

ARTICLE

Hybrid Physics-Informed Graph Neural Network Surrogate for Multi-Physics Optimization of Net-Zero Prefabricated Modular Buildings Driven by BIM Semantic Data

Masood Karamoozian¹, Zarrin Mahdavi pour², Mohammed Ameen³, Faisal Binzagr⁴,
Abdolraheem Khader^{2,*}, Ahmed Hamza Osman³ and Ali Ahmed⁴

¹School of Architecture, Southeast University, Nanjing, China

²School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China

³Department of Information Systems, Faculty of Computing and Information Technology-Rabigh, King Abdulaziz University Rabigh, Jeddah, Saudi Arabia

⁴Department of Computer Science, Faculty of Computing and Information Technology-Rabigh, King Abdulaziz University Rabigh, Jeddah, Saudi Arabia

*Corresponding Author: Abdolraheem Khader. Email: abdolraheem@njust.edu.cn

Received: 11 May 2026; Accepted: 02 July 2026; Published: 28 August 2026

ABSTRACT: Achieving net-zero carbon emissions in the building sector requires computational tools that are simultaneously efficient, physically rigorous, and scalable to complex multi-domain design spaces. Traditional physics-based simulators such as EnergyPlus and finite-element analysis (FEM) packages deliver high fidelity but are computationally prohibitive for large-scale parametric exploration and real-time multi-objective optimization. This study introduces a hybrid Physics-Informed Graph Neural Network (PI-GNN) surrogate that directly leverages semantic Building Information Modeling (BIM) data to enable rapid, accurate, and physically consistent multi-physics prediction and optimization of prefabricated modular buildings. Building components and their physical and functional relationships are encoded as nodes and typed edges in a heterogeneous directed graph, processed by a multi-layer graph attention network (GAT) backbone augmented with physics-informed loss terms. These constraints explicitly enforce energy and mass conservation, transient heat transfer, structural equilibrium, and carbon balance, ensuring physically plausible predictions even under limited training data. Trained on a dataset of more than 852 BIM-derived parametric models coupled with high-fidelity multi-physics simulations (EnergyPlus, OpenSeesPy FEM, One Click LCA), the surrogate achieves mean absolute percentage errors (MAPE) below 7% across energy demand, embodied carbon, operational carbon, and peak structural stress, while delivering inference speedups of up to 1200× over conventional simulation workflows. Integrated with NSGA-II multi-objective evolutionary optimization, the framework identifies Pareto-optimal designs that reduce operational energy by up to 52% and embodied carbon by 28%–45% relative to baseline configurations. Ablation studies confirm that physics-informed regularization reduces prediction errors by 40%–55% in low-data regimes. The proposed PI-GNN offers a generalizable, interpretable, and physics-consistent surrogate modeling paradigm, bridging BIM semantics with deep learning to accelerate sustainable design and support the global transition toward net-zero prefabricated buildings.

KEYWORDS: Building information modeling; physics-informed neural networks; graph neural networks; prefabricated modular buildings; net-zero buildings; multi-physics optimization; surrogate modeling; NSGA-II; sustainable design

1 Introduction

The building sector is among the largest contributors to global energy consumption and greenhouse gas emissions, accounting for approximately 30% of total final energy use and 26% of energy-related CO₂ emissions worldwide [1,2]. Meeting internationally agreed climate targets, most notably the Paris Agreement goal of limiting warming to 1.5°C–2°C above pre-industrial levels, demands a rapid and systemic transition toward net-zero buildings. Prefabricated modular construction has emerged as a particularly promising pathway in this context, offering factory-controlled quality, minimised on-site waste, faster assembly, and substantially improved opportunities for lifecycle carbon reduction and end-of-life material reuse [3].

1.1 Limitations of Conventional Simulation Approaches

Realising true net-zero performance in prefabricated buildings requires holistic optimisation spanning multiple interacting physical domains simultaneously: thermal efficiency, structural safety, embodied and operational carbon, lifecycle cost, and indoor environmental quality. For decades, the primary tools for such evaluations have been high-fidelity physics-based simulators such as EnergyPlus for building energy analysis [4,5], finite element analysis (FEA) platforms for structural assessment [6,7], and dedicated lifecycle assessment (LCA) tools such as One Click LCA [8]. While these simulators deliver high-accuracy predictions grounded in well-established governing equations, they impose severe computational penalties that fundamentally limit their utility in large-scale design exploration.

A single coupled EnergyPlus–FEA evaluation of a multi-module building may require minutes to several hours of computation, rendering iterative parametric studies across thousands of design variants practically infeasible. The challenge is compounded in modular construction contexts, where combinatorial complexity arises from the combinatorial space of module layouts, connection strategies, material selections, and climate scenarios [3,9]. Furthermore, conventional simulation workflows operate on geometric or mesh-based inputs and are poorly equipped to directly consume the semantically rich data encoded in modern Building Information Modeling (BIM) representations, which capture hierarchical assembly relationships, material ontologies, zone-level functional semantics, and multi-physics interaction topologies [10]. The resulting translation overhead—from BIM to simulator-specific input formats—introduces additional time costs and potential information loss. Collectively, these limitations make conventional high-fidelity tools inadequate as the primary computational engine for real-time multi-objective optimisation at the early design stage.

1.2 Advances in Data-Driven and Surrogate Modeling

To alleviate the computational burden of physics-based simulators, the past decade has witnessed rapid growth in data-driven surrogate modeling for building performance prediction [11,12]. Early surrogates employed shallow regression methods (e.g., support vector regression, Gaussian processes) trained on relatively small simulation datasets [13]. The emergence of deep learning substantially elevated surrogate fidelity: convolutional neural networks (CNNs) capture spatial patterns in floor-plan images, recurrent networks model temporal occupancy and HVAC dynamics, and multi-layer perceptrons (MLPs) approximate energy demand from tabular building descriptors [9,14]. However, these approaches share a critical weakness: they are purely data-driven and impose no structural constraints on the predictions. As a consequence, trained models frequently violate fundamental physical laws—such as energy conservation or static equilibrium—particularly when generalising to design configurations outside their training distribution. Prediction errors can propagate silently and mislead optimisation routines toward physically infeasible solutions.

1.3 Emerging Graph-Based and Physics-Informed Approaches

Two complementary research directions have emerged to address these deficiencies. First, Graph Neural Networks (GNNs) provide a natural and expressive representation for buildings, whose components (zones, structural elements, HVAC subsystems, connections) and their physical or functional relationships can be directly encoded as nodes and edges in a heterogeneous graph [10,15]. GNN-based surrogates have demonstrated strong performance in urban-scale energy modeling, construction scheduling, and structural assessment by exploiting relational inductive biases inherent in building topologies [16–18]. Second, Physics-Informed Neural Networks (PINNs), introduced by Raissi et al. [19], embed residuals of governing partial differential equations directly into the training loss, ensuring that predictions remain consistent with thermodynamic, mechanical, and continuity principles even under data scarcity. Extensions of PINNs to graph domains—Physics-Informed Graph Neural Networks (PI-GNNs)—inherit the relational expressiveness of GNNs while enforcing conservation laws at the level of individual edges and nodes [16,20,21]. Recent applications of PI-GNNs have demonstrated superior robustness and generalisability over purely data-driven baselines in building thermal dynamics, fluid networks, and urban energy systems [16,22,23].

1.4 Remaining Gaps and Motivation for the Proposed Framework

Despite these advances, critical gaps persist. Existing surrogate models in the building domain predominantly focus on isolated physical domains—thermal performance alone, or structural response alone—and lack integrated frameworks that jointly predict energy, carbon, structural, and cost metrics from a single shared representation [24,25]. Physics-informed approaches in this domain have largely been demonstrated on single-zone thermal problems or simple network topologies, and their extension to the combinatorial heterogeneity of prefabricated modular assemblies remains largely unexplored [9,22]. Moreover, existing methods rarely consume BIM semantic data directly; instead, they require manual feature engineering or geometric pre-processing, foregoing the rich relational and hierarchical information that BIM encodes [10,26]. Finally, the integration of physics-consistent surrogates into real-time multi-objective optimisation loops for sustainable modular building design has received insufficient attention [13,24].

This study addresses these gaps by introducing a hybrid Physics-Informed Graph Neural Network (PI-GNN) surrogate framework that directly leverages semantic BIM data for rapid, accurate, and physically consistent multi-physics prediction and multi-objective optimisation of net-zero prefabricated modular buildings. The framework converts IFC-format BIM models into heterogeneous directed graphs, processes them through multi-head graph attention layers, and enforces energy conservation, heat transfer, structural equilibrium, and carbon balance constraints via physics-informed loss terms. Integrated with the NSGA-II evolutionary optimiser [27], the surrogate enables real-time exploration of Pareto-optimal design spaces balancing operational energy, embodied carbon, lifecycle cost, and thermal comfort.

The principal contributions of this work are:

1. **BIM-native heterogeneous graph representation:** A systematic conversion pipeline from IFC BIM models to heterogeneous directed graphs that preserves geometric, material, topological, and multi-physics relational information at node and edge levels, enabling end-to-end learning directly from BIM semantics.
2. **Hybrid PI-GNN surrogate with multi-task physics regularisation:** A novel architecture combining multi-head graph attention networks with physics-informed loss terms (thermal dynamics, structural equilibrium, carbon balance, and HVAC continuity), achieving mean absolute percentage errors below 7% across energy, carbon, and structural metrics while delivering 400–1200× speedup over conventional EnergyPlus + FEM workflows.

3. **Integrated multi-objective optimisation:** Embedding of the PI-GNN surrogate within an NSGA-II optimisation loop enables practical early-stage identification of Pareto-optimal prefabricated designs that reduce operational energy by up to 52% and embodied carbon by 28%–45%, computed in near-real-time.
4. **Interpretability:** Attention-weight and saliency-map analyses demonstrate that the learned model emphasises physically meaningful components—thermal bridges for energy predictions, and connection nodes for structural stress—supporting designer trust and decision transparency.

The remainder of this paper is organised as follows. [Section 2](#) reviews related work across four thematic areas. [Section 3](#) details the methodology. [Section 4](#) presents experimental results and ablation studies. [Section 5](#) discusses implications and limitations, and [Section 6](#) concludes with directions for future work.

2 Related Work

This section reviews prior work across four thematic pillars that collectively motivate the proposed framework: (1) BIM-based graph representations for building modeling; (2) surrogate modeling and deep learning for building performance simulation; (3) physics-informed neural and graph network approaches; and (4) multi-objective optimisation for prefabricated and net-zero buildings. For each area, we identify the state of the art, highlight its limitations, and establish how the proposed PI-GNN framework advances beyond existing work.

2.1 BIM-Based Graph Representations for Building Modeling

Building Information Modeling (BIM) provides semantically rich, machine-readable representations of buildings that encode geometry, material properties, assembly hierarchies, zone-level functional semantics, and inter-component relationships in standardised formats such as IFC [10,26]. Early computational uses of BIM were primarily focused on visualisation, clash detection, and quantity take-off, relying on rule-based or relational database approaches [26,28]. More recent research has explored converting IFC models into graph-structured data to enable machine-learning-based reasoning. Gan [10] introduced a BIM-based graph data model for automatic generative design of modular buildings, demonstrating that heterogeneous graph representations can capture containment and spatial adjacency relations between building elements. Li et al. [3] integrated BIM with graph convolutional neural networks for automated generative design and prefabrication of precast buildings. Zhang et al. [17] applied GNN-driven spatial semantic reasoning to BIM-based construction scheduling, exploiting the relational structure of IFC data to improve scheduling optimisation. Nabrotzky [18] surveyed GNN applications for structural engineering workflows, emphasising the role of IFC-derived graphs in improving BIM interoperability. Austern et al. [29] leveraged GNNs to incorporate context into BIM-derived data for building element classification, and Hong et al. [30] demonstrated graph-based construction scheduling without explicit BIM requirements by encoding topological dependencies.

Despite this progress, the majority of BIM-GNN efforts address isolated tasks—classification, scheduling, or geometry reconstruction—and do not extend to multi-physics performance prediction. Furthermore, most methods extract only subset features from BIM (e.g., geometric or connectivity attributes) and do not systematically exploit the full semantic richness of IFC models, including material thermophysical properties, structural parameters, and functional zone semantics simultaneously. The proposed framework addresses this gap by developing a comprehensive IFC-to-heterogeneous-graph conversion pipeline that integrates all physically relevant attributes into a unified representation for multi-task prediction.

2.2 Surrogate Modeling and Deep Learning for Building Performance Simulation

The computational cost of high-fidelity building simulation tools has motivated extensive research into surrogate models as efficient approximations [11,12,14]. Shallow machine learning approaches, including Gaussian processes, random forests, and support vector regression, were among the first surrogates applied to building energy prediction, achieving acceptable accuracy on small, homogeneous datasets but exhibiting limited scalability and generalisation.

The advent of deep learning surrogates substantially raised the performance ceiling. Convolutional neural networks have been applied to spatial floor-plan representations for energy demand estimation, while recurrent architectures capture temporal dynamics in HVAC control and occupancy-driven loads [12,14]. Kerdan and Morillón Gálvez [13] combined artificial neural networks with genetic algorithms in an open-source building energy design optimisation framework, demonstrating early integration of surrogate prediction with evolutionary search. Fellah et al. [14] and Woldegiyorgis et al. [12] reviewed the breadth of machine learning methods deployed for energy efficiency and building performance improvement, confirming strong performance gains over physics-based simulation in terms of computational speed.

Graph-based surrogates have more recently emerged as particularly well-suited to building modeling because they naturally encode the relational and hierarchical structure of building assemblies. Shan et al. [16] developed a physics-informed and explainable GNN surrogate for generalizable urban building energy modeling at district scale. Wang et al. [24] proposed a GNN-enhanced surrogate specifically for optimizing carbon-efficient building thermal insulation systems. Amani [25] demonstrated simulation-based design optimisation for minimising energy consumption through thermal insulation selection, highlighting the continuing relevance of energy-focused surrogate pipelines.

However, a recurring limitation across this body of work is the focus on single-domain prediction. Surrogates for energy, carbon, structural performance, and cost are typically developed in isolation, preventing holistic multi-objective design exploration. Additionally, most surrogates do not enforce physical consistency constraints during training, raising concerns about prediction reliability in out-of-distribution scenarios [11]. The present work directly addresses both limitations by developing a multi-task surrogate with embedded physics regularisation.

2.3 Physics-Informed Neural and Graph Network Approaches

Physics-Informed Neural Networks (PINNs), pioneered by Raissi et al. [19], established the paradigm of embedding residuals of governing partial differential equations into the neural network training loss. This approach produces models whose predictions are constrained to satisfy known physical laws—including heat transfer, mass conservation, and momentum balance—even under limited training data. Early PINN applications focused on fluid mechanics and solid mechanics benchmarks; subsequent work rapidly extended the paradigm to building and energy systems. Gokhale et al. [31] applied PINNs to control-oriented building thermal modeling. Chen et al. [23] employed PINNs for building thermal modeling with demand response control. Jiang et al. [22,32] developed physics-informed modularised neural networks incorporating thermal priors for advanced building control and energy modeling. Ma et al. [11] provided a comprehensive review of physics-informed machine learning for building energy modeling.

The extension of physics-informed learning to graph-structured domains—Physics-Informed GNNs (PI-GNNs)—is more recent but rapidly advancing. Shukla et al. [33] developed scalable algorithms for physics-informed neural and graph networks. Thangamuthu et al. [20] systematically analysed the performance of PI-GNNs for dynamical systems. Ashraf et al. [21] applied PI-GNNs to water distribution systems, enforcing hydraulic conservation laws at the edge level. Chen et al. [34] demonstrated PI-GNNs for predicting fluid flow in porous media. Shan et al. [16] applied physics-informed and explainable GNNs to urban building energy modeling.

Despite this progress, PI-GNN applications in the building domain remain largely fragmented: existing methods address single physical phenomena (predominantly thermal) and have not been demonstrated for joint multi-physics prediction covering energy, structural, and carbon domains simultaneously. Critically, no prior work has developed a PI-GNN that ingests BIM semantic graphs directly and embeds them within a multi-objective optimisation loop. The proposed framework fills this gap by extending PI-GNNs to heterogeneous, multi-relational BIM graphs with multi-task physics-informed regularisation.

2.4 Multi-Objective Optimisation for Prefabricated and Net-Zero Buildings

Optimisation of building performance involves balancing multiple competing objectives—reducing operational energy and embodied carbon while minimising cost and maintaining structural safety and occupant comfort—objectives that frequently trade off against one another [13]. Multi-objective evolutionary algorithms, and NSGA-II in particular [27], have become a standard approach for generating Pareto-optimal design sets in building optimisation. Tian et al. [35] advanced fast clustering-based evolutionary algorithms for large-scale sparse multi-objective problems, while Bian et al. [36] demonstrated archive-based speed-ups for multi-objective evolutionary search.

The primary bottleneck in surrogate-coupled evolutionary optimisation is the computational cost of objective evaluations. When each evaluation requires a full EnergyPlus and FEM simulation run, even small populations over a modest number of generations may demand days of wall-clock time. Surrogate-assisted approaches have therefore attracted considerable attention: Kerdan and Morillón Gálvez [13] integrated neural network surrogates with genetic algorithms for building energy design optimisation; Wang et al. [24] coupled a GNN surrogate with rapid carbon-efficient insulation optimisation; and Amani [25] demonstrated simulation-based optimisation for thermal insulation selection at building scale.

In the domain of prefabricated modular construction, GNNs combined with meta-heuristic search have enabled generative structural design [3], while physics-informed methods have been applied to net-zero energy flexibility quantification [22]. Nevertheless, integrated frameworks that (i) consume BIM semantic graphs as direct input, (ii) perform simultaneous multi-physics prediction under physics constraints, and (iii) embed the resulting surrogate within a full multi-objective evolutionary optimisation pipeline targeting net-zero prefabricated buildings remain notably absent from the literature.

Building upon the advances reviewed above, the proposed hybrid PI-GNN framework directly ingests BIM-derived heterogeneous graphs, enforces multi-domain physics constraints during training, and embeds the surrogate within an NSGA-II optimisation loop—comprehensively addressing the limitations identified across all four thematic areas. The novelty of the contribution lies in the convergence of these four streams within a single, end-to-end, and scalable framework specifically designed for net-zero prefabricated modular buildings.

3 Methodology

This section presents the proposed hybrid Physics-Informed Graph Neural Network (PI-GNN) surrogate framework for rapid, accurate, and physically consistent multi-physics prediction and optimization of net-zero prefabricated modular buildings. The framework proceeds through four tightly coupled stages whose mathematical components build progressively toward the final multi-objective optimization integration: (1) BIM-to-graph conversion, which translates semantically rich IFC data into a heterogeneous directed graph suitable for relational learning; (2) PI-GNN architecture design, which defines the embedding, message-passing, and physics-informed regularization that collectively enable the surrogate to honour governing physical laws; (3) dataset generation and training, which describes how simulation-derived ground-truth labels are used to fit the surrogate; and (4) multi-objective optimization integration, which embeds the trained surrogate in an evolutionary search loop. The overall pipeline is illustrated in Fig. 1.




Figure 1: Workflow of the proposed PI-GNN surrogate framework: BIM input, graph construction, physics-informed training, surrogate prediction, and multi-objective optimization.

3.1 BIM-to-Graph Conversion

3.1.1 Motivation and Graph Definition

Prefabricated modular buildings are inherently relational systems: thermal transfer occurs between adjacent zones, structural forces propagate through connection joints, HVAC flow is distributed across duct networks, and material assemblies are organised in containment hierarchies. These relational properties motivate a graph-theoretic abstraction that directly mirrors the topological and physical structure of the building [3,10,26]. BIM models stored in IFC format are processed using IfcOpenShell [37] to extract semantic, geometric, and topological information. The building is abstracted as a heterogeneous directed graph:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{T}), \quad (1)$$

where $\mathcal{V}$ is the set of nodes representing physical entities—prefabricated modules, structural connections, material layers, thermal zones, and HVAC subsystems—each assigned a type $\mathcal{T}_i \in \mathcal{T}$; $\mathcal{E}$ is the set of directed edges encoding typed multi-physics interactions (structural, thermal, assembly, and fluid-flow relationships); and $\mathcal{T}$ is the finite vocabulary of node types. The directed character of $\mathcal{E}$ is essential: it encodes physical directionality—for instance the direction of gravity loading in structural edges or the positive direction of net heat flux in thermal edges—ensuring that message-passing later propagates information in physically meaningful directions.

3.1.2 Node and Edge Feature Vectors

Each node $v_i \in \mathcal{V}$ carries a heterogeneous feature vector:

$$\mathbf{x}_i = \left[\underbrace{g_i}_{\text{geometry}}, \underbrace{\lambda_i, \rho_i, c_{p,i}, f_i^{\text{ec}}}_{\text{thermal/material}}, \underbrace{E_i, A_i^{\text{cs}}}_{\text{structural}}, \underbrace{z_i, o_i}_{\text{semantic}} \right], \quad (2)$$

where g_i encodes geometric descriptors (volume, surface area, aspect ratio); λ_i , ρ_i , and $c_{p,i}$ are the thermal conductivity, mass density, and specific heat capacity of the element's primary material; f_i^{ec} is the embodied carbon factor (kg CO₂e per unit mass); E_i and A_i^{cs} are Young's modulus and cross-sectional area for structural elements; and z_i , o_i encode zone type and occupancy profile, respectively. This multi-domain feature encoding is one of the key novelties of the framework: by representing all physical attributes of a building component within a single node vector, the GNN can learn cross-domain dependencies—for example, how a material's thermal conductivity jointly influences both heat transfer and the embodied carbon of insulation choices.

Each directed edge $(v_i, v_j) \in \mathcal{E}$ carries a feature vector:

$$\mathbf{e}_{ij} = \left[A_{ij}, R_{ij}, k_{ij}^{\text{mech}}, r_{ij}^{\text{flow}} \right], \quad (3)$$

where A_{ij} is the contact area between elements i and j , R_{ij} is the effective thermal resistance of the interface, k_{ij}^{mech} is the mechanical stiffness of the connection joint, and r_{ij}^{flow} is the fluid-flow resistance for HVAC duct or pipe edges. These edge features directly parameterise the physical interaction laws that are later enforced by the physics-informed loss terms: A_{ij} and R_{ij} appear explicitly in the discretised heat conduction term (Eq. (9)), and k_{ij}^{mech} appears in the structural equilibrium residual (Eq. (11)). An illustrative example is shown in Fig. 2, and the full conversion procedure is formalised in Algorithm 1.

Table 1 summarizes the node types extracted from BIM and their key feature attributes.

This representation preserves the discrete and relational nature of prefabricated systems and provides the structured inductive bias that makes subsequent graph attention-based message passing both efficient and physically interpretable [17,18].




Figure 2: Example of BIM-to-graph conversion: **(left)** BIM model snippet of two connected modules; **(right)** corresponding heterogeneous graph with node types and edge types.

Algorithm 1: BIM-to-graph conversion algorithm

Require: IFC file of prefabricated modular building

Ensure: Heterogeneous directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{T})$

- 1: Initialize empty graph $\mathcal{G}$
- 2: Extract IFC entities: `IfcBuildingStorey`, `IfcSpace`, `IfcBuildingElement`, `IfcRelConnectsElements`, `IfcRelSpaceBoundary`
- 3: **for** each `IfcBuildingElement` **do**
- 4: Create node v_i with type $\mathcal{T}_i$
- 5: Assign feature vector $\mathbf{x}_i$ per Eq. (2)
- 6: Add v_i to $\mathcal{V}$
- 7: **end for**
- 8: **for** each relational entity **do**
- 9: Identify source node v_i and target node v_j

(Continued)

Algorithm 1 (continued)

```

10:   Determine edge type (structural, thermal, assembly, flow)
11:   Compute edge feature vector  $\mathbf{e}_{ij}$  per Eq. (3)
12:   Add directed edge  $(v_i, v_j)$  with features  $\mathbf{e}_{ij}$  to  $\mathcal{E}$ 
13: end for
14: Enforce directionality based on physical rules (gravity direction, positive heat flux)
15: return  $\mathcal{G}$ 

```

Table 1: Node types in the heterogeneous building graph and associated feature dimensions.

Node Type	IFC Entity	Dim.	Key Attributes
Module	IfcBuildingElement	24	Geometry, material, structural, semantic
Connection	IfcRelConnectsElements	12	Contact area, stiffness, joint type
Thermal Zone	IfcSpace	16	Volume, U-values, setpoints, occupancy
Material	IfcMaterial	8	λ, ρ, c_p , ECF, thickness
HVAC Unit	IfcFlowElement	10	COP, rated capacity, flow rate

3.2 Physics-Informed Graph Neural Network Architecture**3.2.1 Architecture Overview**

The proposed surrogate model is a hybrid PI-GNN that processes $\mathcal{G}$ to predict multi-physics performance indicators at both the node level (e.g., zone temperatures, member stresses) and the graph level (e.g., total energy demand, lifecycle carbon). The overall architecture is shown in Fig. 3.

The architecture comprises three principal computational stages: (i) embedding layers that project heterogeneous raw features into a common latent space; (ii) graph attention blocks that propagate and aggregate information across the graph topology; and (iii) physics-informed regularisation that constrains learned representations to honour governing physical laws.

3.2.2 Stage 1: Embedding Layers

Because node and edge features in $\mathcal{G}$ are heterogeneous—spanning geometric scalars, material properties, structural parameters, and functional semantics—they cannot be processed by a single shared linear transformation without sacrificing representational capacity. Separate Multi-Layer Perceptron (MLP) encoders are therefore applied to node and edge features independently, projecting them into a common latent space of dimension d (set to 256 in all experiments):

$$\mathbf{h}_i^{(0)} = \text{MLP}_n(\mathbf{x}_i), \quad \mathbf{e}_{ij}^{(0)} = \text{MLP}_e(\mathbf{e}_{ij}). \quad (4)$$

Here $\mathbf{h}_i^{(0)} \in \mathbb{R}^d$ is the initial latent representation of node v_i , and $\mathbf{e}_{ij}^{(0)} \in \mathbb{R}^d$ is the initial edge embedding. Both MLP_n and MLP_e consist of two fully connected layers with ReLU activations and layer normalisation. Projecting into a shared d -dimensional space is critical: it enables the subsequent graph attention layers to concatenate node and edge embeddings into a unified message, allowing the attention mechanism to consider both the state of connected nodes and the physical properties of the interface between them when computing interaction weights. This is particularly important in multi-physics settings, where the nature of an interaction (thermal, structural, or flow) must influence how information is aggregated.




Figure 3: Hybrid PI-GNN architecture: node/edge embedding → multi-layer graph attention blocks → physics-informed loss regularization → multi-task output heads for energy, carbon, stress, and cost predictions.

3.2.3 Stage 2: Graph Attention Blocks

L layers of multi-head Graph Attention Networks (GAT) [15] are applied to propagate information across the graph topology. At each layer l , the attention coefficient between node v_i and neighbour v_j is computed as:

$$\alpha_{ij}^{(l)} = \frac{\exp\left(\text{LeakyReLU}\left(\mathbf{a}^\top [\mathbf{W}^{(l)} \mathbf{h}_i^{(l)} \parallel \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \parallel \mathbf{e}_{ij}^{(l)}]\right)\right)}{\sum_{k \in \mathcal{N}(i)} \exp\left(\text{LeakyReLU}\left(\mathbf{a}^\top [\mathbf{W}^{(l)} \mathbf{h}_i^{(l)} \parallel \mathbf{W}^{(l)} \mathbf{h}_k^{(l)} \parallel \mathbf{e}_{ik}^{(l)}]\right)\right)}, \quad (5)$$

where $\parallel$ denotes concatenation, $\mathbf{a}$ is a learnable attention vector, $\mathbf{W}^{(l)}$ are layer-specific weight matrices. Critically, the inclusion of edge features $\mathbf{e}_{ij}^{(l)}$ in the attention scoring allows the model to weight interactions differently based on the physical type and magnitude of each interface. The node representation update is:

$$\mathbf{h}_i^{(l+1)} = \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(l)} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)}\right), \quad (6)$$

where σ denotes the ELU activation. With K attention heads, representations are concatenated as:

$$\mathbf{h}_i^{(l+1)} = \left\| \sum_{k=1}^K \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij,k}^{(l)} \mathbf{w}_k^{(l)} \mathbf{h}_j^{(l)} \right) \right\|, \quad (7)$$

where $K = 8$ heads are used in all but the final layer, where averaging is applied. Edge representations are also updated between layers via a residual MLP:

$$\mathbf{e}_{ij}^{(l+1)} = \mathbf{e}_{ij}^{(l)} + \text{MLP}_e^{(l)} \left([\mathbf{e}_{ij}^{(l)} \parallel \mathbf{h}_i^{(l)} \parallel \mathbf{h}_j^{(l)}] \right), \quad (8)$$

enabling edge representations to adapt to evolving node contexts across layers. With $L = 4$ layers and $K = 8$ attention heads, each node ultimately aggregates information from up to its fourth-order neighbourhood, capturing long-range physical dependencies such as the cascading thermal effect of an insulation upgrade across multiple connected zones.

3.2.4 Stage 3: Physics-Informed Regularisation

After L message-passing layers, each node holds a latent representation $\mathbf{h}_i^{(L)}$. A purely data-driven surrogate would proceed directly to output regression; however, without physical constraints the learned mappings may violate conservation laws—particularly when extrapolating to design configurations outside the training distribution [11,19]. The PI-GNN addresses this by augmenting the training objective with physics residual terms derived from the governing equations of each modelled physical domain. The following derivations establish the form, assumptions, boundary conditions, and residual computation for each component.

Thermal Energy Balance ($\mathcal{L}_{\text{energy}}$)

The governing equation for transient heat conduction is derived from the first law of thermodynamics applied to a lumped-capacitance control volume per thermal zone node. The following assumptions are made: (i) uniform temperature within each zone node (lumped thermal mass approximation); (ii) linear Fourier conduction across interfaces; (iii) no explicit zone-to-zone radiation exchange (absorbed into a combined convective-conductive film coefficient); and (iv) prescribed outdoor air temperatures as Dirichlet boundary conditions at envelope-adjacent nodes. Under these assumptions, the discretised heat balance at node v_i at time step t is:

$$\rho_i c_{p,i} V_i \frac{T_i^{t+1} - T_i^t}{\Delta t} = \sum_{j \in \mathcal{N}(i)} \frac{\lambda_{ij} A_{ij}}{d_{ij}} (T_j^t - T_i^t) + Q_i^t + R_{\text{residual},i}, \quad (9)$$

where V_i is the volume of thermal zone node v_i , obtained from the geometric descriptor g_i in the node feature vector (Eq. (2)), and is introduced here so that both sides of Eq. (9) are dimensionally consistent, with each term expressed in units of power (W); λ_{ij} , A_{ij} , d_{ij} are the effective thermal conductivity, contact area, and inter-node distance for edge e_{ij} obtained from $\mathbf{e}_{ij}$; Q_i^t is the internal heat gain; and $R_{\text{residual},i}$ is the physics residual—the quantity that must equal zero if the PI-GNN's predicted temperature field exactly satisfies the governing equation. Deviations from zero indicate violations of thermodynamic consistency. The energy physics loss aggregates these residuals:

$$\mathcal{L}_{\text{energy}} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \|R_{\text{residual},i}^t\|^2. \quad (10)$$

Minimising $\mathcal{L}_{\text{energy}}$ enforces thermodynamic consistency as a soft constraint, guiding predictions toward physically plausible temperature distributions.

Structural Equilibrium ($\mathcal{L}_{\text{structural}}$)

The structural physics constraint is derived from Newtonian static equilibrium. Assumptions: (i) linear elastic material behaviour (Hookean constitutive law); (ii) small-deformation kinematics; and (iii) zero-displacement Dirichlet boundary conditions at foundation nodes. Under quasi-static loading, nodal equilibrium requires:

$$\sum_{j \in \mathcal{N}(i)} k_{ij}^{\text{mech}} (\mathbf{u}_j - \mathbf{u}_i) + \mathbf{b}_i = \mathbf{0} + \mathbf{R}_{\text{structural},i}, \quad (11)$$

where k_{ij}^{mech} is the joint stiffness from $\mathbf{e}_{ij}$; $\mathbf{u}_i, \mathbf{u}_j$ are the PI-GNN-predicted nodal displacement vectors; $\mathbf{b}_i$ is the applied body force. The structural physics loss is:

$$\mathbf{R}_i^{\text{struct}} = \sum_{j \in \mathcal{N}(i)} \hat{\mathbf{f}}_{ij} + \mathbf{b}_i, \quad \mathcal{L}_{\text{structural}} = \frac{1}{|\mathcal{V}_{\text{struct}}|} \sum_{i \in \mathcal{V}_{\text{struct}}} \|\mathbf{R}_i^{\text{struct}}\|^2. \quad (12)$$

Note that $\mathcal{L}_{\text{structural}}$ is computed on a different node subset than $\mathcal{L}_{\text{energy}}$, reflecting the heterogeneous graph's ability to apply domain-specific physics constraints selectively to relevant node types.

Carbon Balance Proxy ($\mathcal{L}_{\text{carbon}}$)

Lifecycle carbon predictions must satisfy an additive consistency constraint: the total predicted lifecycle carbon must equal the sum of embodied and operational contributions. A proxy residual is defined as:

$$R_{\text{carbon}} = \hat{C}_{\text{total}} - \left(\sum_{i \in \mathcal{V}} f_i^{\text{ec}} m_i + \hat{E}_{\text{op}} \cdot \varepsilon_{\text{grid}} \right), \quad (13)$$

where f_i^{ec} is the embodied carbon factor from $\mathbf{x}_i$; m_i is element mass; $\hat{E}_{\text{op}}$ is the PI-GNN-predicted operational energy over a 50-year lifecycle; and $\varepsilon_{\text{grid}}$ is the ASHRAE zone-specific grid carbon intensity. The carbon consistency loss:

$$\mathcal{L}_{\text{carbon}} = \|R_{\text{carbon}}\|^2. \quad (14)$$

This term couples the energy and carbon prediction heads, preventing the PI-GNN from predicting a material system whose implied carbon footprint contradicts the material composition encoded in the node features.

HVAC Flow Continuity ($\mathcal{L}_{\text{continuity}}$)

For flow-type edges, mass conservation under steady-state incompressible flow requires that net mass flux into any interior junction node be zero. With prescribed inlet/outlet mass flow rates as boundary conditions at terminal nodes, the continuity residual is:

$$R_{\text{continuity},i} = \sum_{j \in \mathcal{N}_{\text{flow}}(i)} \dot{m}_{ij} - \dot{m}_{ij}^{\text{pred}}, \quad (15)$$

where $\dot{m}_{ij}$ is the reference mass flow rate inferred from edge flow resistance r_{ij}^{flow} (via the Hagen–Poiseuille relation), and $\dot{m}_{ij}^{\text{pred}}$ is the PI-GNN-predicted flow. The continuity loss:

$$\mathcal{L}_{\text{continuity}} = \frac{1}{N_f} \sum_{i=1}^{N_f} \|R_{\text{continuity},i}\|^2. \quad (16)$$

3.2.5 Total Physics-Informed Training Loss

The four physics residual terms and the data fidelity term are combined into the complete training objective:

$$\mathcal{L} = \underbrace{\frac{1}{N} \sum_{n=1}^N \|\hat{y}^{(n)} - y^{(n)}\|^2}_{\mathcal{L}_{\text{data}}} + \lambda_1 \mathcal{L}_{\text{energy}} + \lambda_2 \mathcal{L}_{\text{structural}} + \lambda_3 \mathcal{L}_{\text{carbon}} + \lambda_4 \mathcal{L}_{\text{continuity}}. \quad (17)$$

The data loss drives statistical accuracy on seen configurations; the four physics losses enforce physical consistency as soft regularisers, penalising physically impossible predictions in regions of the design space where training data are absent. Eq. (17) unifies four independent physical domains—thermodynamics, structural mechanics, lifecycle carbon accounting, and fluid mechanics—within a single differentiable training objective, enabling the PI-GNN to simultaneously learn across all domains while maintaining cross-domain physical consistency.

Logical build-up of Eqs. (1)–(10)

The mathematical development follows a deliberate progression. Eqs. (1)–(3) establish the input representation. Eqs. (4)–(6) define the forward computation of the GNN. Eqs. (9), (11), (13) and (15) derive the physics residuals for each domain. Eq. (17) integrates all residuals into the unified training loss, and Eq. (18) in Section 3.4 completes the methodology by embedding this surrogate in the multi-objective optimisation problem statement.

Role of the loss weights λ_k and sensitivity analysis

Each weight λ_k controls the penalty strength of its corresponding physical domain. The selected values— $\lambda_1 = 0.5$ (energy), $\lambda_2 = 0.3$ (structural), $\lambda_3 = 0.2$ (carbon), $\lambda_4 = 0.15$ (continuity)—were determined by grid search over $[0.1, 1.0]$ for each weight independently, using validation-set MAPE as the selection criterion. Sensitivity analysis conducted by varying each λ_k over ± 0.2 around its selected value while holding the others fixed shows that validation MAPE varies by less than 0.8% across this range, confirming relative robustness to weight perturbation. A linear warmup schedule increases physics loss weights from 0 to their target values over the first 50 epochs, allowing the model to first fit the data before satisfying physical constraints.

3.2.6 Output Heads

Task-specific MLPs regress per-node fields (temperature T_i , stress σ_i) or apply global pooling for whole-building metrics (energy [kWh], carbon [kg CO₂e], cost [\$]). Per-node outputs are substituted back into the physics residual equations during training; graph-level predictions serve as objective function evaluations in the NSGA-II loop (Section 3.4).

The training and inference process is summarized in Algorithm 2.

Algorithm 2: PI-GNN training and inference algorithm

Require: Training graphs $\{\mathcal{G}^{(k)}\}_{k=1}^K$, labels $\{y^{(k)}\}$, hyperparameters $(\lambda_k, d, L, \text{lr}, \text{epochs})$

Ensure: Trained PI-GNN model

1: Initialize parameters Θ (from pre-trained urban GNN weights [16])

(Continued)

Algorithm 2 (continued)

```

2: for each epoch do
3:   for each batch of graphs  $\mathcal{G}_b$  do
4:     Compute embeddings  $\mathbf{h}^{(0)}, \mathbf{e}^{(0)}$  via Eq. (4)
5:     for  $l = 0$  to  $L-1$  do
6:       Compute attention  $\alpha_{ij}^{(l)}$  via Eq. (5)
7:       Apply multi-head GAT:  $\mathbf{h}^{(l+1)}$  via Eq. (7)
8:       Update edge representations via Eq. (8)
9:     end for
10:    Compute predictions  $\hat{y}$  via output heads
11:    Compute losses:  $\mathcal{L}_{\text{data}}, \mathcal{L}_{\text{energy}}, \mathcal{L}_{\text{structural}}, \mathcal{L}_{\text{carbon}}, \mathcal{L}_{\text{continuity}}$ 
12:    Total loss  $\mathcal{L}$  via Eq. (17)
13:    Update  $\Theta$  using AdamW
14:  end for
15:  if validation loss stagnates then
16:    Early stop
17:  end if
18: end for
19: return Trained model
20: Inference: Input  $\mathcal{G} \rightarrow \text{PI-GNN} \rightarrow \text{predictions } \hat{y}$ 

```

The model is implemented in PyTorch Geometric [38], supporting batched graphs, GPU acceleration, and end-to-end differentiability.

3.3 Dataset Generation and Training

3.3.1 Parametric BIM Dataset

The parametric BIM dataset was synthetically generated following established generative-design workflows for modular and prefabricated buildings [3,10], implemented in practice using Revit Dynamo and Grasshopper scripts [39]; full generation methodology and composition statistics are provided in Section 4.1. The dataset covers:

- Module layouts (1–6 modules in width/depth, 1–4 stories),
- Building typologies (residential, office, educational),
- ASHRAE climate zones 1–8 (hot-humid to subarctic),
- Material and insulation levels (low/medium/high embodied carbon variants: concrete, timber, steel hybrid).

High-fidelity simulation labels were obtained through established multi-physics tools:

- **Energy and thermal comfort:** EnergyPlus v9.6 with TMY3 weather files [4,5],
- **Lifecycle carbon:** One Click LCA (embodied + 50-year operational emissions) [8],
- **Structural response:** OpenSeesPy FEM (peak von Mises stress under gravity/wind loads per ASCE 7) [6,7],
- **Lifecycle cost:** RSMeans-based approximations [40].

Table 2 provides a statistical summary of the dataset.

Table 2: Summary statistics of the parametric BIM dataset used for PI-GNN training and evaluation.

Property	Range/Value	Unit
Total models	852	–
Train/validation/test	596/128/128	Models
Climate zones (ASHRAE)	1A–8	–
Typologies	3 (residential, office, educational)	–
Graph nodes per model	50–350	Nodes
Graph edges per model	80–620	Edges
Annual energy demand	18.4–312.7	kWh/(m ² ·yr)
Embodied carbon	42.1–385.6	kg CO ₂ e/m ²
Peak von Mises stress	12.4–241.6	MPa
Simulation time/model	12–180	Minutes

3.3.2 Training Protocol and Hyperparameters

The dataset was split into 70% training (596 models), 15% validation (128 models), and 15% testing (128 models), stratified by climate zone and building typology. Training employed AdamW optimization ($\text{lr} = 5 \times 10^{-4} - 1 \times 10^{-3}$, weight decay = 10^{-5}), batch sizes of 16–32 graphs, and early stopping with patience of 30–50 epochs. Transfer learning from pre-trained urban/building energy GNNs provided initialization for faster convergence and improved generalization [16]. Table 3 summarizes the final hyperparameter configuration selected through grid search on the validation set.

Table 3: PI-GNN training hyperparameters and architecture configuration.

Hyperparameter	Search Range	Selected Value
<i>Architecture</i>		
Hidden dimension d	64, 128, 256	256
GAT layers L	2, 3, 4, 5	4
Attention heads K	4, 8	8
<i>Optimization</i>		
Learning rate	$10^{-4} - 10^{-2}$	5×10^{-4}
Weight decay	$10^{-6} - 10^{-4}$	10^{-5}
Batch size	8, 16, 32	16
Max epochs	300	300
Early stopping patience	20, 30, 50	30
<i>Physics loss weights</i>		
λ_1 (energy)	0.1–1.0	0.50
λ_2 (structural)	0.1–1.0	0.30
λ_3 (carbon)	0.1–0.5	0.20
λ_4 (continuity)	0.05–0.3	0.15

(Continued)

Table 3 (continued)

Hyperparameter	Search Range	Selected Value
Physics warmup epochs	20–100	50
<i>Hardware</i>		
GPU	–	NVIDIA RTX 4090 (24 GB VRAM)

3.4 Multi-Objective Optimization Integration

The trained PI-GNN surrogate was integrated with NSGA-II for Pareto optimization [27]. Design variables included module dimensions, material choices, insulation thickness, window-to-wall ratio (WWR), and orientation. The multi-objective optimization problem is formally stated as:

$$\begin{aligned}
 \min_{\mathbf{d}} \quad & \mathbf{f}(\mathbf{d}) = [E_{\text{op}}(\mathbf{d}), C_{\text{emb}}(\mathbf{d}), C_{\text{lc}}(\mathbf{d}), D_h(\mathbf{d})]^\top \\
 \text{subject to} \quad & \sigma_{\text{max}}(\mathbf{d}) \leq \sigma_{\text{yield}}/\gamma_M, \\
 & D_h(\mathbf{d}) \leq D_h^{\text{max}}, \\
 & \mathbf{d}_{\text{lb}} \leq \mathbf{d} \leq \mathbf{d}_{\text{ub}},
 \end{aligned} \tag{18}$$

where the optimization objectives were: operational energy E_{op} [kWh/year], embodied carbon C_{emb} [kg CO₂e], lifecycle cost C_{lc} [\$], and thermal discomfort hours D_h [h/year]. Constraints enforced structural safety ($\sigma_{\text{max}} < \sigma_{\text{yield}}$) and code compliance. The surrogate enabled second-level per-design evaluations, supporting populations of 100–200 individuals over 150–300 generations—completed in on the order of 1–3 h of wall-clock surrogate time using batched inference (see Section 4.4 for detailed timing).

The methodology thus forms a coherent analytical chain: the BIM-to-graph conversion (Eqs. (1)–(3)) provides a physically complete input representation; embedding and attention stages (Eqs. (4)–(8)) extract relational features; the physics-informed loss (Eqs. (9)–(17)) ensures that the surrogate respects thermodynamic, mechanical, carbon-balance, and fluid-continuity principles; and the multi-objective optimisation (Eq. (18)) translates surrogate predictions into actionable Pareto-optimal design solutions.

The NSGA-II surrogate-assisted optimization procedure is formalised in Algorithm 3.

Algorithm 3: Surrogate-assisted NSGA-II multi-objective optimization

Require: Trained PI-GNN f_Θ , BIM-to-graph converter ϕ , population size N_p , generations G

Ensure: Pareto-optimal design set $\mathcal{P}^*$

- 1: Initialize population $\mathcal{P}^{(0)}$ of N_p random design vectors $\{\mathbf{d}_i\}$
 - 2: **for** $g = 1$ to G **do**
 - 3: **for** each design $\mathbf{d}_i \in \mathcal{P}^{(g-1)}$ **do**
 - 4: Construct graph: $\mathcal{G}_i \leftarrow \phi(\mathbf{d}_i)$
 - 5: Evaluate objectives: $\mathbf{f}_i \leftarrow f_\Theta(\mathcal{G}_i)$ ▷ PI-GNN, $\mathcal{O}(\text{ms})$
 - 6: Check constraints: $\sigma_{\text{max},i}, D_{h,i}$
 - 7: **end for**
 - 8: Apply NSGA-II: non-dominated sorting, crowding distance
 - 9: Apply SBX crossover ($\eta_c = 20$), polynomial mutation ($\eta_m = 20$)
-

(Continued)

Algorithm 3 (continued)

```

10:   Select top  $N_p$  designs  $\rightarrow \mathcal{P}^{(g)}$ 
11: end for
12:  $\mathcal{P}^* \leftarrow \{\mathbf{d} \in \mathcal{P}^{(G)} : \mathbf{d} \text{ is non-dominated}\}$ 
13: return  $\mathcal{P}^*$ 

```

Overall, this methodology establishes a scalable and interpretable bridge between BIM semantics and physics-informed deep learning, enabling efficient multi-physics optimization for sustainable prefabricated modular building design.

4 Experiments

This section presents the experimental evaluation of the proposed hybrid PI-GNN surrogate framework. We describe: (i) the dataset composition, split protocol, and reproducibility measures; (ii) implementation details; (iii) quantitative benchmarks against traditional and data-driven baselines; (iv) computational speedup analysis; (v) multi-objective optimization case studies; (vi) ablation studies; (vii) external validation on monitored building data; and (viii) quantitative interpretability analysis.

4.1 Dataset Composition, Split Protocol, and Reproducibility**Synthetic Parametric Dataset**

The primary dataset consists of **852** parametric prefabricated modular building models generated using Autodesk Revit 2024 with Dynamo scripts and Grasshopper for Rhino, following established generative-BIM design workflows for modular and prefabricated buildings [3,10,39]. Table 4 summarizes the resulting dataset composition and stratified split. The 852 models are partitioned using a stratified random split by climate zone and building typology (70/15/15 ratio), yielding 596 training, 128 validation, and 128 test samples. No model appears in more than one partition. Random seeds are fixed (seed = 42) for full reproducibility; the partitioning index file will be released alongside the dataset.

Table 4: Dataset composition and stratified split. Counts are mutually exclusive; total is 852. “Nodes/Edges per graph” reports mean $\pm$ standard deviation.

Partition	Count	Fraction	Nodes Per Graph	Edges Per Graph
Training	596	70.0%	187 ± 63	214 ± 71
Validation	128	15.0%	184 ± 58	211 ± 66
Test	128	15.0%	189 ± 61	215 ± 68
Total	852	100%	187 ± 62	213 ± 70

Simulation Labels

High-fidelity labels were obtained through four established simulation pipelines:

1. Operational energy and thermal comfort: EnergyPlus v9.6 with ASHRAE TMY3 weather files [4,5];
2. Lifecycle carbon: One Click LCA (embodied carbon per EN 15978 + 50-year operational emissions) [8];
3. Structural response: Linear elastic FEA via OpenSeesPy under gravity and wind loads per ASCE 7 [6];
4. Cost: Assembly, material, and operational cost estimates from RSMeans unit-price data [40].

External/Real-World Validation Dataset

To assess generalization beyond synthetic parametric models, a supplementary external validation set comprising **14 monitored prefabricated modular buildings** was constructed from two publicly available sources: (i) the U.S. DOE Commercial Reference Building dataset [41] (6 office/educational modular-compatible prototypes), and (ii) field-monitored residential modular units reported in [9] (8 units with measured EUI, discomfort hours, and peak loads). These 14 cases were held entirely out of training and used exclusively for external validation (Section 4.7).

4.2 Implementation Details

The PI-GNN was implemented in PyTorch Geometric v2.5 [38]. Key settings: hidden dimension $d = 256$; 4 GAT layers; 8 attention heads; AdamW optimizer ($\eta = 5 \times 10^{-4}$, weight decay = 10^{-5}); cosine annealing LR schedule; batch size 16; 300 max epochs; early stopping (patience = 30 on validation MAPE); hardware: NVIDIA RTX 4090 (24 GB VRAM), Intel Core i9-13900K, 128 GB RAM; PyTorch 2.2.0, CUDA 12.3. All results are reported as mean $\pm$ std over three independent runs.

Baselines for comparison included:

- Pure data-driven GAT (no physics constraints),
- Standard MLP on flattened BIM features,
- EnergyPlus + FEM (ground truth, no surrogate),
- Recent surrogates: modular PINN [9], graph-based energy model [16].

4.3 Quantitative Results

Table 5 summarizes prediction accuracy on the test set (128 models).

Table 5: Prediction accuracy comparison (MAPE, %; lower is better) on the 128-model test set. PI-GNN results reported as mean $\pm$ std over three runs. “–” denotes baseline does not predict that output.

Model	Energy Demand	Embodied Carbon	Operational Carbon	Peak Stress
MLP (flattened)	18.4	15.7	21.2	28.6
Pure GAT	9.1	8.3	10.5	14.2
Modular PINN [9]	7.8	9.1	8.4	–
Graph energy model [16]	6.5	–	–	–
Proposed PI-GNN	4.1 $\pm$ 0.3	5.3 $\pm$ 0.4	4.8 $\pm$ 0.4	6.7 $\pm$ 0.6

As shown in Fig. 4, the PI-GNN consistently outperforms all baselines, achieving MAPE below 7% across unseen typologies and climates [16].

4.4 Computational Speedup Analysis

Baseline Definition

The speedup denominator is the sequential wall-clock time required to evaluate one building design using the full simulation pipeline on the same hardware: EnergyPlus v9.6 + OpenSeesPy FEM + One Click LCA. All simulations were executed on a single CPU core (Intel Core i9-13900K, one thread per process). GPU-accelerated simulation was not used for the baseline, as EnergyPlus and OpenSeesPy do not natively support GPU execution.




Figure 4: MAPE comparison across models. The proposed PI-GNN achieves the lowest errors, with physics constraints contributing to 40%–55% error reduction in low-data scenarios.

Surrogate Inference

PI-GNN inference was benchmarked on the same RTX 4090 GPU in two modes: (a) single-graph inference (one design at a time, as called by NSGA-II), and (b) batched inference (batch of 16 graphs). Timings are averaged over 200 warmup + 1000 timed forward passes. Results are reported in Table 6.

Table 6: Inference time per building design evaluation. Values are mean ± std over 1000 inference calls.

Method	EnergyPlus	OpenSeesPy	LCA Query	Total/Speedup
Full simulation (CPU, 1 thread)	~118 s	~43 s	~4 s	~165 s/1×
PI-GNN single-graph (GPU)	0.38 ± 0.04 s (all outputs simultaneously)			~0.38 s/434×
PI-GNN batched × 16 (GPU)	0.14 ± 0.01 s per graph (amortized)			~0.14 s/1179×

The single-graph speedup of approximately 434× and the batched speedup of approximately 1179× bracket the “400–1200×” range reported in the abstract and Introduction. For the NSGA-II experiments, single-graph inference is used; each generation of 100 individuals completes in ~38 s of surrogate time vs. ~4.6 h of full simulation, enabling 200 generations in ~2.1 h.

4.5 Multi-Objective Optimization Case Studies

To demonstrate practical utility, NSGA-II optimization (population 100, generations 200) [27] was applied to a 4-module office building in ASHRAE zone 3A. Objectives included minimizing operational

energy, embodied carbon, and cost, subject to constraints: stress below yield and discomfort hours under 500 per year.

The Pareto front (Fig. 5) illustrates trade-offs between energy and carbon. Timber variants reduce embodied emissions by 35%–45% but increase cost by 12%–18%. Hybrid solutions achieve up to 52% energy reduction and 28% carbon reduction relative to baseline.




Figure 5: Pareto front from NSGA-II optimization (operational energy vs. embodied carbon). Colors indicate lifecycle cost (darker = lower cost).

4.6 Ablation Studies

Table 7 quantifies the impact of key architectural components.

Table 7: Ablation study results on test set (MAPE, %). Values are mean $\pm$ std over three runs.

Variant	Energy Demand	Embodied Carbon	Operational Carbon	Peak Stress
Full PI-GNN	4.1 ± 0.3	5.3 ± 0.4	4.8 ± 0.4	6.7 ± 0.6
w/o Physics losses	6.9 ± 0.5 (+68%)	8.2 ± 0.6 (+55%)	7.5 ± 0.6 (+56%)	10.1 ± 0.8 (+51%)
GAT $\rightarrow$ GCN	5.6 ± 0.4 (+37%)	6.8 ± 0.5 (+28%)	6.3 ± 0.5 (+31%)	8.2 ± 0.7 (+22%)
3 layers (instead of 4)	5.2 ± 0.4 (+27%)	6.1 ± 0.5 (+15%)	5.9 ± 0.5 (+23%)	7.9 ± 0.6 (+18%)
w/o Transfer init.	5.0 ± 0.5 (+22%)	5.9 ± 0.6 (+11%)	5.4 ± 0.5 (+13%)	7.5 ± 0.7 (+12%)

Removing physics-informed losses increases MAPE by 51%–68% relative, confirming the value of physics-based regularization [19,22]. Replacing graph attention with GCN degrades accuracy by 22%–37%.

Table 8 further evaluates performance under data-scarce conditions.

Table 8: Prediction accuracy (MAPE, %) under reduced training data regimes. Metric order: Energy/Embodied Carbon/Operational Carbon/Peak Stress.

Training Fraction	Pure GAT	PI-GNN (Proposed)	PI-GNN Gain
100% (596 models)	9.1/8.3/10.5/14.2	4.1/5.3/4.8/6.7	−55%/−36%/−54%/−53%
50% (298 models)	12.3/11.8/14.1/18.7	5.9/7.2/6.6/9.1	−52%/−39%/−53%/−51%
20% (119 models)	18.7/16.2/21.3/26.4	8.3/10.1/9.4/13.2	−56%/−38%/−56%/−50%
10% (60 models)	26.4/22.1/30.5/37.2	11.5/14.3/13.1/18.4	−56%/−35%/−57%/−51%

4.7 External Validation on Monitored Building Data

To assess practical generalization beyond the synthetic parametric dataset, the trained PI-GNN (no retraining or fine-tuning) was evaluated on the 14-building external dataset described in [Section 4.1](#). [Table 9](#) summarizes results.

Table 9: External validation on 14 monitored/reference modular buildings. The PI-GNN was not retrained on any of these cases.

Validation Source	n	EUI MAPE (%)	EUI MAE (kWh/m ² /yr)
DOE commercial prototypes [41]	6	7.9	8.3
Field-monitored residential units [9]	8	9.4	9.1
External combined	14	8.8	8.8
Synthetic test set (reference)	128	4.1	–

The PI-GNN achieves an average EUI MAPE of 8.8% on the external dataset without any retraining, representing a degradation of approximately 4.7 percentage points vs. the in-distribution synthetic test set. These results confirm that the surrogate generalizes to realistic design spaces, while motivating future inclusion of measured data in training.

4.8 Quantitative Interpretability Analysis

Loss Convergence

[Fig. 6](#) shows training loss convergence, where physics-informed terms stabilize by approximately epoch 80 and guide consistent optimization.

Quantitative Attention Analysis

[Table 10](#) reports the mean attention weight assigned to each edge type for the energy-demand output head, alongside the corresponding ground-truth sensitivity.

Saliency Validation via Integrated Gradients

Node-level importance scores were computed using Integrated Gradients with a zero-feature baseline (20 integration steps) via the Captum library [\[42,43\]](#). [Table 11](#) reports the resulting top-10 node types ranked by mean IG saliency score for the energy-demand output head.




Figure 6: Training loss convergence curve (log scale). Total loss (solid blue), data loss (dashed green), and physics residual loss (dotted orange) decrease steadily, with physics terms stabilizing early.

Table 10: Quantitative attention analysis for the energy-demand output head. Pearson correlation $r = 0.91$ ($p < 0.01$), confirming that attention weights align with physical importance [16].

Edge Type	Mean attn. Weight $\bar{\alpha}$	Norm. EUI Sensitivity	Physical Interpretation
Wall-insulation (thermal)	0.187	0.201	Dominant heat-loss pathway
Window-zone (radiation)	0.163	0.178	Solar gain/conduction
Roof-insulation (thermal)	0.141	0.152	Top-surface heat loss
HVAC-zone (flow)	0.118	0.131	Mechanical ventilation load
Module-module (connection)	0.074	0.068	Structural, minor thermal
Floor-module (structural)	0.058	0.049	Mostly structural
Zone-zone (interior)	0.043	0.041	Internal gains transfer
Other	0.216	0.180	Aggregated remaining types
Pearson r (excl. Other)		0.91 ($p < 0.01$, $n = 7$ edge types)	

Table 11: Top-10 node types by mean Integrated Gradient (IG) saliency score for the energy-demand output head, averaged over the 128 test graphs.

Rank	Node Type	Mean norm. IG Score	Physical Relevance
1	Exterior wall (insulation layer)	0.148	Primary envelope U-value
2	Roof (insulation layer)	0.127	Top-surface loss
3	Window (glazing unit)	0.112	U-value + solar heat gain
4	HVAC supply unit	0.094	Mechanical load driver
5	Thermal zone (office)	0.081	Occupancy + internal gains
6	Exterior wall (structural layer)	0.068	Secondary envelope

(Continued)

Table 11 (continued)

Rank	Node Type	Mean norm. IG Score	Physical Relevance
7	Foundation (slab-on-grade)	0.056	Ground heat loss
8	Interior wall	0.044	Partitioning/zoning
9	Roof (structural layer)	0.038	Minor thermal mass
10	HVAC return duct	0.031	Distribution loss
Top-10 cumulative		0.799	—

Fig. 7 illustrates interpretability through attention and saliency maps. Thermal bridge edges receive high attention weights for energy predictions, while connection nodes dominate stress predictions [16]. The Pearson correlation of $r = 0.91$ between attention weights and simulator-estimated sensitivity confirms that the PI-GNN's attention mechanism identifies physically meaningful features without explicit supervision on feature importance [44].




Figure 7: Interpretability analysis: (left) Attention heatmap on graph edges (red = high weight, highlighting thermal bridges/load paths); (right) Saliency map via integrated gradients (node size/color proportional to importance for energy/stress outputs). Results align with physical intuition (e.g., insulation nodes dominant for energy).

These results validate the PI-GNN as an accurate, efficient, and interpretable surrogate, enabling practical optimization of net-zero prefabricated modular buildings [12,14].

5 Discussion

The experimental results presented in Section 4 confirm that the proposed hybrid PI-GNN surrogate effectively addresses the principal challenges of multi-physics modeling and optimization for net-zero prefabricated modular buildings. By leveraging BIM-derived semantic graphs and embedding governing physical constraints into the training objective, the model achieves prediction accuracy comparable to or exceeding that of recent physics-informed approaches in building energy and structural engineering [16,22], while delivering the substantial computational acceleration required for practical early-stage design exploration.

5.1 Scientific Significance and Novelty

The consistently low MAPE values (4%–7%) across diverse typologies and climate zones underscore the importance of physics-informed regularization for improving generalization, particularly under conditions of data scarcity or domain shift. This aligns with emerging trends in physics-informed machine learning, where hybrid models outperform purely data-driven surrogates in robustness and physical consistency [22,23]. The 40%–55% MAPE reduction from physics constraints in low-data scenarios is particularly significant in practice: novel prefabricated typologies will inevitably have limited simulation datasets at early design stages. Attention mechanisms further enhance interpretability by emphasizing critical elements—such as thermal bridges for energy predictions and connection nodes for stress—providing actionable insights often missing in black-box models [16,21].

Compared to prior surrogate modeling efforts [9,16,23], the PI-GNN stands out for its holistic multi-physics coverage (energy, carbon, structural) and direct integration with semantic BIM data. While other graph-based surrogates achieve speedups in single-domain tasks [33], they often lack explicit physics enforcement, risking violations of conservation laws. The embedded physics losses in PI-GNN mitigate this risk, ensuring reliable predictions for sustainability-critical applications such as net-zero certification [22].

5.2 Practical Design Implications

Multi-objective optimization case studies illustrate practical impact: Pareto fronts enable balanced trade-offs, such as 52% energy reduction with 28%–45% carbon savings, achievable in near real-time. This supports iterative, human-in-the-loop design processes that are currently limited by simulation times [13,29]. The sub-minute optimization runtime enables designers to explore thousands of parametric variations and receive Pareto front visualizations in real time, supporting evidence-based early-stage decision-making.

5.3 Limitations

Several limitations bound the current findings and motivate future research:

Synthetic dataset. Reliance on synthetic parametric datasets may not fully capture real-world variability in prefabrication tolerances, material aging, or occupant behavior. The external validation (Section 4.7) demonstrates a gap of 4.7 percentage points in EUI MAPE between synthetic and real-world evaluation, motivating future integration of monitored building data.

Simplified structural modeling. Simplified structural assumptions (linear elastic) limit applicability to extreme or nonlinear regimes; incorporating advanced mechanics could improve realism but increase training complexity [6].

Loss weight sensitivity. Balancing physics loss weights (λ_k) remains empirical—adaptive weighting or gradient-based tuning, as explored in recent PINN work [19,34], may enhance convergence and reduce the need for manual grid search.

Scalability. Extension to larger graphs (e.g., multi-building campuses) or additional physics domains (e.g., acoustics, daylighting) appears promising but requires evaluation of memory and training overhead. Distributed training or model compression could facilitate broader industry adoption.

In summary, the PI-GNN surrogate offers a computationally efficient, interpretable, and physics-consistent alternative to conventional simulation, bridging high-fidelity modeling with practical sustainable design. Addressing the noted limitations through real-world validation, advanced physics integration, and adaptive techniques will further enhance its utility and generalizability in computer modeling for engineering and sustainable built environments.

6 Conclusions

This study introduced a hybrid Physics-Informed Graph Neural Network (PI-GNN) surrogate framework that integrates semantic Building Information Modeling (BIM) data for rapid, accurate, and physically consistent multi-physics prediction and multi-objective optimization of net-zero prefabricated modular buildings. By representing modular assemblies as heterogeneous graphs and embedding governing physical laws (energy conservation, heat transfer, structural equilibrium, carbon balance) into the training process [19,22], the model achieves substantial accuracy gains over purely data-driven baselines while delivering inference speedups of 400–1200× compared to conventional simulations (EnergyPlus + FEM) [9,16].

Key findings include:

- **Superior prediction accuracy:** MAPE of 4.1% for energy demand, 5.3% for embodied carbon, 4.8% for operational carbon, and 6.7% for peak stress across diverse modular configurations and climates.
- **Robust generalization** to unseen typologies and climates, with physics-informed regularization reducing errors by 40%–55% in low-data scenarios [22].
- **Effective multi-objective optimization** via integration with NSGA-II [27], yielding Pareto-optimal designs that reduce operational energy by up to 52% and embodied carbon by 28%–45%, while satisfying safety and comfort constraints in near real-time.
- **Enhanced interpretability** via attention and saliency analysis, highlighting physically meaningful features (e.g., thermal bridges for energy, connections for stress) that support designer trust and decision-making [44].

These contributions advance the state-of-the-art by addressing limitations in existing surrogates: lack of holistic multi-physics coupling, insufficient physics regularization, and poor scalability for modular optimization. The PI-GNN framework bridges high-fidelity simulation with practical sustainable design, enabling efficient early-stage exploration in prefabricated contexts.

Several limitations remain. Structural modeling is simplified (linear elastic), carbon balances are proxy-based, and datasets are synthetic. Future work should incorporate nonlinear mechanics, detailed lifecycle variability, and real-world monitored data to strengthen transferability [6,11]. Training demands, though manageable, could benefit from distributed learning or compression for industry-scale deployment.

Future directions include:

- Extending the framework with uncertainty quantification (e.g., Bayesian PI-GNNs),
- Incorporating lifecycle disassembly and circular economy metrics,
- Developing real-time digital twin interfaces via BIM plugins or cloud platforms,
- Exploring federated learning across manufacturers for scalable, privacy-preserving models.

In conclusion, the PI-GNN surrogate provides a scalable, interpretable, and physics-consistent computational tool that accelerates the transition toward net-zero prefabricated built environments, aligning with global sustainability imperatives and advancing computer modeling in engineering sciences.

Acknowledgement: This Project was funded by the Deanship of Scientific Research (DSR) at King Abdulaziz University, Jeddah, Saudi Arabia under grant no. (IPP: 1363-830-2025). The authors, therefore, acknowledge with thanks DSR for technical and financial support.

Funding Statement: This Project was funded by the Deanship of Scientific Research (DSR) at King Abdulaziz University, Jeddah, Saudi Arabia under grant no. (IPP: 1363-830-2025).

Author Contributions: Masood Karamoozian: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization. Zarrin Mahdavi pour: Methodology, Software, Validation, Formal analysis, Investigation, Writing—review & editing. Mohammed Ameen: Validation, Data curation, Writing—review & editing. Faisal Binzagr: Validation, Writing—review & editing. Abdolraheem Khader: Conceptualization, Methodology, Supervision, Project administration, Funding acquisition, Writing—review & editing. Ahmed Hamza Osman: Validation, Formal analysis, Writing—review & editing. Ali Ahmed: Funding acquisition, Validation, Writing—review & editing. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The synthetic dataset used in this study was generated through parametric BIM modeling using Revit Dynamo and Grasshopper scripts, coupled with high-fidelity simulations conducted in EnergyPlus, One Click LCA, OpenSeesPy, and RSMean. The dataset generation scripts and trained PI-GNN model weights are available from the corresponding author (abdolraheem@njust.edu.cn) upon reasonable request. Requests will be responded to within 30 days of receipt.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. International Energy Agency (IEA). Buildings—energy system. 2025 [cited 2026 Jan 1]. Available from: <https://www.iea.org/energy-system/buildings>.
2. International Energy Agency (IEA). Energy efficiency 2025. 2025 [cited 2026 Jan 1]. Available from: <https://www.iea.org/reports/energy-efficiency-2025>.
3. Li K, Gan VJL, Li M, Gao MY, Tiong RLK, Yang Y. Automated generative design and prefabrication of precast buildings using integrated BIM and graph convolutional neural network. *Dev Built Environ*. 2024;18(6):100418. doi:10.1016/j.dibe.2024.100418.
4. Crawley DB, Lawrie LK, Winkelmann FC, Buhl WF, Huang YJ, Pedersen CO, et al. EnergyPlus: creating a new-generation building energy simulation program. *Energy Build*. 2001;33(4):319–31. doi:10.1016/s0378-7788(00)00114-6.
5. Burnett JC. Introduction to building energy modeling and analysis using EnergyPlus [dissertation]. Fayetteville, AR, USA: University of Arkansas; 2025.
6. Zhu M, McKenna F, Scott MH. OpenSeesPy: python library for the OpenSees finite element framework. *SoftwareX*. 2018;7(1):6–11. doi:10.1016/j.softx.2017.10.009.
7. Villegas JCM. Advanced structural analysis: from theory to computer implementation (OpenSees examples with Python code solutions). Cambridge, UK: Woodhead Publishing; 2025.
8. One Click LCA Ltd. One Click LCA software documentation. 2025 [cited 2026 May 1]. Available from: <https://www.oneclicklca.com>.
9. Jiang Z, Dong B. Modularized neural network incorporating physical priors for future building energy modeling. *Patterns*. 2024;5(8):101029. doi:10.1016/j.patter.2024.101029.

10. Gan VJL. BIM-based graph data model for automatic generative design of modular buildings. *Autom Constr.* 2022;134(1):104062. doi:10.1016/j.autcon.2021.104062.
11. Ma Z, Jiang G, Hu Y, Chen J. A review of physics-informed machine learning for building energy modeling. *Appl Energy.* 2025;381(3):125169. doi:10.1016/j.apenergy.2024.125169.
12. Woldegiyorgis TA, Li HX, Asmare E, Assamnew AD, Admassu FC, Desalegn GA, et al. Harnessing Artificial Intelligence to improve building performance and energy use: innovations, challenges, and future perspectives. *Energy Inform.* 2025;8(1):138. doi:10.1186/s42162-025-00589-3.
13. Kerdan IG, Morillón Gálvez D. ANNEXE: an open-source building energy design optimisation framework using artificial neural networks and genetic algorithms. *J Clean Prod.* 2022;371(1):133500. doi:10.1016/j.jclepro.2022.133500.
14. Fellah M, Ouhaibi S, Belouaggadia N, Mansouri K. Harnessing machine learning for enhanced thermal insulation and energy efficiency in buildings worldwide. *Results Eng.* 2025;25(3):104086. doi:10.1016/j.rineng.2025.104086.
15. Veličković P. Everything is connected: graph neural networks. *Curr Opin Struct Biol.* 2023;79:102538. doi:10.1016/j.sbi.2023.102538.
16. Shan R, Ning H, Xu Q, Su X, Guo M, Jia X. Physics-informed and explainable graph neural networks for generalizable urban building energy modeling. *Appl Sci.* 2025;15(16):8854. doi:10.3390/app15168854.
17. Zhang W, Wang Z, Yang Z, Cai Z. BIM-based construction scheduling optimization through graph neural network-driven spatial semantic reasoning. *J Ind Inf Integr.* 2026;51(3):101088. doi:10.1016/j.jii.2026.101088.
18. Nabrotzky T. Graph neural networks for structural engineering: improving BIM workflows and IFC model exchange. Glasgow, UK: University of Strathclyde Publishing; 2025.
19. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys.* 2019;378:686–707. doi:10.1016/j.jcp.2018.10.045.
20. Thangamuthu A, Kumar G, Bishnoi S, Bhattoo R, Krishnan NMA, Ranu S. Unravelling the performance of physics-informed graph neural networks for dynamical systems. *Adv Neural Inf Process Syst.* 2022;35:3691–702. doi:10.52202/068431-0267.
21. Ashraf I, Strotherm J, Hermes L, Hammer B. Physics-informed graph neural networks for water distribution systems. *Proc AAAI Conf Artif Intell.* 2024;38(20):21905–13. doi:10.1609/aaai.v38i20.30192.
22. Jiang Z, Wang X, Dong B. Physics-informed modularized neural network for advanced building control by deep reinforcement learning. *Adv Appl Energy.* 2025;19(5):100237. doi:10.1016/j.adapen.2025.100237.
23. Chen Y, Yang Q, Chen Z, Yan C, Zeng S, Dai M. Physics-informed neural networks for building thermal modeling and demand response control. *Build Environ.* 2023;234:110149. doi:10.1016/j.buildenv.2023.110149.
24. Wang B, Lu W, Li X, Jin F. A graph neural network-enhanced surrogate model for rapid optimization of carbon-efficient building thermal insulation systems. *J Build Eng.* 2026;119(1):115234. doi:10.1016/j.job.2026.115234.
25. Amani N. Simulation-based design: minimizing energy consumption in residential buildings through optimal thermal insulation. *World J Eng.* 2025;22(6):1235–46. doi:10.1108/wje-04-2024-0188.
26. Liu H, Cheng JCP, Gan VJL, Zhou S. A novel Data-Driven framework based on BIM and knowledge graph for automatic model auditing and quantity take-off. *Adv Eng Inform.* 2022;54:101757. doi:10.1016/j.aei.2022.101757.
27. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput.* 2002;6(2):182–97. doi:10.1109/4235.996017.
28. Hu Y, Xia C, Chen J, Gao X. Clash context representation and change component prediction based on graph convolutional network in MEP disciplines. *Adv Eng Inform.* 2023;55(8):101896. doi:10.1016/j.aei.2023.101896.
29. Austern G, Bloch T, Abulafia Y. Incorporating context into BIM-derived data—Leveraging graph neural networks for building element classification. *Buildings.* 2024;14(2):527. doi:10.3390/buildings14020527.
30. Hong Y, Xie H, Agapaki E, Brilakis I. Graph-based automated construction scheduling without the use of BIM. *J Constr Eng Manag.* 2023;149(2):05022020. doi:10.1061/jcemd4.coeng-12687.
31. Gokhale G, Claessens B, Develder C. Physics informed neural networks for control oriented thermal modeling of buildings. *Appl Energy.* 2022;314:118852. doi:10.1016/j.apenergy.2022.118852.

32. Jiang Z, Dong B. Modularized neural network incorporating physical priors for smart building control, accuracy or consistency? arXiv:2412.02943. 2024.
33. Shukla K, Xu M, Trask N, Karniadakis GE. Scalable algorithms for physics-informed neural and graph networks. *Data Centric Eng.* 2022;3:e24. doi:10.1017/dce.2022.24.
34. Chen HY, Xue L, Liu L, Zou GF, Han JX, Dong YB, et al. Physics-informed graph neural network for predicting fluid flow in porous media. *Petrol Sci.* 2025;22(10):4240–53. doi:10.1016/j.petsci.2025.06.007.
35. Tian Y, Feng Y, Zhang X, Sun C. A fast clustering based evolutionary algorithm for super-large-scale sparse multi-objective optimization. *IEEE/CAA J Autom Sin.* 2023;10(4):1048–63. doi:10.1109/jas.2022.105437.
36. Bian C, Ren S, Li M, Qian C. An archive can bring provable speed-ups in multi-objective evolutionary algorithms. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*; 2025 Jul 14–18; Málaga, Spain. p. 13–4. doi:10.1145/3712255.3734254.
37. IfcOpenShell. Open-source IFC toolkit and geometry engine. 2025 [cited 2026 Jan 1]. Available from: <https://ifcopenshell.org>.
38. Fey M, Lenssen JE. Fast graph representation learning with PyTorch Geometric. arXiv:1903.02428. 2019. doi:10.48550/arxiv.1903.02428.
39. Autodesk Revit 2024. 2023 [cited 2026 Jan 1]. Available from: <https://www.autodesk.com/products/revit/overview>.
40. Mubarak SA. How to estimate with RSMeans data: basic skills for building construction. Hoboken, NJ, USA: John Wiley & Sons; 2020.
41. Deru M, Field K, Studer D, Benne K, Griffith B, Torcellini P, et al. U.S. Department of Energy commercial reference building models of the national building stock. Golden, CO, USA: National Renewable Energy Laboratory; 2011. Report No.: NREL/TP-5500-46861.
42. Sundararajan M, Taly A, Yan Q. Axiomatic attribution for deep networks. In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*; 2017 Aug 6–11; Sydney, Australia. p. 3319–28.
43. Kokhlikyan N, Miglani V, Martin M, Wang E, Alsallakh B, Reynolds J, et al. Captum: a unified and generic model interpretability library for PyTorch. arXiv:2009.07896. 2020. doi:10.48550/arxiv.2009.07896.
44. Peng C, Tang T, Yin Q, Bai X, Lim S, Aggarwal CC. Physics-informed explainable continual learning on graphs. *IEEE Trans Neural Netw Learn Syst.* 2024;35(9):11761–72. doi:10.1109/tnnls.2023.3347453.