



MOX-Report No. 41/2026

**Physics-constrained identification of graph-based thermal networks
for spacecraft digital twins**

Sosta, L.; Ciancarelli, C; Marini, L.; Pagani, S.; Regazzoni, F.; Parolini, N.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Physics-constrained identification of graph-based thermal networks for spacecraft digital twins

Luca Sosta^a, Carlo Ciancarelli^b, Leonardo Marini^b, Stefano Pagani^a, Francesco Regazzoni^a, Nicola Parolini^a

^a*MOX-Dipartimento di Matematica, Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy*

^b*Thales Alenia Space Italia, Via Saccomuro 24, Rome, Italy*

Abstract

Reconstructing a thermal model capable of efficiently simulating the behavior of a spacecraft from sparse and localized temperature measurements remains a challenging task. To address this, we introduce a physically-constrained calibration framework for Lumped Parameter Thermal Models (LPTMs), formulated as a trajectory-based inverse problem for graph dynamical systems. The model reconstructs thermal dynamics directly from temperature measurements and known inputs, without relying on a priori parameter values derived from material properties or geometric assumptions. Physical admissibility is enforced at the parameterization level: positivity of nodal coefficients and symmetry of conductive interactions are imposed by construction. This guarantees stable dynamics and restricts the identification problem to a physically meaningful parameter space, improving conditioning without the need of additional regularization. The identification problem is addressed through trajectory matching, ensuring stable rollout over extended time horizons. The methodology is validated on synthetic datasets generated from high-fidelity finite element simulations under progressively complex forcing conditions. The calibrated LPTMs accurately reproduce long-term temperature evolution and exhibit robustness to measurement noise. The proposed framework provides a systematic approach to the calibration of reduced-order thermal models by combining physical structure with data-driven identification. The numerical results show a favorable balance between accuracy and computational efficiency, making the models suitable for integration in spacecraft thermal Digital Twin applications.

Keywords: Inverse problems, Graph dynamical systems, Lumped thermal models, Structure-preserving methods, Continuous-time system identification, Reduced-order modeling, Spacecraft Thermal Modeling, Digital Twin

1. Introduction

Thermal modeling of spacecraft involves the prediction of energy transfer in heterogeneous environments governed by the heat equation and radiative exchange mechanisms [1, 2, 3]. In the absence of atmospheric convection, satellites are subject to strongly time-dependent boundary conditions, including solar irradiation, albedo, and eclipse phases, leading to nonlinear and non-stationary temperature dynamics. As a result, inadequate thermal management under such conditions may compromise subsystem functionality, reduce mission lifetime, or ultimately lead to system failure.

The Digital Twin (DT) paradigm, initially developed in the aerospace sector to support real-time monitoring and predictive maintenance of complex systems [4], aims to combine physics-based models with real-time data for monitoring, prediction, and control. In thermal applications, a DT must remain accurate over long prediction horizons while retaining computational efficiency compatible with real-time deployment and robustness across varying operating conditions.

High-fidelity models based on the finite element method (FEM) provide accurate solutions but are computationally expensive and unsuitable for real-time DT applications. Reduced-order modeling techniques offer a systematic alternative by approximating high-dimensional systems through low-dimensional representations while preserving essential dynamical properties. In this context, projection-based model order reduction techniques derived from FEM discretizations have been proposed to accelerate large-scale simulations while preserving physical fidelity [5, 6, 7, 8]. Lumped Parameter Thermal Models (LPTMs) can be also interpreted as reduced-order models of the underlying heat equation, where the thermal field is approximated by a network of interacting nodes. From a structural perspective,

LPTMs can be interpreted as diffusive dynamical systems defined on graphs, where thermal interactions are governed by Laplacian operators. Such systems belong to the broader class of network dynamical systems, extensively studied in control theory and multi-agent systems [9]. However, unlike projection-based methods, their parameters usually do not directly inherit physical meaning (except in very simple settings, with homogeneous materials, simple geometries, and clearly defined heat transfer paths) and must be inferred from data [10].

Data-driven approaches, including neural networks and operator learning methods [11], have been proposed to approximate complex dynamical systems at reduced computational cost. In recent years, physics-informed machine learning has expanded the range of possible approaches by embedding physical structure into learning architectures to enhance robustness and generalization in dynamical systems [12, 13]. In parallel, system identification techniques for continuous-time dynamics have been revisited through sparse regression methods such as SINDy [14] and Neural Ordinary Differential Equations (Neural ODEs) [15]. Hybrid approaches, such as Universal Differential Equations (UDEs), further combine mechanistic models with data-driven components, enabling the integration of prior physical knowledge within the learning process [16]. Structure-preserving approaches such as Hamiltonian neural networks have also been proposed to enforce physical invariants in learned models [17]. Among these, reduced representations such as Dynamic Mode Decomposition (DMD) have been successfully applied to approximate complex dynamical systems at lower computational cost [18, 19]. Despite their flexibility, these approaches often suffer from poor extrapolation capabilities, lack of physical interpretability, and instability over long prediction horizons, with stability preservation becoming particularly critical in data-scarce regimes [20].

These limitations highlight a gap between physically interpretable reduced-order models and flexible data-driven approaches. In particular, the identification of reduced thermal models from sparse measurements remains an ill-posed inverse problem, due to limited observability and non-uniqueness of the parameters [21, 22].

In this work, we propose a physics-constrained identification framework for lumped thermal models formulated as graph dynamical systems. The approach combines: (i) a graph-based representation of thermal interactions, (ii) a continuous-time system identification formulation inspired to Neural ODEs, (iii) a structure-preserving parameterization enforcing positivity, symmetry, and dissipativity by construction. This formulation ensures that the identified model remains within the class of physically admissible thermal systems, improving stability and generalization without requiring explicit regularization.

The remainder of this paper is organized as follows. Section 2 introduces the theoretical background, including the LPTM formulation, its graph representation, and the connection with Neural ODE models. Section 3 presents the inverse calibration framework and the associated learning problem. Section 4 describes the computational architecture and the enforcement of physical constraints. Section 5 reports numerical results on synthetic test cases, including sensitivity analyses and robustness assessments. Section 6 concludes the paper and outlines future research directions.

2. Background and Problem Formulation

We consider a thermal subsystem equipped with a finite number of temperature probes placed at specific locations within the structure. These sensors provide temperature measurements over time and represent the primary information available for reconstructing the thermal dynamics of the system. Let $\{t_k\}_{k=0}^N$ denote the discrete sampling times. At each time instant, the probes provide a vector of temperature measurements $y_k \in \mathbb{R}^p$, where p is the number of probes. Internal heat sources are assumed to act at the same probing points. These sources account for heater activations or thermal dissipation generated by electronic components, whose power levels are typically known from telemetry. Their values are collected in the input vector $u_k \in \mathbb{R}^p$ and the available dataset can be written as $\mathcal{D} = \{u_k, y_k\}_{k=0}^N$. External thermal interactions, such as radiative exchanges with the surrounding environment, are not directly included in this description and will be introduced separately in the following sections.

We distinguish between the number of measured outputs p and the number of thermal states n used in the model. In general, the graph-based thermal model may include additional nodes beyond sensor locations, corresponding to lumped components or internal thermal masses. Therefore, the state dimension n does not necessarily coincide with the number of measurements p . The observation operator is implicitly assumed to map the full state $T(t) \in \mathbb{R}^n$ to the measured outputs $y(t) \in \mathbb{R}^p$. In the numerical experiments presented in this work, we consider the case $n = p$, that is, one state per sensor, for simplicity.

2.1. Lumped Parameter Thermal Models (LPTMs)

LPTMs are reduced-order models that approximate heat transfer problems by aggregating the domain into a finite set of control volumes (nodes) connected by lumped thermal couplings (edges). The evolution of the temperature T in a domain $\Omega \subset \mathbb{R}^3$ is described by the heat equation arising from conservation of energy and Fourier's law of conduction [1, 2]

$$\rho c_p \frac{\partial T}{\partial t} = p + \nabla \cdot (\lambda \nabla T),$$

where ρ denotes the mass density, c_p is the volume specific heat at constant pressure, p represents any volumetric heat source and λ is the conductivity. The domain Ω is partitioned into disjoint control volumes Ω_i , such that $\Omega = \bigcup_{i=1}^n \Omega_i$. The node temperature is defined as the volume average over a control volume, namely

$$T_i(t) := \frac{1}{|\Omega_i|} \int_{\Omega_i} T(\mathbf{x}, t) dV,$$

while the lumped thermal capacitance is defined as

$$C_i := \int_{\Omega_i} \rho(\mathbf{x}) c_p(\mathbf{x}) dV.$$

Let $\Gamma_{ij} := \partial\Omega_i \cap \partial\Omega_j$ be the interface between nodes i and j , with outward normal \mathbf{n}_{ij} from Ω_i to Ω_j . The net conductive heat flow from i to j is

$$\Phi_{ij}(t) = \int_{\Gamma_{ij}} -\lambda(\mathbf{x}) \nabla T(\mathbf{x}, t) \cdot \mathbf{n}_{ij} dS.$$

Under the usual lumped approximation, this interfacial flow is approximated as

$$\Phi_{ij}(t) \approx G_{ij} (T_i(t) - T_j(t)),$$

where $G_{ij} = 1/R_{ij} > 0$ is an effective conductance that aggregates geometry, materials, and contact properties. External heat inputs acting on node i (heaters, absorbed radiation) are collected in $P_i(t)$. The nodal energy balance becomes

$$C_i \frac{dT_i}{dt} = P_i(t) + \sum_{j \in \mathcal{N}(i)} G_{ij} (T_j - T_i), \quad i = 1, \dots, n,$$

where $\mathcal{N}(i)$ is the set of neighbors of node i .

2.2. Graph-based description: nodes, edges, and matrices

We consider a thermal graph $G = (V, E)$, where the thermal system is represented as a network of interacting components [23, 24]. This graph-based formulation is naturally connected to dynamical systems representations, including structural controllability and observability concepts [25], and is closely related to recent graph-neural and message-passing approaches for multi-component physical systems. In this framework, nodes (vertices) represent control volumes Ω_i associated with physical components or subsystems. Each node is associated with a single temperature state T_i and a capacitance C_i corresponding to the integral of ρc_p over Ω_i . Meanwhile, edges encode lumped thermal couplings between nodes. Each edge (i, j) is weighted by a symmetric conductance $G_{ij} = G_{ji} > 0$ that aggregates conduction paths. We define the weighted adjacency $W = [w_{ij}]$ with $w_{ij} = G_{ij}$ for $i \neq j$ and $w_{ii} = 0$, the degree matrix $D = \text{diag}(d_i)$ with $d_i = \sum_{j \neq i} G_{ij}$, and the (combinatorial) Laplacian L :

$$L := D - W, \quad L_{ij} = \begin{cases} -G_{ij}, & i \neq j, \\ \sum_{k \neq i} G_{ik}, & i = j. \end{cases}$$

For simplicity, we aggregate both controlled fluxes (heaters) and external excitations (orbital fluxes, disturbances) into a single input vector $u(t)$. Thus, we obtain:

$$C \dot{T}(t) = -LT(t) + u(t), \quad C = \text{diag}(C_1, \dots, C_n). \quad (1)$$

2.3. Radiative fluxes and linearization

In addition to conductive exchanges between nodes, satellites experience radiative heat transfer. Two main cases can be distinguished:

- **Radiative exchange with an external sink (deep space)**

The net flux between a node i at temperature T_i and a radiative sink is governed by the Stefan–Boltzmann law [3]:

$$P_i^{\text{rad}} = \epsilon_i \sigma A_i (T_i^4 - T_\infty^4),$$

where ϵ_i is the emissivity, A_i the radiating area, and σ the Stefan–Boltzmann constant. In most satellite applications, the sink is deep space with $T_\infty \approx 0$, so that

$$P_i^{\text{rad}} \approx \epsilon_i \sigma A_i T_i^4.$$

Typically, for satellite thermal monitoring, all quantities A_i , ϵ_i , and T_∞ are known or fixed by design, and the corresponding radiative conductance can be computed analytically without ambiguity. Therefore, this contribution does not require calibration.

- **Radiative exchange between two components**

For two nodes i and j with view factor F_{ij} , emissivities ϵ_i , ϵ_j , and areas A_i , A_j , the net radiative power exchanged is:

$$P_{ij}^{\text{rad}} = \frac{\sigma(T_i^4 - T_j^4)}{\frac{1 - \epsilon_i}{\epsilon_i A_i} + \frac{1}{A_i F_{ij}} + \frac{1 - \epsilon_j}{\epsilon_j A_j}},$$

where σ is the Stefan–Boltzmann constant. We introduce an effective emissivity ϵ_{ij}^* defined by:

$$\frac{1}{\epsilon_{ij}^* A_i F_{ij}} := \frac{1 - \epsilon_i}{\epsilon_i A_i} + \frac{1}{A_i F_{ij}} + \frac{1 - \epsilon_j}{\epsilon_j A_j},$$

so that the net exchange can be compactly written as:

$$P_{ij}^{\text{rad}} = \sigma \epsilon_{ij}^* A_i F_{ij} (T_i^4 - T_j^4),$$

An equivalent conductance between nodes i and j , can be obtained linearizing the radiative exchange around a reference temperature \bar{T} (e.g. the mean of T_i and T_j), namely

$$P_{ij}^{\text{rad}} \approx G_{ij}^{\text{rad}}(\bar{T}) (T_i - T_j),$$

where

$$G_{ij}^{\text{rad}}(\bar{T}) := 4 \sigma \bar{T}^3 \epsilon_{ij}^* A_i F_{ij}.$$

These parameters cannot be fully determined analytically and require calibration based on measurements or high-fidelity simulations. Related data-driven approaches for modeling and correcting radiative thermal couplings in reduced-order spacecraft models have been explored [6, 26].

In the following, radiative exchanges with the external environment that can be reliably estimated from design data are treated as known exogenous inputs. Conversely, unresolved or uncertain radiative interactions between components may be absorbed into the effective conductive coefficients to be identified. This choice is consistent with the interpretation of the model parameters as effective quantities rather than exact physical properties.

2.4. Limitations of direct lumped reductions

The parameters in Lumped Parameter Thermal Models (LPTMs) are often challenging to estimate from nominal material data and geometric specifications. Variations in material properties, uncertainties in geometry, and neglected interactions lead to differences between the idealized model assumptions and the observed effective behavior. The key assumption underlying our approach is that, despite these simplifications, the reduced-order model retains the fundamental structure of thermal physics: heat exchanges between lumped nodes are governed by temperature gradients, energy is conserved, and the dynamics remain consistent with thermodynamic principles. Under this assumption, the LPTM can be interpreted as a physics-constrained surrogate, where the governing equations define the functional form of the dynamics and the parameters are treated as effective lumped coefficients to be identified from observed thermal trajectories. Rather than deriving these coefficients from first principles, we learn them from data while enforcing the structural constraints imposed by the thermal network topology. This motivates the formulation of parameter identification as a data-driven inverse problem, constrained by the physical structure of the thermal network.

3. Inverse Parameter Identification Framework

The objective of the proposed framework is to identify a compact and lumped thermal model from partial observations, which naturally leads to a parameter estimation problem within the theory of inverse problems [21, 22]. Rather than deriving directly the lumped model parameters as defined in Section 2.1, we reinterpret the lumped formulation introduced in Section 2 as a parameterized dynamical system whose coefficients are inferred directly from data.

3.1. Effective Lumped Thermal Model

Let $\phi := \{C_i, G_{ij}\}$ denote the physical parameters that would characterize an ideal lumped thermal model derived from the underlying heat-transfer problem. In practice these quantities are not directly accessible and cannot be reliably determined from geometry and material properties alone. Instead, we introduce a set of effective coefficients $\theta = \{\gamma_i, \delta_{ij}\}$ that parameterize a reduced dynamical system reproducing the observed temperature evolution at the sensor level. These coefficients play the role of nodal (inverse) capacitances and edge conductances within the proxy model, but they should not be interpreted as exact physical quantities.

Let $T(t) \in \mathbb{R}^n$ denote the vector of nodal temperatures. The reduced thermal dynamics are described by the system

$$\dot{T}(t) = f_\theta(T(t), u(t), t),$$

which can be written in the equivalent lumped form

$$C(\theta)\dot{T}(t) = -L(\theta)T(t) + u(t).$$

Here $C(\theta)$ collects the effective nodal capacitances and $L(\theta)$ is the weighted graph Laplacian associated with the thermal network. At the node level, the dynamics can be written as

$$\dot{T}_i = \gamma_i \left(u_i + \sum_{j \in N(i)} \delta_{ij} (T_j - T_i) \right),$$

where $N(i)$ denotes the set of neighbors of node i . The coefficients γ_i and δ_{ij} represent effective thermal properties that must be estimated from data. The coefficient γ_i can be interpreted as an effective inverse thermal capacitance, i.e., $\gamma_i \approx C_i^{-1}$, while δ_{ij} represents an effective thermal conductance between nodes i and j .

3.2. Supervised Learning Perspective

The dynamical system introduced in Section 3.1 naturally inherits the graph structure of the underlying thermal network. Each temperature state corresponds to a node of the graph, while thermal couplings between components are represented by edges. Within this representation, the evolution of each node results from the combination of local power inputs and pairwise heat exchanges with neighboring nodes. The conductive interaction between two nodes i and j can be written as

$$m_{ij}(T_i, T_j; \theta) = \delta_{ij}(T_j - T_i),$$

which represents the heat transfer along the edge connecting the two nodes. The total contribution acting on node i is obtained by aggregating the interactions with all adjacent nodes,

$$\sum_{j \in N(i)} m_{ij}(T_i, T_j; \theta).$$

The nodal dynamics therefore take the form

$$\dot{T}_i = \gamma_i \left(u_i + \sum_{j \in N(i)} m_{ij}(T_i, T_j; \theta) \right).$$

This formulation coincides with the message-passing structure commonly used in graph-based dynamical systems. The same functional form of the interaction law is applied across all edges of the network, while the parameters δ_{ij} remain edge-specific and are learned from data. From this perspective, the lumped thermal model can be interpreted as a structured graph dynamical system in which the interaction functions are fixed by the physical heat transfer law, while the associated coefficients are unknown parameters to be identified.

This observation establishes a direct connection with the framework of Graph Neural Ordinary Differential Equations, which generalize Neural ODE models to structured systems [15, 27]. In that setting, the state evolution is described by a parameterized vector field defined on a graph, and the unknown parameters are identified from trajectory data through gradient-based optimization.

Although the present model does not introduce neural network parametrizations for the interaction functions, its structure matches that of a Graph Neural ODE with physically prescribed message functions. In contrast to fully data-driven approaches such as Physics-Informed Neural Networks, the proposed framework retains an explicit physical structure and focuses on parameter identification rather than function approximation [12, 13]. This reinterpretation allows the calibration problem to be treated within the standard supervised learning setting commonly adopted for Neural ODE models.

3.3. Inverse Problem Formulation

Let $\{t_k\}_{k=0}^N$ denote the sampling times and $\{u_k, y_k\}_{k=0}^N$ the available dataset, where u_k represents the known thermal inputs and y_k the measured temperatures at the sensor locations. Given a parameter vector θ , the effective thermal model introduced in Section 3.1 defines a dynamical system whose evolution is governed by

$$\dot{T}(t) = f_\theta(T(t), u(t), t).$$

For a prescribed initial condition $T(t_0)$ and input trajectory $u(t)$, this system generates a temperature trajectory $\{T_k(\theta)\}_{k=0}^N$, where we denote $T_k(\theta) := T(t_k; \theta)$. The identification task consists in determining the parameter vector θ that best reproduces the observed trajectories. This leads to a trajectory-matching optimization problem, consistent with classical system identification approaches based on time-series data [28],

$$\theta^* = \arg \min_{\theta} \sum_{k=0}^N \|T_k(\theta) - y_k\|_2^2. \quad (2)$$

From the perspective introduced in Section 3.2, this problem can be interpreted as a supervised learning task on dynamical systems defined over a graph. The model parameters θ play the role of trainable weights, while the temperature trajectories represent the training targets.

In practice, solving this optimization problem requires a number of additional design choices, including the parametrization of the coefficients, the strategy used to compute gradients of the trajectory loss, and the numerical structure adopted to enable efficient training. These aspects are addressed in the next section, where we describe the computational architecture used to make the inverse problem numerically tractable while preserving the physical structure of the thermal model.

The trajectory matching differs from standard regression settings, as the supervision is provided through the evolution of the system rather than pointwise targets. In practice, for computational and stability reasons, the trajectory matching loss will be evaluated over multiple shorter time windows, as described in Section 4. The proposed framework can therefore be interpreted as a structured neural differential model on graphs, where physical constraints such as positivity, symmetry, and dissipativity are enforced by construction, while retaining end-to-end trainability from data.

4. Computational Architecture for Stable Calibration

The inverse problem formulated in Section 3 corresponds to the identification of the parameter vector θ of a graph dynamical system from observed trajectories. While conceptually simple, solving this optimization problem in practice requires several architectural choices to ensure numerical stability and efficient training. Stability considerations mirror known challenges in the training of Neural ODE models over long time horizons [27, 29].

In particular, trajectory-based optimization of Neural ODE systems can become unstable when parameters temporarily leave the physically meaningful regime or when gradients are computed over long time horizons. For this reason the calibration framework adopts a number of structural design choices that guarantee stability of both the dynamics and the optimization process.

4.1. Parallel Trajectory Batching

The trajectory-based loss requires repeated integration of the dynamical system over multiple time intervals. For long telemetry sequences evaluating the loss over the entire trajectory at each optimization step would be computationally too expensive and may lead to unstable gradients, since prediction errors accumulate over long horizons.

To address this issue, the available time series is partitioned into multiple shorter temporal segments that are processed simultaneously during training. Let N_B denote the number of batches and T_b the temporal length of each segment. The dataset is partitioned as

$$D_b = \{u_{b,k}, y_{b,k}\}_{k=0}^{N_b}$$

where each batch corresponds to a trajectory initialized from a different time window of the telemetry sequence. Each trajectory segment corresponds to a contiguous time window extracted from the full dataset. In this work, non-overlapping segments are used, although overlapping windows could also be considered.

Instead of evolving each trajectory independently, the system is extended to a larger disconnected graph obtained by replicating the original thermal network N_B times. If $L(\theta)$ denotes the graph Laplacian of the original network, the corresponding operator for the composite system takes a block-diagonal form and corresponding state vector is constructed by concatenating the temperature states of all batches,

$$\bar{L}(\theta) = \text{diag}(L(\theta), \dots, L(\theta)), \quad \bar{T} = [T_1^\top, T_2^\top, \dots, T_{N_B}^\top]^\top.$$

The resulting system evolves according to

$$\dot{\bar{T}} = f_\theta(\bar{T}, \bar{u}, t), \tag{3}$$

where each subgraph evolves independently but shares θ .

This formulation allows all trajectories to be propagated within a single forward pass of the graph-based Neural ODE solver. Since the graph operations involved in the vector field evaluation consist primarily of sparse matrix–vector products and local message computations along edges, the resulting workload can be executed efficiently using parallel computing architectures.

In addition, processing multiple trajectories simultaneously improves the conditioning of the optimization problem. Gradient updates are aggregated across batches corresponding to different initial conditions and forcing scenarios, providing a more informative training signal and reducing the risk of overfitting to individual trajectory segments.

4.2. Bounded Parameter Embedding

The physical parameters of the thermal network are collected in the vector $\theta = \{\gamma_i, \delta_{ij}\}$. These coefficients have two structural properties that make direct optimization very difficult. First, they are all strictly positive. Second, their admissible values may span several orders of magnitude, since they correspond to thermal processes occurring at different characteristic time scales. Very small coefficients correspond to thermally inactive nodes or edges, effectively removing interactions from the network. Conversely, very large coefficients represent increasingly fast heat-transfer processes. The high regime is problematic both numerically and conceptually: extremely large parameters introduce stiff dynamics that require very small integration timesteps, while from a modeling perspective they correspond to resolving thermal processes that fall outside the intended resolution of the lumped approximation.

To enforce physical constraints such as positivity and boundedness of the parameters, we introduce a smooth mapping from unconstrained variables $\tilde{p} \in \mathbb{R}$ to physically admissible parameters p .

Specifically, we define

$$p = \Psi(\tilde{p}) = \frac{s_p}{2} (\tanh(2\tilde{p}) + 1),$$

where $s_p > 0$ is a characteristic scale representing the maximum admissible value of the parameter.

This transformation ensures that $p \in (0, s_p)$ for all $\tilde{p} \in \mathbb{R}$, while preserving smoothness and enabling gradient-based optimization in the unconstrained space. The choice of the hyperbolic tangent provides symmetric saturation and avoids excessively small gradients near the origin, thereby improving numerical stability during training. Similar constrained parameterizations are commonly adopted to enforce positivity and boundedness in neural differential models [30, 31]. The scale s_p is selected based on prior physical knowledge or expected orders of magnitude of the parameters.

4.3. Bounded Evolution During Training

Even with the bounded parameterization introduced in Section 4.2, the optimization process may temporarily explore parameter values far from those corresponding to a stable thermal network. In such situations the predicted trajectories can diverge rapidly during the numerical integration of the dynamical system. This issue is particularly critical when the trajectory loss is evaluated through long forward rollouts, since unstable intermediate dynamics may lead to exploding states and corrupted gradient estimates.

To address this issue, we introduce a stabilization strategy that limits the range of admissible temperatures during the forward propagation of the model. Specifically, after each integration step of the unrolled dynamics, a smooth limiting function is applied to the temperature states. The transformation is designed so that it does not modify the dynamics within the physically relevant operating range, while providing bounded dynamics when the trajectory moves outside this region.

Let T denote a generic nodal temperature. The limited value \hat{T} is defined through a piecewise smooth transformation constructed using hyperbolic tangent functions

$$\hat{T} = \begin{cases} v \tanh(T/v - 1) + 1v, & T < v \\ T, & v \leq T \leq 2v \\ v \tanh(T/v - 2) + 2v, & 2v < T \end{cases} \quad (4)$$

The mapping is linear within the nominal operating interval $200 \text{ K} \leq T \leq 400 \text{ K}$ using $v = 200 \text{ K}$, so that temperatures in this range remain unchanged. Outside this interval the function transitions smoothly toward asymptotic bounds using hyperbolic tangent branches. In particular, the transformation approaches the lower asymptote at 0 K and the upper asymptote at 600 K .

This choice ensures that the transformation behaves as the identity in the physically meaningful temperature range, while preventing unbounded growth of the state variables during unstable intermediate phases of the optimization.

The temperature limiter acts only outside the physically relevant operating range and is primarily used to prevent numerical instabilities during the early stages of training, consistent with the numerical stability issues discussed in [32], when parameters may take unrealistic values. As training progresses, the learned dynamics remain within the linear region of the limiter, making its effect negligible. In all reported experiments, the limiter is removed during evaluation, ensuring that the final model corresponds to the original unconstrained dynamics.

4.4. Energy Conservation and Dissipation

Beyond the numerical stabilization strategies introduced in the previous sections, the thermal model must preserve the structural properties of passive heat-transfer networks. In particular, conductive couplings between nodes are symmetric and cannot generate energy internally. Symmetry of conductive exchanges is enforced by construction through $\delta_{ij} = \delta_{ji}$, while nodal coefficients satisfy $\gamma_i > 0$. Under these conditions the associated graph Laplacian $L(\theta)$ is symmetric positive semidefinite.

Two energy-related quantities can be associated with the thermal network. First, the total thermal content of the system is defined as

$$H(T) = \sum_i C_i T_i = \mathbf{1}^\top C T. \quad (5)$$

This quantity represents the total thermal content of the system relative to a chosen reference temperature. In the absence of external inputs, conductive exchanges only redistribute heat between nodes and therefore preserve this quantity. Indeed, using

$$C(\theta)\dot{T} = -L(\theta)T, \quad (6)$$

one obtains

$$\dot{H} = \mathbf{1}^\top C(\theta)\dot{T} = -\mathbf{1}^\top L(\theta)T = 0, \quad (7)$$

since the Laplacian satisfies

$$L(\theta)\mathbf{1} = 0. \quad (8)$$

In addition to the conserved thermal content, it is convenient to introduce the quadratic functional

$$E(T) = \frac{1}{2} T^\top C(\theta) T. \quad (9)$$

Although not representing the physical thermal energy, this quantity acts as a Lyapunov function for the conductive dynamics. Differentiating along system trajectories yields

$$\dot{E} = T^\top C(\theta)\dot{T} = -T^\top L(\theta)T. \quad (10)$$

Using the structure of the graph Laplacian, the quadratic form can be written explicitly as

$$T^\top L(\theta)T = \frac{1}{2} \sum_{i,j} \delta_{ij} (T_i - T_j)^2, \quad (11)$$

which is always non-negative since $\delta_{ij} > 0$. Such properties are consistent with classical passivity-based formulations of networked physical systems.

Consequently, $\dot{E} \leq 0$. Equality holds only when all connected nodes have identical temperatures, corresponding to a state of internal thermal equilibrium with no conductive heat fluxes.

Embedding symmetry and positivity directly in the parameterization therefore guarantees that the calibrated model preserves the passive and dissipative nature of thermal conduction throughout the entire training process. This ensures that the identified model remains within the class of passive and dissipative systems. Structure-preserving formulations of dynamical systems similarly aim to enforce consistency with underlying physical principles in learned models, including energy conservation and dissipation properties [17, 33].

5. Numerical Results

We first describe the test cases used to assess the proposed hybrid calibration framework. The computational domain represents a simplified yet physically meaningful cross-section of a satellite thermal subsystem. Let $\Omega \subset \mathbb{R}^2$ denote the spatial domain under consideration. Specifically, we consider a square two-dimensional domain Ω , whose boundary $\partial\Omega$ is surrounded by a strongly insulating layer, mimicking the presence of multi-layer insulation and minimizing heat exchange with the external environment. Within Ω , four disjoint regions $\{\Omega_i\}_{i=1}^4$ corresponding to internal subsystems are considered; three of these subsystems, $\Omega_1, \Omega_2, \Omega_3$, are monitored and have different thermal properties, while Ω_4 is thermally isolated and unmonitored. Each region Ω_i is assumed homogeneous, with constant thermal conductivity λ_i and volumetric heat capacity $\rho_i c_i$, as shown in Figure 1.




Figure 1: Computational domain $\Omega = [0, 2]^2 \text{ m}^2$. Thermal conductivity field (left) and volumetric heat capacity field (right).

Ground truth data are generated by solving the transient heat conduction problem for the temperature field $T(x, t)$ on Ω using a Finite Element Method (FEM), a standard discretization approach for heat transfer problems. The temperature field $T(x, t)$ is governed by:

$$c_f \partial_t T - \nabla \cdot (\lambda_f \nabla T) = \sum_j f_j(x, t) \chi_{\Omega_j}(x) \quad \text{in } \Omega, \quad (12)$$

where c_f and k_f denote the thermal capacity and conductivity, respectively. The terms $f_j(x, t)$ represent localized volumetric heat sources acting on subdomains $\Omega_j \subset \Omega$. On the boundary $\partial\Omega = \cup_{i=1}^4 \Gamma_i$, nonlinear Neumann boundary conditions are imposed:

$$-\lambda_f \nabla T \cdot n = q_i(t) + \sigma \varepsilon_{\text{zen}} (T^4 - T_{\text{ext}}^4) \quad \text{on } \Gamma_i, \quad (13)$$

accounting for surface heat fluxes $q_i(t)$ and radiative heat losses linearized towards an external temperature T_{ext} .

5.1. Weak Formulation and Data Generation

The synthetic dataset is generated by solving the two-dimensional heat equation with volumetric sources, prescribed boundary heat fluxes, and radiative cooling. Using an IMEX scheme for time discretization with timestep Δt , the weak formulation reads as follows: find $u \in H^1(\Omega)$ such that for all test functions $v \in H^1(\Omega)$

$$\int_{\Omega} c_f T v dx + \Delta t \int_{\Omega} \lambda_f \nabla T \cdot \nabla v dx - \Delta t \sum_{i=1}^4 \int_{\Gamma_i} (\sigma \varepsilon_{\text{zen}} T_{\text{old}}^3) T v ds = \int_{\Omega} c_f T_{\text{old}} v dx + \Delta t \sum_{i=1}^4 q_i \int_{\Gamma_i} v ds + \Delta t \sum_j f_j \int_{\Omega_j} v dx.$$

Here T denotes the temperature at the current time step and T_{old} the temperature at the previous one. The terms q_i correspond to prescribed heat fluxes on the boundary segments Γ_i , while f_j denote volumetric heat sources acting

on subdomains $\Omega_j \subset \Omega$. Radiative cooling is modeled through a boundary contribution derived from the Stefan–Boltzmann law. The nonlinear T^4 dependence is linearized around the previous state and assuming $T_{ext} \approx 0K$, leading to the coefficient $\sigma \epsilon z_{len} T_{old}^3$ and yielding a stable fully implicit formulation.

The computational domain is $\Omega = [0, 2] \times [0, 2] \text{ m}^2$ and is discretized with a structured 64×64 mesh of rectangular elements. Continuous Galerkin finite elements of degree one (P1) are employed, resulting in $65 \times 65 = 4225$ degrees of freedom. Time integration is solved by direct LU factorization.

5.2. Generation of Time Series and Heat Sources

The resulting FEM solution serves as a surrogate for ground-truth physics, from which synthetic measurements are extracted. All results presented in this work are obtained using simulated data in a controlled setting [5.1](#), where the thermal dynamics, material properties, and boundary conditions are known. High-fidelity FEM simulations provide a reference solution that is not available in real telemetry, allowing us to quantify to what extent the proposed framework is able to recover a consistent lumped thermal model from sparse temperature measurements.

To emulate a realistic operational scenario with sparse sensing, only a limited number of temperature sensors is assumed to be available. All sensors provide pointwise temperature measurements and constitute the only observable data used by the proposed model. From each sensor location $\{\mathbf{x}_k\}_{k=0}^8$, a discrete temperature time series $\{T(\mathbf{x}_k, t_n)\}_{n=0}^{N_t}$ is extracted, and the collection of these sparse observations forms the input dataset for the hybrid model.

In these test cases, 4 sensors are located on the boundary of the square domain, one on each side, and are referred to as boundary nodes. These sensors define the interface through which radiative heat exchange between the corresponding satellite face and deep space is applied. Accordingly, the boundary nodes are endowed with the known geometric and optical properties of the associated satellite face, such as radiating area and emissivity. The remaining sensors are placed inside the monitored subdomains, with some subdomains containing more than one sensor.




Figure 2: Thermal graph on the computational domain. Boundary nodes (1–4) handle radiative exchange with deep space; internal nodes (5–8) correspond to sensor locations within the monitored subdomains. Green edges denote conductive couplings.

The thermal forcing time series were generated using two methodologies to represent deterministic and stochastic phenomena. Periodic forcings, simulating solar irradiation, were modeled through sinusoidal functions with characteristic orbital period, allowing to reproduce the solar exposure and eclipse phases typical of low Earth orbit missions. For stochastic forcings, representing both heating lines and thermal dissipations from internal electronic devices, a Gaussian process approach with radial basis function kernel was employed. This choice is motivated by the intrinsically variable and unpredictable nature of internal thermal loads, which depend on the non-deterministic activation of

electronic components, payload operational cycles, and variable usage conditions during the mission. In the test below, Tests A, B, and C refer to configurations with different characteristic periods of the internal sources, respectively a correlation length of 700, 840 and 560.

5.3. Training Strategy

The training and validation dataset is processed in parallel according to the block-diagonal formulation of Section 4.1. Batch length B is chosen to be sub-orbital, i.e. shorter than the characteristic thermal period of approximately 100 minutes. The sensitivity of the results to the choice of T_b is later investigated in Section 5.6.

To prevent numerical divergence during early optimization phases, we apply a smooth saturation (4) to nodal temperatures at each integration step. The transformation coincides with the identity within the physically admissible operating range [200 K, 400 K], and smoothly compresses values outside this interval through a differentiable mapping with bounded asymptotes. As training progresses and the identified parameters enter a physically meaningful regime, predicted temperatures remain within this interval and the saturation acts as the identity. Consequently, it does not alter the learned dynamics in the operational regime and serves as a stabilization mechanism during early training.

Model parameters are optimized by minimizing the trajectory-based loss defined in (2). Each run is trained for a maximum of 5000 epochs, with early stopping applied when the validation loss ceases to decrease over a fixed patience window. Parameters are optimized using an adaptive first-order gradient method (ADAM [34]), applied to the ODE-constrained trajectory mismatch functional, without learning rate scheduler; runs are conducted independently at fixed learning rates of 10^{-2} and 10^{-3} to explore the sensitivity of convergence to this hyperparameter.

Model performance is evaluated on the long-horizon test segment $\mathcal{D}_{\text{test}}$, which spans several complete orbital cycles and two metrics are reported: the root mean squared error (RMSE) between predicted and reference temperature trajectories and the Pearson correlation coefficient (PCC) (14), which quantifies the linear correlation between predicted and observed time series independently of amplitude offsets.

$$PCC = \frac{\text{cov}(X, Y)}{\sigma_X \cdot \sigma_Y} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 \cdot \sum_{i=1}^n (y_i - \bar{y})^2}} \quad (14)$$

5.4. Computational complexity

Let $G = (V, E)$ denote the thermal graph. Each evaluation of the vector field f_θ consists of local message computations along edges followed by node-wise aggregation. Therefore, the computational cost per timestep scales as

$$O(|E| + |V|),$$

since each edge contributes a constant-cost interaction and each node performs a constant-cost update.

Under the block-diagonal batching strategy of Section 4.1, B independent trajectories are propagated in parallel, leading to an overall cost

$$O(B(|E| + |V|))$$

per timestep, while sharing the same parameter vector ϕ . Memory requirements scale linearly with $B(|E| + |V|)$.

5.5. Test Case Configurations

Three test configurations are considered to assess the proposed framework under progressively complex thermal scenarios. The first (Table 1) involves only external radiative forcing applied at the domain boundary, representing orbital heat fluxes in the absence of internal dissipation. The second (Table 2) features two internal volumetric heat sources acting on distinct subdomains, mimicking the activation of heater lines used in practice to maintain onboard components within their safe operating temperature ranges. The third configuration (Table 3) combines both external and internal forcings simultaneously and constitutes the reference benchmark for all subsequent analyses reported in Section 5.6.

Convergence is consistently achieved within a small number of iterations, often well below the maximum allowed. This is due to three structural properties of the proposed framework. First, the strong physical constraints embedded in the neural field architecture (positivity, symmetry, and energy conservation) substantially reduce the effective

dimensionality of the optimization landscape by confining the search to the physically admissible parameter space from the first iteration. Second, the bounded parameterization introduced in Section (4.2) prevents gradient explosion and ensures smooth differentiability throughout training. Third, the parallelization strategy described in Section 4.1 provides gradient updates averaged over multiple initial conditions and input scenarios simultaneously, yielding a more informative and stable training signal per iteration compared to single-trajectory optimization.

Representative results for the three configurations are reported in Figures 3-4-5. Each panel displays the predicted and reference temperature trajectories over the full time horizon considered. The background shading identifies the three data partitions: the green region corresponds to the training segments, the yellow region to the validation segments, and the red region to the long-horizon test set $\mathcal{D}_{\text{test}}$. In each panel, the reference FEM temperature is shown in blue and the model prediction in orange.

Dataset	Iterations	Max Iter	Learning Rate	Train Loss	Valid Loss	Test RMSE	Test PCC
Test A	2000	2000	1.00×10^{-2}	1.43×10^{-2}	1.66×10^{-2}	2.66×10^{-2}	9.98×10^{-1}
Test A	2073	5000	1.00×10^{-2}	1.43×10^{-2}	1.65×10^{-2}	2.66×10^{-2}	9.98×10^{-1}
Test A	5000	5000	1.00×10^{-3}	1.50×10^{-2}	1.75×10^{-2}	2.71×10^{-2}	9.98×10^{-1}

Table 1: Calibration results for the external forcing configuration.

Dataset	Iterations	Max Iter	Learning Rate	Train Loss	Valid Loss	Test RMSE	Test PCC
Test A	2000	2000	1.00×10^{-2}	2.00×10^{-2}	2.03×10^{-2}	2.16×10^{-2}	9.79×10^{-1}
Test A	4210	5000	1.00×10^{-2}	1.93×10^{-2}	1.96×10^{-2}	2.08×10^{-2}	9.81×10^{-1}
Test A	5000	5000	1.00×10^{-3}	2.31×10^{-2}	2.31×10^{-2}	2.41×10^{-2}	9.75×10^{-1}
Test B	2000	2000	1.00×10^{-2}	2.08×10^{-2}	1.99×10^{-2}	2.40×10^{-2}	9.73×10^{-1}
Test B	4195	5000	1.00×10^{-2}	1.96×10^{-2}	1.88×10^{-2}	2.15×10^{-2}	9.79×10^{-1}
Test B	5000	5000	1.00×10^{-3}	2.27×10^{-2}	2.25×10^{-2}	2.79×10^{-2}	9.64×10^{-1}
Test C	2000	2000	1.00×10^{-2}	1.60×10^{-2}	2.25×10^{-2}	3.96×10^{-2}	9.38×10^{-1}
Test C	4753	5000	1.00×10^{-2}	1.40×10^{-2}	2.30×10^{-2}	3.44×10^{-2}	9.50×10^{-1}
Test C	5000	5000	1.00×10^{-3}	1.82×10^{-2}	3.33×10^{-2}	3.98×10^{-2}	9.34×10^{-1}

Table 2: Calibration results for the internal heater line configuration.

Dataset	Iterations	Max Iter	Learning Rate	Train Loss	Valid Loss	Test RMSE	Test PCC
Test A	2000	2000	1.00×10^{-2}	1.74×10^{-2}	2.18×10^{-2}	2.03×10^{-2}	9.86×10^{-1}
Test A	2456	5000	1.00×10^{-2}	1.72×10^{-2}	2.12×10^{-2}	1.97×10^{-2}	9.87×10^{-1}
Test A	5000	5000	1.00×10^{-3}	1.93×10^{-2}	2.43×10^{-2}	2.23×10^{-2}	9.82×10^{-1}
Test B	2000	2000	1.00×10^{-2}	1.77×10^{-2}	1.91×10^{-2}	2.13×10^{-2}	9.85×10^{-1}
Test B	2336	5000	1.00×10^{-2}	1.76×10^{-2}	1.90×10^{-2}	2.11×10^{-2}	9.85×10^{-1}
Test B	5000	5000	1.00×10^{-3}	1.92×10^{-2}	2.07×10^{-2}	2.37×10^{-2}	9.82×10^{-1}
Test C	2000	2000	1.00×10^{-2}	1.53×10^{-2}	2.18×10^{-2}	2.09×10^{-2}	9.85×10^{-1}
Test C	5000	5000	1.00×10^{-2}	1.28×10^{-2}	1.97×10^{-2}	1.93×10^{-2}	9.88×10^{-1}
Test C	5000	5000	1.00×10^{-3}	1.73×10^{-2}	2.43×10^{-2}	2.32×10^{-2}	9.81×10^{-1}

Table 3: Calibration results for the benchmark configuration.




Figure 3: Predicted (orange) and reference (blue) temperature trajectories for the external source configuration.




Figure 4: Predicted (orange) and reference (blue) temperature trajectories for the internal heater line configuration.




Figure 5: Predicted (orange) and reference (blue) temperature trajectories for the benchmark test, combining external and internal forcings.

The convergence history and the evolution of the learned parameters over training iterations are reported in Figure 6 for the benchmark configuration combining external radiative forcing and internal heater line activations. The training and validation losses decrease monotonically and reach a stable plateau within a limited number of iterations, consistent with the early stopping counts reported in Table 3. The learned physical parameters converge smoothly to well-defined values without oscillations or instabilities, confirming that the bounded parameterization of Section 4.2 effectively conditions the optimization landscape. An analogous convergence behavior is observed across all other test configurations, indicating that the framework is robust with respect to the thermal forcing.




Figure 6: Training diagnostics for the benchmark configuration.

5.6. Sensitivity Analysis and Robustness Assessment

All analyses reported in this section are conducted on the benchmark configuration introduced in Section 5.5, which combines simultaneous external radiative forcing and internal heater line activations. Starting from this common reference setting, the training configuration and the properties of the input data are systematically varied in order to assess how each factor affects the behavior of the calibrated model.

Effect of Training Window Length

In this experiment, the total amount of data available for training and validation is kept fixed, while the number of batches B is varied. Since the total dataset length is held constant, increasing B directly reduces the temporal length T_b of each individual training window, and vice versa. This setup allows the effect of the window length on the identification process to be isolated without altering the overall information content of the dataset. The numerical results are summarized in Table 4.

When using the trajectory-based loss, each window is integrated from a single initial condition, so that prediction errors accumulate over the rollout horizon. Longer windows therefore introduce a more challenging optimization problem in the early training stages, reflected in slower convergence and in a computational cost per epoch that grows approximately linearly with T_b , as confirmed by the training times reported in Table 4. At the opposite extreme, very short windows provide only local supervision, leaving the model insufficiently constrained by long-range thermal dynamics and exposing it to overfitting, as evidenced by the markedly higher test error obtained at $T_b = 500$.

Beyond a minimum threshold, however, further increases in window length yield diminishing returns in predictive accuracy while substantially increasing computational cost. Results show that an intermediate window length provides the best trade-off between long-horizon generalization and computational efficiency. The optimal choice can therefore be identified as the shortest window that still guarantees stable performance on the test set.

T_b	Train Loss	Valid Loss	Test RMSE	Test PCC	Training Time
500	1.67×10^{-2}	1.63×10^{-2}	5.05×10^{-2}	9.04×10^{-1}	5m 56s
1000	1.72×10^{-2}	1.82×10^{-2}	2.00×10^{-2}	9.87×10^{-1}	10m 26s
2000	1.74×10^{-2}	2.18×10^{-2}	2.03×10^{-2}	9.86×10^{-1}	21m 30s
4000	1.97×10^{-2}	2.39×10^{-2}	1.92×10^{-2}	9.87×10^{-1}	35m 54s

Table 4: Effect of the number of training batches B on model performance and computational cost. Total dataset length is held constant.




Figure 7: Effect of training window length on the predicted temperature trajectories. Each panel corresponds to a different value of T_b , with background shading identifying training (green), validation (yellow), and test (red) segments.

Noise Robustness

In this experiment, the robustness of the identification procedure is assessed in the presence of measurement noise. Gaussian white noise is added both to the initial conditions and to the temperature labels used for training and validation. Two noise levels are considered, equal to 1% and 5% of the mean temperature value \bar{x} . More precisely, at each time step a perturbation is sampled from a normal distribution with zero mean and standard deviation equal to $0.01 \bar{x}$ and $0.05 \bar{x}$, respectively. The corresponding results are reported in Table 5 alongside the noiseless baseline.

In the 1% noise regime, model performance remains essentially unchanged with respect to the noiseless case. The trajectory loss, test error, and PCC values are all comparable to the baseline, indicating that the calibration framework is not sensitive to small-amplitude perturbations. When the noise level is increased to 5%, the test error increases accordingly but remains bounded and consistent with the noise magnitude, suggesting that the residual discrepancy is dominated by label uncertainty rather than model mismatch. Importantly, no divergence or instability is observed during training or autoregressive rollout in either regime.

To assess whether errors accumulate over the test horizon, the test segment is partitioned into sequential windows and the distribution of windowed errors is analyzed through box plots. The resulting distributions remain statistically stable across all segments in both noise scenarios, confirming that the learned dynamics do not amplify measurement noise and that long-horizon robustness is preserved even under perturbed observations.

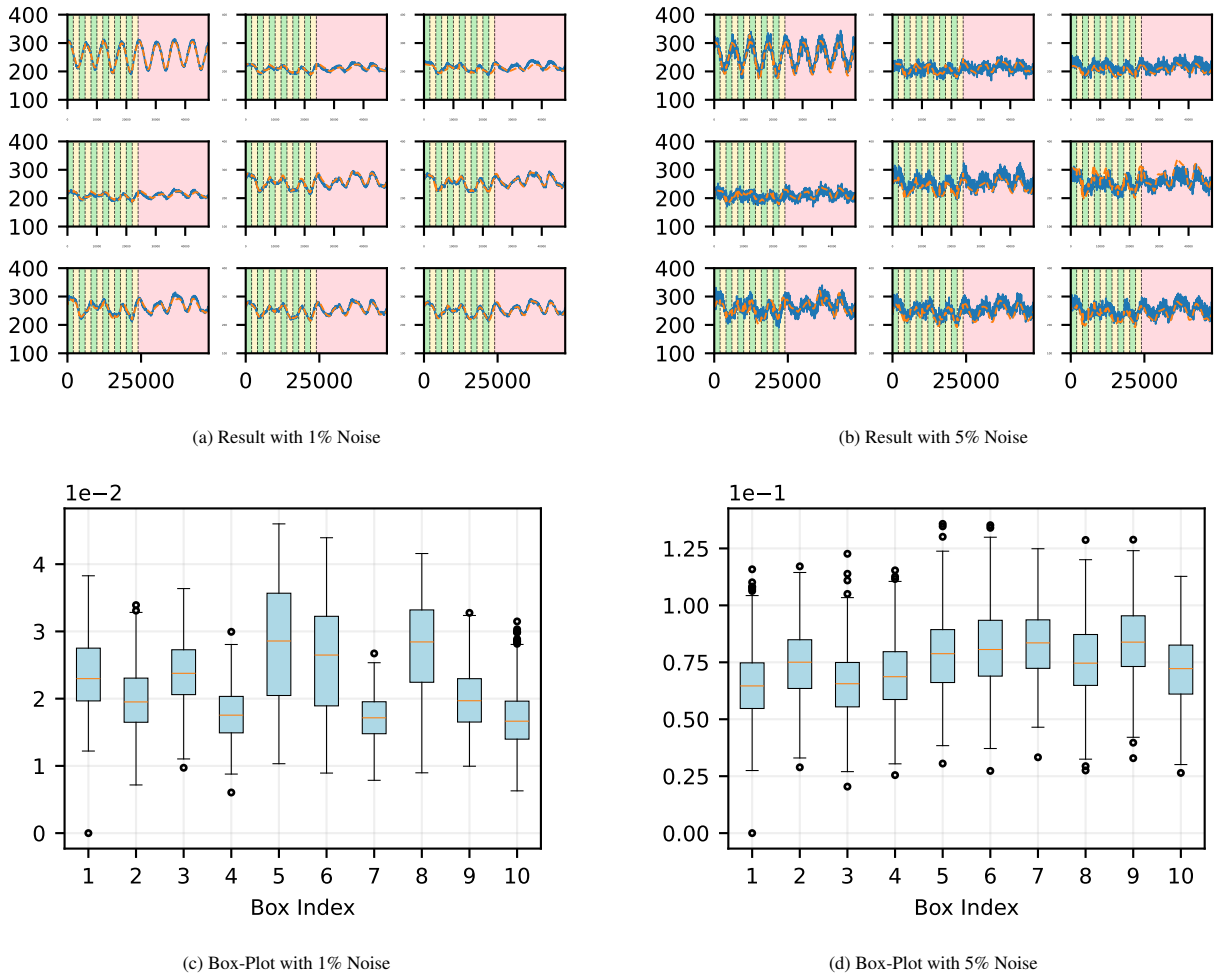


Figure 8: Noise robustness results. Top: predicted (orange) vs. reference (blue) trajectories under 1% (left) and 5% (right) measurement noise. Bottom: box plots of windowed test errors for each noise level, confirming the absence of error accumulation over long horizons.

Noise	Train Error	Valid Error	Test RMSE	Test PCC	Time Total
0%	1.740×10^{-2}	2.177×10^{-2}	2.028×10^{-2}	9.855×10^{-1}	21m 30s
1%	2.013×10^{-2}	2.319×10^{-2}	2.201×10^{-2}	9.832×10^{-1}	18m 15s
5%	5.757×10^{-2}	6.045×10^{-2}	5.679×10^{-2}	8.960×10^{-1}	18m 43s

Table 5: Performance metrics under different levels of measurement noise added to training labels and initial conditions.

5.7. Training Stability Analysis

A reliable calibration framework must produce consistent results across repeated training runs, even in the presence of measurement noise. In this section, training stability is evaluated from two complementary perspectives: the reproducibility of the performance metrics and the consistency of the learned physical parameters across runs.

The benchmark configuration is trained ten times on independently generated datasets, each corrupted by Gaussian white noise at a fixed level of 1% of the mean temperature value \bar{x} as described in Section 5.6. To quantify parameter stability across runs, we adopt a metric based on the signal-to-noise ratio. Such variability analysis is related to identifiability issues in nonlinear system identification [28]. For a given learned parameter θ estimated across K independent training runs, we define

$$\text{SNR}(\theta) = \frac{\mu_\theta}{\sigma_\theta},$$

where μ_θ and σ_θ denote the sample mean and standard deviation of θ across runs, respectively. Large values indicate that the parameter is consistently estimated and well concentrated around its mean, while small values signal that training variability is comparable to or exceeds the identified level. The following thresholds are adopted for interpretation: $\text{SNR} < 1$ indicates that variability dominates and the parameter is not stably identified; $1 \leq \text{SNR} < 2$ denotes marginal stability; $\text{SNR} \geq 2$ corresponds to a structurally stable estimate; and $\text{SNR} \geq 3$ indicates strong concentration around the mean.

The performance metrics across the ten runs exhibit very high reproducibility. Training loss, validation loss, test RMSE, and test PCC all present SNR values well above 30, confirming that the trajectory-based loss leads to consistent convergence regardless of the specific noise realization in the dataset.

At the parameter level, the thermal graph comprises 37 learnable parameters, collecting nodal capacitances γ_i and edge conductances δ_{ij} . Of these, 31 exceed the stability threshold of $\text{SNR} \geq 3$. Six parameters fall below this threshold, of which two present $\text{SNR} < 2$, with a minimum observed value of 1.1.

The larger variability observed for a small subset of parameters is consistent with limited information content in the available excitation/measurement setup. As a consequence, multiple parameter values can yield nearly indistinguishable trajectories on the observed components, leading to larger run-to-run variability without affecting predictive accuracy. This behavior reflects intrinsic practical identifiability limitations of the inverse problem, rather than instability of the training procedure.

6. Conclusions

In this work, we proposed a physics-constrained graph-based identification framework for Lumped Parameter Thermal Models (LPTMs), designed to bridge the gap between physics-based reduced-order models and purely data-driven approaches for spacecraft thermal systems. The methodology formulates thermal dynamics within the formalism of graph-based Neural ODEs, while preserving a physically interpretable message-passing structure and without relying on black-box neural parametrizations, enabling continuous-time parameter identification directly from temperature measurements and known thermal inputs.

A key feature of the proposed approach is the explicit enforcement of the physical structure of the thermal network. The model is constructed so that all admissible parameters satisfy by construction fundamental properties of conductive heat transfer, including positivity of nodal coefficients, symmetry of interactions, and dissipativity of the resulting dynamics. These properties are enforced by construction through a bounded parameterization, without requiring additional regularization terms. The inverse calibration problem is formulated as a trajectory-matching optimization over a graph dynamical system, and solved within a Neural ODE framework using adjoint-based gradient computation. To

ensure numerical stability and computational efficiency, the training pipeline combines bounded parameter embeddings, trajectory batching through block-diagonal graph replication, and controlled forward evolution. This design enables robust training even in the presence of long time horizons and stiff thermal dynamics.

The proposed framework was validated on synthetic datasets generated from high-fidelity FEM simulations under progressively complex forcing conditions. The results demonstrate accurate long-horizon predictions, with high correlation between predicted and reference trajectories and stable autoregressive behavior over multiple orbital cycles. Sensitivity analyses show limited degradation under measurement noise and highlight the role of the training window length in balancing optimization difficulty and generalization. Repeated training runs confirm that most parameters are consistently identified, while residual variability is confined to weakly excited thermal interactions, reflecting intrinsic identifiability limitations rather than instability of the method.

From a modeling perspective, the framework provides a modular, interpretable, and physically grounded approach to the calibration of reduced-order thermal models. The identified parameters retain a clear physical meaning as effective thermal properties, while the graph-based Neural ODE formulation enables flexible integration of data and prior structural knowledge. This combination allows the construction of computationally efficient models suitable for predictive tasks in operational environments.

In the context of spacecraft thermal Digital Twins, the proposed methodology provides a promising basis for continuous calibration from telemetry data, supporting real-time monitoring, anomaly detection, and predictive thermal control. The resulting models can be used to assess deviations from nominal behavior and to evaluate alternative control strategies under varying operational scenarios. The proposed identification framework aligns with current digital-twin initiatives for spacecraft thermal management [4, 35].

Validation of the presented approach on real telemetry data remains an important direction for future work, together with its integration into onboard real-time temperature control systems.

Acknowledgments

This work has been carried out in the framework of the Hybrid Digital Twins for Satellites (HDTS) Work Package of the ASTRAL project, funded by the Research National Center in High Performance Computing, Big Data and Quantum Computing through an Innovation Grant assigned to Thales Alenia Space Italia (TASI). The authors gratefully acknowledge the support of TASI technical staff whose domain expertise was crucial in defining the reference use case, while their continuous support provided essential system-level insights into the spacecraft's thermal subsystem, ensuring the industrial relevance of the proposed framework. SP, NP, FR, LS are members of INdAM-GNCS group. The present research is part of the activities of "Dipartimento di Eccellenza 2023-2027". SP, FR, LS acknowledge support from project FIS, MUR, Italy 2025-2028, Project code: FIS-2023-02228, CUP: D53C24005440001, "SYNERGIZE: Synergizing Numerical Methods and Machine Learning for a new generation of computational models".

References

- [1] H. S. Carslaw, J. C. Jaeger, *Conduction of Heat in Solids*, Oxford University Press, 1959.
- [2] F. P. Incropera, D. P. DeWitt, T. L. Bergman, A. S. Lavine, *Fundamentals of Heat and Mass Transfer*, 6th Edition, John Wiley & Sons, 2007.
- [3] M. F. Modest, *Radiative Heat Transfer*, Academic Press, 2013. [doi:https://doi.org/10.1016/C2010-0-65874-3](https://doi.org/10.1016/C2010-0-65874-3)
- [4] E. H. Glaessgen, D. S. Stargel, The digital twin paradigm for future nasa and u.s. air force vehicles, in: *AIAA Conference*, 2012. [doi:10.2514/6.2012-1818](https://doi.org/10.2514/6.2012-1818)
- [5] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems*, SIAM, 2005. [doi:10.1137/1.9780898718713](https://doi.org/10.1137/1.9780898718713)
- [6] P. Benner, S. Gugercin, K. Willcox, Survey of projection-based model reduction methods for parametric dynamical systems, *SIAM Review* 57 (4) (2015) 483–531. [doi:10.1137/130932715](https://doi.org/10.1137/130932715)
- [7] A. Quarteroni, A. Manzoni, F. Negri, *Reduced Basis Methods for Partial Differential Equations*, Springer, 2016.
- [8] B. S. Jan S Hesthaven, Gianluigi Rozza, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Springer Cham, 2015.
- [9] M. Mesbahi, M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*, Princeton University Press, 2010. [doi:10.1515/9781400835355](https://doi.org/10.1515/9781400835355)
- [10] I. Torralbo, I. Perez-Grande, A. Sanz-Andres, J. Piqueras, Correlation of spacecraft thermal mathematical models to reference data, *Acta Astronautica* 144 (2018) 305–319. [doi:https://doi.org/10.1016/j.actaastro.2017.12.033](https://doi.org/10.1016/j.actaastro.2017.12.033)
- [11] L. e. a. Lu, Learning nonlinear operators via deepnet based on the universal approximation theorem of operators, *Nature Machine Intelligence* (2021). [doi:10.1038/s42256-021-00302-5](https://doi.org/10.1038/s42256-021-00302-5)
- [12] M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics* 378 (2019) 686–707. [doi:10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045)
- [13] G. E. Karniadakis, I. G. Kevrekidis, G. E. Karniadakis, et al., Physics-informed machine learning, *Nature Reviews Physics* 3 (2021) 422–440. [doi:10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5)
- [14] S. L. Brunton, J. L. Proctor, J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proceedings of the National Academy of Sciences* 113 (15) (2016) 3932–3937. [doi:10.1073/pnas.1517384113](https://doi.org/10.1073/pnas.1517384113)
- [15] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, in: *Advances in Neural Information Processing Systems*, Vol. 31, 2018.
- [16] C. e. a. Rackauckas, Universal differential equations for scientific machine learning, *arXiv preprint arXiv:2001.04385* (2020). [doi:10.48550/arXiv.2001.04385](https://doi.org/10.48550/arXiv.2001.04385)
- [17] S. Greydanus, M. Dzamba, J. Yosinski, Hamiltonian neural networks, in: *Advances in Neural Information Processing Systems*, Vol. 32, 2019. [doi:10.48550/arXiv.1906.01563](https://doi.org/10.48550/arXiv.1906.01563)
- [18] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *Journal of Fluid Mechanics* 656 (2010) 5–28. [doi:10.1017/S0022112010001217](https://doi.org/10.1017/S0022112010001217)
- [19] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, J. N. Kutz, On dynamic mode decomposition: Theory and applications, *Journal of Computational Dynamics* 1 (2) (2014) 391–421. [doi:10.3934/jcd.2014.1.391](https://doi.org/10.3934/jcd.2014.1.391)

- [20] T. Reis, T. Stykel, Lyapunov balancing for passivity-preserving model reduction of rc circuits, *SIAM Journal on Applied Dynamical Systems* 10 (1) (2011) 1–34. [doi:10.1137/090779802](https://doi.org/10.1137/090779802).
- [21] A. M. Stuart, Inverse problems: A bayesian perspective, *Acta Numerica* 19 (2010) 451–559. [doi:10.1017/S0962492910000061](https://doi.org/10.1017/S0962492910000061).
- [22] J. Kaipio, E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005. [doi:10.1007/b138659](https://doi.org/10.1007/b138659).
- [23] F. R. K. Chung, *Spectral Graph Theory*, American Mathematical Society, 1997. [doi:10.1090/cbms/092](https://doi.org/10.1090/cbms/092).
- [24] M. Newman, *Networks: An Introduction*, Oxford University Press, 2010. [doi:10.1093/acprof:oso/9780199206650.001.0001](https://doi.org/10.1093/acprof:oso/9780199206650.001.0001).
- [25] C.-T. Lin, Structural controllability, *IEEE Transactions on Automatic Control* 19 (3) (1974) 201–208. [doi:10.1109/TAC.1974.1100557](https://doi.org/10.1109/TAC.1974.1100557).
- [26] K. Willcox, J. Peraire, Balanced model reduction via the proper orthogonal decomposition, *AIAA Journal* 40 (11) (2002) 2323–2330. [doi:10.2514/2.1570](https://doi.org/10.2514/2.1570).
- [27] E. Dupont, A. Doucet, Y. W. Teh, Augmented neural odes, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 32, Curran Associates, Inc., 2019.
- [28] L. Ljung, T. Glad, On global identifiability for arbitrary model parametrizations, *Automatica* 30 (2) (1994) 265–276. [doi:10.1016/0005-1098\(94\)90029-9](https://doi.org/10.1016/0005-1098(94)90029-9).
- [29] P. Kidger, On neural differential equations, arXiv preprint (2022). [doi:10.48550/arXiv.2202.02435](https://doi.org/10.48550/arXiv.2202.02435).
- [30] J. Drgoňa, A. Tuor, D. Vrabie, Constrained neural ordinary differential equations with stability guarantees, arXiv preprint (2020). [doi:10.48550/arXiv.2004.10883](https://doi.org/10.48550/arXiv.2004.10883).
- [31] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, P. Battaglia, Learning to simulate complex physics with graph networks, in: H. D. III, A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning*, Vol. 119 of *Proceedings of Machine Learning Research*, PMLR, 2020, pp. 8459–8468.
- [32] S. Massaroli, M. Poli, J. Park, A. Yamashita, H. Asama, Dissecting neural odes, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems*, Vol. 33, Curran Associates, Inc., 2020, pp. 3952–3963.
- [33] Y. D. Zhong, B. Dey, A. Chakraborty, Symplectic ODE-Net: Learning Hamiltonian dynamics with dissipation, *Physical Review E* 101 (6) (2020) 061301. [doi:10.1103/PhysRevE.101.061301](https://doi.org/10.1103/PhysRevE.101.061301).
- [34] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization (2017). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [35] R. Scigliano, V. De Simone, R. Fusaro, D. Ferretto, N. Viola, Numerical simulation of heat pipe thermal performance for aerospace cooling system applications, *Aerospace* 11 (1) (2024) 85. [doi:10.3390/aerospace11010085](https://doi.org/10.3390/aerospace11010085).

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 40/2026** Marchesin, L.; Menafoglio, A.; Secchi, P.
A Convolution Process for Sea Surface Temperature Hot-Spot Identification in the Mediterranean Sea
- 38/2026** Clemente, A.; Arnone, E.; Mateu, J.; Sangalli, L.M.
Nonparametric estimators over metric graphs
- 39/2026** Patanè, G.; Menafoglio, A.; Krauth, A.; Fechner, P.; Dede', L.; Colosimo, B.M.; Nicolussi, F.
K-Models: a Flexible and Interpretable Method for Ordinal Clustering with Application to Antigen-Antibody Interaction Profiles
- 37/2026** Centofanti, E.; Ziarelli, G.; Scacchi, S.; Pavarino, L.F.
A Neural Latent Dynamics Approach for Solving Inverse Problems in Cardiac Electrophysiology
- 36/2026** Botti, M.; Mascotto, L.; Mosconi, M.
A nonconforming method for a generalized Darcy-Forchheimer model
- 35/2026** Caon, B.; Corti, M.; Bonizzoni, F.; Antonietti, P.F.
High-fidelity and Network-based Spatio-temporal Mathematical Models of Alzheimer's Disease Progression and their Validation Against PET-SUVR Imaging Data
- 34/2026** Mancinelli, F. M.; Torzoni, M.; Maisto, D.; Donnarumma, F.; Corigliano, A.; Pezzulo, G.; Manzoni, A.
Multi-Agent Digital Twins for strategic decision-making using Active Inference
- 33/2026** Franzoni, G.; Mirabella, S.; Dabek, A.; Ferro, N.; Antona, A.; Carlessi, M.; Cinquemani, S.; Matteucci, M.; Cocetta, G.; Perotto, S.
Integrating Environmental Control and Hyperspectral Imaging to Assess Light and Nutrient Effects on Lettuce Post-Harvest Quality in Vertical Farming
- Franzoni, G.; Mirabella, S.; Dabek, A.; Ferro, N.; Antona, A.; Carlessi, M.; Cinquemani, S.; Matteucci, M.; Cocetta, G.; Perotto, S.
Integrating Environmental Control and Hyperspectral Imaging to Assess Light and Nutrient Effects on Lettuce Post-Harvest Quality in Vertical Farming
- 32/2026** Antonietti, P.F.; Bonizzoni, F.; Perugia, I.; Verani, M.
A Multilevel Monte Carlo Virtual Element Method for Uncertainty Quantification of Elliptic Partial Differential Equations