3%**	**90.1%**
EKF	0.224	3.01	66.6%	75.2%	3.9%
UKF	0.232	3.05	69.4%	78.3%	4.0%

5.1.4. Open-loop simulation on low-disturbance segments

Finally, we test whether the identified parameters yield a meaningful generative model by running open-loop simulations (no state updates) on low-disturbance segments of length at least 12 h. Table 6 summarises the open-loop RMSE statistics across all evaluated segments. The RBPF achieves a mean RMSE of 0.22 °C, significantly outperforming the offline LS baseline (0.88 °C) and the EKF/UKF variants (0.91 °C and 0.82 °C, respectively). Notably, the maximum error recorded for the RBPF (0.56 °C) is lower than the mean error of any baseline method.

5.1.5. Sensitivity analysis

We assess sensitivity to the number of particles N_p, the volatility evolution noise σ_η, and the learning-gate thresholds $(F_{\text{low}}, F_{\text{high}})$, with all other settings held at the values reported in Section 4.2.4.

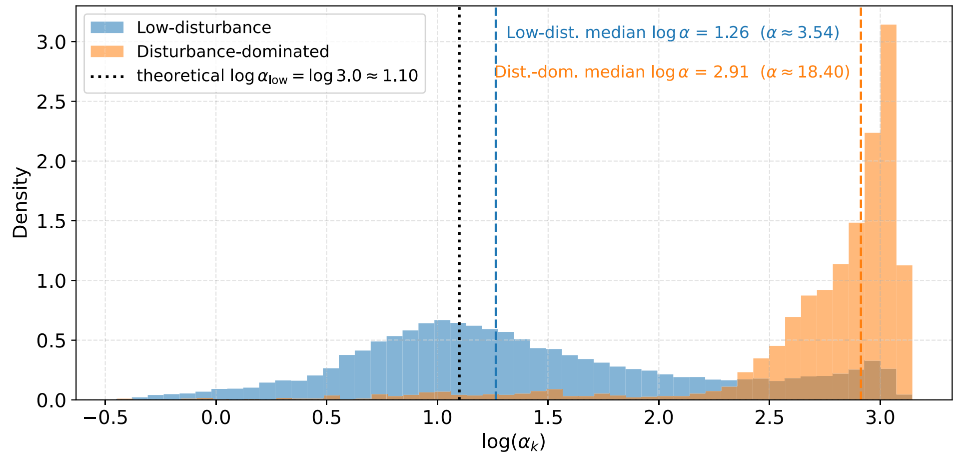

Fig. 5. Semi-synthetic testbed: distribution of $\log(\alpha_k)$ in low-disturbance and disturbance-dominated periods. Dashed coloured lines show the regime medians in both the logarithmic and original scales. The black dotted line shows the theoretical low-disturbance baseline $\log(\alpha_{\text{low}}) = \log(3.0) \approx 1.10$, obtained from $\mathbf{Q}_{\text{true}}$ and $\mathbf{Q}_0$.

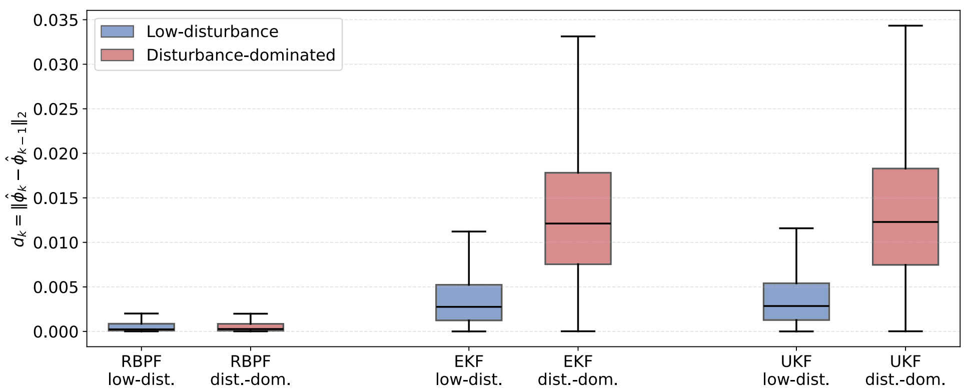

Fig. 6. Semi-synthetic testbed: per-step log-parameter update magnitude d_k split by low-disturbance and disturbance-dominated regimes. The RBPF shows nearly regime-invariant adaptation (PLI = 1.16), whereas EKF and UKF show disturbance-amplified parameter drift (PLI = 4.39 and PLI = 4.31, respectively).

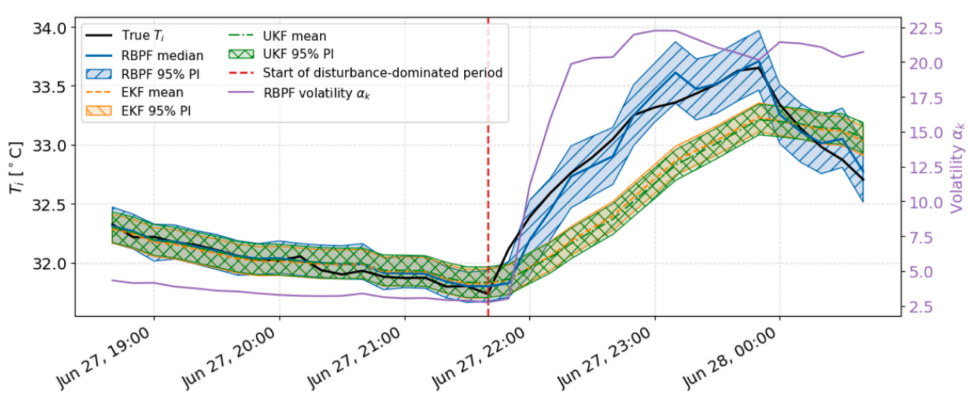

Fig. 7. The RBPF widens its prediction intervals at the start of a disturbance-dominated period through increased volatility, while EKF/UKF remain overconfident and fail to cover transient deviations. In this semi-synthetic testbed, unmeasured heat-gain events are injected independently of the diurnal cycle (Section 4.2), so the displayed start time is incidental.

The particle count exhibits clear diminishing returns: the median open-loop RMSE across low-disturbance segments drops from 0.25 °C at $N_p = 3{,}000$ to 0.20 °C at $N_p = 9{,}000$ and remains essentially unchanged at $N_p = 15{,}000$, while runtime scales linearly with N_p. We therefore recommend $N_p \in [6{,}000, 9{,}000]$.

The volatility evolution noise shows the clearest sensitivity. Setting $\sigma_\eta = 0.05$ is too restrictive: the volatility state cannot track regime transitions, raising the median open-loop RMSE across low-disturbance segments to 0.28 °C and yielding 90.9% coverage in disturbance-dominated periods. At $\sigma_\eta = 0.20$, overall coverage drifts to 94.6%, below the nominal 95% level. $\sigma_\eta \approx 0.10$ provides the best balance.

The histogram-calibrated volatility-based learning gate, by contrast, is robust to threshold choice. Across $F_{\text{low}} \in [0.15, 0.40]$ and $F_{\text{high}} \in [0.50, 0.90]$ (with $F_{\text{low}} < F_{\text{high}}$), the median open-loop RMSE across low-disturbance segments remains within 0.20–0.26 °C, overall 95% coverage within 95.3–95.7%, and median 95% prediction-interval width

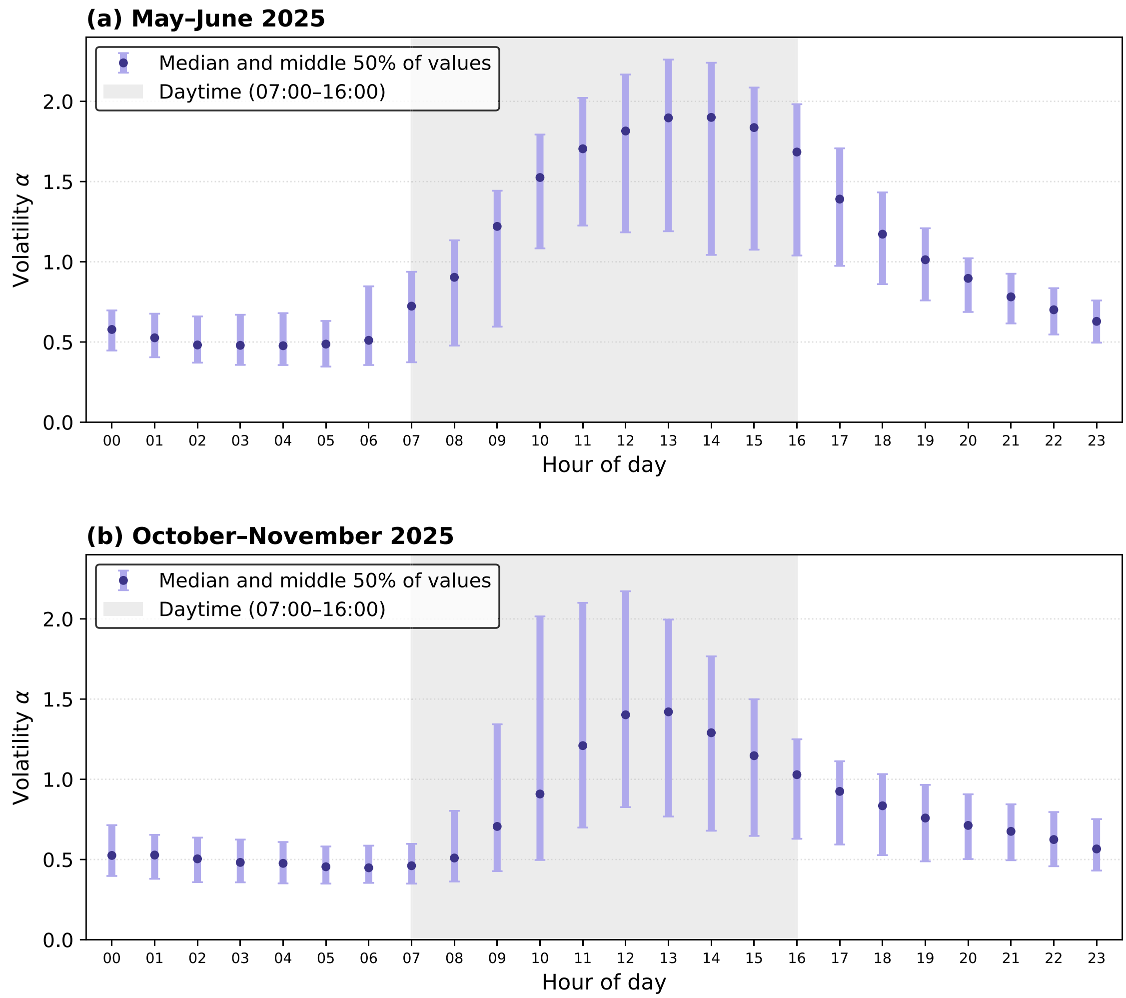

Fig. 8. Ecolar Home: diurnal pattern of the RBPF stochastic-volatility state α_k for (a) May–June 2025 and (b) October–November 2025. Points show hourly medians and bars show the middle 50% of values. The panels share the same axis limits and are shown separately because the periods are non-contiguous and the RBPF is run independently on each.

Table 6
Open-loop RMSE on low-disturbance segments (state updates disabled). Summary statistics over all evaluated low-disturbance segments.

Method	RMSE (°C)			
	Mean	Median	Min	Max
RBPF	**0.2203**	**0.1984**	**0.0519**	**0.5595**
LS (offline)	0.8752	0.7797	0.0771	2.1632
EKF	0.9064	0.8235	0.1477	1.9492
UKF	0.8216	0.7615	0.1286	1.8018

within 0.338–0.340 °C. Any pair satisfying $F_{\text{high}} - F_{\text{low}} \gtrsim 0.3$ yields comparable performance. The full sensitivity analysis, including the complete set of additional results, is available in the accompanying repository [40].

5.2. Real-building testbed results: ecolar home

Because ground-truth parameters are unavailable for Ecolar Home, we evaluate the RBPF through three complementary diagnostics: the diurnal volatility pattern, one-step calibration with measurement updates, and open-loop simulation over low-disturbance night-time windows. Both measurement periods and their measured trajectories are described in Section 4.3 and illustrated in Fig. 3.

Table 7
Ecolar Home: one-step RBPF performance over the two real-building measurement periods. Coverage is reported for nominal 95% prediction intervals.

Metric	RBPF
Coverage (95% nominal)	95.23%
Mean prediction-interval width (°C)	0.178
One-step RMSE (°C)	0.045

Table 8
Ecolar Home: open-loop simulation error during night-time windows, with state updates disabled. Values summarize the distribution of night-block RMSE.

Period	Method	Median RMSE (°C)	Std. RMSE (°C)
03 May–01 Jun 2025	LS (offline)	0.305	0.115
03 May–01 Jun 2025	**RBPF**	**0.175**	**0.091**
03 Oct–20 No. 2025	LS (offline)	0.550	0.606
03 Oct–20 No. 2025	**RBPF**	**0.193**	**0.134**

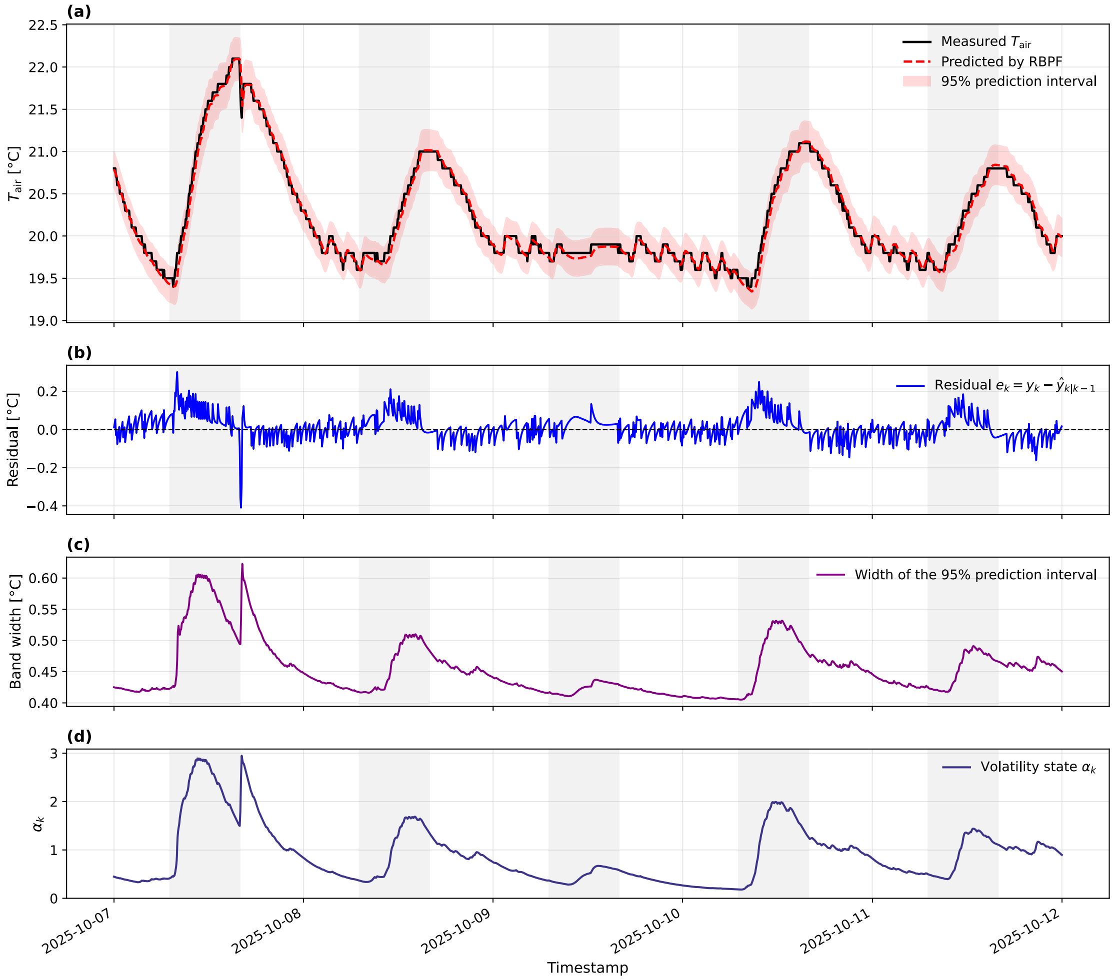

Fig. 9. Ecolar Home (7–12 October 2025): one-step RBPF diagnostics over a representative window. (a) Measured indoor-air temperature and one-step prediction with its 95% prediction interval; (b) one-step residual; (c) width of the 95% prediction interval; and (d) stochastic-volatility state α_k. Daytime (07:00–16:00) is shaded consistently across all panels. The rise in α_k coincides with wider prediction intervals, while coverage stays near nominal (95.6% daytime, 97.1% outside daytime).

Table 9
Representative computational cost of the RBPF implementation using the semi-synthetic particle count (N_p = 9,000) and the runtime-benchmark setting $n_{sub} = 10$. This benchmark is distinct from the Ecolar configuration (N_p = 20,000, $n_{sub} = 5$; Section 4.3.6).

Quantity	Value
Implementation	Python 3.12, CPU-only
Processor	Intel Xeon CPU @ 2.20 GHz
Particles N_p	9,000
Sampling period Δt	300–600 s
Euler sub-steps n_{sub}	10
Mean update time	0.05–0.06 s
Equivalent update rate	16–20 updates/s
Persistent memory footprint	2.9 MB

5.2.1. Diurnal volatility behaviour

Fig. 8 shows the hourly distribution of the inferred stochastic-volatility state α_k for both Ecolar measurement periods. In both periods, the RBPF assigns lower volatility during night-time and early-morning hours, and higher volatility during daytime hours. The increase starts around 07:00, remains elevated through the day, and decreases again after approximately 16:00. The two periods differ mainly in the evening: in the May–June period the volatility remains elevated later into the evening than in October–November, consistent with the longer daylight (the sun is still up) and the correspondingly later solar gains. Because the two measurement periods are non-contiguous and the RBPF is run independently on each, we present them as separate panels, which also makes this seasonal contrast visible.

This pattern is consistent with the intended interpretation of α_k as an inverse model-trust indicator. During daytime periods, solar gains, occupancy effects, ventilation, and other unmeasured heat gains are more likely. The RBPF responds by increasing α_k, which indicates lower trust in the deterministic RC model, inflates the effective process covariance, and reduces parameter rejuvenation through the volatility-based learning gate g_k. Conversely, lower night-time volatility permits parameter learning primarily during periods in which the deterministic RC input structure is more reliable.

5.2.2. Closed-loop one-step calibration

Table 7 reports the one-step RBPF performance over the two Ecolar measurement periods with measurement updates enabled. The empirical coverage of the nominal 95% prediction intervals is 95.23%, close to nominal calibration under real operation. The average prediction-interval width is 0.178 °C and the one-step RMSE is 0.045 °C.

Fig. 9 shows the one-step behaviour over a representative window (7–12 October 2025). Panel (a) confirms that the prediction mean $\hat{y}_{k|k-1}$ follows the measured indoor-air temperature T_{air}, the variable the filter predicts. Panel (b) shows the one-step residual $e_k = y_k - \hat{y}_{k|k-1}$. The residual stays small and near zero at night but becomes systematic during the day, reflecting unmeasured heat gains that the deterministic RC model cannot explain. Panels (c) and (d) show that the rise in α_k coincides with wider one-step prediction intervals. Over this window the 95% interval stays close to nominal coverage: 95.6% during the 07:00–16:00 daytime hours and 97.1% outside them.

5.2.3. Open-loop validation of the learned RC parameters on low-disturbance nights

For the open-loop validation defined in Section 4.3.7, we freeze the learned parameters and simulate the RC model through each night-time window without measurement updates.

Table 8 reports the distribution of night-block open-loop RMSE for the offline LS baseline and the RBPF model. In the May–June period, the median night-block RMSE decreases from $0.305\,^{\circ}\mathrm{C}$ for LS to $0.175\,^{\circ}\mathrm{C}$ for RBPF, corresponding to a reduction of approximately 43%. In the October–November period, the median RMSE decreases from $0.550\,^{\circ}\mathrm{C}$ to $0.193\,^{\circ}\mathrm{C}$, corresponding to a reduction of approximately 65%. The RBPF also yields lower variability across night blocks in both periods.

The lower night-block errors show that the RBPF parameters remain useful for simulation after learning from the full data stream.

5.3. Computational cost and real-time feasibility

We evaluated the computational cost of the proposed RBPF using a CPU-only Python 3.12 implementation on an Intel Xeon CPU @ 2.20 GHz. Table 9 summarises the runtime-benchmark settings.

For a fixed RC state dimension, the dominant cost is the per-particle conditional Kalman prediction/update and likelihood evaluation. Therefore, the runtime scales approximately linearly with the number of particles N_p. The Liu–West rejuvenation step adds the computation of the weighted covariance in the log-parameter space and sampling from the rejuvenation kernel; for the low-dimensional RC parameter vectors considered here, this cost is small compared with the per-particle propagation and likelihood evaluation.

More generally, if n_x denotes the number of RC states, d_θ the number of parameters, and N_p the number of particles, the cost per update can be summarized as

$$\mathcal{O}\big(N_p\, C_{\mathrm{KF}}(n_x) + N_p\, d_\theta^2\big),$$

where $C_{\mathrm{KF}}(n_x)$ denotes the cost of the conditional Kalman prediction/update and RC model propagation for one particle. In the single-zone models used here, $n_x \le 5$ and $d_\theta \le 10$, so these matrix operations are small. For independent multi-zone deployment, the cost scales approximately linearly with the number of zones and can be parallelised across zones. For strongly coupled multi-zone RC models, the state dimension increases and the Kalman propagation cost becomes the main scaling factor.

Relative to the 5–10 min sampling intervals used in the experiments, the measured runtime uses less than 0.02% of the available time budget per update. This indicates that the method is compatible with real-time online estimation at typical building-control sampling rates. The present implementation is not optimised for embedded deployment; it is a vectorised Python prototype. A compiled or hardware-specific implementation would likely reduce runtime further, but deployment on a specific low-cost controller should still be benchmarked on the target device.

6. Conclusion

This paper addresses two gaps in the online identification of building thermal models: the *parasitic learning* of unmeasured heat gains into RC parameters and the lack of probabilistic calibration during disturbance-dominated regimes. By proposing a Rao–Blackwellized Particle Filter (RBPF) augmented with a stochastic-volatility state, we demonstrated that stochastic volatility can limit the absorption of structural mismatch into the RC parameters while preserving the identified input–output dynamics used for prediction.

Our results on both the semi-synthetic and Ecolar Home testbeds show that the proposed RBPF outperforms standard EKF and UKF approaches in terms of dynamics preservation and uncertainty calibration. Finally, the computational requirements of the method were found to be compatible with typical building sampling intervals, making the approach practically deployable for online estimation.

Several limitations bound these results. The real-building evaluation covers two multi-week measurement periods on a single building, so the findings should be read as a focused case study rather than a validation across building types and climates. In the absence of on-site solar and occupancy instrumentation, the learned RC parameters are assessed on real data only indirectly—through one-step calibration and open-loop simulation on low-disturbance night-time windows—whereas a parameter-level ground truth is available only in the semi-synthetic testbed. Finally, the estimator is evaluated in open loop; using the stochastic-volatility state as an inverse model-trust indicator inside a closed-loop MPC remains to be demonstrated.

Future work will focus on embedding the estimator in a Model Predictive Control (MPC) loop. There, the RBPF plays two roles: it adapts the RC parameters to slowly varying building conditions, and the stochastic-volatility state α_k provides an online inverse model-trust indicator: larger α_k means that the deterministic model is currently less reliable. Since MPC requires a sufficiently accurate model and can be outperformed by simpler rule-based control once model uncertainty exceeds a threshold [7], the volatility—or its running statistics—is a natural online proxy for that threshold: it could act as a supervisory “knob” that arbitrates between MPC and a simpler fallback controller, or that widens the MPC constraints when trust is low. We present this as an outlook to be validated in closed loop.

The present study addresses a multi-input, single-output (MISO) configuration, in which only the indoor-air temperature is measured; extending the estimator to multi-output (MIMO) settings with several measured zone temperatures is left for future work. Likewise, the stochastic volatility used here is a single scalar state that scales the covariance of the Gaussian process noise; future work will consider a richer parametrisation, with several stochastic-volatility states scaling different blocks of that covariance.

Finally, the inferred volatility state offers potential for anomaly detection: distinct patterns in the volatility signal could be classified to diagnose specific malfunctions or identify structural shortcomings.

CRediT authorship contribution statement

Mohamed Ramadane: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Doris Bohnet:** Writing – review & editing, Visualization, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Large Language Models to refine the language and formatting. The authors take full responsibility for the content.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research is funded by the IoT Sustainability Lab under the Interreg VI Alpenrhein-Bodensee-Hochrhein program, supported by the

European Regional Development Fund (ERDF) and the Swiss Confederation. Open Access funding enabled and organized by HTWG Hochschule Konstanz.

Data availability

The data and code are available in a public Github repository linked in the manuscript.

References

[1] C. Maduta, T. Letnik, A. Anthrakidis, M. Economidou, Towards climate neutrality within the European union: assessment of the energy performance of buildings directive implementation in member states, Energy Build. 301 (2023) 113716. https://doi.org/10.1016/j.enbuild.2023.113716
[2] Eurostat, Energy use in EU households down second year in a row, 2025, (Eurostat News Release). Accessed: 2025-11-20, https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250625-2.
[3] Q. Hamp, F. Levihn, Model predictive control for dynamic indoor conditioning in practice, Energy Build. 277 (2022) 112548. https://doi.org/10.1016/j.enbuild.2022.112548
[4] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E.P. Ollé, J. Oravec, M. Wetter, D. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annu. Rev. Control 50 (2020) 190–232. https://doi.org/10.1016/j.arcontrol.2020.09.001
[5] S. Freund, G. Schmitz, Implementation of model predictive control in a large-sized, low-energy office building, Build. Environ. 197 (2021) 107830. https://doi.org/10.1016/j.buildenv.2021.107830
[6] S. Prívara, J. Cigler, Z. Váňa, O. Šálek, L. Ferkl, Building modeling as a crucial part for building predictive control, Energy Build. 56 (2013) 8–22. https://doi.org/10.1016/j.enbuild.2012.10.024
[7] M. Maasoumy, M. Razmara, M. Shahbakhti, A. Sangiovanni-Vincentelli, Handling model uncertainty in model predictive control for energy efficient buildings, Energy Build. 77 (2014) 377–392. https://doi.org/10.1016/j.enbuild.2014.03.057
[8] F. Bünning, B. Huber, P. Heer, A. Aboudonia, J. Lygeros, Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings, Energy Build. 211 (2020) 109792. https://doi.org/10.1016/j.enbuild.2020.109792
[9] A. Kathirgamanathan, M. De Rosa, E. Mangina, D.P. Finn, Data-driven predictive control for unlocking building energy flexibility: a review, Renew. Sustain. Energy Rev. 135 (2021) 110120. https://doi.org/10.1016/j.rser.2020.110120
[10] J. Yan, G. Lin, Comparison of day-matching and machine-learning models in predicting building peak electric consumption, 2025. https://doi.org/10.63044/s25com25
[11] G. Pinto, Z. Wang, A. Roy, T. Hong, A. Capozzoli, Transfer learning for smart buildings: a critical review of algorithms, applications, and future perspectives, Adv. Appl. Energy 5 (2022) 100084. https://doi.org/10.1016/j.adapen.2022.100084
[12] M. Montazeri, C. Remlinger, B. Bejar Haro, P. Heer, Fully data-driven and modular building thermal control with physically consistent modeling, Appl. Energy 390 (2025) 125770. https://doi.org/10.1016/j.apenergy.2025.125770
[13] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy Build. 43 (7) (2011) 1511–1522. https://doi.org/10.1016/j.enbuild.2011.02.005
[14] R. Serasinghe, N. Long, J. Clark, Parameter identification methods for low-order gray box building energy models: a critical review, Energy Build. 311 (2024) 114123. https://doi.org/10.1016/j.enbuild.2024.114123
[15] S. Rouchier, M.J. Jiménez, S. Castaño, Sequential Monte Carlo for on-line parameter estimation of a lumped building energy model, Energy Build. 187 (2019) 86–94. https://doi.org/10.1016/j.enbuild.2019.01.045
[16] G. Lin, A. Casillas, M. Sheng, J. Granderson, Performance evaluation of an occupancy-based HVAC control system in an office building, Energies 16 (20) (2023). https://doi.org/10.3390/en16207088
[17] L. Raillon, C. Ghiaus, Sequential Monte Carlo for states and parameters estimation in dynamic thermal models, in: Proceedings of Building Simulation 2017: 15th Conference of the International Building Performance Simulation Association, IBPSA, San Francisco, USA, 2017, pp. 75–84. https://doi.org/10.26868/25222708.2017.020
[18] S.F. Fux, A. Ashouri, M.J. Benz, L. Guzzella, EKF based self-adaptive thermal model for a passive house, Energy Build. 68 (2014) 811–817. https://doi.org/10.1016/j.enbuild.2012.06.016
[19] Y. Li, J. Castiglione, R. Astroza, Y. Chen, Real-time thermal dynamic analysis of a house using RC models and joint state–parameter estimation, Build. Environ. 188 (2021) 107184. https://doi.org/10.1016/j.buildenv.2020.107184
[20] D. Cheng, W. Zhang, J. Liu, Window-varying particle filter for parameter identification of space thermal model, IEEE Trans. Instrum. Meas. 66 (1) (2017) 165–176. https://doi.org/10.1109/TIM.2016.2619986
[21] S. Ham, P. Karava, Online building energy model to evaluate heating and cooling-related behavior changes for eco-feedback in a multifamily residential building, 2019, pp. 2379–2387. https://doi.org/10.26868/25222708.2019.211234
[22] S. Rouchier, M. Rabouille, P. Oberlé, Calibration of simplified building energy models for parameter estimation and forecasting: stochastic versus deterministic modelling, Build. Environ. 134 (2018) 181–190. https://doi.org/10.1016/j.buildenv.2018.02.043
[23] J. Leprince, C. Miller, H. Madsen, K. Basu, R. van der Vlist, W. Zeiler, Grey-brick buildings, an open data set of calibrated RC models of dutch residential building heat dynamics, in: Proceedings of the Twentieth ACM Conference on Embedded Networked Sensor Systems, 2022, pp. 1067–1071.
[24] G. Reynders, J. Diriken, D. Saelens, Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals, Energy Build. 82 (2014) 263–274. https://doi.org/10.1016/j.enbuild.2014.07.025
[25] J. Leprince, H. Madsen, C. Miller, J.P. Real, R. van der Vlist, K. Basu, W. Zeiler, Fifty shades of grey: automated stochastic model identification of building heat dynamics, Energy Build. 266 (2022) 112095. https://doi.org/10.1016/j.enbuild.2022.112095
[26] P. Radecki, B. Hencey, Online building thermal parameter estimation via unscented Kalman filtering, in: Proceedings of the American Control Conference (ACC), IEEE, 2012, pp. 3056–3062. https://doi.org/10.1109/ACC.2012.6315699
[27] A. Martinčević, A. Starčić, M. Vašak, Parameter estimation for low-order models of complex buildings, in: IEEE PES Innovative Smart Grid Technologies, Europe, 2014, pp. 1–6. https://doi.org/10.1109/ISGTEurope.2014.7028767
[28] S. Baldi, S. Yuan, P. Endel, O. Holub, Dual estimation: constructing building energy models from data sampled at low rate, Appl. Energy 169 (2016) 81–92. https://doi.org/10.1016/j.apenergy.2016.02.019
[29] V. Zamani, S. Abtahi, Y. Chen, Y. Li, Parameter-input estimation of RC thermal models of buildings using unscented Kalman filter and nonlinear least square method, Indoor Built Environ. 34 (1) (2024) 1–15. https://doi.org/10.1177/1420326X241270775
[30] D. Cheng, X. Chu, W. Zhang, Y. Liu, Y. Hou, Application of particle filter to parameter identification of space thermal model, in: Proceedings of the 2014 International Conference on Neural Networks and Computing (ICNC), 2014, pp. 589–593. https://doi.org/10.1109/ICNC.2014.6975901
[31] A.R. Coffman, P. Barooah, Simultaneous identification of dynamic model and occupant-induced disturbance for commercial buildings, Build. Environ. 128 (2018) 153–160. https://doi.org/10.1016/j.buildenv.2017.10.020
[32] A.H.J. (ed), Stochastic Processes and Filtering Theory, vol. 64 of *Mathematics in Science and Engineering*, Elsevier, 1970. https://doi.org/10.1016/S0076-5392(09)60365-9
[33] D. Simon, Nonlinear Kalman Filtering, John Wiley & Sons, Ltd, 2006, pp. 393–431. https://doi.org/10.1002/0470045345.ch13
[34] S.J. Julier, J.K. Uhlmann, New extension of the Kalman filter to nonlinear systems, in: I. Kadar (Ed.), Signal Processing, Sensor Fusion, and Target Recognition VI, Vol. 3068, International Society for Optics and Photonics, SPIE, 1997, pp. 182 – 193. https://doi.org/10.1117/12.280797
[35] E.A. Wan, R. Van Der Merwe, The unscented Kalman filter for nonlinear estimation, in: Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373), 2000, pp. 153–158. https://doi.org/10.1109/ASSPCC.2000.882463
[36] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, İ. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat. Methods 17 (2020) 261–272. https://doi.org/10.1038/s41592-019-0686-2
[37] J.S. Liu, M. West, Combined parameter and state estimation in simulation-based filtering, in: A. Doucet, N. de Freitas, N. Gordon (Eds.), Sequential Monte Carlo Methods in Practice, Springer, New York, 2001, pp. 197–223. https://doi.org/10.1007/978-1-4757-3437-9_10
[38] R.L. Winkler, A decision-theoretic approach to interval estimation, J. Am. Stat. Assoc. 67 (337) (1972) 187–191.
[39] Deutscher Wetterdienst, Climate Data Center (CDC) – overview of climate data, 2025, (https://www.dwd.de/EN/ourservices/cdc/cdc_ueberblick-klimadaten_en.html). Accessed: 2025-11-23.
[40] M. Ramadane, Research code, 2025, (GitHub repository). https://github.com/mramadane/Research_code.