

Tuning performance of grey-box models for thermal building applications[☆]

Savvas Panagi^{a,b}^{*}, Lazaros Aresti^b, Paul Christodoulides^c, Chrysovalantis Spanias^a,
Petros Aristidou^b

^a Distribution System Operator, Electricity Authority of Cyprus, Nicosia, Cyprus

^b Department of Electrical & Computer Engineering, & Informatics, Cyprus University of Technology, Limassol, Cyprus

^c Faculty of Engineering and Technology, Cyprus University of Technology, Limassol, Cyprus

ARTICLE INFO

Keywords:

Optimal tuning
Buildings
Thermal modeling
Grey-box models

ABSTRACT

The electrification of heating and cooling systems alters energy consumption patterns, presenting both challenges and opportunities for the operation of power systems. To this end, RC (grey-box) models have emerged as promising modeling structures for forecasting the thermal needs of buildings. This paper presents a tuning algorithm and performance evaluation of the grey-box model using realistic measurement data. It provides insights into the sensitivity of the learning process to sub-optimal solutions and the computational burden. Four different grey-box structures with varying complexity are evaluated. The proposed methodology is critical for integrating heating and cooling systems into future power systems. The results reveal that while the 4R3C model is the most detailed model structure among the evaluated structures, the 3R2C model proves to be the most stable, offering the best trade-off between computational burden and model complexity. Moreover, the physical representation of the building through this training structure can be challenging, as the optimization process does not consistently converge to a unique set of parameter values, indicating the presence of multiple local optima. The tuning framework is provided as an open-source modeling tool, aiming to support further research on grey-box models.

1. Introduction

1.1. Motivation

The urgent need to address climate change is accelerating the transition towards clean energy technologies, particularly in the heating sector. Conventional heating systems, which are still largely dependent on fossil fuels, are being increasingly replaced by electricity-based alternatives, such as heat pumps. In the European Union, heating and cooling account for approximately 50% of total energy consumption, with over 70% still relying on fossil fuels, while around 80% of residential energy use is dedicated to space and water heating [1,2]. As a result, there is a growing need for accurate and adaptable modeling approaches that can capture the thermal behavior of buildings and predict their future heating and cooling demands.

The thermal-electrical interface typically used for modeling thermal needs and demand-side flexibility has been made rigorous by building thermodynamics, resulting in oversimplified models that cannot fully capture real-world performance dynamics. This limitation arises from the gap between disciplines: building engineers are often unfamiliar

with the technical details of energy networks, and conversely, power system engineers may lack expertise in building dynamics. Therefore, adopting simple, intuitive methodologies for quantifying flexibility [3] for energy-based systems such as phase change materials [4] and hydrogen or fuel cells [5] can greatly enhance mutual understanding and collaboration between the two communities.

1.2. Building thermal needs

A notable contribution towards the modeling of thermal needs and demand-side flexibility has been carried out in the context of the AMBIENCE project [6], which explores the activation of energy-efficient services in buildings through innovative business models and control mechanisms for active building energy performance contracting. A significant challenge in this domain lies in the limited availability of automated methodologies for generating detailed thermodynamic models and the lack of representative building stock data [7,8]. Projects such as TABULA [9], Hotmaps [10], and BuiltHub [11] have been developed to address the challenge of limited building stock data by

[☆] This article is part of a Special issue entitled: 'SESAAU2025' published in Smart Energy.

^{*} Corresponding author at: Department of Electrical & Computer Engineering, & Informatics, Cyprus University of Technology, Limassol, Cyprus.
E-mail address: savvas.panagi@cut.ac.cy (S. Panagi).

providing comprehensive, open-access datasets. These datasets support various applications, from national energy planning to local building-level simulations, and serve as foundational elements for developing and testing energy management methodologies.

Building thermal modeling methodologies are broadly categorized into three high-level approaches: (i) white-box, (ii) grey-box, and (iii) black-box models [7], each with distinct advantages and limitations. White-box models, based on detailed physical principles, are primarily utilized by simulation tools such as TRNSYS [12], EnergyPlus [13], and IDA ICE [14] to perform comprehensive multizone simulations of building energy flows. These models require extensive inputs, including precise specifications of building envelope materials, and detailed definitions of construction elements (e.g., windows, doors, roofs), considering their area and orientation [15]. While white-box models offer high accuracy due to their detailed physical foundations, they need a significant amount of data and computational resources, making them complex, building-specific, and time-intensive to develop [8,16]. Moreover, they lack adaptability, as generating a white-box model typically requires deep knowledge of the building's physical structure, making it impractical to derive such models for each building.

Black-box models leverage large datasets collected from a representative sample of buildings to infer thermal demand profiles. These models rely on statistical or machine learning (ML) techniques to capture complex patterns and dependencies among multiple input features, such as the approach used in [17]. A significant limitation of such models is their inability to describe the physical characteristics of a given building, such as construction materials, thermal mass, or specific usage and occupancy schedules. As a result, they provide generalized approximations that may not reflect the variability of real-world scenarios at the individual level. Moreover, the predictive performance is sensitive to the availability and quality of the training data. In many practical cases, especially in the residential sector, the necessary volume of historical data is either limited or unavailable, thereby constraining the applicability of purely black-box solutions. However, this type of method is simple during the operational phase; it does not require prior knowledge of the building's physical structure and can be easily implemented using widely available ML libraries [7].

Grey-box models provide a hybrid modeling approach that combines physical principles with mathematical identification techniques to represent the thermal dynamics of buildings. These models analyze a building's thermal behavior by representing its components, such as walls, windows, and air volumes, as thermal resistances and capacitances, analogous to an electrical RC network [18]. This structure enables the derivation of state-space equations governing temperature evolution, utilizing a relatively small set of input parameters. Positioned between the entirely physical white-box and the black-box approaches, grey-box models utilize physical knowledge to define the model structure, while estimating parameters using measurement data. This duality enables a balance between interpretability, accuracy, and computational efficiency. Grey-box models are an ideal choice for applications that require mathematical modeling, such as model predictive control (MPC) structures [8,16,19]. Its implementation will become easier in the near future, when smart meters are expected to be installed in almost all buildings. Compared to black-box models, grey-box methods require less training data and offer more physical insight. Compared to white-box models, they are simpler to implement, less data-intensive, and more adaptable to different building types [8].

However, grey-box models still present some drawbacks. The individual implementation requires tuning/identifying a new set of parameters for each building. Moreover, for forecasting, the model may be sensitive to the initial values, which are more observable in models with multiple nodes. Finally, the usage of the building or even its building parameters, such as insulation quality, may change over time; this could lead to an inaccurate model, which may need to be adjusted/retuned [16].

1.3. Grey-box models

Numerous implementations have been proposed in the literature to identify effective grey-box modeling approaches for building thermal behavior. These studies typically focus on a series of fundamental questions, including: (i) the selection of an appropriate model structure, (ii) the required quantity of training data (duration, frequency), (iii) the origin of this data—either from measurements or simulations, (iv) assumptions about building occupancy during training, (v) the specification of physical bounds for parameters to ensure plausibility and interoperability, and (vi) the best practices for the training process.

The literature has not converged on a standard grey-box configuration that can be applied across different building types and use cases. For instance, Harb et al. [7] evaluated four distinct grey-box structures and concluded that the 4R2C model provided the best trade-off between accuracy and complexity. This model exhibited superior dynamic performance by capturing temperature variations and transient responses during system modulation. Regarding the learning strategy, the GlobalSearch algorithm was employed to enhance the likelihood of finding the global optimum. During the learning process, if the optimal parameter hits the bounds, the simulation is repeated by relaxing the bound limits only for this parameter. Similar grey-box models have been investigated in [15], with the 6R2C model being considered the most suitable model.

Moreover, in [15], improvements related to occupancy heat gains and variable ventilation modeling were given. Another critical aspect of the study was the investigation of how different training data selection ranges influence forecast performance. Regarding the optimization procedure, an effort was made to guide the solution space by initializing physically meaningful values and applying appropriate bounds to each variable. However, no modeling or learning strategy was employed to ensure that the solution corresponds to a global optimum, meaning that the final results depend strongly on the solver behavior and the local convergence landscape. It is also worth noting that the 4R2C structure proposed in [15] differs significantly from that presented in [7]. The key difference is that the first 4R2C model explicitly models a single ambient node and also uses a supplementary resistance to characterize variable airflow ventilation. Whereas the second 4R2C model separates indoor air from an interior node area and uses two different but correlated ambient representations. This is a common place when comparing all proposed structures.

In [16], a passive house equipped with a building-integrated photovoltaic facade is modeled using grey-box approaches. The only deviation from conventional models is the separate treatment of one exterior wall, which is thermally decoupled due to elevated temperatures caused by solar gains behind the BIPV structure. In [20], a multi-room grey-box model was developed using a second-order model representation, where each room was divided into distinct thermal zones to capture both short-term and long-term temperature dynamics. In [8], a large-scale implementation of grey-box modeling was applied to a comprehensive set of 2245 representative buildings, including residential and non-residential types. This practical application demonstrates the scalability and flexibility of grey-box models in capturing the dynamic thermal behavior of the entire European building stock.

In [19], the authors focused on evaluating the performance of a parameter estimation algorithm for grey-box building thermal models. Through extensive Monte Carlo simulations, they demonstrated that the identified thermal parameters, R and C , exhibit high sensitivity to both the initial guess and the input dataset. As a result, the parameter estimates exhibit significant variability, which undermines their physical interpretability.

A systematic methodology for identifying suitable grey-box models for building thermal dynamics is proposed in [21], emphasizing a forward selection strategy that incrementally increases model complexity based on likelihood-ratio tests, ensuring both physical interpretability and statistical validation. In [22], the authors investigated the quality and robustness of grey-box building models as a function of the accuracy and type of input and observation signals, using detailed white-box simulations as a controlled reference.

Table 1
Overview of grey-box models, training durations, and data sources used in related studies.

Ref.	Models (Best Model)	Number of buildings	Training duration	Data source	Problem addressed
[15]	2nd–3rd order (6R2C)	1	14 days	White-box	Modeling level
[7]	1st–3rd order (4R2C)	3	26 days	Measurements	Modeling level
[16]	1st, 2nd, 4th order (6R4C)	1	12 days	Measurements	self-adapting model
[20]	2nd order	1	17 days	Measurements	Stochastic
[8]	2nd order	1	60 days	White-box	Modeling level
[19]	2nd order	1	8 days	Measurements	Physical interpretation
[23]	2nd order	1	5 days	Measurements	Modeling
[21]	1st–5th order (4R4C)	1	6 days	Measurements	Model selection strategy
[22]	1st–5th order	2	6 days	White-box	Data sensitivity

1.4. Limitations

Despite the progress, grey-box models still face several limitations that hinder their wider adoption. A significant issue is the lack of transparency, particularly regarding parameter estimation procedures, solar gain treatment, and the absence of open-source implementations. Many works omit critical details, such as the tuning algorithm, initialization strategy, and the associated computational cost of the calibration process, which can be substantial depending on the model complexity. The treatment of solar irradiance and heating inputs is often inconsistent across studies, with each model distributing gains differently among internal zones, the envelope, and internal contents.

To enhance clarity and emphasize the main limitations identified in the literature, these challenges can be summarized as follows:

1. Lack of transparency in parameter estimation procedures, including insufficient reporting of tuning algorithms, initialization strategies, and computational cost.
2. Inconsistent treatment of solar irradiance and heating inputs across different model structures.
3. Absence of standardized practices regarding the selection, duration, and granularity of training datasets.
4. Limited discussion on retraining frequency and long-term model robustness.
5. Lack of clear physical bounds and interoperability guidelines for estimated parameters.

These gaps affect both the reliability and generalization of grey-box models. The most common model structures, data sources, quantity of training data, and problems addressed by each work are summarized in Table 1.

1.5. Novelty and contributions

This paper addresses these challenges through a grey-box tuning algorithm that balances accuracy and flexibility. Moreover, it provides a sensitivity analysis for all critical decisions. The objective is to deliver a robust tuning and modeling framework that is suitable for integration into optimization-based control schemes. The proposed framework primarily focuses on practical deployment and can support power system planning, short-term scheduling, and operational decision-making in active distribution grids, particularly in the coordinated management of electrified heating systems.

The contributions of the paper are threefold:

- A systematic framework to model and forecast the thermal energy needs of individual buildings, based on validated and interpretable grey-box structures.

- Insights into the sensitivity of the models to different initial values, and a computational sensitivity analysis assessing how tuning complexity and parameter initialization affect model performance.
- An open-source implementation that enables stakeholders to efficiently use, adjust, and extend the proposed models for broader studies on electrification and grid integration.

The remainder of the paper is structured as follows. In Section 2, the modeling framework and mathematical modeling of building thermodynamics are given. The tuning process is discussed in Section 3. Section 4 provides a brief explanation of the case study. Then, in Section 5, the simulation results are discussed. Finally, conclusions are drawn in Section 6.

2. Thermal modeling framework

2.1. Electric analogy modeling

In grey-box modeling, the building is represented through a lumped-parameter thermal network inspired by analogies with electrical circuits. In such analogy, temperature corresponds to voltage, heat flow to electric current, thermal capacitance to electric capacitance (representing thermal mass), and thermal resistance to electrical resistance (representing insulation and heat transfer limitations).

The model structure is typically described using the $xRyC$ notation, where x and y denote the number of thermal resistances (R) and thermal capacitances (C), respectively. The number of states in the model defines its order, and it directly influences the trade-off between computational complexity and predictive accuracy. First-order, second-order, or higher-order models can be formulated depending on the desired fidelity and the availability of training data. These models are based on the following set of assumptions that users should be aware of when applying them in practice [7,15]:

1. Thermal resistance and capacitance values (R , C) are assumed constant over time, representing average building behavior.
2. Building components (walls, windows, etc.) are considered to have uniform material properties.
3. Most of the grey-box models consider a single thermal zone per building model. Also, multi-thermal zone grey box models are proposed; however, they are out of the subject of this paper.
4. Convective heat transfer is treated as a linear process.
5. Occupancy profiles are assumed to be uniform and constant.
6. Air infiltration is assumed constant unless explicitly modeled.

Capturing the thermal response of a building requires knowledge of four main variables: solar irradiance (Q_{irrad}), outdoor air temperature (T_o), heating power input from the HVAC system (Q_h), and the resulting indoor air temperature (T_{in}). In practice, three of these inputs are

sufficient to estimate the fourth. Depending on the application, the focus may be on forecasting indoor temperature based on external conditions and heating input or, conversely, estimating the required heating power to maintain thermal comfort (T_{in}).

2.2. Mathematical modeling

Building upon the electrical–thermal analogy (See Fig. A.6 in the Appendix for illustrative models), the thermal dynamics of a building zone are formulated using the principle of energy conservation [20]. The heat balance at each thermal node can be expressed as:

$$C_i \frac{dT_i}{dt} = \sum Q_{in} - \sum Q_{out} + d\omega_i$$

where C_i [J/K] is the thermal capacitance of node i , T_i [°C] is the temperature at node i , and Q [W] denotes thermal power exchange due to either active sources (e.g., heating systems, solar gains) or passive heat transfer through conduction. The conductive heat transfer between two thermal nodes is typically modeled by:

$$Q_w = \frac{1}{R_w} (T_j - T_i), \quad (1)$$

where R_w [K/W] is the thermal resistance of element w , and T_i , T_j are the temperatures on either side of the surface. In practice, building thermal behavior is affected by various uncertainties, including unknown internal heat gains, occupant-driven variability, and measurement noise. To account for these effects, a disturbance term $d\omega_i$ [W] is added to the heat balance. This stochastic representation, as discussed in [23], helps capture:

- Model simplifications and neglected physical phenomena,
- Unmeasured or time-varying internal gains (e.g., from occupants or appliances),
- Sensor and measurement noise.

The full set of ordinary differential equations (ODEs) is derived from the selected RC network structure, where each node corresponds to a thermal element in the building (e.g., indoor air, wall mass). These equations constitute the foundation for simulation and parameter estimation in the grey-box modeling framework. The continuous-time dynamics can be expressed in compact form as:

$$\frac{d\mathbf{T}}{dt} = A(\theta) \cdot \mathbf{T}(t) + B(\theta) \cdot \mathbf{U}(t) \quad (2)$$

where \mathbf{T} is the state vector and \mathbf{U} is the input vector. A and B represent the matrices of parameters, that are dependent on the grey-box model structure. For instance, in the 1R1C model, the state, input, and parameters vectors are defined as:

$$\mathbf{T}(t) = [T_{in}(t)], \quad \mathbf{U}(t) = [Q_h(t) \quad Q_s(t) \quad T_a(t)]^T \quad (3)$$

$$A(\theta) = \left[-\frac{1}{R_{in,a} \cdot C_{in}} \right], \quad B(\theta) = \left[\frac{1}{C_{in}} \quad \frac{A_{in}}{C_{in}} \quad \frac{1}{R_{in,a} \cdot C_{in}} \right] \quad (4)$$

To enable discrete-time simulation and optimization, the system is discretized using the forward Euler discretization:

$$\mathbf{T}(k+1) = A_d(\theta) \cdot \mathbf{T}(k) + B_d(\theta) \cdot \mathbf{U}(k) \quad (5)$$

with $A_d(\theta) = I + \Delta T \cdot A(\theta)$ and $B_d(\theta) = \Delta T \cdot B(\theta)$, where I is the identity matrix of appropriate dimension (matching the size of $A(\theta)$), and ΔT is the discrete time step used for the numerical simulation and control horizon. The Appendix presents a compact formulation of the ODEs used in all grey-box models, alongside a detailed description of their structural components and parameter definitions.

The heat input from solar irradiance and the heating system (Q_h , Q_s) constitutes the most critical heat transfer affecting the model behavior. In the modeling structure presented in this study, the solar irradiance heat flux is expressed as $A_i Q_s$, and the heating system heat flux as $f_i Q_h$, where i denotes node i of the grey-box model. The heating

system coefficients must satisfy the $\sum f_i = 1$, as the power inputs are distributed across the nodes of the modeling structure. The solar heat gain Q_s is defined as $Q_s = Irr$ [W/m²] $\cdot B_w$ [m²], where B_w is the effective window area and Irr is the measured solar irradiance. In most studies, the effective window area is considered to be approximately 60% of the measured window size [20]. However, in our study, B_w is treated as a tunable parameter, constrained only by logical bounds.

3. Tuning methodology

3.1. Optimization framework

The identification of the unknown physical parameters $\theta \in \Theta$ in the grey-box model is formulated as a nonlinear optimization problem because of bilinear equations. The objective is to minimize the deviation between the predicted indoor temperature $T_{in}(k)$ and the corresponding measured values $T_{in}^{meas}(k)$ over a time horizon of N steps which may correspond from a few days up to some weeks:

$$\min_{\theta \in \Theta} \frac{1}{N+1} \sum_{k=0}^N (T_{in}(k) - T_{in}^{meas}(k))^2 \quad (6)$$

$$\text{s.t. } \mathbf{T}(k+1) = A_d(\theta) \cdot \mathbf{T}(k) + B_d(\theta) \cdot \mathbf{U}(k), \quad \forall k \quad (7)$$

$$\theta_{min} \leq \theta \leq \theta_{max}, \quad \forall \theta \in \Theta, \quad \forall k \quad (8)$$

where:

- Θ is the vector of parameters to be tuned (e.g., thermal resistances and capacities),
- $[\theta_{min}, \theta_{max}]$ denotes the set of admissible physical bounds,
- $T_{in}(k) \in \mathbf{T}$ is the simulated temperature at time step k ,
- $T_{in}^{meas}(k)$ is the measured indoor air temperature.

The θ limits were selected based on building physics to ensure the chosen model correlates directly with the modeled building's physical attributes.

3.2. Tuning framework

The overall identification and validation process is illustrated in Fig. 1, which involves three main stages:

1. **Data Acquisition and Filtering:** Generate synthetic data through detailed simulation or obtain on-site data. These raw data are preprocessed (testing and validation) and resampled to a common frequency interval.
2. **Parameter tuning:** Optimal parameters θ^* are found by solving the proposed optimization problem. In each iteration, the initial values are randomly selected within the bounds.
3. **Validation:** The identified model is tested on unseen datasets to verify the performance. From the entire set of tests, the parameters of the most accurate run are selected.

Initially, the user selects the model structures for tuning. Then a repetitive tuning process is executed for a predefined number of iterations, denoted as n_{iter} . The model is initialized in each iteration with randomly selected parameter values within the specified bounds. After completing all iterations, the set of parameters that yielded the lowest mean squared error (MSE) in the validation dataset is retained as the optimal configuration for that model.

3.3. Metric

Performance metrics are used to compare the model's performance and assess its accuracy. The metric concerns the model's ability to predict indoor air temperature. This is quantified by the RMSE, which measures the disparity between simulated indoor temperatures and actual measurements as defined in Eq. (9).

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{t=1}^N (T_{in}[t] - \hat{T}_{in}[t])^2} \quad (9)$$

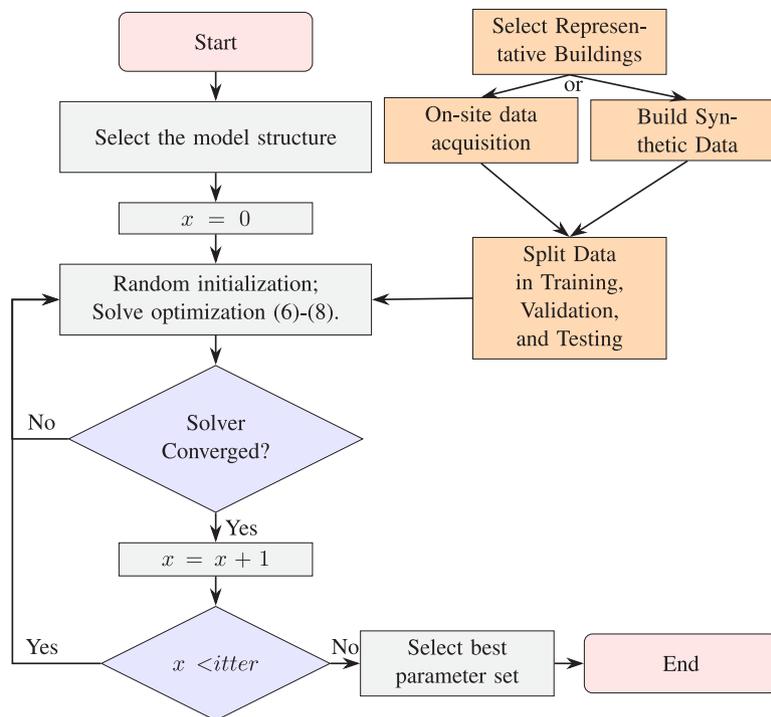


Fig. 1. Grey-box model tuning and validation framework.

Table 2
Overview of case study buildings.

Case Study	Location	Type	Floor area	Available Data
1	Trondheim	Residential	100 m ²	22 Days
2	Canada	Library	2100 m ²	60 Days
3 & 4	California	Office	57 m ²	20 Days

4. Case study buildings and input data

Measurement data collected from four buildings with varying usage types were initially used to validate and assess the robustness of the proposed modeling framework under real data. To ensure a comprehensive evaluation, the buildings were selected to represent a range of structural characteristics, including building size (57–2100 m²), and diverse climate conditions across different locations (Norway, Canada, United States). The measurement data used in this study are available online at Mendeley Data [24]. A short description of the building characteristics is given in Table 2.

4.1. Buildings

Case Study 1 (CS1) is a zero-emission single-family house, located on the NTNU campus in Trondheim, which is a 100 m² single-family house. Constructed with a wooden frame, the building is highly insulated, featuring 35–40 cm of mineral wool, and maintains a window-to-wall ratio of about 20%. The layout is divided into four zones: two bedrooms, a bathroom, and a living area.

Case Study 2 (CS2) is the first Net-Zero Energy Library in Canada, built in 2014. Situated in Varennes, 30 km northeast of Montreal within ASHRAE Climate Zone 6, this two-story building encompasses approximately 2100 m² of floor space. It features a sizable basement housing the mechanical room and a mezzanine dividing the two levels. The two floors are mainly open-plan areas designated for library use, with a few closed office spaces on the first floor. A 110.5 kWp photovoltaic array, integrated into the building's architecture, is mounted on the south-facing roof.

Case Studies 3 and 4 (CS3, CS4), named FLEXLAB, located at Lawrence Berkeley National Laboratory in Berkeley, California, are well-equipped experimental test centers. The facility comprises four testbeds, each with two identical rooms to enable side-by-side comparisons of a new technology versus a baseline scenario. The data was gathered from the testbed “XR” which includes two identical adjoining rooms, Room Cell A and Room Cell B, each simulating a small commercial office space with a floor area of 57 m² and featuring a large south-facing window.

4.2. Thermal data

The heating requirements are relative to all zones, with the building's average air temperature calculated by weighting each zone's temperature against its volume. Moreover, each simulation study case obtained heat consumption, outdoor air temperature, and global solar irradiance on a horizontal surface. Obtaining solar irradiance data is challenging. To simplify model implementation, it is best to avoid on-site measurements and instead use global horizontal irradiance data from nearby meteorological stations. Bathrooms are excluded from the calculations due to the very high temperature fluctuations caused by hot-water use.

5. Results and discussions

The optimization problem is formulated as a nonlinear programming problem (due to bilinear terms) and solved using the IPOPT solver, a primal–dual interior-point method well-suited to nonlinear programming. The implementation used a laptop equipped with an M4 Pro Max chip, featuring a 12-core CPU and 36 GB of unified

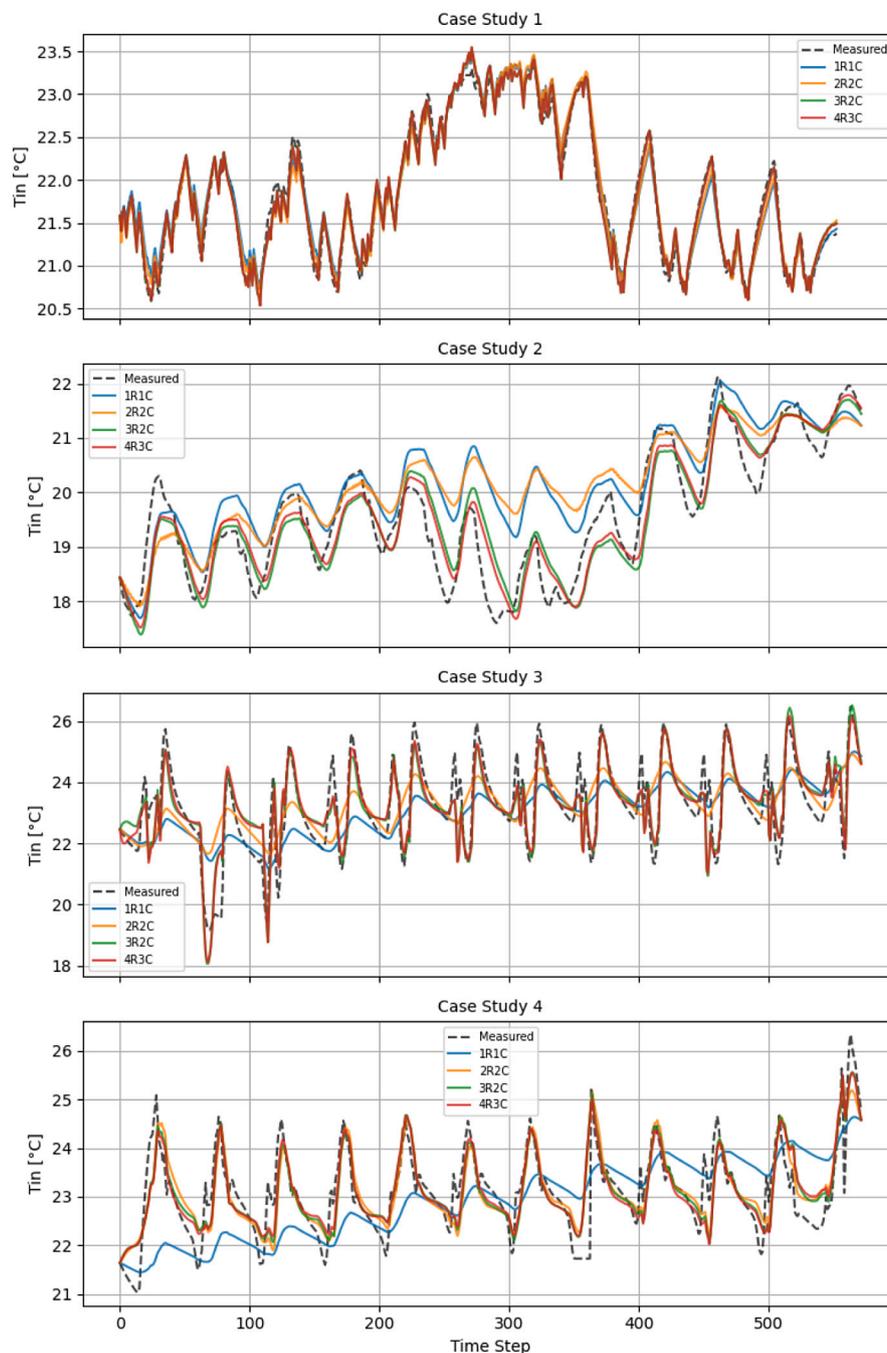


Fig. 2. Comparison of measured and simulated indoor temperatures for all grey-box models across the four case studies.

memory. The solver successfully converged to a feasible solution in almost all scenarios tested. A 30-minute temporal resolution is used for all studies. The modeling and optimization framework is implemented in an open-source Python-based toolkit, publicly available on GitHub.¹

5.1. Tuning performance evaluation

To assess the accuracy of different models and the computational time, real-world data from case studies was utilized. The first 12 days of each dataset are used for model training, the next 5 days for validation, and the last 3 days are used for testing purposes. As the methodology

suggests that repeated runs are required to ensure convergence of the optimization process, 30 iterations were chosen for each test case in this analysis. Fig. 2 shows a comparison of measured and simulated indoor temperatures across all grey-box model structures and case studies for the training dataset. The parameters after the tuning are presented in Table 4. The performance metric, RMSE, is summarized in Table 3, for training, validation, and testing phases.

The 1R1C model shows clear limitations in capturing indoor temperature dynamics in some test cases. While it performs reasonably well under stable conditions (e.g., case studies 1 and 2), it struggles in more dynamic cases (e.g., case studies 3 and 4), where it underestimates peak values and produces overly damped responses. This is due to its simple structure, which includes only one thermal resistance and one capacitance, capturing just a single dominant time constant. As a result,

¹ <https://github.com/SPS-L/HeatPump-FlexModeling-Toolkit>

Table 3

RMSE metrics for different grey-box thermal models during training, validation, and testing phases. “CS” denotes the case study identifier used for each building scenario.

	Training RMSE [K]				Validation RMSE [K]				Testing RMSE [K]			
	1R1C	2R2C	3R2C	4R3C	1R1C	2R2C	3R2C	4R3C	1R1C	2R2C	3R2C	4R3C
CS 1	0.15	0.14	0.10	0.10	0.28	0.22	0.18	0.18	0.36	0.32	0.33	0.31
CS 2	0.87	0.93	0.51	0.44	1.34	1.42	0.86	0.48	0.34	0.37	0.36	0.74
CS 3	1.16	1.12	0.65	0.62	1.06	0.97	0.86	0.76	0.67	0.58	0.49	0.48
CS 4	0.93	0.51	0.48	0.48	1.35	0.60	0.52	0.53	0.85	0.64	0.56	0.69

Table 4

Summary of results across varying grey-box models and buildings for the training dataset. Different buildings denoted with Case Study 1 / Case Study 2 / Case Study 3 / Case Study 4.

Variable	1R1C	2R2C	3R2C	4R3C
$R_{int,in}$	–	–	–	0.016 / 0.025 / 0.004 / 0.003
$R_{in,e}$	–	0.048 / 0.0005 / 0.0016 / 0.023	–	0.06 / 0.002 / 0.093 / 0.0027
$R_{in,a}$	0.013 / 0.00043 / 0.012 / 0.072	–	–	0.049 / 0.002 / 0.039 / 0.035
$R_{e,a}$	–	0.0014 / 0.0005 / 0.001 / 0.0051	–	0.0029 / 1 / 0.0097 / 0.019
C_{int}	–	–	–	4.2e6 / 2.7e8 / 1.2e6 / 2.6e6
C_{in}	2.3e7 / 2.5e9 / 5.1e7 / 2.5e7	6.3e6 / 1.3e9 / 1.1e8 / 1.1e6	–	4.2e6 / 2.7e8 / 1.2e6 / 2.6e6
C_e	–	1e6 / 3.9e7 / 4.5e8 / 3.3e8	–	2.7e7 / 5.0e8 / 1.2e6 / 2.8e6
A_{int}	–	–	–	0.00 / 0.00 / 1.60 / 2.35
A_{in}	3.47 / 630.00 / 5.35 / 3.06	0.39 / 226.05 / 14.51 / 0.37	–	0.86 / 76.07 / 1.02 / 1.60
A_e	–	12.37 / 0.00 / 0.00 / 8.55	–	3.59 / 0.00 / 3.87 / 0.00
$f_{h,int}$	–	–	–	0.17 / 0.81 / 0.54 / 0.35
$f_{h,in}$	1 / 1 / 1 / 1	0.26 / 0.00 / 1.00 / 0.13	–	0.22 / 0.13 / 0.46 / 0.49
$f_{h,e}$	–	0.74 / 1 / 0 / 0.87	–	0.61 / 0.06 / 0.00 / 0.16

Table 5

Summary of tuning sensitivity analysis in terms of random initializations.

Metric	1R1C	2R2C	3R2C	4R3C
No. of Optimal Solutions	43	7	9	2
Mean RMSE	0.28	0.32	0.12	0.13
RMSE Std. Dev.	0.24	0.21	0.15	0.15

it cannot differentiate between fast internal gains and slower thermal effects through the building envelope, leading to poor alignment with measured data.

In contrast, the model’s performance increases as the complexity of the model rises. The 3R2C model shows some improvement, which is mainly caused by the extra ventilation parameter, making the 3R2C model slightly better than the 2R2C. The 4R3C structure demonstrates the highest performance among all models evaluated during the training and validation phase. However, in the testing phase, it does not present significant improvement over the other models, and it was even worse in CS2 and CS4.

Therefore, we cannot clearly distinguish a universally superior model, as each dataset exhibits different sensitivities under the various models. However, we can confidently conclude that the most stable model is the 3R2C. It consistently outperformed the others, most likely due to its higher complexity compared to the 1R1C and 2R2C models, allowing it to capture both fast and slow thermal dynamics, while still maintaining lower complexity than the 4R3C model. The latter’s higher degrees of freedom make it more difficult to tune effectively.

5.2. Computational analysis

The boxplots in Fig. 3 illustrate the expected trend in computational time across the different thermal models. The reported values correspond to the computation time for each iteration. The increasing model complexity (from 1R1C to 4R3C) generally leads to longer computation times. This is evident in most case studies, especially in CS3 and CS4, where the 4R3C model exhibits significantly higher computation times compared to simpler models, as expected due to its increased number of parameters and degrees of freedom.

High computational times are observed for some iterations (longer outliers), and this becomes more evident as the complexity of the

model increases. This behavior can be attributed to the fact that some repetitions required significantly more iterations to converge because of poor initialization, which delays convergence, distorting the expected performance. As a result, the computational time for a specific iteration can appear higher than the median value, even though its typical convergence time remains lower.

5.3. Tuning sensitivity to initial parameter values

To systematically evaluate the sensitivity of the parameter tuning process to random initialization, we conduct 100 independent optimization runs using different initial parameter values. Sensitivity is quantified using three indicators: (i) the variability of the final RMSE across runs, (ii) the frequency of convergence to the best-performing solution, and (iii) the dispersion of the optimized parameter values. Case study 1 was used for this sensitivity analysis; however, the results were very similar across all test cases. As before, 17 days are used for model training and tuning. We assess the impact of initialization variability on the final model performance. The RMSE and the corresponding optimized parameter values are recorded for every trial. The summary key results presented in Table 5 underscore the critical influence of initialization on the training of grey-box models.

The 1R1C model includes three degrees of freedom: $R_{in,a}$, C_{in} , and A_{in} . As shown in Fig. 4, the model exhibits low sensitivity to the initial values of R_{in} , which consistently converge to approximately 0.013. In contrast, A_{in} and C_{in} are identified as the most sensitive parameters. The optimized value of C_{in} varies between $2e^7$ and $4e^8$, depending on the initialization. Despite this variation, the optimization process converges towards the same optimal solution across 43% of the trials. In the remaining 57% of the cases, slightly elevated RMSE values were observed. Therefore, even if the optimum solution cannot be reached, good model performance can still be gained.

In contrast, the 2R2C model introduces seven degrees of freedom, significantly increasing the nonlinearity and complexity of the optimization problem. As shown in Fig. 5, this results in a greater tendency to converge to suboptimal solutions, with multiple local minima observed across different initializations. Interestingly, the tuning variables do not converge to a single parameter value, indicating the difficulty in plausibly describing the building characteristics.

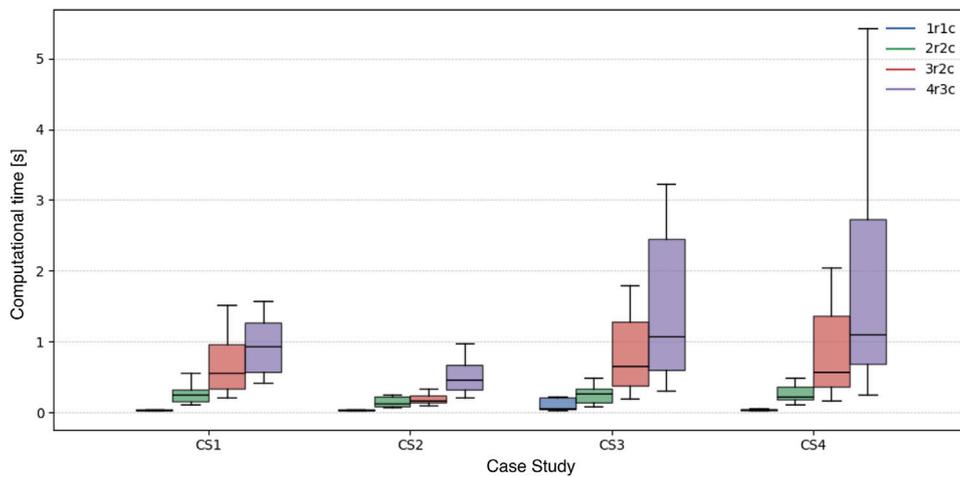


Fig. 3. Computational time for each thermal model across all case studies. Each boxplot represents the optimization time over 30 iterations per tuning process.

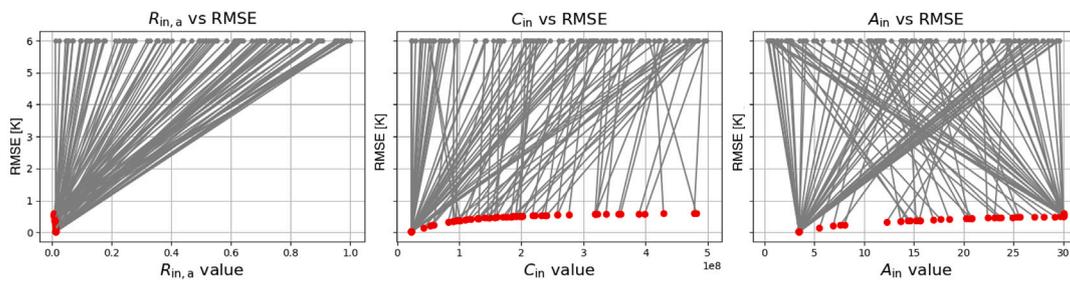


Fig. 4. Parameter tuning values and training RMSE for the 1R1C model. Each line connects the initial randomized value (grey marker) to the final optimized value (red marker).

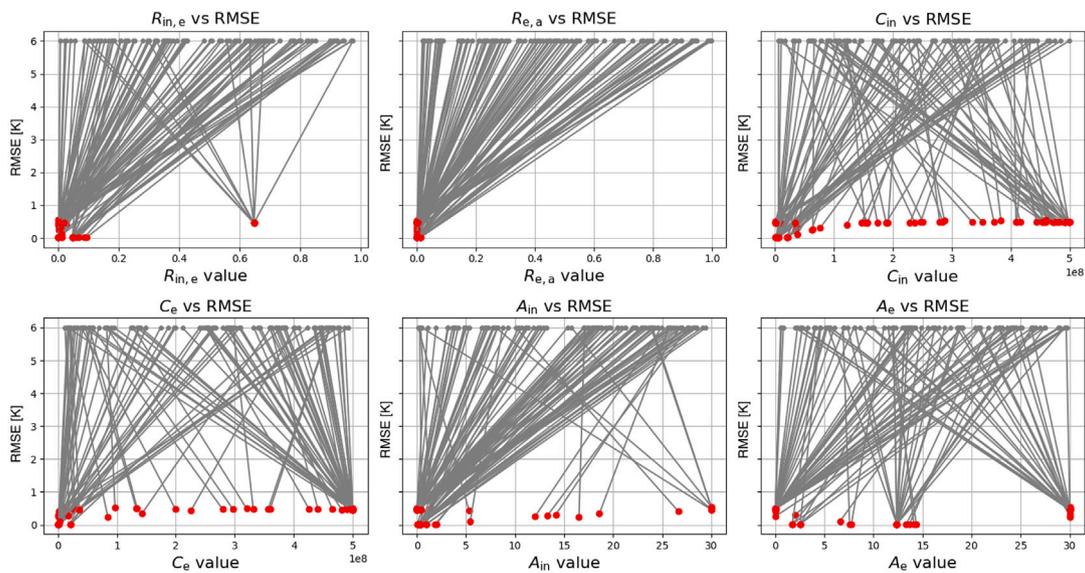


Fig. 5. Parameter tuning values and training RMSE for the 2R2C model. Each line connects the initial randomized value (blue marker) to the final optimized value (red marker).

The 3R2C model, although more complex than the previous two structures, as shown in Table 5, nearly all initializations converge to a similar optimal value. As a result, the standard deviation of the RMSE values is reduced, indicating highly stable and consistent performance across all trials. This suggests that the additional state may better capture the dominant thermal dynamics of the building, effectively

reducing model bias without excessively increasing parameter redundancy, validating the improved results of this model as explained in Section 5.1.

For the 4R3C model, although convergence towards similar RMSE levels is frequently observed, a noticeable dispersion in the optimized parameter values remains. This indicates that the increased number of

degrees of freedom introduces parameter redundancy, allowing multiple parameter configurations to produce similar model performance. While predictive accuracy is preserved, the interpretability and physical reliability of individual parameter estimates may be compromised.

Finally, the number of times reaching the optimal solution is significantly reduced for 2R2C, 3R2C, and 4R3C models compared to the 1R1C model. These results emphasize the necessity of employing multiple training runs with randomized initializations to enhance parameter estimation's reliability and reduce the likelihood of entrapment in local minima.

6. Conclusions

In conclusion, this work presents a systematic methodology for identifying and tuning grey-box thermal models. While the 4R3C model achieved the highest accuracy during the training and validation phases, its performance during the testing phase was less encouraging. Additionally, its computational cost increases exponentially, making it unsuitable for large-scale implementations. Simpler models, such as the 1R1C, demonstrated high computational efficiency but failed to capture the dynamic thermal behavior under varying external conditions. Overall, the 3R2C model emerged as the most suitable option, offering a balanced trade-off between accuracy and computation time.

Moreover, the tuning process was found to be sensitive to the initial parameter values, indicating the presence of multiple local minima in the solution space. Repeating the optimization with randomized initializations proved essential to ensure robustness and to avoid convergence to suboptimal solutions. While the primary advantage of grey-box models lies in maintaining the physical interpretability of model parameters, the experimental results indicate that this is not always guaranteed. The final solution depends heavily on the solver's algorithm and the input data, which may limit interpretability when convergence to the global optimum is not achieved.

Future work will adopt the proposed framework to optimally operate heat pumps in active distribution grids. More specifically, the proposed linear model equations of any model structure, together with the tuning algorithm, could be integrated within optimization frameworks and combined with different heating and cooling systems. By associating these models with the corresponding network elements, the operator could properly control such systems to support operational decision-making and ensure the safe and optimal operation of electrical network grids.

CRedit authorship contribution statement

Savvas Panagi: Writing – original draft, Visualization, Software, Methodology, Investigation, Data curation, Conceptualization. **Lazaros Aresti:** Writing – review & editing, Methodology, Formal analysis. **Paul Christodoulides:** Writing – review & editing, Supervision, Project administration, Conceptualization. **Chrysovalantis Spanias:** Writing – review & editing, Supervision, Methodology. **Petros Aristidou:** Writing – review & editing, Validation, Supervision, Project administration, Methodology, Conceptualization.

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.

Acknowledgments

This project has received funding from the European Union's Framework Programme for Research and Innovation *Horizon Europe* (HORIZON) under the Marie Skłodowska-Curie Actions Doctoral Networks (MSCA-DN), Grant Agreement No. 101120278 - DENSE.

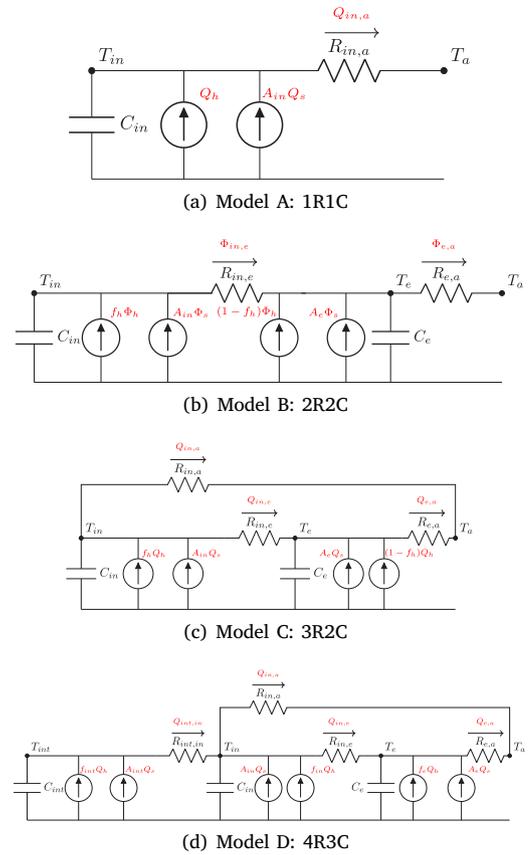


Fig. A.6. Thermal RC network representations of the four grey-box model structures evaluated in this study.

Appendix. Model structures

A.1. Model 1R1C – single-node

The most basic configuration is the 1R1C model, shown in Fig. A.6(a), and widely used in early-stage estimation studies [7,8,16]. It consists of a single thermal node with capacitance C_i representing the thermal capacity of the whole building's indoor environment. This node is connected to the external environment through a single thermal resistance $R_{in,ext}$; a higher resistance means better insulation. Heat inputs from the heating system (\dot{Q}_h) and solar irradiance (\dot{Q}_{sun}) are applied directly to the internal air node. This model provides a lumped estimation of building thermal dynamics, with minimal computational burden. The product $A_{in}Q_s$ is a key component of this model structure. According to [7], A_{in} is a tunable parameter typically ranging from 0 to 0.25, while [8] defined this variable without further explaining how this parameter is estimated, tuned, or bounded.

$$C_{in} \frac{dT_{in}}{dt} = \frac{T_a - T_{in}}{R_{in,a}} + Q_h + A_{in}Q_s \quad (A.1)$$

A.2. Model 2R2C – envelope split

A 2R2C model structure — illustrated in Fig. A.6(b) — represents the thermal mass of the building envelope (T_e), distinguishing it from the indoor air temperature T_{in} [8]. The key difference in the model presented here compared with the proposed one by [8] is that the thermal input powers (Q_s, Q_h) were modeled in both indoor and envelope temperatures, with the corresponding tuning factor instead of

only indoor input thermal powers.

$$C_{in} \frac{dT_{in}}{dt} = \frac{T_e - T_{in}}{R_{in,e}} + f_h Q_h + A_{in} Q_s \quad (\text{A.2})$$

$$C_e \frac{dT_e}{dt} = \frac{T_{in} - T_e}{R_{in,e}} + \frac{T_a - T_e}{R_{a,e}} + (1 - f_h) Q_h + A_e Q_s \quad (\text{A.3})$$

A.3. Model 3R2C – ventilation-aware

The 3R2C model — depicted in Fig. A.6(c) — introduces an additional resistance to represent heat transfer due to ventilation, distinguishing it from internal conduction ($R_{in,e}$) and external envelope resistance ($R_{e,a}$). This model extends the 2R2C structure, enabling more accurate predictions in scenarios with constant airflow rates. This model is also adopted from [15,19]. Again, the heat input flux varies between models in the literature. Here, a superset relies on the tuning process to define which is more important.

$$C_{in} \frac{dT_{in}}{dt} = \frac{T_e - T_{in}}{R_{in,e}} + \frac{T_a - T_{in}}{R_{in,a}} + f_h Q_h + A_{in} Q_s \quad (\text{A.4})$$

$$C_e \frac{dT_e}{dt} = \frac{T_{in} - T_e}{R_{in,e}} + \frac{T_a - T_e}{R_{e,a}} + (1 - f_h) Q_h + A_e Q_s \quad (\text{A.5})$$

A.4. Model 4R3C – four-node structure

The 4R3C model illustrated in Fig. A.6(d) extends the previously mentioned 3R2C model by introducing a second thermal node to represent internal contents, such as furniture or indoor walls, improving the accuracy of capturing thermal inertia. The internal components temperature, the indoor air temperature, and the temperature of the envelope are included. The heat flux from the heating system is partly transferred to each node as described in (A.9).

$$C_{int} \frac{dT_{int}}{dt} = \frac{T_{in} - T_{int}}{R_{int,in}} + f_{h,int} Q_h + A_{int} Q_s \quad (\text{A.6})$$

$$C_{in} \frac{dT_{in}}{dt} = \frac{T_e - T_{in}}{R_{in,e}} + \frac{T_a - T_{in}}{R_{in,a}} + \frac{T_{int} - T_{in}}{R_{int,in}} + f_{h,in} Q_h + A_{in} Q_s \quad (\text{A.7})$$

$$C_e \frac{dT_e}{dt} = \frac{T_a - T_e}{R_{e,a}} + \frac{T_{in} - T_e}{R_{in,e}} + f_{h,e} Q_h + A_e Q_s \quad (\text{A.8})$$

$$f_{h,int} + f_{h,in} + f_{h,e} = 1 \quad (\text{A.9})$$

Data availability

Throughout the manuscript, references to the data and code are provided via the corresponding links and the GitHub repository.

References

- [1] Gaur AS, Fitiwi DZ, Curtis J. Heat pumps and our low-carbon future: A comprehensive review. *Energy Res Soc Sci* 2021;71:101764.
- [2] European Commission. Heat pumps. 2023, https://energy.ec.europa.eu/topics/energy-efficiency/heat-pumps_en. [Accessed 11 April 2025].
- [3] Panagi S, Spanias C, Aristidou P. A thermal-electrical co-optimization framework for active distribution grids with electric vehicles and heat pumps. 2026, arXiv preprint arXiv:2601.21651.
- [4] Emami A, Chitsaz A, Nouri A, Hasanazadeh A. Smart demand-side management with paraffin-based phase change materials for an energy flexible and decarbonized building in London. *Appl Therm Eng* 2026;292:129815.
- [5] Emami A, Chitsaz A, Nouri A. IoT-driven smart energy management with a closed PEMFC-PEMEC loop: A sustainable approach to decarbonizing flexible buildings in London. *Smart Energy* 2025;19:100191.
- [6] AMBIENCE Consortium. AMBIENCE Project: Active managed buildings with energy performance contracting. 2019, <https://www.ambience-project.eu/>. Funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 847054.
- [7] Harb H, Boyanov N, Hernandez L, Streblov R, Müller D. Development and validation of grey-box models for forecasting the thermal response of occupied buildings. *Energy Build* 2016;117:199–207.
- [8] Jankovic I, Fernandez X, Diriken J. Database of grey-box model parameter values for EU building typologies. *Ambience EU Proj Rep* 2021.
- [9] TABULA Project Team. TABULA: Typology approach for building stock energy assessment. 2012, <https://episcopes.eu/welcome/>. Funded by the European Union's Intelligent Energy Europe programme.
- [10] Hotmaps Project Consortium. Hotmaps: The open source mapping and planning tool for heating and cooling. 2020, <https://www.hotmaps-project.eu/>. Funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 723677.
- [11] BuiltHub Consortium. BuiltHub: Building stock data flow hub. 2020, <https://builthub.eu/>. Funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 957026.
- [12] University of Wisconsin-Madison. TRNSYS: Transient system simulation tool. 2023, <https://www.trnsys.com/>. [Accessed 11 April 2025].
- [13] US Department of Energy. EnergyPlus energy simulation software. 2023, <https://energyplus.net/>. [Accessed 11 April 2025].
- [14] EQUA Simulation AB. IDA indoor climate and energy (IDA ice). 2023, <https://www.equa.se/en/ida-ice>. [Accessed 11 April 2025].
- [15] Berthou T, Stabat P, Salvazet R, Marchio D. Development and validation of a gray box model to predict thermal behavior of occupied office buildings. *Energy Build* 2014;74:91–100.
- [16] Fux SF, Ashouri A, Benz MJ, Guzzella L. EKF based self-adaptive thermal model for a passive house. *Energy Build* 2014;68:811–7.
- [17] Navarro-Espinoso A, Mancarella P. Probabilistic modeling and assessment of the impact of electric heat pumps on low voltage distribution networks. *Appl Energy* 2014;127:249–66.
- [18] Déqué F, Ollivier F, Poblador A. Grey boxes used to represent buildings with a minimum number of geometric and thermal parameters. *Energy Build* 2000;31(1):29–35.
- [19] Brastein OM, Perera DWU, Pfeifer C, Skeie N-O. Parameter estimation for grey-box models of building thermal behaviour. *Energy Build* 2018;169:58–68.
- [20] Andersen KK, Madsen H, Hansen LH. Modelling the heat dynamics of a building using stochastic differential equations. *Energy Build* 2000;31(1):13–24.
- [21] Bacher P, Madsen H. Identifying suitable models for the heat dynamics of buildings. *Energy Build* 2011;43(7):1511–22.
- [22] Reynders G, Diriken J, Saelens D. Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals. *Energy Build* 2014;82:263–74.
- [23] Madsen H, Holst J. Estimation of continuous-time models for the heat dynamics of a building. *Energy Build* 1995;22(1):67–79.
- [24] Sartori I, Walnum HT, Skeie KS, Georges L, Knudsen MD, Bacher P, Candanedo J, Sigounis A-M, Prakash AK, Pritoni M, et al. Sub-hourly measurement datasets from 6 real buildings: Energy use and indoor climate. *Data Brief* 2023;48:109149.