

Article

DT-LCAF: Digital Twin-Enabled Life Cycle Assessment Framework for Real-Time Embodied Carbon Optimization in Smart Building Construction

Naif Albelwi

Department of Architectural Engineering, College of Engineering at Yanbu, Taibah University,
Yanbu Al-Bahr 46425, Saudi Arabia; nbelwi@taibahu.edu.sa

Abstract

The construction sector contributes approximately 39% of global carbon emissions, with embodied carbon—emissions from material extraction, manufacturing, transportation, and construction—representing a systematically underestimated yet increasingly critical component of building life cycle environmental impacts. Traditional Life Cycle Assessment (LCA) methods suffer from static database dependencies, delayed feedback cycles, and limited integration with active construction decision-making, creating a fundamental gap between environmental assessment and construction operations. This paper presents the Digital Twin-Enabled Life Cycle Assessment Framework (DT-LCAF), a dynamic construction-phase embodied carbon accounting system aligned with the EN 15978 standard (stages A1–A5) that integrates Building Information Modeling (BIM), Internet of Things (IoT) sensor networks, and machine learning designed to support real-time sustainability decision-making during smart building construction, with computational performance validated through the offline processing of historical datasets. The framework introduces two enabling mechanisms: (1) a Multi-Scale Carbon Prediction Network (MSCPN) employing hierarchical graph attention networks to capture material interdependencies across component, system, and building scales; and (2) a Reinforcement Learning-based Carbon Optimization Engine (RL-COE) that generates constraint-aware recommendations for material substitution, supplier selection, and construction sequencing while respecting structural, economic, and temporal constraints. Experimental evaluation employs two complementary validation strategies using proxy embodied carbon labels (not ground-truth construction measurements): embodied carbon prediction accuracy is assessed using proxy carbon labels derived from the CBECS dataset (5900 commercial buildings) combined with the ICE Database v3.0 emission factors, achieving a 10.24% MAPE, representing a 23.7% improvement over the best-performing baseline in predicting these proxy estimates; temporal responsiveness and streaming data ingestion capabilities are validated using the Building Data Genome Project 2 (1636 buildings, 3053 m). The RL-COE optimization engine demonstrates an 18.4% mean carbon reduction rate within the proxy label framework across building types while maintaining cost and schedule feasibility. A BIM-based case study illustrates the framework's construction-phase update loop, showing how embodied carbon estimates evolve dynamically as construction progresses. The limitations regarding the proxy-based nature of embodied carbon labels and the absence of ground-truth construction-phase measurements are explicitly discussed. The framework contributes to smart city sustainability by enabling scalable, data-driven embodied carbon intelligence across building portfolios. All quantitative results are based on proxy embodied carbon estimates derived from building characteristics and standard

Academic Editor: Antonio Caggiano

Received: 29 January 2026

Revised: 14 February 2026

Accepted: 24 February 2026

Published: 27 February 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

emission factor databases, rather than measured project data. The reported performance therefore demonstrates a proof-of-concept within the proxy system, and real-project, measurement-based validation remains future work.

Keywords: embodied carbon assessment; digital twin; life cycle assessment; construction-phase optimization; BIM-LCA integration; graph neural networks; reinforcement learning; EN 15978; smart building construction; sustainable construction

1. Introduction

The buildings and construction sector represents the largest contributor to anthropogenic greenhouse gas emissions, accounting for approximately 39% of global energy-related CO₂ emissions [1–3]. Within this footprint, embodied carbon emissions arising from material extraction, manufacturing, transportation, and on-site construction activities constitutes approximately 11% of global energy-related CO₂ emissions and is projected to exceed 50% of whole-life building emissions by 2050 as operational emissions decline through grid decarbonization and energy efficiency improvements [4–7]. Unlike operational carbon, which accumulates over decades of building use, embodied carbon is substantially determined during the design and construction phases, creating a narrow but critical window for meaningful intervention [8–10]. This temporal characteristic makes construction-phase optimization essential for achieving sectoral decarbonization objectives [11–14].

Traditional Life Cycle Assessment (LCA) methodologies exhibit fundamental limitations that impede their integration with active construction decision-making [15–17]. Conventional approaches rely on static emission factor databases that fail to capture supply chain variability, produce results only after design completion when modification costs are prohibitive, and lack mechanisms for incorporating real-time construction progress data [18–21]. This temporal misalignment between LCA outputs and construction decision points represents a critical barrier to effective embodied carbon optimization [22–25]. Project managers and sustainability practitioners require actionable carbon intelligence during construction when material substitutions, supplier changes, and sequencing modifications remain feasible rather than retrospective assessments that only serve documentation purposes [26–29]. Figure 1 illustrates the research motivation and the specific gaps addressed by the proposed framework.




Figure 1. Research motivation and framework positioning. The DT-LCAF framework addresses all four gaps through integrated digital twin architecture.

Digital twin technology offers transformative potential for bridging this gap by enabling the continuous synchronization between physical construction progress and virtual building models [30–32]. Recent advances in BIM-LCA integration have demonstrated the feasibility of automated quantity extraction and environmental impact calculation [33–35]. Machine learning approaches have achieved promising results in carbon prediction tasks [36–38], while IoT-enabled monitoring systems have improved real-time data acquisition in building management contexts. Multi-objective optimization algorithms have addressed carbon–cost–time trade-offs in construction planning. However, no existing framework comprehensively integrates these capabilities into a unified system that supports dynamic, construction-phase embodied carbon assessment and optimization.

1.1. Research Hypothesis and Objectives

The central hypothesis is that a dynamically updated, Digital Twin-Enabled Life Cycle Assessment Framework can provide accurate embodied carbon estimates and actionable optimization recommendations during active construction, overcoming the temporal limitations of traditional static BIM-LCA approaches, operationalized through three testable propositions:

Proposition 1.

- *Prediction Accuracy: A hierarchical graph-based neural network architecture modeling material interdependencies across component, system, and building scales will achieve a superior embodied carbon prediction accuracy compared to methods ignoring the hierarchical structure.*
- *Test Method (Section 4.2): MSCPN was compared against 10 baselines using 5-fold cross-validation on CBECS dataset (5900 buildings) with proxy embodied carbon labels.*
- *Test Outcome: CONFIRMED with caveats—MSCPN achieves a 10.24% MAPE (23.7% improvement over best baseline at 13.42%), $R^2 = 0.912$ vs. 0.856, $p < 0.001$, but validates the prediction of proxy labels rather than ground-truth construction measurements.*

Proposition 2.

- *Optimization Effectiveness: A reinforcement learning-based optimization engine learning from digital twin feedback will identify carbon reduction opportunities, outperforming traditional metaheuristics while respecting structural, economic, and temporal constraints.*
- *Test Method (Section 4.4): RL-COE was compared against NSGA-II and ACO-Opt baselines across building types using simulated construction scenarios with defined constraints.*
- *Test Outcome: CONFIRMED with caveats—RL-COE achieves an 18.4% mean carbon reduction vs. NSGA-II (14.1%) and ACO-Opt (12.3%), <2% constraint violations, $p < 0.001$, but demonstrates internal consistency within simulation rather than validated real-world effectiveness.*

Proposition 3.

- *Temporal Responsiveness: A framework designed for streaming data ingestion and incremental updates will support construction-phase update frequencies suitable for practical decision support.*
- *Test Method (Section 4.6): Computational performance was measured across varying building complexities (100–10,000 components) and action spaces (50–2000 actions) using the offline processing of BDGP2 streaming data characteristics.*

- *Test Outcome: CONFIRMED with caveats—sub-100 ms prediction latency for buildings up to 10,000 components, <100 ms RL-COE decision time for action spaces up to 2000 alternatives, 45.2 ms end-to-end median latency, but validates algorithmic performance under controlled conditions rather than operational deployment with system integration and network communication.*

1.2. Contributions

This paper presents the Digital Twin-Enabled Life Cycle Assessment Framework (DT-LCAF), a dynamic embodied carbon accounting and optimization system aligned with the EN 15978 [8] standard for construction-phase assessment (stages A1–A5). The technical and methodological contributions of this research directly address the three testable propositions articulated in Section 1.1. Each contribution provides specific capabilities that enable hypothesis validation. The contributions are framed in terms of sustainability outcomes enabled by the proposed technical mechanisms.

- **Dynamic Construction-Phase Embodied Carbon Accounting:** A comprehensive framework that continuously updates embodied carbon estimates as construction progresses, integrating BIM-extracted quantities, IoT-derived progress data, and supplier-specific emission factors. This addresses the fundamental limitation of static LCA approaches that cannot respond to construction-phase changes.
- **Data-Driven Decision Support for Sustainable Material and Supplier Selection:** The Multi-Scale Carbon Prediction Network (MSCPN), a hierarchical graph attention architecture, enables accurate carbon prediction across component, system, and building granularities, providing practitioners with multi-level carbon hotspot identification to prioritize intervention efforts.
- **Constraint-Aware Optimization for Practical Carbon Reduction:** The Reinforcement Learning-based Carbon Optimization Engine (RL-COE) generates ranked recommendations for material substitution, supplier selection, and construction sequencing that explicitly balance carbon reduction against cost, schedule, and structural performance constraints—reflecting the trade-offs practitioners face in real projects.
- **Comprehensive Validation with Explicit Scope Delineation:** Experimental evaluation employs two complementary validation strategies with clearly distinguished purposes: (a) embodied carbon prediction accuracy assessment using proxy labels derived from CBECS building characteristics combined with ICE Database emission factors; and (b) temporal responsiveness validation using Building Data Genome Project 2 streaming data. A supplementary BIM-based case study demonstrates the construction-phase update mechanism.

1.3. Scope and Limitations

This study is subject to several scope limitations. Most critically, all embodied carbon labels used for model training and evaluation are proxy estimates derived from building characteristics and standardized emission factor databases (ICE v3.0 and Ecoinvent 3.9.1), rather than ground-truth measurements from real construction projects, due to the absence of large-scale datasets with detailed material, supplier, transport and site-energy records. Consequently, the reported prediction accuracy (e.g., 10.24% MAPE) and the 18.4% carbon reduction rate demonstrate only proof-of-concept performance within the proxy-label framework and do not confirm real-world construction accuracy or engineering-level reliability. All quantitative results in this paper therefore refer exclusively to proxy labels, and validation against measured construction-phase embodied carbon remains a key future research direction. In addition, the current implementation is limited to cradle-to-practical completion stages (A1–A5), excludes operational and end-of-life

stages, and evaluates the framework using batch-processed historical data rather than real-time site deployment. Finally, applicability depends on the availability of BIM models with at least LOD 300 for reliable quantity extraction.

1.4. Smart City Relevance

Beyond individual building optimization, DT-LCAF contributes to smart city sustainability objectives by enabling scalable embodied carbon intelligence across building portfolios. The framework's modular architecture supports the aggregation of building-level digital twins to district and city scales, allowing municipalities to monitor construction-sector carbon contributions, identify urban carbon hotspots, and evaluate the embodied carbon implications of alternative material regulations or procurement policies. This positions the framework as infrastructure for evidence-based urban sustainability governance.

The remainder of this paper is organized as follows. Section 2 reviews related work across digital twin technology, BIM-LCA integration, machine learning for carbon prediction, and IoT-enabled monitoring. Section 3 presents the proposed methodology, including system architecture, the mathematical modeling of embodied carbon, the MSCPN prediction network, and the RL-COE optimization engine. Section 4 presents the results and analysis. Section 5 discusses the findings, practical implications, and limitations. Section 6 concludes with future research directions.

2. Literature Review

This section critically examines the existing literature across five research domains that inform the proposed DT-LCAF framework: digital twin technology in construction, BIM-LCA integration, urban digital twins for smart city carbon management, machine learning for carbon prediction, and IoT-enabled monitoring and optimization. For each domain, we identify the current state of knowledge, methodological limitations, and specific gaps that motivate the present work. Table 1 provides a systematic comparison of representative frameworks against the capabilities required for construction-phase embodied carbon optimization.

Table 1. Systematic comparison of DT-LCAF with existing frameworks.

Capability	Static LCA	BIM-LCA [4]	ML-Carbon [11]	DT-Blockchain [28]	ACO-Opt [17]	DT-LCAF
Real-time updating	X	X	X	Partial	X	✓
Dynamic emission factors	X	X	X	X	X	✓
Hierarchical building model	X	Partial	X	X	X	✓
Construction-phase focus (A1–A5)	✓	✓	X	✓	✓	✓
IoT integration	X	X	X	X	X	✓
Constraint-aware optimization	X	X	X	X	Partial	✓
Adaptive learning	X	X	X	X	X	✓
Uncertainty quantification	X	X	Partial	X	X	✓
Multi-scale prediction	X	X	X	X	X	✓
Large-scale validation	X	X	✓	X	X	✓

Capability ratings are assessed as ✓ (fully implemented), Partial (implemented with limitations), or X (not implemented).

2.1. Digital Twin Technology in Construction

Digital twin technology has emerged as a transformative paradigm in the architecture, engineering, and construction (AEC) industry, enabling bidirectional data flow between physical assets and their virtual counterparts. Zhang et al. [1] provided a comprehensive review of digital twin applications for sustainability in construction, documenting accelerating adoption rates while noting that integration with embodied carbon assessment remains largely unexplored. The review identified that most implementations focus on operational phase monitoring rather than construction-phase optimization.

Subsequent work has examined specific dimensions of digital twin deployment. Omrany et al. [2] investigated the social sustainability implications of digital twin adoption, highlighting workforce and community considerations that complement environmental objectives. Zahedi et al. [3] proposed a five-level maturity model for digital twin implementation in sustainable construction, providing a roadmap for organizational capability development. Zhang et al. [25] developed residential carbon accounting frameworks using digital twin principles, though their scope remained limited to post-occupancy assessment rather than construction-phase intervention.

At the urban scale, Peldon et al. [31] examined digital twin applications in city development, demonstrating the potential for coordinated infrastructure management but without addressing embodied carbon tracking during construction phases. This pattern—sophisticated digital twin architectures applied primarily to operational performance—represents a systematic gap in the literature. Construction-phase embodied carbon, which is determined during a narrow temporal window when intervention remains feasible, has received insufficient attention despite its growing contribution to whole-life building emissions.

2.2. BIM-LCA Integration

The integration of Building Information Modeling with Life Cycle Assessment has attracted substantial research interest as a means of automating environmental impact calculation. Zhao et al. [4] demonstrated BIM-LCA application to healthcare facilities, achieving automated quantity extraction but relying on static emission factors that cannot capture supply chain variability. Ma et al. [5] conducted a BIM-based LCA of high-rise building structures, finding a 30–40% variation in embodied carbon depending on the material selection—a finding that underscores the optimization potential but that was not translated into a decision support mechanism.

Chen et al. [6] reviewed the technological evolution of BIM-LCA integration, identifying five distinct integration approaches ranging from manual data transfer to fully automated assessment. Their analysis revealed that even advanced implementations remain design-phase oriented, producing results after major decisions have been finalized. Gao et al. [7] developed LCA-BIM models specifically for prefabricated buildings, addressing an important construction typology but maintaining the static assessment paradigm. Parece et al. [8] created BIM-based tools using standardized classification systems to improve interoperability, while Hemmati et al. [23] compared mass timber and steel carbon footprints using BIM-integrated methods. Zhai et al. [9] investigated lightweight strategies for wooden structure buildings based on embodied carbon calculations, demonstrating that material selection optimization in timber construction can achieve substantial carbon reductions while maintaining structural performance. Their work highlights the importance of material-specific assessment approaches that DT-LCAF incorporates through its multi-scale carbon prediction capabilities.

Three systematic weaknesses characterize the existing BIM-LCA approaches. First, the dependence on generic emission factor databases (e.g., Ecoinvent, ICE Database) that represent industry averages rather than supplier-specific values introduces uncertainty that compounds across building components. Second, the design-phase orientation of

existing tools means that LCA results arrive too late to inform construction-phase decisions such as supplier selection, transportation mode choice, or installation sequencing. Third, the absence of feedback mechanisms prevents learning from actual construction outcomes to improve future predictions. These limitations motivate the dynamic emission factor modeling and construction-phase update capabilities central to DT-LCAF.

2.3. Urban Digital Twins and Smart City-Scale Carbon Management

Digital twins have increasingly become foundational infrastructure for smart city development, enabling the integrated analysis, monitoring, and optimization of complex urban systems. Villani et al. [32] demonstrated how urban digital twins support sustainability and resilience objectives by integrating multi-layer data sources and enabling scenario simulation for city-scale environmental decision-making. Omrany et al. [34] examined urban digital twin adoption in the built environment, outlining applications for resilient and sustainable urban frameworks through continuous data-driven feedback cycles.

Comprehensive reviews have further established digital twins as central to smart city research agendas. Huzzat et al. [33] surveyed digital twin technologies across intelligent city applications, identifying sustainability assessment, infrastructure optimization, and policy support as emerging application domains. Mazzetto [35] conducted a systematic literature review addressing integration challenges and future prospects, emphasizing the need for scalable frameworks that connect building-level analytics with city-scale planning objectives.

Recent research has specifically examined the intersection of digital twins with climate change mitigation and urban resilience. Mabrouk et al. [36] reviewed digital twin city frameworks for climate-resilient urban planning, finding that dynamic environmental assessment and scenario evaluation capabilities are essential for effective climate risk management. Ghorbany et al. [37] proposed a decision support tool combining carbon emission and cost optimization for sustainable urban development, representing the growing application of data-driven technologies to city-level sustainability planning. Yessef et al. [38] positioned digital twins as enablers of intelligent urban management, emphasizing their potential for targeted governance and strategic planning.

Despite these advances, a critical examination reveals that current urban digital twin research predominantly addresses operational performance, energy efficiency, and aggregated urban indicators. Construction-phase embodied carbon—which represents a substantial and growing share of urban carbon footprints—remains largely absent from smart city digital twin implementations. This gap is significant because construction activities are spatially and temporally concentrated, creating opportunities for coordinated intervention that operational emissions do not afford. The proposed DT-LCAF framework addresses this gap by enabling construction-phase embodied carbon tracking at the building level while providing a natural scalability to district and city scales through digital twin aggregation.

2.4. Machine Learning for Carbon Prediction

Machine learning approaches have demonstrated promising results for building carbon prediction, though predominantly for operational rather than embodied carbon. Zheng et al. [10] developed operational carbon prediction models achieving an R^2 of 0.847, establishing the viability of data-driven approaches for building energy assessment. Zhang et al. [11] applied stacking ensemble methods (combining XGBoost, LightGBM, and CatBoost) to provincial-level carbon prediction, achieving an 8.7% MAPE and demonstrating the effectiveness of ensemble architectures for carbon modeling tasks.

Explainability has emerged as an important consideration for practical deployment. Zhu et al. [12] proposed explainable machine learning frameworks incorporating SHAP

(SHapley Additive exPlanations) values to identify feature contributions to carbon predictions, addressing the interpretability requirements of sustainability practitioners. ElArwady et al. [30] combined digital twin concepts with machine learning for indoor thermal comfort modeling, demonstrating the potential for integrated approaches, though not addressing embodied carbon. Figueiredo et al. [29] proposed blockchain-enabled digital twin architectures for LCA data management, addressing data provenance and integrity concerns.

The current machine learning approaches for carbon prediction exhibit two significant limitations relevant to embodied carbon assessment. First, the existing models treat buildings as flat feature vectors, ignoring the hierarchical structure (components → systems → buildings) through which embodied carbon actually accumulates. This architectural limitation prevents models from learning the compositional relationships that characterize building assemblies. Second, most approaches address operational carbon prediction, where temporal patterns in energy consumption provide rich training signals; embodied carbon prediction receives less attention despite its distinct data characteristics and modeling requirements. The Multi-Scale Carbon Prediction Network (MSCPN) proposed in this work directly addresses these limitations through hierarchical graph attention mechanisms that explicitly model building structure.

2.5. IoT-Enabled Monitoring and Optimization

Internet of Things technologies have transformed building monitoring capabilities, enabling real-time data acquisition at unprecedented granularity. Poyyamozhi et al. [13] surveyed IoT applications for building energy management, documenting sensor technologies, communication protocols, and data analytics approaches. Franco et al. [14] demonstrated synergistic IoT integration achieving 12.8% energy savings through coordinated monitoring and control, establishing the practical value of real-time building intelligence.

Specific monitoring applications have addressed diverse building performance dimensions. Chaudhari et al. [15] reviewed occupancy detection algorithms using IoT sensor networks, an application relevant to operational carbon but not directly addressing construction-phase monitoring. Fay et al. [26] proposed green IoT architectures for environmental monitoring, including carbon tracking, though focused on ambient monitoring rather than construction process integration.

Multi-objective optimization approaches have addressed construction sustainability trade-offs. Zhang et al. [16] applied optimization algorithms to building energy performance, while Xu et al. [17] employed ant colony optimization for construction carbon minimization, demonstrating metaheuristic approaches to carbon-aware construction planning. However, these optimization approaches operate on static problem formulations and do not adapt to real-time construction progress data.

The IoT literature reveals mature capabilities for real-time data acquisition, but with limited application to construction-phase embodied carbon monitoring. The existing implementations predominantly serve operational building management rather than construction process tracking. The DT-LCAF framework addresses this gap by specifying IoT integration for construction progress monitoring, enabling the dynamic updates that distinguish the proposed approach from static assessment methods.

2.6. Research Gaps and Positioning

The preceding review identifies five critical gaps that collectively motivate the proposed framework:

Gap 1: Absence of Real-Time Digital Twin-LCA Integration for Embodied Carbon:

The existing digital twin implementations in construction focus predominantly on operational-phase optimization. No current framework provides continuous, construction-

phase embodied carbon assessment through integrated digital twin architecture. DT-LCAF addresses this gap through its five-layer architecture connecting IoT data acquisition, digital twin state management, carbon prediction, optimization, and visualization.

Gap 2: Static Emission Factor Dependence in BIM-LCA Tools: The current BIM-LCA integration relies on fixed emission factor databases that cannot capture supplier-specific efficiency, transportation variability, or temporal changes in manufacturing processes. DT-LCAF addresses this gap through dynamic emission factor modeling that incorporates real-time supplier and logistics data.

Gap 3: Neglect of Hierarchical Building Structure in ML-Based Carbon Prediction: The existing machine learning models for carbon prediction treat buildings as unstructured feature collections, ignoring the component–system–building hierarchy through which embodied carbon accumulates. DT-LCAF addresses this gap through the Multi-Scale Carbon Prediction Network (MSCPN), which employs hierarchical graph attention to explicitly model building structure.

Gap 4: Non-Adaptive Optimization Strategies: The current optimization approaches for construction carbon operate on static problem formulations established at project initiation. They cannot incorporate real-time feedback from construction progress to adjust recommendations. DT-LCAF addresses this gap through the Reinforcement Learning-based Carbon Optimization Engine (RL-COE), which learns adaptive policies from digital twin interactions.

Gap 5: Limited Large-Scale Validation: Many proposed frameworks demonstrate concepts on single buildings or small datasets, limiting generalizability assessment. DT-LCAF addresses this gap through validation on two large-scale public datasets comprising over 7500 buildings, enabling a statistically robust comparison with baseline methods.

2.7. Comparison with Existing Frameworks

Table 1 systematically compares the proposed DT-LCAF framework against representative existing approaches across key capability dimensions required for construction-phase embodied carbon optimization.

BIM-LCA [4] has building hierarchy but treats components independently without modeling interdependencies and calculates A1–A5 stages statically at design completion rather than updating dynamically during construction; DT-Blockchain [28] architecturally supports real-time recording but implements periodic batch updates and uses hierarchical storage with simple summation rather than dependency modeling; ACO-Opt [17] includes cost/schedule constraints as hard limits checked post hoc rather than integrated into optimization; ML-Carbon [11] provides ensemble variance estimates capturing only model uncertainty, not data uncertainty from emission factors or measurement error. Assessment criteria require <1 min latency, supplier-specific time-varying emission factors, explicit component–system–building relationship modeling, direct sensor stream ingestion, multi-objective soft constraint optimization, online learning, aleatoric + epistemic uncertainty separation, multi-scale predictions with consistency enforcement, and 1000+ building validation with statistical significance.

3. Methodology and Analysis

This section presents the complete methodology of the DT-LCAF framework. Section 3.1 provides the system architecture overview. Section 3.2 details the mathematical modeling of embodied carbon aligned with EN 15978. Section 3.3 describes the construction of proxy embodied carbon labels and associated scope limitations. Section 3.4 presents the Multi-Scale Carbon Prediction Network (MSCPN). Section 3.5 details the Reinforcement Learning-based Carbon Optimization Engine (RL-COE). Section 3.6 provides a complexity analysis. Figure 2 illustrates the overall framework architecture.




Figure 2. DT-LCAF framework architecture comprising five integrated layers. (a) Data acquisition layer integrating BIM models (IFC format), IoT sensor networks (MQTT protocol), supplier databases (EPDs), and environmental impact databases (Ecoinvent 3.9.1, ICE v3.0); (b) digital twin engine maintaining synchronized virtual representation with geometric models, material properties, construction program, and real-time state using Extended Kalman Filtering; (c) Multi-Scale Carbon Prediction Network (MSCPN) employing hierarchical graph attention networks for component-, system-, and building-level predictions with uncertainty quantification; (d) Reinforcement Learning-based Carbon Optimization Engine (RL-COE) generating constraint-aware recommendations for material substitution, supplier selection, and construction sequencing; (e) visualization dashboard providing 3D carbon hotspot overlay, optimization trajectory display, scenario comparison, and ranked recommendations for multiple stakeholder roles.

3.1. System Architecture Overview

The DT-LCAF framework comprises five interconnected layers that collectively enable real-time embodied carbon monitoring and optimization throughout construction stages A1–A5. Each layer performs distinct functions while maintaining a bidirectional data flow with adjacent layers, creating an integrated system capable of continuous carbon assessment and adaptive optimization.

3.1.1. Layer 1: Data Acquisition

The data acquisition layer integrates heterogeneous data sources with varying fidelity, update frequencies, and uncertainty characteristics. BIM models in IFC format provide geometric quantities and material specifications extracted at minimum Level of

Development (LOD) 300, including volumes in cubic meters, masses in kilograms, surface areas in square meters, and linear dimensions in meters for all building components. Material specifications encompass density, composition, and classification codes mapped to emission factor databases.

IoT sensor networks transmit real-time construction progress indicators via MQTT protocol, including material delivery confirmations, installation completion signals, equipment utilization metrics, and environmental conditions such as temperature and humidity that affect curing times. Sensor data arrives at frequencies ranging from sub-second for equipment telemetry to daily for progress milestones. Supplier databases provide product-specific Environmental Product Declarations (EPDs) where available, manufacturing location coordinates for transportation distance calculation, and production efficiency indicators. When supplier-specific data is unavailable, industry benchmark values from the ICE Database v3.0 are used with appropriate uncertainty inflation.

Environmental impact databases supply baseline emission factors from Ecoinvent 3.9.1 and ICE Database v3.0, providing stage-specific emission factors for stages A1 through A5 for materials classified by type, manufacturing process, and geographic region. Construction progress tracking incorporates Earned Value Management (EVM) data, including planned versus actual progress, schedule variance, and activity completion status.

Data streams are categorized by origin and fidelity into three types. Measured data comprise IoT-derived construction progress indicators and environmental conditions representing the highest fidelity. Estimated data include BIM-extracted quantities and database emission factors at medium fidelity. Simulated data encompass schedule projections, dependency graphs, and scenario-based what-if analyses with variable fidelity depending on modeling assumptions.

Missing data are handled through hierarchical imputation where supplier-specific values are preferred, followed by material category benchmarks, followed by industry averages. Imputation uncertainty is propagated through the framework via the probabilistic modeling described in Section 3.2.3. Data ingestion employs Apache Kafka for scalable, fault-tolerant streaming with exactly once semantics, supporting ingestion rates exceeding 10,000 events per second for large construction sites.

3.1.2. Layer 2: Digital Twin Engine

The digital twin engine maintains a continuously synchronized virtual representation of the construction project comprising four interconnected elements. Geometric models provide three-dimensional building geometry at an appropriate LOD, updated as construction progresses from design at LOD 300 to as-built representations at LOD 500. Material properties include component-level attributes such as density in kilograms per cubic meter, carbon intensity in kilograms of CO₂ equivalent per kilogram, recyclability parameters, thermal conductivity, and structural properties. The construction program maintains the activity network with predecessor and successor dependencies, planned durations, resource assignments, and critical path identification. The real-time state captures the current construction progress, material inventory status, supplier commitments, and environmental conditions.

The digital twin supports bidirectional data flow where physical-to-virtual synchronization updates the virtual model based on sensor data, while virtual-to-physical simulation enables scenario evaluation before physical implementation. State synchronization employs Extended Kalman Filtering for sensor fusion, providing filtered state estimates and anomaly detection when sensor readings deviate significantly from predicted values.

Unlike static BIM models that serve primarily as design documentation, the DT-LCAF digital twin functions as a dynamic computational environment. It continuously

updates to reflect the actual construction state, supports real-time scenario simulation, and enables predictive carbon assessment for decisions not yet implemented. This capability of evaluating alternative materials, schedules, and logistics routes before physical commitment distinguishes the framework from fixed LCA approaches and visualization-focused digital models.

3.1.3. Layer 3: Carbon Prediction Module

The Multi-Scale Carbon Prediction Network (MSCPN), detailed in Section 3.4, consumes digital twin state data to produce embodied carbon estimates at three granularities. Component-level predictions address individual building elements such as beams, columns, panels, and fixtures. System-level predictions aggregate building systems including structural, envelope, MEP, and finishes. Building-level predictions provide whole-building embodied carbon estimates. Uncertainty quantification through Monte Carlo dropout provides prediction confidence intervals, typically at 90% coverage, enabling risk-informed decision-making.

MSCPN recalculation is triggered by construction progress events rather than fixed time intervals. Specific triggers include material delivery confirmation, which updates inventory and A4 transportation emissions; installation completion signals, which update A5 construction emissions and remaining work; supplier change notifications, which update A1 through A3 emission factors; schedule revisions, which update projected completion and sequencing options; and threshold-based change detection when a greater than 5% change occurs in any input feature since the last prediction. In practice, the update frequency ranges from hourly during active construction phases to daily during planning or low-activity periods, depending on sensor resolution and project requirements.

While the framework architecture is designed for real-time operation with event-driven updates, the experimental validation in Section 4 employs the offline processing of historical datasets to enable a controlled comparison with baseline methods. True real-time operation requires deployment on active construction sites with live IoT data streams—a validation approach not undertaken in the current research. The computational performance results (Section 4.6) demonstrate that the framework achieves a sub-100 ms latency suitable for real-time deployment, establishing architectural feasibility, but online deployment under construction site conditions with network latency, sensor failures, and concurrent user access remains to be validated.

3.1.4. Layer 4: Optimization Engine

The Reinforcement Learning-based Carbon Optimization Engine (RL-COE), detailed in Section 3.5, generates adaptive recommendations across three decision categories. Material substitution recommendations identify alternative materials meeting structural and functional requirements with lower embodied carbon. Supplier selection recommendations identify preferred suppliers based on carbon intensity, transportation distance, and reliability. Construction sequencing recommendations optimize activity ordering to minimize waste, rework, and logistics emissions.

Recommendations are constrained by structural performance requirements, budget limits, and schedule deadlines. The engine learns optimal policies through an interaction with the digital twin environment, continuously improving as construction progresses and new data becomes available.

3.1.5. Layer 5: Visualization Dashboard

The interactive decision support interface provides four primary visualization capabilities. The 3D carbon hotspot overlay presents BIM model visualization with color-coded carbon intensity, enabling the rapid identification of high-impact components. The

optimization trajectory display shows carbon reduction achieved through construction phases, comparing actual performance against baseline and optimized scenarios. Scenario comparison enables a side-by-side evaluation of alternative material and sequencing choices with quantified carbon, cost, and schedule implications. Ranked recommendations present actionable suggestions ordered by carbon reduction potential, implementation feasibility, and constraint satisfaction.

The dashboard serves multiple stakeholder roles where project managers focus on schedule and cost implications, sustainability consultants examine carbon hotspots and reduction pathways, and procurement teams evaluate supplier alternatives.

3.1.6. Scalability to District and City Levels

While presented at an individual building scale, the DT-LCAF architecture supports aggregation to larger scales. Multiple building-level digital twins can be federated to provide district-wide or city-wide embodied carbon intelligence, enabling portfolio-level carbon monitoring across concurrent construction projects, urban carbon hotspot identification at a neighborhood scale, policy impact assessment for material regulations or procurement standards, and coordinated optimization across projects sharing suppliers or logistics infrastructure. This scalability positions DT-LCAF as an infrastructure for smart city sustainability governance, though city-scale deployment introduces additional challenges, including data interoperability, privacy, and coordination, that are addressed in Section 6.

3.2. Mathematical Modeling of Embodied Carbon

3.2.1. Life Cycle Carbon Emission Framework

Embodied carbon is modeled in accordance with EN 15978, which defines the system boundary for construction-phase assessment. Total embodied carbon EC_{total} is computed as the cumulative sum across building components and life cycle stages according to Equation (1).

$$EC_{total} = \sum_{i=1}^N \sum_{j \in M} EC_{i,j} \quad (1)$$

where

EC_{total} = Total embodied carbon of the building measured in kilograms of CO₂ equivalent (kgCO_{2e}), representing cumulative emissions across all components and life cycle stages.

N = Number of building components extracted from BIM model, typically ranging from 50 (simple warehouse) to 10,000 (complex hospital).

M = Set of life cycle stages within system boundary {A1, A2, A3, A4, A5} according to EN 15978.

$EC_{i,j}$ = Embodied carbon contribution of component i in stage j , measured in kgCO_{2e}.

i = Component index, ranging from 1 to N .

j = Life cycle stage index, corresponding to stages A1 through A5.

In this formulation, N represents the number of building components extracted from BIM, $M = \{A1, A2, A3, A4, A5\}$ defines the life cycle stages within the system boundary, and $EC_{i,j}$ denotes the embodied carbon contribution of component i in stage j .

Component-stage embodied carbon is calculated according to Equation (2).

$$EC_{i,j} = Q_i \times EF_{i,j} \times \omega_{i,j} \quad (2)$$

where

Q_i = Quantity of component i extracted from BIM, measured in kilograms (kg) for materials, cubic meters (m³) for volumetric elements, or square meters (m²) for surface elements.

$EF_{i,j}$ = Emission factor for component i in stage j , measured in kgCO₂e per unit (kgCO₂e/kg, kgCO₂e/m³, or kgCO₂e/m²), obtained from LCA databases or supplier EPDs. $\omega_{i,j}$ = Waste factor accounting for material losses during stage j , a dimensionless ratio typically ranging from 1.02 to 1.15 (representing 2% to 15% waste).

The quantity Q_i represents the quantity of component i extracted from BIM measured in kilograms for materials, cubic meters for volumetric elements, or square meters for surface elements. The emission factor $EF_{i,j}$ provides the emission factor for component i in stage j measured in kilograms of CO₂ equivalent per unit. The waste factor $\omega_{i,j}$ accounts for material losses during stage j . Table 2 shows the life cycle stage definitions and data sources according to EN 15978

Table 2. Life cycle stage definitions and data sources according to EN 15978.

Stage	Description	Data Source	Update Mechanism
A1	Raw material extraction	LCA database with supplier efficiency scaling	Supplier EPD updates
A2	Transport to manufacturer	Supplier location and transport mode data	Supplier change notification
A3	Manufacturing	LCA database with manufacturing efficiency factors	Supplier EPD updates
A4	Transport to construction site	Supplier and site locations with transport mode	Delivery confirmation signals
A5	Construction and installation	Equipment factors and waste rates	Installation completion signals

Stages A1 through A3, representing the cradle-to-gate boundary, are primarily determined by material selection and supplier choice, with baseline values from LCA databases adjusted by supplier-specific efficiency factors when available. Stages A4 and A5, representing the gate-to-practical completion boundary, are dynamically modeled using real-time IoT data, including actual transportation distances, modes, equipment utilization, and observed waste rates.

The current implementation focuses on stages A1 through A5 because these stages are determined during the construction phase when intervention remains feasible, are supported by established emission factor databases, and are addressable through the material, supplier, and sequencing decisions that RL-COE optimizes. Operational stages B1 through B7 and end-of-life stages C1 through C4 are excluded from the current scope, though the framework architecture could accommodate their inclusion in future extensions.

3.2.2. Dynamic Emission Factor Modeling

Traditional LCA employs static emission factors from generic databases, failing to capture supply chain variability that significantly affects actual embodied carbon. The DT-LCAF dynamic emission factor model incorporates real-time supplier and logistics information according to Equation (3).

$$EF_{i,j}^{dyn(t)} = EF_{i,j}^{base} \times \alpha_{sup(t)} \times \beta_{trans(d_i)} \times \gamma_{mfg(t)} \quad (3)$$

where

$EF_{i,j}(t)$ = Time-varying dynamic emission factor at time t , measured in kgCO₂e per unit.
 $EF_{i,j}^{base}$ = Baseline emission factor from environmental databases (Ecoinvent 3.9.1, ICE v3.0), measured in kgCO₂e per unit.

$\alpha_{sup}(t)$ = Time-varying supplier efficiency factor, dimensionless, typically ranging from 0.7 to 1.5, where values < 1.0 indicate above-average efficiency and values > 1.0 indicate below-average efficiency.

$\beta_{trans}(d_i)$ = Transportation distance modifier for component i as a function of distance d_i , a dimensionless multiplier typically ranging from 1.0 to 1.8.

$\gamma_{mfg}(t)$ = Manufacturing process efficiency modifier at time t , dimensionless, ranging from 0.8 to 1.2.

t = Current time period measured in days since project commencement.

The baseline emission factor $EF_{i,j}^{base}$ is obtained from environmental databases. The time-varying supplier efficiency factor $\alpha_{sup}(t)$ captures deviations from industry-average manufacturing efficiency. The transportation distance modifier $\beta_{trans}(d_i)$ adjusts for actual logistics distances for component i . The manufacturing process efficiency modifier $\gamma_{mfg}(t)$ accounts for production technology variations. The variable t denotes the current time period.

The transportation distance modifier is calculated according to Equation (4).

$$\beta_{trans(d_i)} = 1 + \frac{(d_i \times ef_{mode})}{EF_{i,AA}^{base}} \quad (4)$$

The transportation distance d_i is computed in kilometers from the supplier location to the construction site. The mode-specific emission factor ef_{mode} is measured in kilograms of CO₂ equivalent per tonne-kilometer with values of 0.089 for road transport by truck, 0.016 for rail transport, and 0.012 for ship transport. Combined modes use weighted averages by distance proportion.

The supplier efficiency factor is modeled according to Equation (5).

$$\alpha_{(sup,t)} = \frac{EF_{sup,actual}(t)}{EF_{sup,reference}} \quad (5)$$

where $EF_{sup,actual}$ = Actual emission factor for the specific supplier at time t [kgCO₂e/unit] obtained from supplier-provided EPDs or real-time production monitoring. $EF_{sup,reference}$ = Reference emission factor from standard database for material category [kgCO₂e/unit] obtained from ICE Database v3.0 or Ecoinvent v3.9.1. Table 3 shows the parameter ranges and calibration assumptions for dynamic emission factor modeling.

Table 3. Parameter ranges and calibration assumptions for dynamic emission factor modeling.

Parameter	Symbol	Typical Range	Default Value	Calibration Source
Supplier efficiency factor	α_{sup}	0.7–1.5	1.0	Supplier EPD or industry benchmark
Transportation modifier	β_{trans}	1.0–1.8	Distance-dependent	Supplier location and transport mode
Manufacturing efficiency	γ_{mfg}	0.8–1.2	1.0	Manufacturing efficiency reports
Waste factor	ω	0.02–0.15	Material-dependent	Industry waste rate studies

When supplier-specific data is unavailable, default values are used with inflated uncertainty bounds to reflect the additional estimation uncertainty. This approach ensures that predictions appropriately convey lower confidence when relying on generic rather than supplier-specific information. Table 4 shows the variable ranges and units for dynamic emission factor modeling.

Table 4. Variable ranges and units for dynamic emission factor modeling.

Variable	Symbol	Unit	Typical Range	Data Source
Total embodied carbon	EC_total	kgCO ₂ e	200,000–5,000,000	Calculated
Component count	N	count	50–10,000	BIM model
Component quantity	Q _i	kg, m ³ , m ²	Varies by type	BIM extraction
Baseline emission factor	EF ^{base}	kgCO ₂ e/unit	0.05–3.5	ICE Database v3.0
Supplier efficiency	α _{sup}	dimensionless	0.7–1.5	EPD/industry data
Transport modifier	β _{trans}	dimensionless	1.0–1.8	Calculated from distance
Manufacturing efficiency	γ _{mfg}	dimensionless	0.8–1.2	Supplier reports
Waste factor	Ω	dimensionless	1.02–1.15	Industry benchmarks
Transport distance	d _i	Km	10–5000	Supplier location
Transport emission factor	ef _{mode}	kgCO ₂ e/(t·km)	0.012–0.089	Transport mode data

3.2.3. Uncertainty Quantification

Embodied carbon estimates incorporate multiple uncertainty sources including emission factor variability, quantity measurement imprecision, waste rate fluctuation, and model uncertainty. Total uncertainty is propagated using first-order Taylor series approximation according to Equation (6).

$$\sigma_{EC}^2 = \sum_{i=1}^N \sum_{j \in M} \left(\frac{\partial EC_{total}}{\partial EC_i} \right)^2 \sigma_{i,j}^2 + 2 \sum_{i < k} \rho_{ik} \sigma_i \sigma_k \quad (6)$$

The variance $\sigma_{i,j}^2$ represents the uncertainty of the emission factor for component i in stage j , derived from EPD uncertainty ranges typically spanning $\pm 10\%$ to $\pm 30\%$. The correlation coefficient ρ_{ik} captures the correlation between components sharing common materials or suppliers.

Components sharing the same material type exhibit a positive correlation with ρ of approximately 0.6 to 0.8 because emission factor errors affect all instances. Components from the same supplier exhibit a correlation with ρ of approximately 0.4 to 0.6 through shared manufacturing efficiency uncertainty. Cross-material and cross-supplier correlations are assumed negligible, with ρ approximately zero.

For MSCPN predictions, epistemic uncertainty representing model uncertainty is quantified through Monte Carlo dropout. A total of T forward passes with dropout enabled at inference time produce a distribution of predictions according to Equations (7) and (8). During training, dropout randomly disables network connections (setting weights to zero with probability $p = 0.2$) to prevent overfitting. During inference, dropout is typically disabled. However, Monte Carlo dropout keeps dropout enabled during inference and performs multiple forward passes to capture model uncertainty

$$\hat{EC} = \left(\frac{1}{T} \right) \sum_{t=1}^T f_{\psi(G; dropout_t)} \quad (7)$$

where $\hat{EC}^t$ = Predicted embodied carbon from dropout sample t [kgCO₂e]. s = Input state (building component graph, features). θ = Trained model parameters (weights, biases). $\xi(t)$ = Random dropout mask for pass t (different random connections disabled each pass).

$$\sigma_{model} = \sqrt{\left[\left(\frac{1}{T-1} \right) \sum_{t=1}^T (f_{\psi(G; dropout_t)} - \hat{EC})^2 \right]} \quad (8)$$

The total prediction uncertainty combines aleatoric and epistemic components according to Equation (9).

$$\sigma_{total} = \sqrt{\sigma_{EC}^2 + \sigma_{model}^2} \quad (9)$$

3.3. Construction of Proxy Embodied Carbon Labels and Dataset Scope Limitations

A critical methodological consideration is the construction of embodied carbon labels for model training and evaluation. This section explicitly documents the proxy label construction process, its assumptions, and its associated limitations to ensure the appropriate interpretation of experimental results.

3.3.1. The Ground-Truth Challenge

No large-scale public dataset provides ground-truth embodied carbon measurements from actual construction projects. Measuring true embodied carbon requires the comprehensive tracking of actual material quantities, supplier-specific emission factors, realized transportation distances and modes, observed waste rates, and equipment energy consumption. Such data are rarely collected systematically and are almost never made publicly available due to commercial sensitivity and the substantial effort required for comprehensive tracking.

This data availability constraint necessitates the use of proxy embodied carbon labels, which are estimates derived from building characteristics and standardized emission factor databases rather than direct measurements. While proxy labels enable large-scale model training and evaluation, they introduce systematic uncertainty that must be explicitly acknowledged when interpreting results.

3.3.2. Proxy Label Construction Process

For the CBECS dataset, proxy embodied carbon labels are constructed through a four-step process. In the first step of building characteristic extraction, the following attributes are obtained from CBECS building records: gross floor area in square meters, number of floors, building age, principal construction type including steel frame, concrete, masonry, and wood options, wall materials, roof materials, and building activity type.

In the second step of material intensity estimation, material intensity factors measured in kilograms per square meter or cubic meters per square meter are applied based on construction type and building activity, derived from published benchmarks and industry studies. For example, steel frame office buildings typically require structural steel at 45 to 65 kg per square meter and concrete at 0.15 to 0.25 cubic meters per square meter. Concrete frame healthcare facilities typically require reinforced concrete at 0.35 to 0.50 cubic meters per square meter and rebar at 8 to 12 kg per square meter.

In the third step of emission factor application, estimated material quantities are mapped to emission factors from ICE Database v3.0, applying stage-specific factors for stages A1 through A5. Default waste factors are applied by material type based on industry benchmarks.

In the fourth step of aggregation, component-stage contributions are summed to produce building-level embodied carbon estimates measured in kilograms of CO₂ equivalent and kilograms of CO₂ equivalent per square meter.

3.3.3. Explicit Assumptions and Limitations

The proxy label construction process embeds several assumptions that limit the validity of claims that can be made from experimental results.

The first assumption is that material intensity benchmarks are representative. Material quantities are estimated from building characteristics using published intensity factors rather than extracted from actual construction documents. Buildings deviating from

typical construction practices will have misestimated quantities, and this deviation is not captured in the proxy labels.

The second assumption is that generic emission factors apply uniformly. ICE Database emission factors represent UK and European industry averages. The application to U.S. buildings from CBECS introduces a geographic bias, and actual supplier-specific factors may differ by $\pm 30\%$ or more from database values.

The third assumption is that the default waste factors are accurate. Waste factors are applied uniformly by material type, but actual waste rates vary significantly by contractor capability, site conditions, and project management quality. This variation is not captured in the proxy labels.

The fourth assumption is that construction process variations are negligible. The proxy labels assume typical construction processes, and unusual construction methods, challenging site conditions, and non-standard equipment are not captured in the estimation.

These limitations mean that the experimental results demonstrate the framework's ability to predict proxy embodied carbon labels with stated accuracy, not its ability to predict true embodied carbon from actual construction. The 23.7% improvement over baselines reflects the superior modeling of the relationship between building characteristics and proxy labels. While we hypothesize that improved proxy prediction correlates with improved true prediction, this hypothesis requires validation on projects with ground-truth measurements, which represents an important direction for future research.

3.3.4. Validation Task Separation

To maintain methodological rigor, the experimental validation in Section 4 employs two datasets for distinct purposes. The CBECS dataset is used for embodied carbon prediction accuracy assessment, where proxy labels constructed as described above are used to train and evaluate MSCPN prediction accuracy. This validates the hierarchical graph attention architecture's ability to capture building structures for carbon prediction.

The BDGP2 dataset is used exclusively to validate temporal responsiveness, specifically the framework's ability to ingest streaming data, update predictions incrementally, and maintain computational performance under continuous data arrival. The Building Data Genome Project 2 contains hourly energy meter readings representing operational energy rather than embodied carbon. BDGP2 does not validate embodied carbon prediction accuracy, and no claims regarding embodied carbon prediction are made based on BDGP2 results.

This separation ensures that claims about prediction accuracy and claims about temporal responsiveness are supported by appropriate evidence from suitable datasets.

3.4. Multi-Scale Carbon Prediction Network (MSCPN)

The Multi-Scale Carbon Prediction Network addresses the fundamental limitation of existing machine learning approaches for carbon prediction, namely their failure to model the hierarchical structure through which embodied carbon accumulates in buildings. MSCPN employs graph neural networks with hierarchical attention mechanisms to learn representations at component, system, and building scales, enabling predictions that respect the compositional nature of building assemblies.

3.4.1. Graph Representation of Building Components

Buildings are represented as attributed heterogeneous graphs $G = (V, E, X, A)$ where V is the set of nodes representing building components, E is the set of edges representing relationships between components, $X \in \mathbb{R}^{N \times d}$ is the node feature matrix with d -dimensional feature vectors, and $A \in \mathbb{R}^{N \times N}$ is the weighted adjacency matrix encoding

relationship strengths. Figure 3 illustrates the MSCPN architecture with its hierarchical processing stages.




Figure 3. Multi-Scale Carbon Prediction Network (MSCPN) architecture illustrating input feature encoding combining geometric, material, process, and temporal attributes; four graph attention layers with hierarchical attention mechanisms; multi-scale aggregation producing component, system, and building-level representations; and prediction heads with uncertainty quantification through Monte Carlo dropout.

Node features x_i for component i concatenate four feature categories according to Equation (10).

$$x_i = [f_{geo}^i \parallel f_{mat}^i \parallel f_{proc}^i \parallel f_{temp}^i] \quad (10)$$

where

x_i = Feature vector for component i , dimension $d = 41$.

f_{geo}^i = Geometric feature vector, dimension $d_g = 12$, containing volume (m^3), surface area (m^2), length (m), width (m), height (m), position coordinates x, y, z (m), and orientation angles (radians).

f_{mat}^i = Material feature vector, dimension $d_m = 15$, containing density (kg/m^3), carbon intensity for stages A1–A5 ($kgCO_2e/kg$), recyclability index (0–1), thermal conductivity ($W/(m \cdot K)$), and structural grade (categorical).

f_{proc}^i = Process feature vector, dimension $d_p = 8$, containing prefabrication level (0–1), quality grade (1–5), supplier tier (1–3), and manufacturing method (categorical encoding).

f_{temp}^i = Temporal feature vector, dimension $d_t = 6$, containing installation phase (1–5), predecessor count (integer), planned duration (days), and critical path indicator (0–1).

$\parallel$ = Concatenation operator combining vectors sequentially.

Table 5 shows the node feature categories and dimensions for MSCPN graph representation. The geometric feature vector f_{geo}^i with dimension $d_g = 12$ contains volume, surface area, length, width, height, position coordinates in three dimensions, and orientation angles. The material feature vector f_{mat}^i with dimension $d_m = 15$ represents density, carbon intensity values for stages A1 through A5, recyclability index, thermal conductivity, and structural grade. The process feature vector f_{proc}^i with dimension $d_p = 8$ captures prefabrication level, quality grade, supplier tier, and manufacturing method encoding. The temporal feature vector f_{temp}^i with dimension $d_t = 6$ stores installation phase, predecessor count, planned duration, and critical path indicator. The total feature dimension is $d = 41$.

Table 5. Node feature categories and dimensions for MSCPN graph representation.

Category	Notation	Dimension	Features Included
Geometric	f_{geo}^i	$d_g = 12$	Volume, surface area, length, width, height, position coordinates (x, y, z), orientation angles
Material	f_{mat}^i	$d_m = 15$	Density, carbon intensity for stages A1–A5, recyclability index, thermal conductivity, structural grade
Process	f_{proc}^i	$d_p = 8$	Prefabrication level, quality grade, supplier tier, manufacturing method encoding
Temporal	f_{temp}^i	$d_t = 6$	Installation phase, predecessor count, planned duration, critical path indicator

Edge relationships encode three types of connections between components according to Equation (11).

$$e_{ij} = 1 \text{ if } spatial(i,j) \vee functional(i,j) \vee material(i,j), \text{ otherwise } 0 \quad (11)$$

Spatial adjacency is indicated when the Euclidean distance between component centroids is less than a threshold distance, with the default set to 5 m. A functional relationship is indicated when components belong to the same building system or have structural or MEP dependency. Material similarity is indicated when components share the same material category or supplier.

When IFC models are available, graph construction proceeds by parsing the IFC file to extract IfcBuildingElement entities and their geometric representations, computing bounding boxes and centroids for spatial relationship detection, extracting material associations from IfcMaterialLayerSetUsage and IfcMaterialConstituentSet, identifying functional relationships from IfcRelContainedInSpatialStructure and IfcRelAggregates, and constructing the adjacency matrix based on relationship criteria.

When IFC models are unavailable, as in the CBECS validation, abstract graphs are constructed from building characteristics by generating a synthetic component inventory based on building type and construction system, applying typical spatial layouts for the building type, assigning material properties from building characteristic codes, and constructing relationships based on typical system configurations. The experimental validation uses abstract graphs constructed from CBECS characteristics, while real deployment would use IFC-derived graphs with a higher fidelity. Figure 4 illustrates the graph representation using a 20-component office building example, demonstrating the node categorization by building system, the three edge types encoding spatial, functional, and material relationships, the structure of the 41-dimensional node feature vector, and the BIM-to-graph construction pipeline that transforms IFC models into the attributed graph representation consumed by MSCPN.




Figure 4. Graph representation of a 20-component building for MSCPN. Nodes represent building components colored by system: structural (blue), envelope (green), MEP (orange), and finishes (purple). Edges encode three relationship types: spatial adjacency (dashed), functional dependency (solid), and material similarity (dotted).

3.4.2. Hierarchical Graph Attention Mechanism

MSCPN employs multi-head graph attention to learn importance weights for message passing between adjacent components. At layer l , the attention coefficients are computed according to Equation (12).

$$\alpha_{ij}^l = \exp(\text{LeakyReLU}(a^{(l)T} [w^{(l)} h_i^{l-1} \| w^l h_j^{l-1}])) / \sum_{k \in \mathcal{N}_i} \exp(\text{LeakyReLU}(a^{(l)T} [w^{(l)} h_i^{l-1} \| w^l h_k^{l-1}])) \quad (12)$$

The hidden representation $h_i^l \in \mathbb{R}^{d_h}$ represents node i at layer l . The learnable weight matrix $W^{(l)} \in \mathbb{R}^{d_h \times d_h}$ transforms node representations. The attention parameter vector $a^l \in \mathbb{R}^{2d_h}$ computes attention scores. The neighborhood $\mathcal{N}_i$ denotes the set of nodes adjacent to node i . The symbol $\|$ denotes vector concatenation.

Node representations are updated through attention-weighted aggregation according to Equation (13).

$$h_i^l = \sigma(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^l W^l h_j^{l-1}) \quad (13)$$

The activation function σ is the Exponential Linear Unit (ELU), which provides smooth gradients for negative inputs and avoids the dying neuron problem associated with ReLU activations.

Multi-head attention with K heads improves representational capacity according to Equation (14).

$$h_i^l = \parallel_k^{=1K} \sigma(\Sigma_i \in \mathcal{N}_i \alpha_{ij}^{l,k} W^{l,k} h_j^{l-1}) \quad (14)$$

MSCPN uses $K = 8$ attention heads and $L = 4$ graph attention layers, producing the hidden dimension $d_h = 128$ with 16 dimensions per head.

3.4.3. Multi-Scale Aggregation

To capture carbon accumulation across hierarchical granularities, MSCPN employs hierarchical pooling operations that aggregate component representations to system and building levels. Let $\mathcal{C}_k$ denote the set of components belonging to building system $k \in \{\text{structural, envelope, MEP, finishes}\}$. System-level representations are computed according to Equation (15).

$$h_{sys}^k = \text{Readout}(\{h_i^l: i \in \mathcal{C}_k\}) \quad (15)$$

The Readout function is a permutation-invariant aggregation combining mean, max, and attention-weighted sum according to Equation (16).

$$\text{Readout}(H) = W_{mean} \cdot \left(\frac{1}{|H|}\right) \Sigma_h \in H h + W_{max} \cdot \max_h \in H(h) + W_{att} \Sigma_h \in H \beta_h h \quad (16)$$

Attention weights β_h are computed via a learned attention mechanism over system components according to Equation (17).

$$\beta_h = \frac{\exp(w_{att}^T h)}{\Sigma_{h' \in H} \exp(w_{att}^T h')} \quad (17)$$

Building-level representation aggregates all system representations according to Equation (18).

$$h_{bld} = \text{MLP}(\text{Concat}(\{h_{sys}^k\}_{k=1}^{K_{sys}})) \quad (18)$$

The number of building systems $K_{sys} = 4$ corresponds to structural, envelope, MEP, and finishes systems. The multi-layer perceptron (MLP) provides a non-linear transformation of the concatenated system representations.

3.4.4. Prediction Heads and Training Objective

The framework includes separate prediction heads for each granularity level. Component-level predictions are computed according to Equation (19).

$$\hat{EC}_i = \text{MLP}_{comp}(h_i^l)_{for all i \in V} \quad (19)$$

System-level predictions are computed according to Equation (20).

$$\hat{EC}_{sys}^k = \text{MLP}_{sys}(h_{sys}^k)_{for all k \in \{1, \dots, K_{sys}\}} \quad (20)$$

Building-level predictions are computed according to Equation (21).

$$\hat{EC}_{bld} = \text{MLP}_{bld}(h_{bld}) \quad (21)$$

The training objective integrates multi-scale prediction losses with hierarchical consistency regularization according to Equation (22).

$$\mathcal{L} = \mathcal{L}_{comp} + \lambda^1 \mathcal{L}_{sys} + \lambda^2 \mathcal{L}_{bld} + \lambda^3 \mathcal{L}_{cons} \quad (22)$$

Individual losses use the Huber loss function for robustness to outliers according to Equation (23).

$$\mathcal{L}_{comp} = \left(\frac{1}{N}\right) \Sigma_i^{=1N} \mathcal{L}_{\delta}(EC_i, \hat{EC}_i) \quad (23)$$

The consistency loss imposes a hierarchical coherence between prediction levels according to Equation (24).

$$\mathcal{L}_{cons} = \|\hat{E}C_{bld} - \sum_k^{=1Ksys} \hat{E}C_{sys}^k\|^2 + \sum_k^{=1Ksys} \|\hat{E}C_{sys}^k - \sum_i \in \mathcal{C}_k \hat{E}C_i\|^2 \quad (24)$$

This consistency regularization ensures that building-level predictions equal the sum of system-level predictions, and system-level predictions equal the sum of component-level predictions within each system, reflecting the physical reality that embodied carbon is additive across the building hierarchy. Table 6 shows the MSCPN hyperparameter configuration.

Table 6. MSCPN hyperparameter configuration.

Hyperparameter	Symbol	Value	Justification
Graph attention layers	L	4	Sufficient receptive field for building graphs
Attention heads	K	8	Balance between expressiveness and computation
Hidden dimension	d_h	128	Adequate capacity for building-scale graphs
Learning rate	η	0.001	Standard for graph neural networks
Dropout rate	p	0.2	Regularization and uncertainty quantification
Component loss weight	-	1.0	Base scale for loss balancing
System loss weight	λ_1	0.5	Reduced weight for aggregated predictions
Building loss weight	λ_2	0.3	Reduced weight for highest aggregation
Consistency loss weight	λ_3	0.1	Soft constraint for hierarchical coherence

3.5. Reinforcement Learning-Based Carbon Optimization Engine (RL-COE)

The Reinforcement Learning-based Carbon Optimization Engine is designed to generate practical construction recommendations rather than abstract numerical solutions. The optimization process focuses on material substitution between available alternatives, supplier and transportation logistics selection, construction activity sequencing, and waste minimization during installation stages. Optimization results are delivered to practitioners as ranked, constraint-aware suggestions that explicitly trade embodied carbon savings against structural performance, cost, and schedule feasibility.

3.5.1. Markov Decision Process Formulation

The carbon optimization problem is formulated as a Markov Decision Process (MDP) defined by the tuple (S, A, P, R, γ) comprising state space, action space, transition probability, reward function, and discount factor.

The state s_t at construction step t encodes the current project status according to Equation (25).

$$s_t = [h_{bld(t)} \| p_t \| b_t \| c_t \| m_t] \quad (25)$$

The building representation $h_{bld(t)} \in \mathbb{R}^{d_h}$ from MSCPN captures the current construction state. The progress vector $p_t \in [0,1]^{N_{act}}$ provides normalized construction progress across activities. The remaining budget $b_t \in \mathbb{R}^+$ indicates the available financial resources. The cumulative carbon $c_t \in \mathbb{R}^+$ represents the embodied carbon to date. The material availability vector $m_t \in \{0,1\}^{N_{mat}}$ indicates which materials are currently available.

Actions encompass three decision categories according to Equation (26).

$$a_t = (a_t^{mat}, a_t^{sup}, a_t^{seq}) \quad (26)$$

The material selection action $a_t^{mat} \in \{0,1\}^{(N \times M_{alt})}$ indicates the material alternative selection for components with available substitutes, where M_{alt} is the maximum number of alternatives per component. The supplier selection action $a_t^{sup} \in \{0,1\}^{(N \times$

S) indicates supplier selection from S qualified suppliers. The sequencing action $a_t^{\text{seq}} \in \Pi(N_{\text{pending}})$ encodes construction sequence permutation for pending activities.

The reward function balances carbon reduction with constraint satisfaction according to Equation (27).

$$r_t = -\Delta EC_t - \lambda_{\text{cost}} \varphi_{\text{cost}}(\Delta \text{Cost}_t) - \lambda_{\text{time}} \varphi_{\text{time}}(\Delta \text{Time}_t) - \lambda_{\text{struct}} \varphi_{\text{struct}}(a_t) \quad (27)$$

The carbon change ΔEC_t represents the change in predicted embodied carbon from action a_t . The penalty functions φ softly penalize budget overruns, schedule slippage, and structural performance degradation. The trade-off coefficients λ regulate preferences among objectives.

The cost penalty function uses a quadratic soft constraint according to Equation (28).

$$\varphi_{\text{cost}}(\Delta \text{Cost}_t) = \max(0, \Delta \text{Cost}_t - \varepsilon_{\text{cost}})^2 \quad (28)$$

The tolerance threshold $\varepsilon_{\text{cost}}$ defines the acceptable cost increase before penalties apply. Similar formulations apply to schedule and structural penalties. Table 7 shows the constraint parameters and feasibility thresholds for RL-COE optimization.

Table 7. Constraint parameters and feasibility thresholds for RL-COE optimization.

Constraint	Symbol	Default Value	Interpretation
Cost tolerance	$\varepsilon_{\text{cost}}$	5% of component cost	Acceptable cost premium for low-carbon alternatives
Schedule tolerance	$\varepsilon_{\text{time}}$	2 days per activity	Acceptable delay for sequencing optimization
Structural threshold	$\varepsilon_{\text{struct}}$	$0.95 \times \text{baseline}$	Minimum structural performance ratio
Cost penalty weight	λ_{cost}	0.3	Relative importance of cost constraint
Time penalty weight	λ_{time}	0.2	Relative importance of schedule constraint
Structural penalty weight	λ_{struct}	0.5	Relative importance of structural constraint

3.5.2. Constraint Enforcement and Feasibility Analysis

A critical consideration for practical deployment is ensuring that RL-COE recommendations are implementable in real construction projects. The framework employs a two-stage approach to constraint enforcement.

In the first stage of soft constraint learning, penalty terms in the reward function encourage the policy to avoid constraint violations during training. This approach allows the policy to learn the trade-off structure between carbon reduction and constraint satisfaction, potentially accepting minor violations when carbon benefits are substantial.

In the second stage of hard constraint filtering, at inference time, generated actions are validated against hard constraints before presentation to users. Actions violating structural requirements with a performance below $\varepsilon_{\text{struct}} \times \text{baseline}$ are rejected. Actions exceeding budget limits beyond a hard threshold of 10% above $\varepsilon_{\text{cost}}$ are flagged with warnings. Actions causing critical path delays beyond a hard threshold of 5 days are identified for user review. Table 8 shows the constraint violation analysis from RL-COE training on the CBECS validation set.

Table 8. Constraint violation analysis from RL-COE training on CBECS validation set.

Constraint Type	Actions Generated	Violations Before Filter	Violations After Filter	Feasibility Rate
Structural performance	10,000	847 (8.5%)	0 (0%)	100%

Cost threshold	10,000	1234 (12.3%)	156 (1.6%)	98.4%
Schedule deadline	10,000	523 (5.2%)	89 (0.9%)	99.1%
Combined feasibility	10,000	2145 (21.5%)	203 (2.0%)	98.0%

This analysis demonstrates that RL-COE generates predominantly feasible recommendations, with the filtering stage catching edge cases where the learned policy produces marginally infeasible actions.

3.5.3. Policy Network Architecture

The policy network $\pi_\theta(a|s)$ employs a Transformer-based architecture to model construction decision sequences with long-range dependencies. State encoding produces query, key, and value representations according to Equation (29).

$$Q = W_Q \cdot z_s, K = W_K \cdot z_s, V = W_V \cdot z_s \quad (29)$$

The encoded state representation z_s is projected by learned matrices W_Q , W_K , and W_V .

Self-attention computes contextualized representations according to Equation (30).

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (30)$$

Action probabilities are computed through type-specific heads according to Equations (31) and (32).

$$\pi_\theta^{mat}(a^{mat}|s) = \text{softmax}(W_{mat} \cdot z_L + b_{mat}) \quad (31)$$

$$\pi_\theta^{sup}(a^{sup}|s) = \text{softmax}(W_{sup} \cdot z_L + b_{sup}) \quad (32)$$

The final Transformer layer output z_L provides the representation for action prediction.

3.5.4. Proximal Policy Optimization Training

Policy learning employs Proximal Policy Optimization (PPO) for stable training dynamics. The surrogate objective is defined according to Equation (33).

$$\mathcal{L}^{PPO}(\theta) = \mathbb{E}_t[\min(r_{t(\theta)\hat{A}_t}, \text{clip}(r_{t(\theta)}, 1-\epsilon, 1+\epsilon)\hat{A}_t)] \quad (33)$$

The probability ratio $r_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{old}}(a_t|s_t)$ compares new and old policies. The advantage estimate $\hat{A}_t$ measures the benefit of the taken action. The clipping parameter $\epsilon = 0.2$ constrains policy updates.

Advantages are computed using Generalized Advantage Estimation (GAE) (see Appendix A) according to Equation (34).

$$\hat{A}_t = \sum_{l=0}^{\infty} (\gamma \lambda_{GAE})^l \delta_{t+l} \quad (34)$$

Temporal difference errors are defined according to Equation (35).

$$\delta_t = r_t + \gamma V_{\varphi(s_{t+1})} - V_{\varphi(s_t)} \quad (35)$$

Algorithm 1 shows the RL-COE training algorithm and Algorithm 2 shows the DT-LCAF realtime inference.

Algorithm 1 RL-COE Training Algorithm

- Input:** Digital twin environment $\mathcal{E}$, trained MSCPN f_ψ , initial policy π_θ , value function V_ϕ , hyperparameters $\gamma, \lambda_{GAE}, \epsilon, K$
 - Output:** Optimized policy parameters θ^*
-

```

3. Initialize replay buffer  $\mathcal{D} \leftarrow \emptyset$ 
4. for episode  $e = 1$  to  $E_{max}$  do
5.     Reset environment:  $s_0 \leftarrow \mathcal{E}.reset()$ 
6.     Initialize trajectory  $\tau \leftarrow \emptyset$ 
7.     for  $t = 0$  to  $T_{max}$  do
8.         Compute building representation:  $\mathbf{h}_{bld}(t) \leftarrow f_{\psi}(\mathcal{G}_t)$ 
9.         Encode state:  $s_t \leftarrow \text{encode}(s_t, \mathbf{h}_{bld}(t))$ 
10.        Sample action:  $a_t \sim \pi_{\theta}(\cdot | s_t)$ 
11.        Execute in environment:  $s_{t+1}, r_t \leftarrow \mathcal{E}.step(a_t)$ 
12.        Append to trajectory:  $\tau \leftarrow \tau \cup \{(s_t, a_t, r_t, s_{t+1})\}$  break
13.        if  $s_{t+1}$  is terminal then
14.            break
15.        end if
16.    end for
17.    Compute advantages  $\{\hat{A}_t\}$  using GAE (Equation (31))
18.    Add trajectory to buffer:  $\mathcal{D} \leftarrow \mathcal{D} \cup \tau$ 
19.    if  $|\mathcal{D}| \geq \text{batch size}$  then
20.        for epoch  $k = 1$  to  $K$  do
21.            Sample mini-batch  $\mathcal{B}$  from  $\mathcal{D}$ 
22.            Compute PPO loss:  $\mathcal{L}^{PPO} \leftarrow \text{Equation (30)}$ 
23.            Update policy:  $\theta \leftarrow \theta + \alpha_{\pi} \nabla_{\theta} \mathcal{L}^{PPO}$ 
24.            Compute value loss:  $\mathcal{L}^V \leftarrow \frac{1}{|\mathcal{B}|} \sum (V_{\phi}(s) - R)^2$ 
25.            Update value:  $\phi \leftarrow \phi - \alpha_V \nabla_{\phi} \mathcal{L}^V$ 
26.        end for
27.        Clear buffer:  $\mathcal{D} \leftarrow \emptyset$ 
28.    end if
29. end for
30. return  $\theta^*$ 

```

Algorithm 2 DT-LCAF Real-Time Inference

```

1. Input: BIM model  $\mathcal{B}$ , IoT sensor stream  $\mathcal{J}$ , trained MSCPN  $f_{\psi}^*$ , trained policy  $\pi_{\theta^*}$ , con-
   straint parameters
2. Output: Carbon predictions  $\{\widehat{EC}_t\}$ , optimization recommendations  $\{a_t^*\}$ 
3. Initialize digital twin:  $DT \leftarrow \text{DigitalTwin}(\mathcal{B})$ 
4. while construction in progress do
5.     Receive sensor data:  $\mathbf{z}_t \leftarrow \mathcal{J}.receive()$ 
6.     Update digital twin state:  $s_{DT}(t) \leftarrow DT.sync(\mathbf{z}_t)$ 
7.     Extract component graph:  $\mathcal{G}_t \leftarrow \text{BIM2Graph}(\mathcal{B}, s_{DT}(t))$ 
8.     Predict embodied carbon at all scales:
9.          $\widehat{EC}_{comp}, \widehat{EC}_{sys}, \widehat{EC}_{bld} \leftarrow f_{\psi}^*(\mathcal{G}_t)$ 
10.    Encode optimization state:  $s_t \leftarrow \text{encode}(DT, \widehat{EC}_{bld})$ 
11.    Generate optimal action:  $a_t^* \leftarrow \text{argmax}_a \pi_{\theta^*}(a | s_t)$ 
12.    Validate against constraints:
13.    if satisfies_constraints( $a_t^*, s_t$ ) then
14.        Present recommendation to user via dashboard
15.    else
16.        Generate constrained alternative:  $a_t^* \leftarrow \text{constrained\_search}(s_t)$ 
17.    end if
18.    Update dashboard visualizations
19.    Log predictions and recommendations
20. end while

```

3.6. Complexity Analysis

The MSCPN forward pass computational complexity is $O(|E| \cdot d_h + N \cdot L \cdot d_h^2)$, where $|E|$ is the number of edges, N is the component count, L is the layer count, and d_h

is the hidden dimension. For sparse building graphs where $|E| = O(N)$, this simplifies to $O(N \cdot L \cdot d_h^2)$, enabling linear scaling with building size.

The RL-COE policy evaluation complexity per step is $O(d_s \cdot d_h \cdot L_{\text{trans}} + N \cdot M_{\text{alt}})$, where d_s is the state dimension, L_{trans} is the number of Transformer layers, and M_{alt} is the maximum number of material alternatives. Both components support GPU parallelization through batched operations, achieving real-time performance as demonstrated in Section 4. Table 9 shows the computational complexity summary for DT-LCAF components.

Table 9. Computational complexity summary for DT-LCAF components.

Component	Operation	Complexity	Typical Latency
MSCPN	Forward pass	$O(N \cdot L \cdot d_h^2)$	15–45 ms
MSCPN	Uncertainty (T passes)	$O(T \cdot N \cdot L \cdot d_h^2)$	150–450 ms
RL-COE	Policy evaluation	$O(d_s \cdot d_h \cdot L_{\text{trans}})$	5–15 ms
RL-COE	Constraint validation	$O(N \cdot M_{\text{alt}})$	2–8 ms
End-to-end	Full inference	Sum of above	25–520 ms

4. Results

This section involves an extensive experimental analysis of the DT-LCAF framework, dataset descriptions, implementation information, baseline analysis, ablation, and prediction accuracy and optimization performance. Every experiment will be aimed at confirming that the framework is efficient in the real-time optimization of the construction of smart buildings' embodiment of carbon.

4.1. Experimental Setup

4.1.1. Datasets

To test the presented DT-LCAF, we rely on two massive, publicly accessible databases that contain more detailed building features and energy/carbon-related indicators. Table 10 is a summary of the important attributes of the two databases.

Table 10. Summary of experimental datasets.

Characteristic	[12]	BDGP2
Number of buildings	5900	1636
Number of features	200+	150+
Temporal resolution	Annual	Hourly
Total data points	5900 records	53.6 million
Geographic scope	United States	International
Building type categories	16	6
Primary use case	Carbon modeling	Energy prediction
Data availability	Public domain	Open source (MIT)
Year of publication	2022	2020

Dataset 1: Commercial Buildings Energy Consumption Survey (CBECS) 2018

The CBECS dataset is the largest national survey of commercial building characteristics and energy consumption in the United States, which is provided by the Energy Information Administration (EIA) of the United States. The 2018 survey consists of 5900 commercial buildings that have more than 200 characteristics that comprise building size, age, type of constructions, materials, energy, and operation patterns as per 16 patterns of activities in buildings.

- **URL:** <https://www.eia.gov/consumption/commercial/data/2018/> (accessed on 25 November 2025).
- **Buildings:** 5900 commercial buildings that have statistical weights that represent 5.9 million U.S. buildings.
- **Features:** There are 200+ building features, such as the type of structure, envelope material, floor area, multi-story, and year of construction.
- **Building types:** There are 16 categories (office, retail, education, healthcare, warehouse, hotel, etc.).
- **Geographic coverage:** Each of the U.S. census regions and climate zones.
- We complete the use of CBECS data with embodied carbon computation reliant on the Inventory of Carbon and Energy (ICE) Database v3.0 emission elements by encompassing material intensity factors contingent upon the categories of building age and construction type to estimate the component-level embodied carbon to train and evaluate the model.
- While CBECS provides large-scale coverage (5900 buildings representing 5.9 million U.S. commercial buildings), it lacks component-level construction details, specific site locations, construction process documentation, and temporal consistency—requiring embodied carbon estimates based on aggregated characteristics and assumed intensity factors rather than actual construction data. This means that the validation results demonstrate that the hierarchical graph attention architecture (MSCPN) outperforms baselines when predicting carbon estimates derived from building statistics and achieves a consistent performance across building types and regions but does not validate absolute prediction accuracy against ground-truth construction carbon or readiness for engineering deployment. The dataset enables architectural comparison and scalability assessment rather than the verification of real-world prediction precision.

Dataset 2: Building Data Genome Project 2 (BDGP2)

The data of BDGP2 are high-resolution hourly data on energy meters of non-residential buildings in various countries, which allow us to explore the patterns of building performance over time. Building characteristics and systems have comprehensive metadata in the dataset.

- **URL:** <https://github.com/buds-lab/building-data-genome-project-2> (accessed on 25 November 2025).
- **Buildings:** North American, European, and Asian non-residential buildings, numbered 1636.
- **Energy meters:** 3053 m of electricity, gas, steam and chilled water.
- **Measurements:** The amount of data was 53.6 million hourly (17,544 measurements per meter in two years).
- **Time range:** 2016–2017 (complete two-year period).
- **Building types:** There are six major categories (education, office, public, residential, retail, other).

The Building Data Genome Project 2 (BDGP2) validates only the framework's temporal responsiveness—its ability to ingest streaming data, update predictions incrementally, and maintain computational performance under continuous data arrival—not embodied carbon prediction accuracy, as BDGP2 contains operational energy meter data (hourly electricity, gas, steam, chilled water consumption over 2 years) unrelated to construction materials. Like CBECS, BDGP2 lacks BIM models, material specifications, and construction documentation, requiring abstract component graphs based on building type templates, meaning validation demonstrates streaming data processing capabilities and acceptable prediction latency but does not validate embodied carbon accuracy or actual IoT sensor integration with construction monitoring systems.

4.1.2. Implementation Details

The DT-LCAF model is executed in Python 3.10 based on PyTorch 2.1 to handle deep learning and graph neural network operations and Stable-Baselines3 to handle the reinforcement learning algorithms. Table 11 shows important hyperparameter values.

Table 11. Hyperparameter configuration.

Parameter	MSCBN	RL-COE
Hidden dimension d_h	128	256
Number of layers	4 (GNN)	6 (Transformer)
Attention heads	8	8
Learning rate	3×10^{-4}	1×10^{-4}
Batch size	64	256 trajectories
Training epochs/episodes	200	1000
Dropout rate	0.1	0.1
Weight decay	1×10^{-5}	1×10^{-5}
PPO clip parameter ϵ	–	0.2
Discount factor γ	–	0.99
GAE parameter λ	–	0.95
Loss weights $\lambda_1, \lambda_2, \lambda_3$	0.5, 0.3, 0.2	–

The training is done on a workstation with two NVIDIA A100 80 GB cards, an AMD EPYC 7763 64-core Processor and 512 GB RAM. It is estimated that MSCBN training needs an average of 4 h in the CBECS dataset; RL-COE training needs an average of 12 h on the 1000 episodes of the simulation of a digital twin environment.

4.1.3. Baseline Methods

We put DT-LCAF into competition with 12 state-of-the-art algorithms based on traditional LCA, machine learning strategies, and optimization algorithms:

1. **Single Static LCA:** The conventional method of carrying out a Life Cycle Assessment based on the ICE Database using emission factors that are not used dynamically.
2. **BIM-LCA [4]:** BIM-based LCA application that provides an automatic extraction of quantities.
3. **RF-Carbon:** 500-tree chaos random forest regression of carbon.
4. **XGBoost-Carbon:** Early stopping and 1000-estimator Gradient boosting.
5. **LSTM-Carbon:** Multiplex neural network consisting of two LSTM lines (128 units in layers).
6. **GCN-Carbon:** Baseline of Graph Convolutional Network (GCN) (4 layers, 128 dimensions).
7. **GAT-Carbon:** Baseline graph attention network (four layers, eight heads).
8. **ml-provincial [11]:** Ensemble of XGBoost, LightGBM, and CatBoost.
9. **CatBoost-SHAP [12]:** SHAP-based feature selection given CatBoost.
10. **ACO-Opt [17]:** Ant colony optimization construction for carbon minimization.
11. **NSGA-II:** Pareto-optimal solution genetic algorithm with multi-objectives.
12. **DT-Blockchain [28]:** Cryptocurrency blockchain-based data management of LCA.

4.1.4. Evaluation Metrics

We use full performance regarding prediction accuracy and optimization:

Prediction Accuracy Metrics

$$\text{MAPE} = \frac{1}{N} \sum_{i=1}^N \left| \frac{EC_i - \widehat{EC}_i}{EC_i} \right| \times 100\% \quad (36)$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (EC_i - \widehat{EC}_i)^2} \quad (37)$$

$$R^2 = 1 - \frac{\sum_{i=1}^N (EC_i - \widehat{EC}_i)^2}{\sum_{i=1}^N (EC_i - \bar{EC})^2} \quad (38)$$

Optimization Metrics

$$\text{CRR} = \frac{EC_{\text{baseline}} - EC_{\text{optimized}}}{EC_{\text{baseline}}} \times 100\% \quad (39)$$

4.2. Carbon Prediction Results

Figure 5 provides a comparison of the accuracy of the prediction of embodied carbon for all the specified baseline methods using the CBECS dataset and 5-fold cross-validation.




Figure 5. Embodied carbon prediction accuracy (MAPE) comparison across 11 methods on CBECS dataset. MSCPN achieves 10.24% MAPE, representing 23.7% improvement over the best baseline CatBoost-SHAP (13.42%). Error bars indicate standard deviation across 5-fold cross-validation.

Table 12 provides detailed quantitative comparison across all prediction metrics.

Table 12. Carbon prediction performance comparison.

Method	MAPE (%)	RMSE (kgCO ₂ e/m ²)	R ²	MAE (kgCO ₂ e/m ²)	Time (s)
Static LCA	28.45	156.32	0.612	112.45	2.1
BIM-LCA	22.18	128.76	0.695	98.34	8.4
RF-GBDT	18.92	108.43	0.754	76.28	6.5
XGBoost	17.64	96.29	0.782	68.92	6.2

LSTM	16.58	101.28	0.771	72.14	8.3
GCN	15.47	89.45	0.813	63.27	7.4
GAT	14.89	85.12	0.832	61.38	7.8
ML-Physical	13.67	92.14	0.805	66.73	8.2
DT-KNN-SHAP	12.45	78.65	0.865	54.82	9.1
MSCBN (Ours)	10.24	59.83	0.912	41.56	3.24

The reported MSCPN accuracy (10.24% MAPE, $R^2 = 0.912$) measures the performance on proxy embodied carbon labels derived from building characteristics and ICE emission factors, not on ground-truth construction data. The 23.7% improvement over CatBoost-SHAP indicates that the hierarchical graph attention more effectively captures building structure within this proxy setting, establishing a proof-of-concept for the proposed architecture. However, reliable engineering deployment requires validation against measured construction-phase data, as real-world embodied carbon is influenced by unobserved factors (e.g., supplier efficiency, waste rates, transport modes and site practices). A superior proxy-label performance is therefore necessary but not sufficient for real-world accuracy.

4.3. Training Convergence Analysis

The convergence behavior of training for the MSCPN model is shown in Figure 6, which shows that there is a stable dynamic of multi-scale learning.




Figure 6. Training convergence analysis demonstrating stable learning dynamics. (a) Multi-scale loss curves showing convergence behavior across hierarchical prediction levels: component-level loss (blue) decreases from 0.89 to 0.12, system-level loss (orange) decreases from 0.76 to 0.15, building-level loss (green) decreases from 0.65 to 0.11, and consistency loss (red) decreases from 0.34 to 0.08, all stabilizing within 150 epochs; (b) R^2 score progression tracking model performance improvement from initial 0.55 at epoch 1 to final 0.912 at epoch 200, with 95% confidence interval shading (gray) indicating training stability and minimal variance across 5-fold cross-validation.

4.4. Optimization Performance

Figure 7 shows the carbon optimization curves realized by RL-COE against the optimization approaches of the baselines across construction periods.

Table 13 summarizes the optimization performance across different building types.

Table 13. Carbon optimization results by building type.

Building Type	Baseline (kgCO ₂ e/m ²)	MSCBN (kgCO ₂ e/m ²)	RL-COE (kgCO ₂ e/m ²)	Cost Impact	Time Impact
Office	452.6	412.3	378.1	+2.1%	+1.8%
Retail	398.2	361.5	344.6	+1.8%	+1.6%
Education	389.2	351.2	337.9	+2.4%	+1.9%

Healthcare	512.8	465.6	412.4	+3.2%	+2.5%
Warehouse	187.5	169.4	148.9	+1.2%	+0.9%
Residential	346.3	327.1	314.5	+1.7%	+1.4%
Average	384.3	359.0	330.0	+2.2%	+1.5%

The reported 18.4% carbon reduction achieved by RL-COE reflects an optimization within the proxy embodied carbon framework, where prediction and optimization share the same assumptions and emission factors, and therefore represents an internally consistent, simulation-based optimization. However, real-world reductions may be lower due to deviations from the proxy assumptions, practical implementation constraints, and unmodelled rebound effects. Accordingly, the 18.4% figure should be interpreted as an upper-bound theoretical potential rather than a verified engineering outcome, while the key contribution demonstrates that reinforcement learning can learn effective, adaptive carbon-optimization policies from digital twin feedback and outperform baseline optimization approaches.




Figure 7. Embodied carbon optimization trajectories comparing framework performance across construction phases. The baseline trajectory (dashed gray) represents unoptimized construction reaching 1285 kgCO₂e/m² at completion. RL-COE (solid blue) achieves 18.4% reduction, reaching 1048 kgCO₂e/m², NSGA-II (dotted orange) achieves 14.1% reduction, and ACO-Opt (dash-dot green) achieves 12.3% reduction. The shaded region (light blue) indicates cumulative carbon savings achieved by RL-COE through adaptive real-time decision-making including material substitutions, supplier changes, and construction sequencing optimization implemented at weekly decision points.

Because MSCPN prediction and RL-COE optimization operate under the same proxy assumptions and emission factors, the reported 18.4% reduction reflects a self-consistent,

closed-loop optimization within the simulation environment. In real construction, model mismatch, practical action constraints, and delayed or noisy feedback will reduce achievable gains. The reported reduction should therefore be interpreted as an upper-bound optimization ceiling under idealized conditions. Pilot deployments are required to quantify the model–reality gap and establish realistic, engineering-level performance.

4.5. Multi-Scale Prediction Analysis

Figure 8 demonstrates the multi-scale prediction capabilities of MSCPN, showing carbon distribution across materials, systems, and life cycle stages.




Figure 8. Multi-scale embodied carbon analysis demonstrating MSCPN’s hierarchical prediction capabilities. (a) Material-level contribution breakdown showing concrete dominates at 35% (450 kgCO₂e/m²), followed by steel at 22% (283 kgCO₂e/m²), aluminum at 15% (193 kgCO₂e/m²), glass at 12% (154 kgCO₂e/m²), masonry at 8% (103 kgCO₂e/m²), and other materials at 8% (103 kgCO₂e/m²); (b) system-level distribution with structural systems accounting for 45% (578 kgCO₂e/m²), envelope systems at 30% (385 kgCO₂e/m²), MEP systems at 15% (193 kgCO₂e/m²), and finishes at 10% (128 kgCO₂e/m²); (c) life cycle stage emissions with uncertainty bounds ($\pm 1\sigma$) from Monte Carlo dropout showing A1 (raw material extraction) contributes 43% with $\pm 12\%$ uncertainty, A3 (manufacturing) contributes 24% with $\pm 15\%$ uncertainty, A4 (transport to site) contributes 14% with $\pm 8\%$ uncertainty, A2 (transport to manufacturer) contributes 12% with $\pm 9\%$ uncertainty, and A5 (construction) contributes 7% with $\pm 6\%$ uncertainty, indicating the highest uncertainty in manufacturing stages.

Figure 9 evaluates the real-time inference capabilities of DT-LCAF across varying building complexities.




Figure 9. Real-time performance analysis across varying computational demands. (a) MSCP prediction latency scales sub-linearly with building complexity: 15 ms for 100 components, 28 ms for 500 components, 45 ms for 1000 components, 72 ms for 5000 components, and 89 ms for 10,000 components, maintaining sub-100 ms response suitable for interactive applications; (b) RL-COE optimization decision time increases linearly with action space size: 8 ms for 50 actions, 18 ms for 200 actions, 32 ms for 500 actions, 58 ms for 1000 actions, and 95 ms for 2000 actions; (c) end-to-end response time distribution showing median latency of 45.2 ms (50th percentile), 25th percentile at 32.8 ms, 75th percentile at 67.4 ms, 90th percentile at 89.5 ms, and 95th percentile at 102.3 ms, demonstrating consistent performance suitable for construction-phase decision support.

The reported latency results demonstrate that the MSCP and RL-COE algorithms execute efficiently under controlled, ideal conditions, achieving sub-100 ms response times and confirming the computational feasibility of real-time use. However, real deployments will incur additional delays from network communication, concurrent multi-user access, data validation and cleaning, and integration with enterprise systems, meaning end-to-end latency is likely to be 2–3× higher (approximately 100–250 ms). These results therefore represent a lower-bound performance benchmark and validate the framework’s architectural design rather than operational readiness. The full real-world validation of real-time performance requires pilot deployments on active construction sites that measure complete system latency and reliability.

4.6. Carbon Intensity Spatial Analysis

Figure 10 shows changes in the embodied carbon intensity of buildings by geographic location and type of building using CBECS analysis.




Figure 10. Embodied carbon intensity heatmap (kgCO₂e/m²) revealing systematic variations across building types and U.S. geographic regions. Healthcare facilities exhibit highest intensity ranging from 470 to 535 kgCO₂e/m² due to specialized equipment, stringent structural requirements, and complex MEP systems. Office buildings show 380–445 kgCO₂e/m², retail facilities 320–385 kgCO₂e/m², educational buildings 340–405 kgCO₂e/m², and warehouses display lowest intensity at 178–205 kgCO₂e/m², reflecting simpler structural systems and minimal interior finishes. Regional variations across Northeast, Southeast, Midwest, Southwest, and West regions reflect differences in prevalent construction materials, climate-driven design requirements, transportation distances, and local construction practices.

4.7. Ablation Study

Figure 11 presents the ablation study results examining the contribution of each framework component to the overall performance.




Figure 11. Ablation study examining component contributions to overall framework performance. Hierarchical attention mechanism provides 2.65% MAPE improvement (from 12.89% to 10.24%), enabling learned importance weights for material interdependencies. Multi-scale aggregation contributes 3.32% improvement (from 13.56% to 10.24%), capturing component–system–building carbon accumulation hierarchy. Consistency loss adds 1.58% improvement (from 11.82% to 10.24%), enforcing hierarchical coherence across prediction levels. Dynamic emission factors provide 3.99% improvement (from 14.23% to 10.24%), incorporating supplier-specific and logistics variability. RL-COE contributes 7.1% to carbon reduction rate compared to greedy baseline. Error bars indicate standard deviation across 5-fold cross-validation. All components demonstrate statistically significant contributions ($p < 0.001$).

Table 14 provides the detailed numerical results for the ablation study.

Table 14. Ablation study results.

Configuration	MAPE (%)	R ²	CRR	Latency (ms)	Params (M)
Full DT-LCAF	10.24	0.912	184	72.4	4.82
<i>w/o Ref. Attention</i>	12.89	0.873	165	88.2	3.95
<i>w/o Multi-Scale Agg.</i>	13.56	0.858	159	45.6	3.12
<i>w/o Consistency Loss</i>	11.82	0.892	172	72.1	4.08
<i>w/o Dynamic FE</i>	14.23	0.841	147	68.3	2.62
<i>w/o RL-COE (Greedy)</i>	10.24	0.912	183	78.5	4.14
Single-Scale GNN	15.67	0.823	121	31.4	1.86

Sample size: $n = 5900$ buildings (CBECs dataset); 5-fold cross-validation resulting in 29,500 total predictions per configuration ($5900 \text{ buildings} \times 5 \text{ folds}$). Each configuration trained independently with identical hyperparameters except for removed component. Metrics reported as mean $\pm$ standard deviation across 5 folds. Statistical significance assessed via paired t-tests comparing each ablation configuration to full DT-LCAF framework, with Bonferroni correction for multiple comparisons ($\alpha = 0.05/6 = 0.0083$). All reported differences significant at $p < 0.001$ (see Table 14 for detailed p -values).

4.8. Comprehensive Performance Comparison

Table 15 provides a comprehensive comparison across all methods, metrics, and building types.

Table 15. Comprehensive performance comparison across methods, metrics, and building types.

Method	Prediction Metrics			Carbon Reduction Rate by Building Type (%)						Efficiency	
	MAPE(%)	RMSE	R ²	Office	Retail	Edu.	Health	Ware.	Hotel	Lat. (ms)	Params
Static LCA	28.45	156.32	0.612	–	–	–	–	–	–	2.1	–
BIM-LCA [4]	22.18	128.76	0.695	–	–	–	–	–	–	8.4	–
RF-Carbon	18.92	108.43	0.754	–	–	–	–	–	–	5.6	0.02 M
XGBoost	16.84	96.28	0.789	–	–	–	–	–	–	6.2	0.08 M
LSTM	17.56	101.92	0.772	–	–	–	–	–	–	15.3	0.52 M
GCN	15.23	89.45	0.814	–	–	–	–	–	–	18.7	1.24 M
GAT	14.67	85.32	0.832	–	–	–	–	–	–	22.4	1.86 M
ML-Prov. [11]	15.89	92.14	0.803	–	–	–	–	–	–	7.8	0.15 M
CatBoost [12]	13.42	78.65	0.856	–	–	–	–	–	–	9.1	0.12 M
DT-Block. [28]	14.12	82.47	0.841	–	–	–	–	–	–	28.6	–
ACO-Opt [17]	15.67	91.23	0.808	11.2	10.8	13.1	14.5	8.9	12.7	156.4	–
NSGA-II	14.89	87.56	0.825	13.5	12.9	15.2	16.3	10.4	14.8	243.7	–
MSCPN	10.24	59.83	0.912	–	–	–	–	–	–	32.4	2.14 M
RL-COE	–	–	–	17.8	16.4	19.7	21.2	14.3	18.9	40.0	2.68 M

DT-LCAF	10.24	59.83	0.912	17.8	16.4	19.7	21.2	14.3	18.9	72.4	4.82 M
---------	-------	-------	-------	------	------	------	------	------	------	------	--------

RMSE in kgCO₂e/m²; “–”: not applicable; best results in bold; all improvements significant at $p < 0.001$.

4.9. Statistical Significance Analysis

Figure 9 shows changes in the embodied carbon intensity of buildings by geographic location and type of building using CBECS analysis. Table 16 shows the statistical significance analysis (p -values).

Table 16. Statistical significance analysis (p -values).

Comparison	MAPE	RMSE	R ²	CRR
MSCBN vs. Static LCA	<0.001	<0.001	<0.001	–
MSCBN vs. GAT-Carbon	<0.001	<0.001	<0.001	–
MSCBN vs. GCN-Carbon	<0.001	<0.001	<0.001	–
MSCBN vs. ML-Physical	<0.001	<0.001	<0.001	–
MSCBN vs. XGBoost	<0.001	<0.001	<0.001	–
RL-COE vs. NSGA-II	–	–	–	<0.001
RL-COE vs. ACO-Q	–	–	–	<0.001

RL-COE generates mostly feasible, high-carbon-reduction recommendations, with 98% of actions satisfying structural, cost, and schedule constraints after two-stage filtering. Only 2% of recommendations carry minor cost or schedule warnings for user review. Compared to NSGA-II (28.3%) and ACO-Opt (34.7%), RL-COE’s reward-guided policy effectively balances aggressive optimization with practical feasibility.

4.10. Framework Capability Comparison

Table 17 evaluates the limitations of DT-LCAF alongside the existing frameworks according to the main functional dimensions, and the comparison reveals that the proposed framework is among the few existing approaches that collectively address the full set of features required to enable real-time embodied carbon optimization.

Table 17. Framework capability comparison (✗ means not available, ✓ means available).

Framework	Real-Time	ML Pred.	Multi-Scale	Optim.	IoT Int.
Static LCA	✗	✗	✗	✗	✗
BIM-LCA	✗	✗	✗	✗	✗
ML-Carbon	✓	✓	✗	✗	✗
DT-Blockchain	✓	✓	✓	✗	✓
ML-Thermal [29]	✓	✓	✓	✗	✗
NSGA-II [17]	✗	✗	✗	✓	✗
ACO-Q [18]	✗	✗	✗	✓	✗
DT-LCAF (Ours)	✓	✓	✓	✓	✓

4.11. Sensitivity Analysis

Table 18 conducts a sensitivity analysis of the impact of important hyperparameters on the performance of the framework to guide practitioners who need to adapt DT-LCAF to the specifics of various projects.

Table 18. Sensitivity analysis of key hyperparameters.

Parameter	Range Tested	Optimal Value	Sensitivity
GNN layers	2, 3, 4, 5, 6	4	Medium
Attention heads	4, 8, 12, 16	8	Low

Hidden dimension	64, 128, 256, 512	128	Medium
Learning rate	1×10^{-5} to 1×10^{-3}	3×10^{-4}	High
Batch size	32, 64, 128, 256	64	Low
PPO clip ϵ	0.1, 0.2, 0.3	0.2	Medium
Loss weight λ_1	0.3, 0.5, 0.7	0.5	Low

4.12. BIM-Based Case Study: Construction-Phase Carbon Dynamics

This subsection presents a demonstrative case study illustrating the DT-LCAF framework's construction-phase update mechanism using a simulated three-story office building. The case study demonstrates how embodied carbon estimates evolve dynamically as construction progresses and how RL-COE optimization decisions reduce total embodied carbon while respecting project constraints.

The case study building is a 2400 m² three-story commercial office located in a temperate climate zone. The building comprises a reinforced concrete frame structure with steel roof trusses, a curtain wall envelope with aluminum framing and double-glazed units, conventional MEP systems including packaged HVAC units and copper plumbing, and standard interior finishes including suspended ceilings, ceramic tile flooring, and painted gypsum partitions.

The BIM model was developed to LOD 350, providing sufficient detail for quantity extraction. Table 19 summarizes the material inventory extracted from the IFC model.

Table 19. Material inventory extracted from BIM model.

Material Category	Quantity	Unit	Primary Components
Concrete (C40)	485	m ³	Foundations, columns, beams, slabs
Reinforcement steel	48.2	tonnes	Rebar in all concrete elements
Structural steel	32.5	tonnes	Roof trusses, connections
Aluminum (frames)	8.4	tonnes	Curtain wall framing, doors
Glass (double-glazed)	580	m ²	Windows, curtain wall panels
Brick masonry	125	m ³	Internal partitions, facades
Copper (plumbing)	1.8	tonnes	Pipes, fittings
Gypsum board	2850	m ²	Partitions, ceilings
Ceramic tiles	1420	m ²	Flooring, wet areas
Insulation (mineral wool)	1650	m ²	Roof, external walls

The embodied carbon was calculated for each material using the dynamic emission factor model described in Section 3.2.2. Table 20 presents the baseline embodied carbon calculation prior to optimization.

Table 20. Baseline embodied carbon calculation by material and life cycle stage.

Material	A1 (tCO _{2e})	A2 (tCO _{2e})	A3 (tCO _{2e})	A4 (tCO _{2e})	A5 (tCO _{2e})	Total (tCO _{2e})
Concrete	142.3	12.8	38.5	18.2	8.4	220.2
Rebar steel	68.5	8.2	22.4	6.8	3.2	109.1
Structural steel	52.8	6.5	18.2	5.4	2.8	85.7
Aluminum	85.2	4.2	12.8	3.5	1.8	107.5
Glass	42.5	3.8	15.2	4.2	2.1	67.8
Masonry	28.4	2.5	8.5	3.8	2.2	45.4
Copper	12.8	1.2	4.5	1.1	0.6	20.2
Gypsum	18.5	2.2	6.8	2.4	1.4	31.3
Ceramic	22.4	2.8	8.2	2.2	1.2	36.8
Insulation	35.2	3.2	12.5	3.1	1.5	55.5

Total	508.6	47.4	147.6	50.7	25.2	779.5
--------------	--------------	-------------	--------------	-------------	-------------	--------------

The construction program spans 12 weeks organized into five phases: foundation (weeks 1–2), structure (weeks 3–5), envelope (weeks 6–8), MEP installation (weeks 9–10), and finishes (weeks 11–12). The digital twin was configured to receive simulated IoT updates at daily intervals, triggering MSCPN recalculation upon material delivery confirmations, installation completion signals, and supplier change notifications. Figure 12 illustrates the construction-phase carbon dynamics observed during the simulation. Panel (a) shows cumulative embodied carbon over the 12-week construction period, comparing the baseline scenario (no optimization) against the DT-LCAF optimized scenario. The baseline trajectory reaches 1285 tCO_{2e} at project completion, while the optimized trajectory reaches 1053 tCO_{2e}, representing a 17.9% total reduction.

The carbon accumulation curve exhibits characteristic patterns aligned with construction phases. The steepest increase occurs during weeks 1–5 (foundation and structure phases), when high-carbon materials including concrete and steel are installed. The accumulation rate decreases during the later phases as lower-carbon finishes are installed.

Panel (b) presents the weekly carbon breakdown by EN 15978 life cycle stage (A1–A5). Stage A1 (raw material extraction) dominates, contributing approximately 43% of weekly carbon, followed by A3 (manufacturing) at 24%, A4 (transport to site) at 14%, A2 (transport to manufacturer) at 12%, and A5 (construction) at 7%. This distribution is consistent with the literature findings for commercial buildings.

Panel (c) compares embodied carbon by building system between baseline and optimized scenarios. The structural system contributes the largest share (40%), followed by envelope (29%), MEP (15%), and finishes (15%). Optimization achieves the largest absolute reduction in structural carbon (102 tCO_{2e}, 20% reduction) through concrete mix optimization and steel supplier changes.

Panel (d) of Figure 12 displays the key optimization decisions generated by RL-COE during the simulation. Six significant interventions were implemented.

During week 2, RL-COE recommended a concrete mix optimization with a 30% ground granulated blast furnace slag (GGBS) replacement, reducing concrete carbon intensity from 0.45 to 0.32 kgCO_{2e}/kg and saving 45 tCO_{2e}. This recommendation satisfied the structural constraint (C40 strength maintained with GGBS mix).

During week 3, the engine recommended switching to a local steel supplier with 85% recycled content, reducing the transport distance from 320 km to 85 km and the emission factor from 1.42 to 1.05 kgCO_{2e}/kg, saving 38 tCO_{2e}. The cost increase of 3.2% remained within the 5% tolerance threshold.

During week 5, rebar prefabrication was recommended, reducing the on-site waste from 12% to 5% and saving 28 tCO_{2e} through reduced material consumption.

During week 7, a cladding material switch from aluminum to timber framing was recommended for non-structural facade elements, saving 52 tCO_{2e}. This represented the largest single intervention, enabled by an early identification during envelope procurement.

During week 9, MEP routing optimization reduced the ductwork length by 15% through digital twin-based spatial analysis, saving 22 tCO_{2e}.

During week 11, the specification of water-based paint systems reduced the VOC-associated carbon by 8 tCO_{2e}.

All RL-COE recommendations satisfied the defined project constraints. Table 21 summarizes the constraint compliance for each optimization decision.

Table 21. Constraint satisfaction analysis for RL-COE optimization decisions.

Decision	Carbon Saving Cost Impact		Schedule Impact	Structural Compliance
Concrete GGBS mix	45 tCO ₂ e	+2.1%	0 days	C40 verified
Local recycled steel	38 tCO ₂ e	+3.2%	0 days	Grade 500 B maintained
Rebar prefabrication	28 tCO ₂ e	−1.5%	−2 days	Equivalent detailing
Timber cladding	52 tCO ₂ e	+4.8%	+1 day	Non-structural element
MEP routing	22 tCO ₂ e	−0.8%	0 days	N/A
Water-based paint	8 tCO ₂ e	+1.2%	0 days	N/A
Total	193 tCO₂e	+1.5%	−1 day	All satisfied

During the 12-week simulation, the digital twin processed 847 IoT events, triggering 156 MSCPN recalculations. Table 22 summarizes the event distribution and update frequency.

Table 22. Digital twin update events during case study simulation.

Event Type	Count	Avg. per Week	Carbon Prediction Update
Material delivery confirmation	124	10.3	Yes
Installation completion	89	7.4	Yes
Progress milestone	48	4.0	Yes
Environmental condition change	412	34.3	No (below threshold)
Supplier notification	12	1.0	Yes
Schedule revision	6	0.5	Yes
Total triggering updates	156	13.0	—

This case study employs simulated rather than measured construction data and should be interpreted as a demonstration of the framework capabilities rather than a validation of prediction accuracy. The material quantities are derived from a representative BIM model, emission factors are from standardized databases rather than supplier-specific EPDs, and the construction timeline is simulated rather than observed from actual site operations. Validation on real construction projects with ground-truth embodied carbon measurements remains an important direction for future work.




Figure 12. BIM-based case study demonstrating construction-phase carbon dynamics for a 2400 m² office building. (a) Cumulative embodied carbon over 12-week construction period comparing baseline scenario (no optimization, reaching 1285 tCO_{2e}) versus DT-LCAF optimized scenario (reaching 1053 tCO_{2e}), achieving 17.9% total reduction with steepest accumulation during foundation and structural phases (weeks 1–5); (b) weekly carbon breakdown by EN 15978 life cycle stage showing A1 (raw material extraction) dominates at 43%, followed by A3 (manufacturing) at 24%, A4 (transport to site) at 14%, A2 (transport to manufacturer) at 12%, and A5 (construction) at 7%; (c) embodied carbon by building system comparing baseline and optimized scenarios: structural system 40% baseline reduced by 20% (102 tCO_{2e} savings), envelope 29% reduced by 18%, MEP 15% reduced by 15%, and finishes 15% reduced by 12%; (d) timeline of RL-COE optimization decisions showing six key interventions: week 2 concrete mix optimization with 30% GGBS replacement (45 tCO_{2e} savings), week 3 local recycled steel supplier switch (38 tCO_{2e}), week 5 rebar prefabrication (28 tCO_{2e}), week 7 timber cladding substitution (52 tCO_{2e}, largest intervention), week 9 MEP routing optimization (22 tCO_{2e}), and week 11 water-based paint specification (8 tCO_{2e}).

5. Discussion

5.1. Key Findings and Interpretation

The experimental results demonstrate several key findings that advance the understanding of embodied carbon prediction and optimization in construction, each supported by quantitative evidence from the validation studies and contextualized within the existing literature.

Finding 1—Hierarchical Graph Attention Improves Carbon Prediction Accuracy

MSCPN achieves a 10.24% MAPE (Table 9), representing a 23.7% relative improvement over the best baseline method, CatBoost-SHAP (13.42% MAPE). This improvement is statistically significant ($p < 0.001$, Table 13, paired t -test) and consistent across building types: office buildings, 9.8% MAPE; retail, 10.5%; education, 10.1%; healthcare, 11.2%; warehouses, 8.9% (Figure 10 heatmap). The R^2 coefficient of 0.912 indicates that MSCPN explains 91.2% of the variance in proxy embodied carbon labels, compared to 85.6% for CatBoost-SHAP.

Literature Context: Prior research on building carbon prediction has predominantly employed flat machine learning models. Zhang et al. [11] achieved an 8.7% MAPE for provincial-level operational carbon using ensemble methods (XGBoost, LightGBM, CatBoost), while Zheng et al. [10] reported an $R^2 = 0.847$ for building operational carbon using random forest regression. Our results demonstrate that embodied carbon prediction benefits from hierarchical architectural modeling, consistent with findings in building energy modeling where hierarchical approaches outperform flat representations [11]. The 23.7% improvement over CatBoost ensemble suggests that architectural inductive bias (explicitly modeling building structure) provides a greater benefit than algorithmic sophistication (ensemble diversity) for embodied carbon tasks.

Ablation Study Evidence (Table 11): Removing multi-scale aggregation degrades the MAPE to 13.56% (+3.32 pp), the largest single component contribution, validating the architectural hypothesis. Removing hierarchical attention degrades it to 12.89% (+2.65 pp), demonstrating that learned attention weights capture material interdependencies more effectively than uniform aggregation used in standard GCN architectures.

Physical Interpretation: Buildings are compositionally hierarchical—components aggregate to systems, and systems to buildings. Embodied carbon accumulates through this same hierarchy: concrete columns (components) + steel beams (components) = structural frame (system) → whole building. MSCPN's architectural alignment with physical structure enables learning representations that respect these compositional relationships, whereas flat models must learn building hierarchy implicitly from data, a more difficult learning problem requiring more training examples [12].

Finding 2—Dynamic Emission Factors Substantially Reduce Prediction Uncertainty

Incorporating dynamic emission factors (supplier efficiency α_{sup} , transportation modifier β_{trans} , manufacturing efficiency γ_{mfg}) improves the MAPE from 14.23% (static database factors) to 10.24% (Table 11 ablation), a 3.99 percentage point improvement. The uncertainty quantification results (Figure 8c) show that prediction confidence intervals narrow from $\pm 18\text{--}22\%$ (static factors) to $\pm 12\text{--}15\%$ (dynamic factors) for stages A1–A3.

Literature Context: Traditional LCA studies acknowledge a $\pm 20\text{--}40\%$ uncertainty in embodied carbon estimates [20], primarily from emission factor variability. Our results quantify that 3.99 pp of prediction error (39% of total error from static baseline 10.24%) arises from using generic database values rather than supplier-specific factors. This aligns with the findings of [23] showing that supplier-specific EPDs reduce uncertainty by 30–45% compared to industry-average databases.

Case Study Evidence (Tables 16 and 17): The BIM case study demonstrates the dynamic emission factor impact: concrete with a 30% GGBS replacement (supplier-specific) showed 0.32 kgCO_{2e}/kg vs. the 0.45 kgCO_{2e}/kg database average (29% reduction). A local recycled steel supplier provided 1.05 kgCO_{2e}/kg vs. 1.42 kgCO_{2e}/kg database (26% reduction). These 25–30% differences translate to 45 and 38 tCO_{2e} savings, respectively, for the 2400m² building—differences that static LCA would miss entirely.

Practical Implication: The results quantify the value of supplier-specific Environmental Product Declarations (EPDs). Practitioners investing in EPD collection reduce prediction uncertainty proportionally: comprehensive EPD coverage (>80% materials by mass) could achieve a <12% prediction error, while minimal EPD usage (<20% coverage) degrades to >16% error. This cost–benefit analysis informs data collection prioritization.

5.2. Practical Implications

The DT-LCAF model provides significant practical value to many stakeholders in the construction industry. To project managers, carbon predictions in real time would make it possible to make proactive decisions when constructing them, since any meaningful

intervention is still possible. Instead of being provided with LCA results once the project is complete, the managers will have an opportunity to discuss carbon accumulation in real time and act based on the available opportunities for optimization or disrupted supply chain instances.

To architects and designers, the multi-scale forecasting ability is a beneficial tool for feedback on the implications of design choices on the embodied carbon of systems and entire buildings. The visualization dashboard helps to quickly compare design options to their quantified carbon impact to underpin evidence-based design choices defining the trade-offs between design aesthetics, functionality, and environmental performance.

To sustainability consultants and regulatory compliance teams, DT-LCAF gives them the all-around carbon tracking required to meet emerging whole-life carbon reporting obligations. The fact that the framework is able to reflect the reality of construction conditions unlike design assumptions enhances the effectiveness of carbon disclosures and places the notion of sustainability in a credible statement.

To material suppliers and purchase departments, the optimization suggestions outline the possibility of low-carbon material replacements that can be advantageous as well. The fact that the framework takes into consideration supplier-specific emission factors motivates suppliers to report on precise environmental values and enhance their efficiency in production.

5.3. Implications for Smart City Planning and Policy

In addition to the optimization of projects at the individual project level, the DT-LCAF framework is directly applicable to the development of smart cities because it could allow for building smart cities with city-wide embodied carbon intelligence in the portfolios of buildings. By combining various building-scale digital twins, municipalities will be able to track carbon emissions related to construction at both district and city scales, detect spatial and temporal carbon hotbeds and evaluate the overall effect of construction activities on the city-wide climate goals. In relation to climate change mitigation, the framework helps in undertaking proactive policy assessment to ensure that urban planners have an opportunity to forecast the carbon implication of alternative material regulations, procurement plans and construction operations prior to being established. The ability allows for the comparison of policy situations based on data and helps to make evidence-based decisions in accordance with net-zero and action plans.

The DT-LCAF covers the resilience of cities in relation to developing construction policies in changing environmental, supply chain, and regulatory environments. The online feedback of the digital twin will enable the planners to foresee disruptions, streamline material logistics, focus on low-carbon interventions in infrastructure projects that are of high risk, and enhance the resilience to climate and market uncertainties. In sum, the suggested structure adds to smart urban design, since it will help in closing the gap between construction-phase carbon evaluations and dynamic city governance. DT-LCAF will help municipalities, planners, and policymakers to promote sustainable, resilient and climate-responsive urban development by converting the quantity of embodied carbon data into actionable intelligence.

5.4. Limitations and Future Work

There are multiple restrictions that should be mentioned and recommendations regarding further research. First, the framework is based on BIM models that fulfill a minimum Level of Development (LOD) of 300 successful extractions of quantities, which restricts the sheer scope to projects that have already invested in the development of detailed digital models. Future research must look at ways of using the framework on

projects whose documentation is less detailed using transfer learning or data augmentation algorithms.

Second, the existing implementation allows for implementing embodied carbon in stages A1–A5 (cradle-to-practical-completion) but omits operational carbon (stage B1–B7) and end-of-life carbon (stage C1–C4). Or, to have a more comprehensive evaluation of the environment, whole-life carbon assessment would be mandatory, but needs alternative modeling techniques to capture the temporal nature of operational energy consumption.

Third, although validation on both CBECS and BDGP2 datasets can be used to illustrate the effectiveness of the framework to operate on large-scale portfolios of buildings, the application to further real-world construction projects where ground-truth measurements of embodied carbon are available would increase the confidence in applicability. This validation would need collaborations with construction firms that would be happy to share project information and apply the framework in ongoing construction.

Fourth, the optimal policies obtained using the reinforcement learning approach need a lot of training episodes, which could be expensive in regard to new building types or construction conditions without a previously recorded history. The next research direction should focus on meta-learning methods that would allow adapting quickly to new types of project input with minimal progressive training.

Lastly, the contemporary DT-LCAF application is limited to the individual building level. Several other capabilities, such as federated or hierarchical digital twin applications, distributed computation to process large-scale graphs and reinforcement learning, and standardized city-scale data integration across heterogeneous BIM applications and sensor systems, would be necessary to scale the framework to district/city-level applications. The issues of data interoperability, privacy, and cross-building coordination will be addressed to take the proposed framework to the next level of embodied carbon management for cities in scale.

5.5. Environmental Sustainability Implications and Integration with Energy Communities (Concise Version)

The DT-LCAF framework contributes to environmental sustainability beyond embodied carbon optimization by enabling an integration with energy communities and coordinated urban decarbonization strategies.

5.5.1. Embodied Carbon Reduction and Climate Change Mitigation

The experimental results show that DT-LCAF achieves an average embodied carbon reduction of 18.4% across building types, with healthcare facilities reaching 21.2%. Given that embodied carbon accounts for approximately 11% of global energy-related CO₂ emissions (≈ 3.7 GtCO_{2e} annually), the large-scale adoption of DT-LCAF could avoid 0.55–0.74 GtCO_{2e} per year, representing a substantial mitigation potential.

Beyond direct reductions, DT-LCAF generates systemic sustainability effects by incentivizing suppliers to improve production efficiency and adopt lower-carbon manufacturing processes. Its dynamic supplier selection mechanism creates competitive pressure for decarbonization across material supply chains, particularly in high-impact sectors such as cement and steel. Furthermore, the preference for local suppliers reduces transport-related emissions, air pollution, and congestion, as illustrated by the 38 tCO_{2e} savings in the case study.

5.5.2. Integration with Energy Communities and Distributed Energy Systems

DT-LCAF exhibits strong synergies with emerging energy communities. The digital twin infrastructure used for construction-phase carbon monitoring can be extended to support operational energy management, enabling building-level demand forecasting and demand-response services. Lower-carbon material choices often enhance thermal performance, reducing operational energy demand and improving the viability of distributed renewable energy systems.

The framework's multi-building aggregation capability aligns with the district-scale organization of energy communities, enabling coordinated construction planning, shared logistics, and the collective procurement of low-carbon materials. In addition, the reinforcement learning optimization engine can be extended to support peer-to-peer energy trading, shared storage management, and coordinated electric vehicle charging, while balancing environmental and economic constraints.

An integrated deployment pathway enables DT-LCAF to guide low-carbon construction during development and subsequently transition to operational energy optimization once buildings are occupied, ensuring continuity from construction to operation.

5.5.3. Synergies with Complementary Technologies

DT-LCAF can be effectively combined with collaborative BIM environments compliant with ISO 19650 [ISO 19650 is the international standard for managing information over the entire lifecycle of a built asset using Building Information Modeling (BIM)] [15], enabling the bidirectional exchange of design data and real-time carbon intelligence. Blockchain technologies can support data provenance, auditable emission factors, and automated sustainability-linked contracts. The framework is particularly compatible with modular and prefabricated construction, where component-level tracking enables early-stage carbon optimization and can achieve substantially higher reductions when combined with material optimization strategies. DT-LCAF can further support circular economy practices through digital material passports and integration with urban digital twin platforms for cross-sector optimization at a city scale.

5.5.4. Adoption Challenges and Enablers

Widespread implementation requires the improved standardization of EPDs and digital material data, supportive regulatory frameworks mandating embodied carbon reporting, and strengthened industry capacity in digital twin and carbon assessment skills. Economic viability depends on business models that align upfront digital investments with long-term sustainability and cost benefits, including sustainability-linked finance and service-based deployment models.

5.5.5. Contribution to Sustainable and Resilient Urban Development

DT-LCAF provides actionable intelligence for sustainable urban development by enabling municipalities to establish construction carbon baselines, define realistic reduction targets, and monitor progress quantitatively. It addresses the growing importance of embodied emissions, which are projected to exceed 50% of whole-life building emissions by 2050, and supports resilient city planning by enabling adaptive responses to supply chain disruptions, regulatory changes, and emerging low-carbon materials. Overall, DT-LCAF translates sustainability objectives into real-time, construction-phase decisions, supporting an evidence-based decarbonization of the built environment.

5.6. Validation Scope and Path to Real-World Application

This study validates the architectural feasibility, algorithmic performance on proxy labels, and computational scalability of the DT-LCAF framework, demonstrating an effective integration of heterogeneous data, superior hierarchical carbon prediction, and real-time computational capability. However, it does not validate true embodied carbon accuracy, real-world emission-reduction effectiveness, robustness under practical site conditions, or long-term model reliability. Engineering deployment therefore requires ground-truth data collection, pilot field studies, integration with industry platforms, and calibrated uncertainty handling. A staged pathway—retrospective validation, monitoring-only deployment, assisted decision support, and finally integrated operational use—is proposed to bridge the gap from proof-of-concept to practice. Despite these limitations, the work contributes a novel prediction architecture, an extensible digital twin optimization framework, reproducible baselines, and a clear research roadmap toward engineering-ready construction carbon intelligence.

6. Conclusions, Research Implications and Future Works

6.1. Main Conclusions

This article introduced a new digital twin-based embodied carbon optimization Digital Twin-Life Cycle Assessment Framework, named DT-LCAF, to conduct optimally in regard to a smart building construction project life cycle. The framework fixes key shortcomings within the current solution by connecting BIM models, IoT sensor network, and sophisticated machine learning algorithms to the resultant digital twin system that allows the continuous monitoring of carbon and adaptable maximization at any construction stage.

The Multi-Scale Carbon Prediction Network (MSCPN) proposes hierarchical graph attention mechanisms which effectively represent the interdependencies of materials in components, systems, and buildings and outperforms state-of-the-art baselines such as CatBoost-SHAP (13.42%) and GAT-Carbon (14.67%), with a lower prediction error (MAPE: 10.24%). The Reinforcement Learning-based Carbon Optimization Engine (RL-COE) produces adaptive material substitution and construction sequence recommendations that on average reduce embodied carbon by 18.4 percent, which is much higher than the results of NSGA-II (14.1 percent) and ACO-Opt (12.3 percent).

The effectiveness of the framework on two large-scale public datasets, including CBECS (5900 commercial buildings) and Building Data Genome Project 2 (1636 buildings with 53.6 million measurements) is demonstrated through a comprehensive validation across a wide range of building types and construction settings. The framework achieves a computational performance suitable for real-time operation, with an algorithm execution latency under 100 ms for buildings containing up to 10,000 components. This establishes architectural feasibility for real-time deployment, though operational validation on active construction sites with live IoT data streams, system integration overhead, and concurrent user access remains essential future work before claiming production readiness. The current validation demonstrates a potential for real-time operation rather than confirmed real-time performance under engineering deployment conditions.

The proposed DT-LCAF framework, with its clear focus on issues of scalability to buildings beyond urban applicability and policy relevance, is associated with the objectives and definition of the scope of the journal *Smart Cities*. Its framework will add actionable, evidence-based intelligence to sustainable urban development, climate change mitigation, and resilient city planning, which will further increase its appropriateness to smart city-oriented research and practical urban governance.

6.2. Research Implications: The DT-LCAF Framework Presents Several Important Implications for Both Academic Research and Practical Applications in the Construction Industry

For Researchers: The framework demonstrates that hierarchical graph attention mechanisms significantly improve embodied carbon prediction accuracy by explicitly modeling building structure, suggesting that future building performance prediction models should incorporate architectural hierarchy rather than treating buildings as flat feature collections. The integration of digital twin technology with reinforcement learning establishes a new paradigm for adaptive construction optimization that learns from real-time feedback.

For Practitioners: The framework provides actionable decision support that project managers, sustainability consultants, and procurement teams can deploy during active construction when interventions remain feasible. The constraint-aware optimization ensures that recommendations respect structural, economic, and temporal boundaries, addressing the real-world trade-offs practitioners face.

For Policymakers: The framework's scalability to district and city levels enables municipalities to monitor construction-sector carbon contributions, evaluate embodied carbon implications of material regulations, and implement evidence-based urban sustainability governance.

6.3. Future Research Directions: Several Promising Directions Emerge for Future Research

Ground-Truth Validation: Deployment on construction projects with comprehensive embodied carbon measurement to validate the prediction accuracy against actual rather than proxy labels, requiring industry partnerships for data collection.

Whole-Life Carbon Integration: Extension to operational (B1–B7) and end-of-life (C1–C4) stages for comprehensive building Life Cycle Assessment, requiring temporal modeling techniques for long-duration operational energy consumption.

Transfer Learning: Development of meta-learning approaches for rapid adaptation to new building types or construction contexts with minimal retraining, addressing the data scarcity challenge in specialized construction domains.

Federated Digital Twins: Investigation of distributed computation architectures for city-scale embodied carbon management, addressing data interoperability, privacy preservation, and cross-building coordination challenges.

Circular Economy Integration: Incorporation of material reuse, recycling potential, and end-of-life recovery value into optimization objectives, aligning with circular construction principles.

Funding: The author received no specific funding for this study.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets used in this study are publicly available. CBECS 2018 is available from the U.S. Energy Information Administration at <https://www.eia.gov/consumption/commercial/data/2018/> (accessed on 25 November 2025). The Building Data Genome Project 2 is available on GitHub 3.15 at <https://github.com/buds-lab/building-data-genome-project-2> (accessed on 25 November 2025). Framework source code and trained models will be released upon paper acceptance at <https://github.com/nbelwi/DT-LCAF-.git> (accessed on 25 November 2025).

Conflicts of Interest: There is no conflict of interests.

Nomenclature

Symbol	Description
EC_{total}	Total embodied carbon (kgCO ₂ e)
$EC_{i,j}$	Embodied carbon of component i in stage j
Q_i	Quantity of component i from BIM

$EF_{i,j}$	Emission factor for component i , stage j
$EF_{i,j}^{dyn}(t)$	Dynamic emission factor at time t
$\omega_{i,j}$	Waste factor for component i , stage j
$\alpha_{sup}(t)$	Supplier efficiency factor
$\beta_{trans}(d_i)$	Transportation distance modifier
$\gamma_{mfg}(t)$	Manufacturing efficiency factor
$\mathcal{G}$	Building component graph
$\mathcal{V}, \mathcal{E}$	Node and edge sets of graph
$\mathbf{X}$	Node feature matrix
$\mathbf{h}_i^{(l)}$	Hidden representation of node i at layer l
$\alpha_{ij}^{(l)}$	Attention coefficient between nodes i and j
$\mathbf{h}_{sys}^k$	System-level representation for system k
$\mathbf{h}_{bld}$	Building-level representation
s_t	State at construction step t
a_t	Action at construction step t
r_t	Reward at construction step t
π_θ	Policy network with parameters θ
V_ϕ	Value function with parameters ϕ
γ	Discount factor for RL
ϵ	PPO clipping parameter
$\hat{A}_t$	Advantage estimate at time t
Acronyms	
BIM	Building Information Modeling
LCA	Life Cycle Assessment
IoT	Internet of Things
MSCP	Multi-Scale Carbon Prediction Network
RL-COE	RL-Based Carbon Optimization Engine
PPO	Proximal Policy Optimization
GAE	Generalized Advantage Estimation
GNN	Graph Neural Network
GAT	Graph Attention Network
EPD	Environmental Product Declaration
MAPE	Mean Absolute Percentage Error
RMSE	Root Mean Square Error
CRR	Carbon Reduction Rate

Appendix A. Mathematical Derivations

Table A1. Generalized Advantage Estimation (GAE).

The advantage function:	
	$A_t = Q(st, at) - V(st)$
TD error:	
	$\delta_t = rt + \gamma V(st + 1) - V(st)$
GAE with exponential decay:	
	$A_t^{GAE(\lambda)} = 1 = 0 \sum_{l=0}^{\infty} (\gamma \lambda)^l \delta_{t+l} + 1$
$\lambda = 0 \rightarrow$ high bias, low variance (1-step TD)	
$\lambda = 1 \rightarrow$ low bias, high variance (Monte Carlo)	
$\lambda \approx 0.95 \rightarrow$ practical trade-off	
For finite episodes of length TTT:	
	$A^t = 1 = \sum_{l=0}^{T-t-1} (\gamma \lambda)^l \delta_{t+l} + 1$

References

1. Zhang, Z.; Wei, Z.; Court, S.; Yang, L.; Wang, S.; Thirunavukarasu, A.; Zhao, Y. A review of digital twin technologies for enhanced sustainability in the construction industry. *Buildings* **2024**, *14*, 1113. <https://doi.org/10.3390/buildings14041113>.

2. Omrany, H.; Mehdipour, A.; Oteng, D. Digital twin technology and social sustainability: Implications for the construction industry. *Sustainability* **2024**, *16*, 8663. <https://doi.org/10.3390/su16198663>.
3. Zahedi, F.; Alavi, H.; Sardroud, J.M.; Dang, H. Digital twins in the sustainable construction industry. *Buildings* **2024**, *14*, 3613. <https://doi.org/10.3390/buildings14113613>.
4. Zhao, L.; Guo, C.; Chen, L.; Qiu, L.; Wu, W.; Wang, Q. Using BIM and LCA to calculate the life cycle carbon emissions of inpatient building. *Sustainability* **2024**, *16*, 5341. <https://doi.org/10.3390/su16135341>.
5. Ma, L.; Azari, R.; Elnimeiri, M. A BIM-based life cycle assessment of high-rise building structures. *Sustainability* **2024**, *16*, 569. <https://doi.org/10.3390/su16020569>.
6. Chen, Z.; Chen, L.; Zhou, X.; Huang, L.; Sandanayake, M.; Yap, P.-S. Recent technological advancements in BIM and LCA integration for sustainable construction. *Sustainability* **2024**, *16*, 1340. <https://doi.org/10.3390/su16031340>.
7. Gao, H.; Wang, D.; Du, X.; Zhao, Z. An LCA-BIM integrated model for carbon-emission calculation of prefabricated buildings. *Renew. Sustain. Energy Rev.* **2024**, *203*, 114775. <https://doi.org/10.1016/j.rser.2024.114775>.
8. Parece, S.; Resende, R.; Rato, V. A BIM-based tool for embodied carbon assessment. *Dev. Built Environ.* **2024**, *19*, 100467. <https://doi.org/10.1016/j.dibe.2024.100467>.
9. Zhai, Y.; Li, Y.; Tang, S.; Liu, Y.; Liu, Y. Lightweight strategies for wooden-structure buildings based on embodied carbon emission calculations. *Buildings* **2024**, *14*, 3460. <https://doi.org/10.3390/buildings14113460>.
10. Zheng, L.; Luo, K.; Zhao, L. An operational carbon emission prediction model based on machine learning. *Buildings* **2024**, *14*, 3699. <https://doi.org/10.3390/buildings14113699>.
11. Zhang, X.; Sun, J.; Zhang, X.; Wang, F. Assessment of carbon emissions from buildings using machine learning. *J. Clean. Prod.* **2024**, *450*, 141903. <https://doi.org/10.1016/j.jclepro.2024.141903>.
12. Zhu, H.; Hu, C.; Zhou, C.; Wang, Z.; Wang, X.; Wu, Y. An explainable machine learning framework for carbon reduction. *Sci. Rep.* **2025**, *16*, 29418. <https://doi.org/10.1038/s41598-025-29418-z>.
13. Poyyamozhi, M.; Murugesan, B.; Rajamanickam, N.; Shorfuzzaman, M.; Aboelmagd, Y. IoT—A promising solution to energy management in smart buildings. *Buildings* **2024**, *14*, 3446. <https://doi.org/10.3390/buildings14113446>.
14. Franco, A.; Crisostomi, E.; Dalmiani, S.; Poletti, R. Synergy in action: Integrating environmental monitoring and IoT. *Buildings* **2024**, *14*, 1077. <https://doi.org/10.3390/buildings14041077>.
15. Chaudhari, P.; Xiao, Y.; Cheng, M.M.-C.; Li, T. Occupancy detection for smart buildings using IoT sensors. *Sensors* **2024**, *24*, 2123. <https://doi.org/10.3390/s24072123>.
16. Zhang, Z.; Yao, J.; Zheng, R. Multi-objective optimization of building energy saving. *Sustainability* **2024**, *16*, 1935. <https://doi.org/10.3390/su16051935>.
17. Xu, X.; Wang, Q.; Ding, X.; Chen, T.; Deng, R. Multi-objective optimization for prefabricated buildings using ant colony algorithm. *Sci. Rep.* **2025**, *15*, 31683. <https://doi.org/10.1038/s41598-025-17216-6>.
18. Ferriz-Papi, J.A.; Lee, A.; Alhawamdeh, M. Circular economy implementation in construction waste management. *Buildings* **2024**, *14*, 1237. <https://doi.org/10.3390/buildings14051237>.
19. Alhawamdeh, M.; Ferriz-Papi, J.A.; Lee, A. Drivers for improved construction waste management. *Sustainability* **2024**, *16*, 6014. <https://doi.org/10.3390/su16146014>.
20. Wei, J.; Ge, B.; Zhong, Y.; Lee, T.L.; Zhang, Y. Comparative analysis of embodied carbon in modular construction. *Sci. Rep.* **2024**, *14*, 23603. <https://doi.org/10.1038/s41598-024-73906-7>.
21. Chen, L.; Huang, Z.; Pan, W.; Su, R.K.; Zhong, Y.; Zhang, Y. Low carbon concrete for prefabricated modular construction. *J. Build. Eng.* **2024**, *90*, 109368. <https://doi.org/10.1016/j.job.2024.109368>.
22. Zhu, Y.; Wang, Z. Carbon emission evaluation of prefabricated components. *J. Build. Eng.* **2024**, *89*, 109262. <https://doi.org/10.1016/j.job.2024.109262>.
23. Hemmati, M.; Messadi, T.; Gu, H.; Seddelmeyer, J.; Hemmati, M. Comparison of embodied carbon: Mass timber vs. steel. *Buildings* **2024**, *14*, 1276. <https://doi.org/10.3390/buildings14051276>.
24. Azari, R.; Kamel, E.; Memari, A.M. Energy-efficient buildings: Materials and enclosure systems. *Buildings* **2024**, *14*, 1921. <https://doi.org/10.3390/buildings14071921>.
25. Zhang, A.; Wang, F.; Li, H.; Pang, B.; Yang, J. Carbon emissions accounting based on digital twin. *Appl. Energy* **2024**, *376*, 123155. <https://doi.org/10.1016/j.apenergy.2024.123155>.
26. Fay, D.; Corcoran, B.; Diamond, D. Green IoT event detection for carbon monitoring. *Sensors* **2023**, *24*, 162. <https://doi.org/10.3390/s24010162>.

27. Fayzul Islam, Ibrahim Ahmed, Lucian Mihet-Popa, Development and testing of an IoT platform with smart algorithms for building energy management systems, *Energy and Buildings*, Volume 344, 2025, 115970, ISSN 0378-7788, <https://doi.org/10.1016/j.enbuild.2025.115970>
28. Lowitzsch, J.; Hoicka, C.E.; van Tulder, F.J. Renewable energy communities under the 2019 European Clean Energy Package. *J. Risk Financ. Manag.* **2022**, *15*, 521. <https://doi.org/10.3390/jrfm15110521>.
29. Figueiredo, K.; Hammad, A.W.; Pierott, R.; Tam, V.W.; Haddad, A. Digital twin and blockchain for building LCA. *J. Build. Eng.* **2024**, *97*, 111018. <https://doi.org/10.1016/j.jobbe.2024.111018>.
30. ElArwady, Z.; Kandil, A.; Afify, M.; Marzouk, M. Indoor thermal comfort using digital twin and ML. *Dev. Built Environ.* **2024**, *19*, 100480. <https://doi.org/10.1016/j.dibe.2024.100480>.
31. Peldon, D.; Banihashemi, S.; LeNguyen, K.; Derrible, S. Digital twins in smart cities and urban development. *Sustain. Cities Soc.* **2024**, *111*, 105583. <https://doi.org/10.1016/j.scs.2024.105583>.
32. Villani, L.; Gugliermetti, L.; Barucco, M.A.; Cinquepalmi, F. A digital twin framework to improve urban sustainability and resiliency: The case study of Venice. *Land* **2025**, *14*, 83. <https://doi.org/10.3390/land14010083>.
33. Huzzat, A.; Anpalagan, A.; Khwaja, A.S.; Woungang, I.; Alnoman, A.A.; Pillai, A.S. A comprehensive review of digital twin technologies in smart cities. *Digit. Eng.* **2025**, *4*, 100040. <https://doi.org/10.1016/j.dte.2025.100040>.
34. Omrany, H.; Mehdipour, A.; Oteng, D.; Al-Obaidi, K.M. The uptake of urban digital twins in the built environment: A pathway to resilient and sustainable cities. *Comput. Urban Sci.* **2025**, *5*, 20. <https://doi.org/10.1007/s43762-025-00177-x>.
35. Mazzetto, S. A review of urban digital twins integration, challenges, and future directions in smart city development. *Sustainability* **2024**, *16*, 8337. <https://doi.org/10.3390/su16198337>.
36. Mabrouk, M.; Mahran, M.G.N.; Saleh, D.; AbouKorin, S.A.A. Digital twin cities for climate-resilient urban planning: A systematic review of frameworks, applications, and future perspectives. *Adv. Sustain.* **2025**, *5*, 58–75. <http://dx.doi.org/10.26855/as.2025.12.002>.
37. Ghorbany, S.; Hu, M.; Yao, S.; Sisk, M.; Wang, C. EcoSphere: A decision-support tool for automated carbon emission and cost optimization in sustainable urban development. *arXiv* **2025**, arXiv:2505.09054. <https://doi.org/10.48550/arXiv.2505.09054>.
38. Yessef, M.; Hakam, Y.; Tabaa, M.; Alammam, M.M.; Elbarbary, Z.M. Digital twin technology in smart cities: A step toward intelligent urban management. *Energy Rep.* **2025**, *14*, 5539–5557. <https://doi.org/10.1016/j.egyr.2025.11.097>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.