EAI Endorsed Transactions
on Scalable Information Systems
Research Article EAI.EU

Real-Time Scheduling Mechanisms for Heterogeneous Distributed Systems in Edge-Enabled Urban Renewal Digital Twin Platforms

Xiang Li[1], Xiang Wang[2] and Wei Peng[3,*]

[1] School of Architecture and Urban Planning, Nanjing University, Nanjing, 210000, China
[2] Jiangsu Provincial Academy of Building Research Co., Ltd., Nanjing, 210000, China
[3] Jiangsu Provincial Architectural Design Institute Co., Ltd., Nanjing, 210000, China

Abstract

INTRODUCTION: Urban renewal digital twin systems must support high-fidelity rendering and millisecond-level interactive control across geographically distributed, heterogeneous edge computing infrastructures. Conventional scheduling approaches, often assuming hardware homogeneity and static task graphs, struggle with dynamic service-chain reconfigurations driven by urban spatial entropy fluctuations, leading to resource misallocation and service instability in large-scale distributed environments.
OBJECTIVES: This paper aims to develop a real-time, scalable scheduling framework for heterogeneous distributed edge systems that jointly accounts for hardware diversity, evolving directed acyclic graph (DAG)-based workloads, and decentralized decision-making while preserving data locality and model privacy.
METHODS: We propose a hierarchical scheduling architecture integrating physical and logical coordination: (1) a resource affinity mask encodes hardware constraints as a prior to prune infeasible placements; (2) a spatio-temporal graph neural network captures critical path dynamics in non-stationary task DAGs; and (3) a federated policy distillation mechanism enables knowledge transfer across structurally diverse cloud-edge-end agents without sharing raw models or data.
RESULTS: Experiments on the Alibaba Cluster Trace and Shanghai Telecom datasets show that the proposed method reduces average latency to 43.8 ms, achieving a 25.6% reduction compared with CO-MARL, sustains a 94.7% task completion rate under long-tail traffic surges, and achieves an edge inference latency of 2.1 ms.
CONCLUSION: The "physical priors plus topology awareness" paradigm demonstrates that heterogeneous-aware, distributed coordination is essential for real-time digital twin services at scale.

Keywords: Urban Digital Twin; Heterogeneous Edge Computing; Distributed Task Scheduling; Federated Policy Distillation; Deep Reinforcement Learning

Received on 11 February 2026, accepted on 17 June 2026, published on 29 June 2026

doi: 10.4108/eetsis.11900

1. Introduction

As urban renewal enters a phase focused on improving existing infrastructure, digital twins (DTs) are evolving from static Building Information Modeling (BIM)-based visualization toward dynamic virtual-physical interaction[1]. Practical deployments, such as those in Xiong'an New Area and Shanghai Lingang, have begun to adopt real-time interactive platforms[2][3]. This evolution gives rise to pronounced bi-modal computational workloads, characterized by the coexistence of Graphics Processing Unit (GPU)-intensive high-fidelity rendering and latency-sensitive

*Corresponding author. Email: weipengjs11@163.com

control tasks[4]. In resource-constrained edge environments, coordinating heterogeneous and competing workloads to simultaneously ensure visual continuity and millisecond-level control responsiveness has become a fundamental bottleneck limiting the scalability of digital twin platforms[5]. Addressing this challenge is critical to system availability and safety in complex urban scenarios, such as preventing cybersickness or failures in emergency control.

Although substantial progress has been made in computing continuum architectures, several fundamental limitations remain. Most existing solutions adopt hardware-agnostic assumptions, overlooking acceleration heterogeneity across edge nodes and consequently causing resource mismatches for digital twin-specific tasks, such as rendering subtasks[6]. In addition, many deep reinforcement learning (DRL)-based approaches implicitly assume task independence, ignoring the directed acyclic graph (DAG) nature of digital twin services and often leading to service logic breakdowns due to violated dependencies[7]. Similarly, federated learning based on parameter averaging fails to address structural heterogeneity among intelligent agents, thereby limiting effective knowledge transfer across nodes with disparate computational capacities[8]. Moreover, urban renewal projects typically involve multiple stakeholders, such as design institutes, construction companies, and regulatory agencies, necessitating distributed architectures to protect data privacy, which renders traditional centralized cloud-based scheduling inadequate[9].

To bridge these gaps, this paper proposes an edge-enabled, scalable real-time scheduling mechanism tailored for urban renewal digital twin platforms. The main contributions are threefold. First, a heterogeneous-aware multi-agent DRL algorithm is developed by introducing a resource affinity mask to mitigate computational mismatches. By explicitly embedding hardware constraints into the decision space, the mechanism automatically rejects infeasible scheduling actions and significantly improves resource matching efficiency. Second, a spatio-temporal dependency-aware DAG orchestration mechanism is proposed, leveraging graph neural network (GNN) embeddings to model service topology and penalize constraint violations, thereby ensuring end-to-end service-chain integrity in distributed environments. Finally, a federated policy distillation framework is designed to enable knowledge transfer across heterogeneous agents through logit exchange, avoiding the structural incompatibility introduced by parameter aggregation while preserving data privacy.

This work departs from traditional assumptions of homogeneous resources and independent tasks, offering a new perspective on heterogeneous distributed optimization. From an application standpoint, the proposed mechanism substantially enhances the robustness of digital twin systems under non-stationary workloads and effectively mitigates failure risks induced by long-tail latency. Beyond urban renewal scenarios, it also establishes a new paradigm for infrastructure orchestration in future metaverse systems.

2. Related Work

2.1 Application Scenarios and Challenges

In urban renewal DT systems, such as real-world deployments in Xiong'an New Area and Singapore's Virtual Singapore, task workloads have exhibited pronounced bi-modal characteristics. One category consists of computation-intensive, high-fidelity 3D rendering and holographic visualization, typically implemented through Unity or Unreal Engine rendering pipelines and heavily reliant on high-performance GPU clusters[10]. The other category comprises latency-sensitive real-time physical feedback and emergency control tasks, which depend on fast logical inference at the edge[11]. This coexistence of "rendering–control" demands renders traditional BIM-based static models inadequate.

Although existing benchmark datasets cover urban traffic flows or energy consumption measurements, they generally lack annotations that capture the nonlinear relationship between task "spatial entropy" and computational complexity, making it difficult to quantify cognitive load under dynamic environments[12]. Moreover, current evaluation methodologies predominantly rely on average response time as a single metric. Such coarse-grained measurements obscure high-latency anomalies in the long-tail distribution, which are precisely the critical factors leading to cybersickness or control failures in digital twin systems[13]. Consequently, in resource-constrained and highly heterogeneous edge environments, simultaneously satisfying visual fidelity and millisecond-level control responsiveness remains a fundamental challenge for digital twin platforms.

2.2 Overview of Mainstream Approaches

To address the above challenges, existing research has primarily explored two directions: infrastructure architectures and scheduling algorithms.

From an infrastructure perspective, recent studies have focused on constructing IoT-Edge-Cloud computing continua. By integrating 5G network slicing technologies, these architectures effectively alleviate uplink transmission bottlenecks caused by massive sensor data, offering advantages in instantaneous access, bandwidth isolation, and reduced transmission jitter[14]. However, their resource orchestration layers typically treat edge nodes as homogeneous resource pools, neglecting substantial differences in hardware acceleration capabilities[15]. Such hardware-agnostic assumptions easily lead to resource mismatches when handling digital twin workloads that depend on specialized acceleration, such as high-fidelity rendering, thereby causing severe performance degradation, including rendering frame-rate collapse[16].

From the perspective of task scheduling strategies, traditional mixed-integer nonlinear programming (MINLP) and heuristic methods are often computationally prohibitive for real-time scenarios. More recent efforts have shifted toward data-driven approaches, employing Autoregressive

Integrated Moving Average (ARIMA) or Long Short-Term Memory (LSTM) models to predict traffic loads, which perform well under workloads with strong periodicity[18]. However, these methods heavily rely on stationarity assumptions and tend to fail in the presence of abrupt, non-stationary event streams during urban renewal processes, such as emergency fire simulations[17]. Although recent 2024–2025 studies [13], [15] incorporate federated learning to preserve data privacy in edge networks, most adopt parameter averaging-based aggregation schemes that inadequately account for structural heterogeneity across edge devices. As a result, global models suffer significant inference degradation on resource-constrained nodes[18]. This limitation arises because Federated Averaging (FedAvg) requires identical parameter dimensions across participants, whereas heterogeneous computational capabilities necessitate structurally different models, rendering direct parameter averaging mathematically infeasible.

2.3 Closely Related Studies

The studies most closely related to this work focus on digital twin task scheduling based on DRL.

The first category investigates collaborative computing mechanisms in smart construction sites and building engineering scenarios. Existing works propose multi-agent DRL-based resource scheduling frameworks for construction sites, enabling edge nodes to make independent decisions and alleviating single-point-of-failure issues in large-scale deployments. Other studies apply distributed reinforcement learning to optimize component transportation sequences in prefabricated building digital twins[19].These efforts are representative in reducing signaling overhead and improving system robustness. However, most of them assume hardware homogeneity and consider only CPU utilization as the system state. In contrast, this work focuses on heterogeneous urban renewal digital twin environments, where rendering, simulation, and control tasks may require different hardware accelerators in edge – cloud infrastructures[20]. In our experimental setting, approximately 30% of edge nodes are configured as GPU-enabled nodes to emulate such heterogeneity. Based on this setting, the proposed resource affinity mask allows agents to reject hardware-incompatible assignments, such as assigning rendering subtasks to CPU-only nodes, thereby reducing performance degradation caused by resource mismatches.

The second category concentrates on real-time scheduling in BIM-based collaborative design and mobile construction equipment scenarios. Related studies employ DRL to handle topology dynamics induced by equipment mobility or leverage edge computing to reduce simulation transmission latency in bridge digital twins. Furthermore, while recent 2024 works have explored spatio-temporal graph networks for edge prediction [9] or dependent task offloading [19], these approaches generally lack dynamic topology awareness or assume hardware homogeneity [21]. In practice, urban renewal digital twin services inherently form DAGs composed of data acquisition, safety verification, and rendering stages, where structural simulation must precede visualization to avoid presenting unsafe designs. Unlike existing studies, this work explicitly addresses heterogeneous constraints and DAG dependencies in urban renewal digital twin scenarios by designing a scheduling mechanism that integrates physical priors with topology awareness.

2.4 Summary and Research Gaps

In summary, although computing continuum architectures address connectivity issues and generic multi-agent DRL algorithms enable large-scale collaboration, there remains a lack of real-time scheduling mechanisms specifically tailored to the dual characteristics of heterogeneous computational constraints and DAG-structured task flows in digital twin platforms. To explicitly highlight our key differences and advantages over recent state-of-the-art (SOTA) works, a concise core design comparison is summarized in Table 1.

Table 1. Core Design Comparison with Recent SOTA Works

Methodology / Ref.	Heterogeneous Edge Support	Dynamic DAG Topology	Low-Comm Knowledge Transfer	Urban Digital Twin Focus
Adaptive-FL [15]	✓	×	×(Transmits Params)	×
ST-GCN SOTA [9]	✓	×	×	×
DAG-DRL [23]	×(Homogeneous)	✓	×	✓
Ours (Proposed)	✓	✓	✓ (Transmits Logits)	✓

While existing literature demonstrates the potential of edge computing for latency reduction, understanding of how to achieve policy-level, rather than parameter-level, knowledge transfer among heterogeneous agents remains limited. Current solutions either suffer from rendering stalls due to ignored hardware heterogeneity or experience service logic failures due to neglected task dependencies.

To fill these gaps, this study introduces three key innovations. (1) Resource affinity masking, which encodes GPU/CPU heterogeneity as hard constraints and forcibly eliminates mismatched task–resource assignments, thereby removing invalid exploration. (2) Spatio-temporal GNN embeddings, which leverage graph attention networks to capture DAG dependency structures and ensure scheduling decisions satisfy service-chain temporal constraints. (3) Federated policy distillation, which replaces parameter averaging with logit-based knowledge transfer, enabling decision-level alignment among heterogeneous agents. Together, these innovations address the core technical challenges faced by existing approaches in urban renewal digital twin scenarios.

3. Methodology

3.1 Problem Formulation

In the context of urban renewal digital twin platforms, the underlying environment is abstracted as a heterogeneous computing network, whose inputs consist of static resources and dynamic task streams. Computing nodes are modeled as heterogeneous resource pools. Each node j, in addition to its baseline CPU computing capability, is characterized by an attribute vector $C_j = [f_j^{gpu}, M_j]$, which quantifies its hardware acceleration capabilities for high-fidelity graphics rendering and AI inference.

The task stream is composed of dynamically arriving service requests. To capture the logical coupling inherent in digital twin applications, each request is formalized as a directed acyclic graph (DAG) $G_k = (\mathcal{V}, \mathcal{E})$, where each node $v_i \in \mathcal{V}$ represents a subtask with specific resource affinity requirements, and each edge $\mathcal{E}$ denotes an irreversible precedence dependency. From a data distribution perspective, the task arrival process and channel states are assumed to follow an unknown joint distribution $\mathcal{P}$, reflecting the spatio-temporal dynamics of urban entities.

Based on this formulation, the system output is defined as a spatio-temporal resource mapping policy π. At each time step, the policy generates a binary decision matrix that specifies the target execution node j for each subtask v_i, subject to the following constraints.

First, under the atomicity constraint, each subtask must be assigned to exactly one node:

$$\sum_{j\in N} x_{ij} = 1, \forall i. \tag{1}$$

Second, under the topological precedence constraint, the start time of any subtask must not precede the completion and data transmission of all its predecessor tasks:

$$T_{start}(v) \geq \max_{u\in Pred(v)} (T_{finish}(u) + T_{trans}(u, v)). \tag{2}$$

Third, under the resource capacity constraint, the aggregate resource demand of tasks assigned to any node must not exceed its available capacity:

$$\sum_{i:x_{ij}=1} r_i \leq C_j, \forall j. \tag{3}$$

where r_i denotes the computational resource demand of subtask v_i.Consequently, the scheduling policy must jointly address two tightly coupled challenges: spatial resource allocation and temporal execution ordering.

Overall, the optimization objective is to identify an optimal policy π^* that minimizes the long-term expected system cost while satisfying heterogeneous resource and DAG topology constraints. The cost function is defined as a weighted sum of end-to-end latency and energy consumption, aiming to endow the digital twin platform with adaptive scheduling capabilities under non-stationary data streams. Formally, the objective is to minimize

$$\mathbb{E}_{\mathcal{P}}[Cost(\pi)], \tag{4}$$

thereby achieving a Pareto-optimal balance between visual interaction smoothness and system energy efficiency.

3.2 Overall Framework

To enable real-time scheduling of urban digital twin tasks in heterogeneous edge environments, this paper proposes a hierarchical and cooperative intelligent framework, as illustrated in Figure 1. The framework is designed to transform dynamic DAG-based service flows (inputs) into optimal computation allocation decisions (outputs), covering the entire pipeline from perception and orchestration to distributed decision-making. At the lowest level, heterogeneous nodes collect local resource states, such as GPU utilization, via local agents, which are then fused with DAG task structures to form state representations. These representations are subsequently fed into three tightly coupled core modules that jointly complete the scheduling loop.

First, the data flow enters the heterogeneity-aware multi-agent DRL scheduling module, which serves as the decision engine. Independent agents deployed on edge nodes leverage a resource affinity mask to filter the state space, eliminating hardware-incompatible options and producing preliminary offloading actions that satisfy the atomicity constraint. Next, to handle task dependencies, the decision flow is passed to the spatio-temporal dependency-aware DAG orchestration module. This module employs GNN embeddings to encode task topology, performing temporal validation and reordering of preliminary decisions to ensure strict compliance with DAG precedence constraints and to prevent service-chain disruption. Finally, the system periodically activates the federated policy distillation-based collaborative training module. Under privacy-preserving settings, this module enables knowledge transfer among heterogeneous agents through the exchange of policy logits, allowing weaker nodes to imitate the decision patterns of stronger ones and continuously optimize long-term rewards. These three modules operate in a tightly coupled manner, collectively achieving Pareto-optimal trade-offs between latency and energy efficiency in complex urban environments.

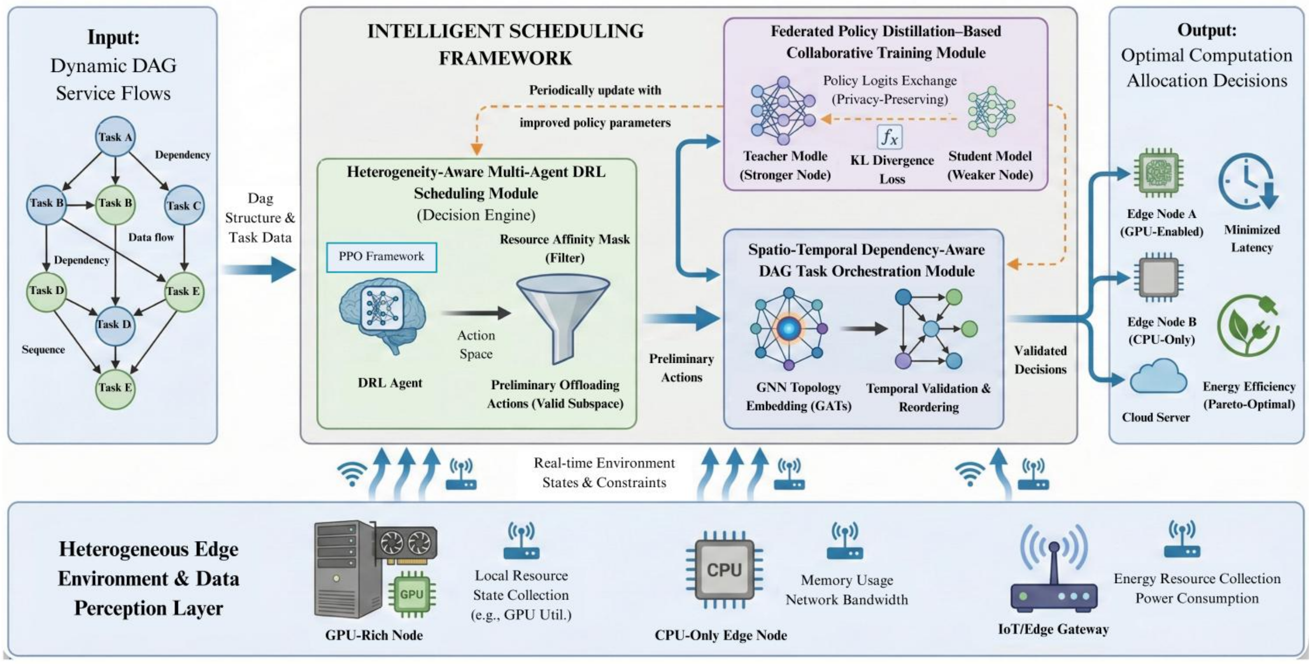

Figure 1. Overview of the proposed hierarchical scheduling framework for urban digital twins

3.3 Module Description

Heterogeneity-Aware Multi-Agent DRL Scheduling Module

Motivation.

In digital twin platforms, edge nodes exhibit extreme physical heterogeneity, such as GPU/CPU disparities. Conventional DRL methods typically explore within homogeneous action spaces, which frequently leads to invalid decisions, for example, assigning high-fidelity rendering tasks to CPU-only nodes, resulting in frame-rate collapse and slowed convergence. This module is designed to address feature-resource mismatches and invalid exploration by explicitly incorporating physical constraints to ensure scheduling feasibility.

Principle.

The core mechanism is a Proximal Policy Optimization (PPO) framework augmented with a resource affinity mask. By modifying the action distribution, actions that violate physical constraints are assigned zero probability, thereby pruning the original action space $\mathcal{A}$ into a valid subspace $\mathcal{A}_{\text{valid}}$. Leveraging prior knowledge of resource affinity, this mechanism preemptively blocks invalid gradient propagation, ensuring that the policy network is optimized strictly within the feasible solution space and significantly improving sample efficiency in heterogeneous environments.

Implementation.

As illustrated in Figure 2, a resource affinity mask layer is inserted before the Softmax layer of the Actor–Critic network. The mask is defined as

$$M \in \{0, -\infty\}^{|\mathcal{N}|}. \tag{5}$$

If node jprovides the hardware acceleration required by subtask v_i, then $M_j = 0$; otherwise, $M_j = -\infty$. The resulting action probability is computed as

$$\pi_\theta(a_t \mid s_t) = \mathrm{Softmax}(l(s_t) + M(v_i)), \tag{6}$$

where extremely negative values force the probabilities of invalid actions toward zero. The optimization objective maximizes the clipped surrogate function to guarantee monotonic policy improvement:

$$L^{CLIP}(\theta) = \mathbb{E}_t[\min\,(r_t(\theta)\widehat{A}_t, \mathrm{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\widehat{A}_t)], \tag{7}$$

where r_t denotes the policy ratio, $\widehat{A}_t$ the advantage estimate, and ϵ the clipping parameter.

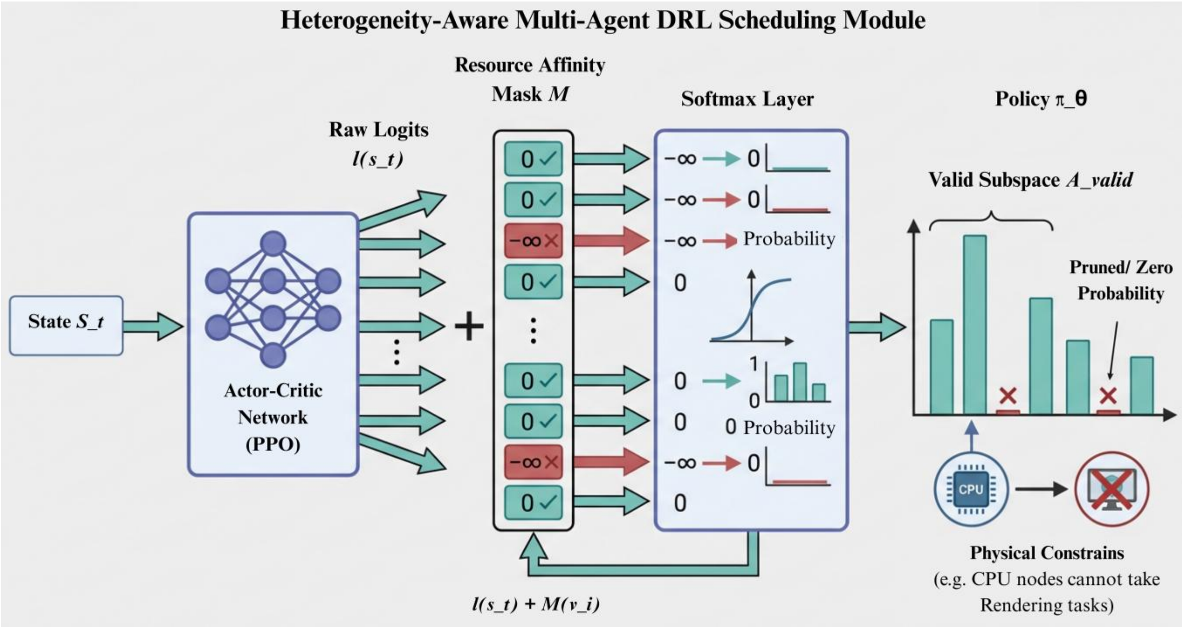

Figure 2. Resource affinity mask mechanism in the Actor-Critic network

Spatio-Temporal Dependency-Aware DAG Task Orchestration Module

Motivation.

Urban digital twin services inherently form strict DAGs comprising data acquisition, simulation, and rendering stages. Traditional schedulers often assume task independence and ignore precedence constraints, leading to temporal violations, such as scheduling successor tasks before predecessors complete, and ultimately causing virtual-physical synchronization failures. This module aims to model long-range topological dependencies by transforming graph constraints into learnable representations.

Principle.

The core mechanism is topology embedding based on Graph Attention Networks (GATs). Through message passing, the model aggregates features from DAG neighbors and enables each node to dynamically attend to the states of its predecessors (e.g., completion times). Multi-head attention captures critical-path bottlenecks and generates embeddings that encode implicit spatio-temporal positional information, guiding scheduling decisions to respect topological constraints.

Implementation.

As shown in Figure 3, the inputs consist of a task adjacency matrix and node features. The adjacency matrix A is constructed from DAG dependencies: if task uis a direct predecessor of task v (i.e., $(u,v) \in \mathcal{E}$), then $A_{uv} = 1$; otherwise, $A_{uv} = 0$. This matrix is dynamically updated at each time step based on the subgraph of unfinished tasks, reflecting the real-time topology of the service chain. The $(l+1)$-th layer embedding of task vaggregates information from its predecessor neighborhood $\mathcal{N}(v)$ as

$$h_v^{(l+1)} = \sigma\left(\sum_{u \in \mathcal{N}(v) \cup \{v\}} \alpha_{vu} W^{(l)} h_u^{(l)}\right), \tag{8}$$

where $W^{(l)}$ denotes the learnable weight matrix and α_{vu} the attention coefficient. To penalize temporal violations, DAG constraints are incorporated into the reward function:

$$R_{topo} = -\lambda \sum_{v \in \mathcal{V}} \max\left(0, T_{start}(v) - \max_{u \in Pred(v)} (T_{finish}(u) + T_{trans}(u,v))\right) \tag{9}$$

which forces the agent to schedule coupled tasks toward low-latency transmission configurations.

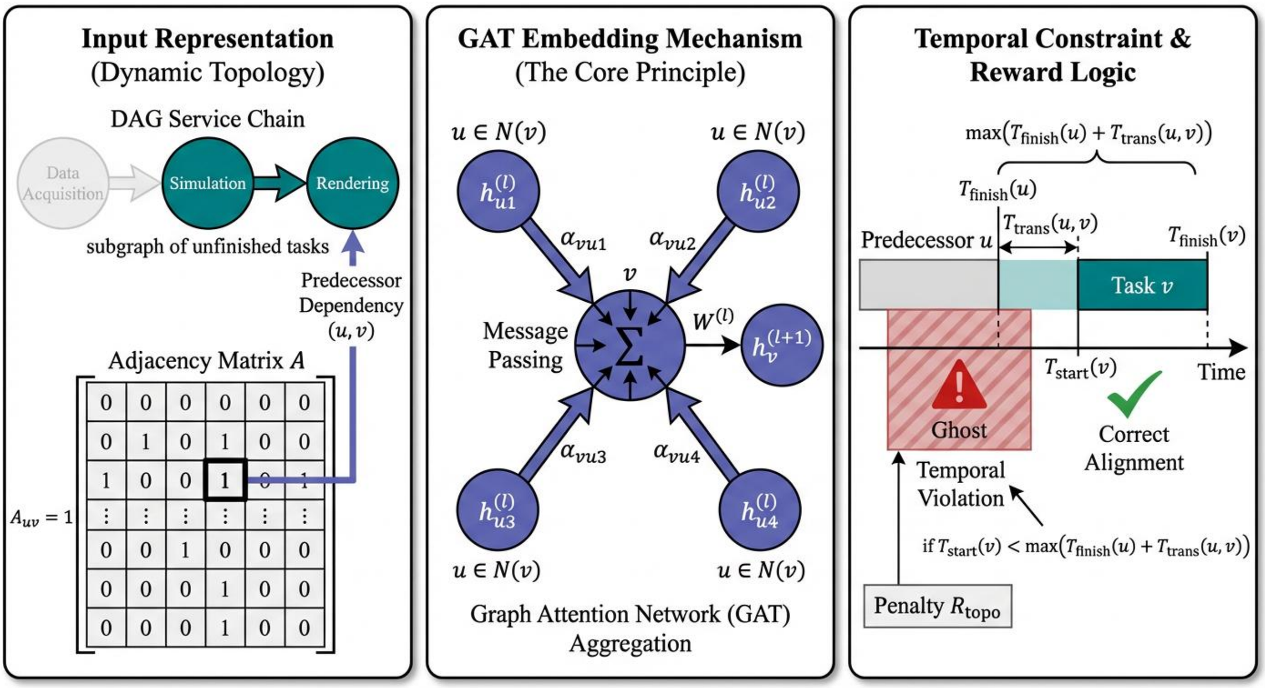

Figure 3. Spatio-Temporal Dependency-Aware DAG Task Orchestration Module. This module leverages graph attention networks to capture dynamic topological dependencies and output valid scheduling actions

Federated Policy Distillation–Based Collaborative Training Module

Motivation.

Digital twin platforms typically adopt cloud-edge-end hierarchical deployments, where model architectures differ across nodes due to computational constraints. Conventional FedAvg requires homogeneous model structures and is therefore unsuitable for such heterogeneous settings. This module aims to enable cross-model knowledge exchange, allowing weaker nodes to imitate stronger ones without synchronizing parameters.

Principle.

The core mechanism is policy distillation based on Kullback–Leibler (KL) divergence. Theoretically, the logits produced by teacher models contain rich class-level information. In the federated setting, only policy distributions are exchanged, rather than gradients, thus preserving privacy. Knowledge transfer is achieved by minimizing the KL divergence between local and global policy distributions, enabling lossless migration of decision knowledge.

Implementation.

As illustrated in Figure 4, the global average policy is defined as

$$\bar{\pi} = \frac{1}{K}\sum_{k=1}^{K} \pi_k. \tag{10}$$

During local updates of θ_k, each agent minimizes both the RL loss and the distillation loss:

$$LKD(\theta k) = DKL(\pi \parallel \pi\theta k) = a \in \mathcal{A}\bar{\pi}(a \mid s)\log\ \pi(a \mid s)\pi\theta k(a \mid s) \tag{11}$$

The final gradient is computed as:

$$\nabla\theta k = \nabla LCLIP + \beta\nabla LKD \tag{12}$$

allowing heterogeneous nodes to absorb global knowledge while maintaining lightweight local models.

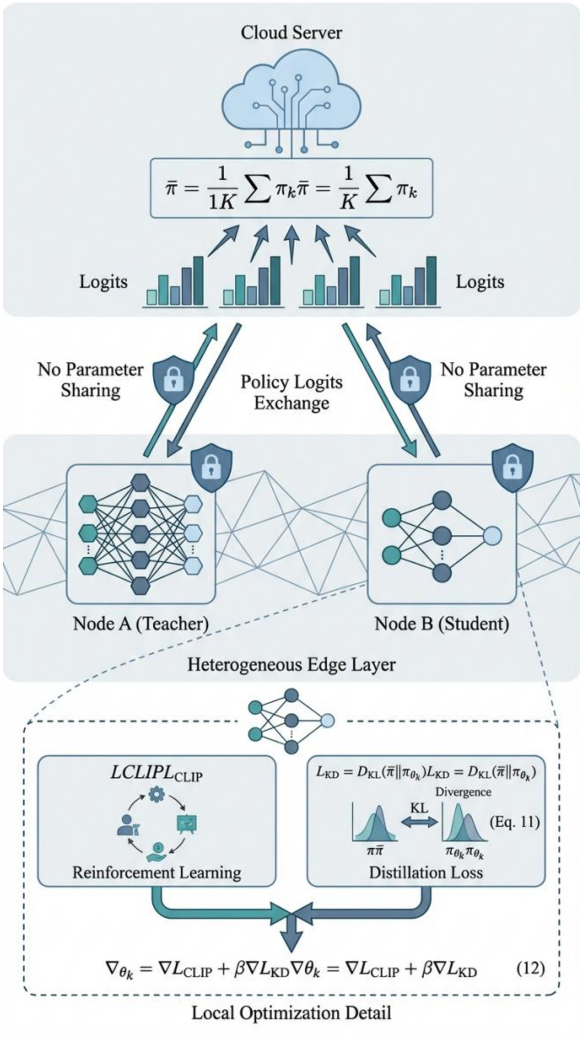

Figure 4. Federated Policy Distillation-Based Collaborative Training Module

3.4 Objective Function and Optimization

To achieve Pareto-optimal task scheduling in heterogeneous edge computing environments, the training process is formulated as a multi-objective joint optimization problem. The core objective is to identify an optimal set of neural network parameters θ^* such that each agent can simultaneously maximize long-term cumulative rewards derived from environmental feedback while minimizing deviations from the global policy distribution. Accordingly, the overall objective is defined as a composite loss function $L_{total}(\theta)$, as shown in Eq. (13), which consists of a weighted combination of the primary reinforcement learning loss and an auxiliary federated distillation loss:

$$L_{total}(\theta) = L_{PPO}(\theta) + \eta \cdot L_{KD}(\theta) \quad (13)$$

Here, $L_{PPO}(\theta)$ denotes the reinforcement learning loss induced by PPO, which aims to maximize the agent's own decision rewards, while $L_{KD}(\theta)$ represents the loss introduced by federated policy distillation, encouraging the local model to align with global knowledge. The coefficient η serves as a balancing factor that regulates the contribution of knowledge distillation during training. In this study, η is set to 0.5, as determined through grid search, to ensure that agents maintain personalized exploration without deviating excessively from the globally optimal policy.

To compute $L_{PPO}(\theta)$, we first define the environment feedback reward function R_t. Considering the dual requirements of digital twin workloads, a hybrid reward function incorporating latency, energy consumption, and topological constraints is designed, as shown in Eq. (14):

$$R_t = -\left(\alpha \cdot T_{sys}^{(t)} + \beta \cdot E_{sys}^{(t)} + \gamma \cdot P_{DAG}^{(t)}\right) \quad (14)$$

The negative sign is introduced to conform to the reinforcement learning objective of reward maximization. In this formulation, $T_{sys}^{(t)}$ denotes the total system latency, including both computation and transmission delays, while $E_{sys}^{(t)}$ represents the total energy consumption incurred by task execution. The weighting coefficients α and β correspond to latency and energy, respectively, and are both set to 0.5 in this work to pursue balanced performance. The term $P_{DAG}^{(t)}$ is a topology penalty that quantifies the degree of DAG dependency violation. Its associated coefficient $\gamma$is assigned a relatively large value of 10.0 to force the agent to prioritize logical integrity of service chains. Based on this reward formulation, $L_{PPO}(\theta)$ is further composed of a linear combination of the clipped policy loss, value function loss, and entropy regularization, with coefficients following standard PPO practice.

The second term in the total loss, $L_{KD}(\theta)$, is designed to facilitate collaboration among heterogeneous agents. As shown in Eq. (15), the KL divergence is employed to measure the discrepancy between the local policy and the global average policy:

$$L_{KD}(\theta) = \sum_{a \in \mathcal{A}} \bar{\pi}(a \mid s) \log \left(\frac{\bar{\pi}(a \mid s)}{\pi_\theta(a \mid s)}\right) \quad (15)$$

where $\pi_\theta(a \mid s)$ and $\bar{\pi}(a \mid s)$ denote the probability distributions of the local agent and the global average policy, respectively. By minimizing the divergence between these distributions, the local model is guided to incorporate consensus experience from other agents during parameter updates, thereby accelerating convergence.

Regarding optimization settings, the Adam optimizer is adopted to minimize the total objective $L_{total}(\theta)$, due to its adaptive handling of sparse gradients, which is particularly suitable for graph neural network features. The specific hyperparameters are configured as follows. The learning rate is set to 3×10^{-4} to avoid gradient explosion or vanishing. The batch size is fixed at 64, meaning that 64 trajectories are sampled per update. The discount factor Γ is set to 0.99, indicating a strong emphasis on long-term cumulative rewards. The PPO clipping parameter ϵ is set to 0.2 to constrain policy updates and prevent training collapse. With these settings, the system is able to adaptively learn scheduling policies that approach the Pareto frontier under complex heterogeneous constraints, providing reliable support for real-time virtual-physical interaction in urban digital twin platforms.

4. Experiment and Results

4.1 Experimental Setup

To comprehensively evaluate the real-time scheduling performance of the proposed method in heterogeneous edge computing environments, particularly under the bi-modal workloads and virtual-physical interaction requirements of urban renewal digital twin platforms, we adopt the Alibaba Cluster Trace v2018 as the primary benchmark dataset to enable fair comparisons with existing state-of-the-art (SOTA) approaches. In addition, the more challenging Shanghai Telecom Dataset is introduced to validate the generalization capability and robustness of the proposed method under realistic urban dynamics, such as topology variations induced by user mobility and long-tail workload distributions (see Table 2).

Table 2. Dataset Overview

Dataset	Primary Purpose	Task Type	Data Modality	Data Scale	Key Features / Target Variables	Source
Alibaba Cluster Trace v2018	Core evaluation and ablation study	Dynamic task offloading and scheduling	Structured logs (CSV/Parquet)	3M / 500K / 500K	Task arrivals; heterogeneous resource demands (CPU/memory + GPU demand derived/simulated from CPU traces); DAG dependencies; end-to-end latency (avg. 40 – 60 ms)	https://github.com/alibaba/clusterdata/tree/master/cluster-trace-v2018
Shanghai Telecom Dataset	Generalization and robustness evaluation	Edge task allocation and real-time scheduling	Spatio-temporal sequences (GPS trajectories + traffic logs)	1M / 200K / 200K	User mobility; heterogeneous node loads; long-tail traffic dynamics; end-to-end latency and utilization	https://www.kaggle.com/datasets/mexwell/telecom-shanghai-dataset

The Alibaba Cluster Trace v2018 dataset is split into training, validation, and test sets in chronological order using a fixed random seed of 42. Preprocessing includes Min-Max normalization to $[0,1]$, extraction of DAG dependency structures (average depth 3–5), and extension of CPU-related fields to simulate heterogeneous GPU computing demands measured in floating-point operations per second (FLOPs). The Shanghai Telecom Dataset is processed following the official pipeline: GPS trajectories are discretized into latitude–longitude grids with 100 m resolution, traffic logs are converted into 1-minute time series, and data are partitioned by user ID to prevent information leakage.

The two datasets are complementary. The Alibaba dataset emphasizes large-scale heterogeneous resource management, while the Shanghai dataset captures urban mobility and long-tail workload dynamics, enabling evaluation under both static and highly dynamic scenarios. Through preprocessing, the Alibaba dataset is extended to emulate heterogeneous digital twin rendering pipelines by incorporating GPU affinity and DAG dependencies, whereas the Shanghai dataset is mapped to non-stationary workloads driven by urban spatial entropy fluctuations, reflecting millisecond-level interaction bottlenecks at the edge.

Due to the lack of publicly available real-world datasets for urban renewal digital twin systems, a proxy-mapping strategy is adopted. Specifically, tasks with high memory and computational demands in the Alibaba dataset are labeled as BIM rendering-type tasks, while user mobility trajectories in the Shanghai dataset are analogized to construction equipment dynamics. This mapping preserves heterogeneous workloads, dynamic topology, and DAG dependencies, though it cannot fully capture all real-world constraints, such as multi-stakeholder coordination.

All experiments are conducted on a platform equipped with an NVIDIA RTX 3090 GPU (24 GB) and an Intel Xeon Gold 6240 CPU (36 cores). A heterogeneous environment is constructed using PyTorch 2.0.1 and Gym, with 50-500 nodes and approximately 30% GPU-enabled nodes. A three-layer multilayer perceptron (MLP) Actor-Critic architecture with hidden dimensions [256, 128, 64] and ReLU activations is used. Key hyperparameters include the Adam optimizer (learning rate 3×10^{-4}), batch size 64, discount factor 0.99, and 1000 training epochs with 2048 steps each. The total loss is defined as $L_{total} = L_{PPO} + 0.5 \cdot L_{KD}$, where the distillation weight is determined via grid search. Four-head GATs are employed for topological embedding, and invalid actions are filtered using the proposed masking mechanism. For robustness evaluation, Gaussian noise $\mathcal{N}(0, \sigma^2)$ with $\sigma \in [0.1, 0.5]$ is injected to simulate bandwidth fluctuations. All results are reported as mean ± standard deviation over five runs with a fixed random seed of 42.

The performance of the proposed framework is evaluated using a comprehensive set of metrics, including average latency, energy consumption, resource utilization balance, and task completion rate, as defined in Table 3.

Table 3. Evaluation Metric

Metric	Formula	Interpretation	Reference Values
Average Latency	$\frac{1}{N}\sum(T_{finish,i} - T_{start,i})$	Average task completion time, reflecting real-time interaction capability	<100 ms is excellent; 20 – 30% improvement over baseline is significant
Energy Consumption	$\sum P_j \cdot T_{exec,j}$	Total energy cost, indicating sustainability and operational efficiency	Unit: kJ; 15 – 25% reduction vs. baseline is considered efficient
Resource Utilization Balance	$Var(U_k)$	Variance of node utilization, measuring fairness of heterogeneous allocation	<0.05 indicates balanced; >0.1 is suboptimal
Task Completion Rate	$\frac{\#tasks\ completed\ on\ time}{N} \times 100\%$	Reliability under dependency constraints, ensuring continuous service	>95% is excellent; <90% under long-tail scenarios indicates insufficient robustness

4.2 Baseline Methods

To comprehensively evaluate the effectiveness of the proposed mechanism, we compare our framework with four SOTA methods.

First, Yang et al.[22] employ multi-agent deep reinforcement learning to address large-scale collaborative scheduling problems (hereafter CO-MARL); however, their method relies on a hardware-homogeneous assumption. In contrast, the proposed method introduces a resource affinity mask that explicitly avoids invalid exploration caused by CPU/GPU mismatches, which commonly degrade the performance of this baseline. Second, Guo et al.[23] apply DRL to handle task sequencing constraints with the goal of optimizing execution order (hereafter DAG-DRL). While effective at modeling temporal dependencies, this approach lacks the capability to capture deep topological features. By comparison, our method leverages a spatio-temporal GNN embedding layer to extract critical-path information, enabling more reliable preservation of service-chain integrity. Third, Kamdjou et al.[24] design a resource management framework specifically for digital twin-enabled extended reality systems (hereafter DT-Res), focusing on optimizing particular resource dimensions. Unlike this single-objective design, our method adopts a multi-objective Pareto optimization mechanism, which demonstrates stronger robustness under non-stationary bi-modal workload fluctuations. Finally, Peng et al.[25] propose SCOF, which utilizes parameter averaging to enable federated offloading, but this strategy strictly requires homogeneous model structures across participants. The proposed federated policy distillation mechanism overcomes this limitation by enabling knowledge transfer across heterogeneous neural network architectures, without enforcing parameter-level consistency.

Overall, these baseline methods are constrained by idealized hardware assumptions or enforced model homogeneity, which prevents them from simultaneously handling complex DAG topologies and adapting to non-stationary workloads. Such limitations are key factors underlying long-tail latency and service disruptions observed in real-world urban digital twin platforms. Accordingly, we expect the proposed framework, by integrating heterogeneity awareness, topology embedding, and policy distillation, to significantly outperform the above SOTA methods in supporting high-fidelity rendering and millisecond-level control interactions.

4.3 Quantitative Results

Table 4 reports the performance comparison of all methods on the Alibaba dataset (static heterogeneous resources) and the Shanghai dataset (dynamic urban workloads). On the computation-intensive Alibaba dataset, our method reduces the average latency to 43.8 ms, achieving a 25.6% improvement over CO-MARL. Intuitively, this substantial latency reduction directly benefits on-site structural engineers. When using AR/ VR headsets for high-fidelity holographic inspections of building renovations, this improvement effectively prevents rendering stalls and eliminates cybersickness, achieving seamless virtual-physical interaction. In contrast to DT-Res, which sacrifices latency to reduce energy consumption, the proposed approach attains a more favorable Pareto balance, enabling smooth rendering and real-time control simultaneously at the edge.

On the highly dynamic Shanghai dataset, our method maintains a 94.7% task completion rate despite long-tail workload perturbations, significantly outperforming all baselines. In practical urban renewal contexts, this result suggests that the proposed scheduler can better support time-sensitive services, such as structural warning updates or emergency evacuation-related digital twin tasks, under workload congestion.

Table 4. Cross-Dataset Performance Comparison

(Mean ± Std, N = 5, significance level † $p < 0.01$ vs. best baseline)

Datasets / Metrics	CO-MARL	DAG-DRL	DT-Res	SCOF	Ours
Alibaba Cluster Trace					
Latency (ms) ↓	58.9 ± 4.2	49.3 ± 3.5	61.4 ± 2.9	55.6 ± 5.4	43.8 ± 3.2†
Energy (kJ) ↓	145.3 ± 5.1	132.6 ± 4.8	118.2 ± 3.5	139.7 ± 6.2	121.5 ± 2.8
Comp. Rate (%) ↑	91.2 ± 2.4	93.5 ± 1.8	89.8 ± 2.1	90.5 ± 3.5	96.4 ± 1.8†
Shanghai Telecom Dataset					
Latency (ms) ↓	82.1 ± 6.8	65.4 ± 5.5	78.3 ± 4.9	74.2 ± 8.1	58.2 ± 5.7†
Energy (kJ) ↓	168.4 ± 8.2	152.1 ± 6.1	135.8 ± 4.5	158.9 ± 9.4	141.2 ± 5.2
Comp. Rate (%) ↑	84.5 ± 3.1	87.4 ± 4.2	82.1 ± 3.8	85.6 ± 5.2	94.7 ± 2.5†

Figure 5 illustrates the reward convergence curves during training. By masking invalid action spaces, our method converges stably within approximately 165 to 180 epochs, over 2× faster than CO-MARL (320 epochs) and SCOF (>400 epochs). Furthermore, while SCOF suffers from severe oscillations due to gradient conflicts from parameter averaging across heterogeneous devices, our logit-based policy distillation enables smooth knowledge transfer, substantially reducing training variance.

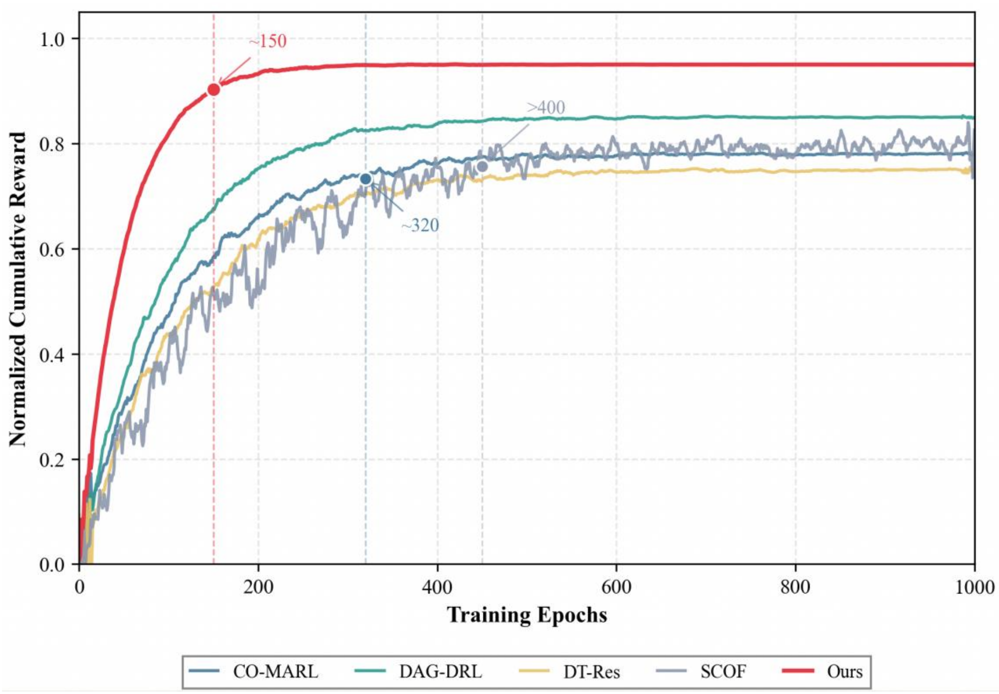

Figure 5. Reward convergence curves during training

Although the proposed method reaches performance saturation around 170 epochs (with reward fluctuation margins ⩽0.04), all methods are trained for 1000 epochs to ensure fair comparison and to examine long-term stability. Subsequent epochs are used to assess potential overfitting or performance degradation. Experimental results show that the proposed method remains reasonably stable between 170 and 1000 epochs, demonstrating convergence quality rather than merely faster convergence speed.

Table 5 quantifies computational efficiency. Specifically, the total computational cost of 1.15 GFLOPs is primarily distributed between the DRL agent (~0.80 GFLOPs) and the lightweight GNN embedding module (~0.35 GFLOPs), while the affinity masking computation is negligible. Despite the integration of GNN embeddings and policy distillation, the proposed method requires only 2.1 ms per inference, which is well below the 10.0 ms decision threshold for real-time digital twin scheduling. Compared with the heavy parameter synchronization overhead of SCOF, the proposed approach minimizes computational cost while maintaining high accuracy, achieving a favorable trade-off between efficiency and performance.

Table 5. Inference Efficiency and Computational Overhead

Method	Model Size (MB)	FLOPs (G)	Training Time (h)	Inference Latency (ms)	Communication Cost per Round (KB)
CO-MARL	4.2	0.85	12.5	1.8	N/A (Centralized)
DAG-DRL	5.8	1.20	14.2	2.5	N/A (Centralized)
DT-Res	3.5	0.65	11.8	1.5	N/A (Centralized)
SCOF	12.4	2.55	28.4	5.2	12687.0 (Params)
Ours	4.8	1.15	13.1	2.1	16.5 (Logits)

The inference latency of 2.1 ms reported in Table 5 ensures that scheduling decisions do not become a bottleneck for real-time digital twin interactions, satisfying the 10.0 ms decision window requirement and supporting imperceptible virtual-physical feedback for end users. To further evaluate scalability under high-concurrency scenarios, Figure 6 illustrates the framework's performance across varying edge node scales (50 to 500). The scheduling inference latency scales near-linearly ($\mathcal{O}(N)$), remaining under 8.4 ms even with 500 nodes, thereby avoiding exponential computational overhead. Furthermore, transmitting policy logits rather than complete model parameters maintains an ultra-low communication cost (e.g., merely 16.5 KB per node per round), demonstrating strong engineering feasibility for large-scale deployments.

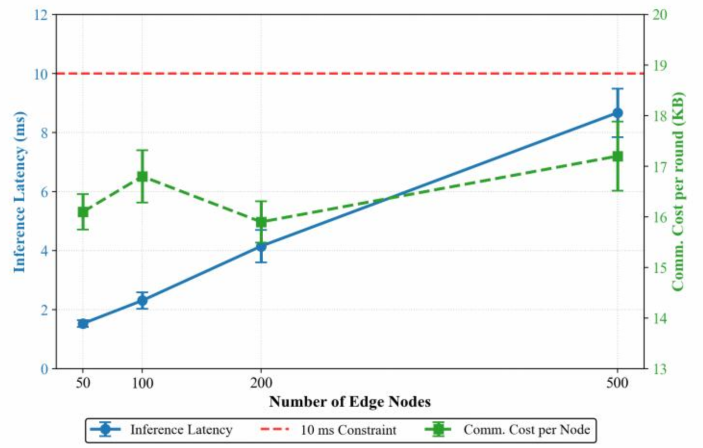

Figure 6. Scalability analysis of the proposed framework under varying edge node scales (50 to 500)

4.4 Qualitative Results

Figures 7 and 8 present representative scheduling visualizations on the Alibaba Cluster Trace dataset and the Shanghai Telecom dataset, respectively. These cases qualitatively show how scheduling decisions affect virtual–physical interaction under static heterogeneous resources and dynamic urban mobility.

Successful Case: Precise Adaptation to Heterogeneous Resources and DAG Topology

Figure 7 illustrates a representative urban renewal scenario constructed from the Alibaba Cluster Trace dataset. It presents a Gantt chart for scheduling a BIM-based building façade renovation task, where CPU-based preprocessing must be completed before GPU-based rendering. At t = 15 ms, only one CPU node is available. CO-MARL assigns the rendering subtask to this CPU node, increasing the execution time to 78 ms and delaying subsequent tasks. In contrast, the proposed method filters out this hardware-incompatible option through the resource affinity mask and waits for 3 ms until a GPU node becomes available. As a result, the task is completed within 43 ms due to GPU acceleration. This "strategic waiting" behavior shows that the proposed method can better coordinate resource compatibility and DAG precedence constraints in heterogeneous urban digital twin scheduling.

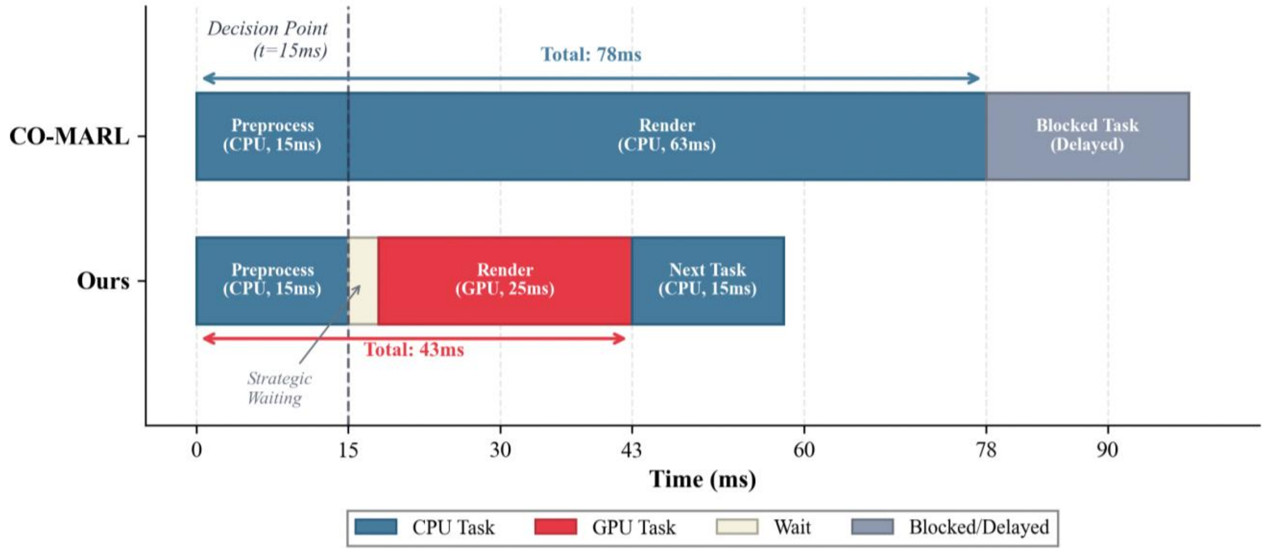

Figure 7. Gantt chart comparison of scheduling decisions for a BIM façade renovation task. The figure shows how the proposed method avoids hardware-incompatible assignment and maps the rendering subtask to a suitable GPU node

Failure Case: Topology Disruption under Extreme Mobility

Figure 8 presents a construction-site mobility scenario based on the Shanghai Telecom dataset, simulating an engineer wearing AR glasses and moving rapidly within a historical building renovation site. The device makes an abrupt turn near the edge of base-station coverage, such as entering a deep foundation pit, causing the actual trajectory to deviate sharply from the predicted trajectory. In this case, the model fails to complete the real-time BIM model synchronization task within the required time window, resulting in a mismatch between the 3D model displayed on the AR device and the physical environment.

The main cause is the delayed topology updating of the spatio-temporal GNN embedding layer. Since the GNN aggregates historical states from neighboring nodes to infer topology evolution, abrupt movement may invalidate the learned neighborhood structure before task migration is completed. This case indicates that the current framework remains limited under sudden link disruption and ultra-high-frequency spatial changes.

This failure case directly motivates the future work. Event-triggered graph updating can be introduced to detect abrupt topology changes earlier; dynamic graph pruning can remove obsolete dependencies after physical links become invalid; and adaptive attention thresholding can reduce the influence of outdated neighboring states. Therefore, the future improvement should focus on reducing the delay between physical topology disruption and scheduling adaptation, rather than simply replacing the current model with a larger GNN.

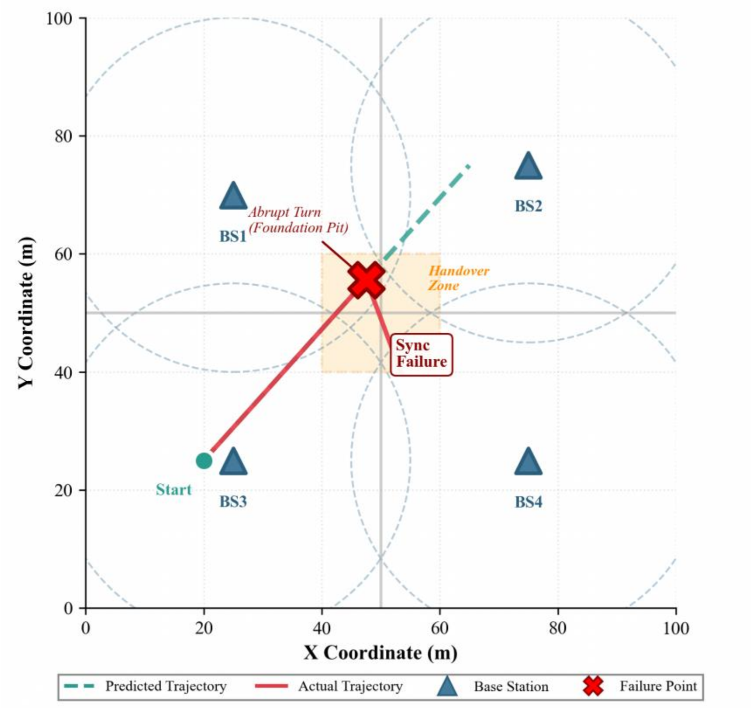

Figure 8. Trajectory map illustrating a BIM synchronization failure caused by abrupt mobility on the Shanghai Telecom dataset

4.5 Robustness Analysis

To evaluate the system's resilience to disturbances under realistic network fluctuations, we conduct robustness experiments on the most challenging Shanghai Telecom Dataset by injecting Gaussian noise $\mathcal{N}(0, \sigma^2)$ with intensity $\sigma \in [0.1, 0.5]$ into communication bandwidth, thereby simulating the impact of signal jitter on transmission latency in urban environments. Figure 9 illustrates the variation in task completion rates of different methods as the noise intensity increases.

As shown in Figure 9, performance degradation is observed for all methods as the noise level rises from 0.1 to 0.5; however, the proposed method (Ours) exhibits a pronounced "noise-resilient plateau." Specifically, under the extreme condition of $\sigma = 0.5$, the performance curve of DAG-DRL experiences a cliff-like drop, with the task completion rate falling below 75%. Its collapse point appears at approximately $\sigma = 0.3$, revealing the fragility of purely sequential prediction approaches: once noise-induced delays occur on a single link, DAG dependencies trigger a domino effect, ultimately causing the entire service chain to exceed its deadline. In contrast, the proposed method shows only a gradual decline and maintains a task completion rate above 91.5% even at $\sigma = 0.5$. This robustness can be directly attributed to the synergistic effect of the spatio-temporal GNN embedding layer and the resource affinity mask.

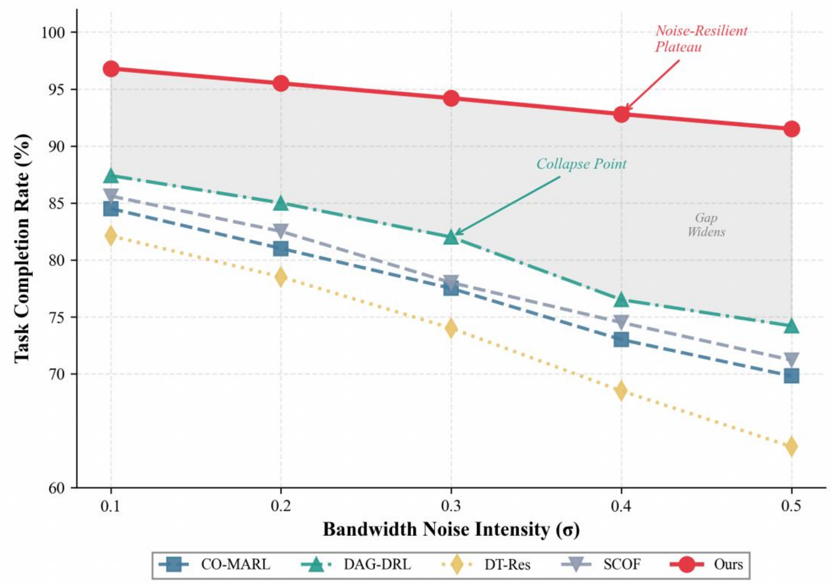

Figure 9. Robustness analysis under bandwidth noise perturbation. The results demonstrate that the proposed framework maintains stable performance and graceful degradation despite severe network fluctuations, outperforming baseline methods

First, the GNN attention mechanism functions as a form of "soft buffer." By aggregating neighborhood states, it dynamically identifies and avoids high-latency risk paths, rather than rigidly relying on static topology. Second, the resource affinity mask enforces the allocation of computation-intensive tasks to GPU-enabled nodes, substantially reducing computation time (T_{comp}) and thereby reserving greater time margins to absorb uncertainty in transmission latency (T_{trans}).

Compared with DT-Res, although it achieves reasonable energy efficiency under low-noise conditions, its task completion rate drops by 18.5% under high noise levels, demonstrating that isolated resource optimization lacks elasticity against stochastic disturbances. By explicitly penalizing DAG violations in the reward function and embedding topology awareness into the decision loop, the proposed method successfully preserves service continuity under non-ideal channel conditions. These results indicate that the proposed approach is not only effective in stable laboratory environments but also possesses practical engineering value for deployment in noisy urban settings.

4.6 Ablation Study

To quantitatively disentangle the individual contributions of each proposed component, we design three ablated variants, without resource affinity mask (w/o Mask), without GNN embedding (w/o GNN), and without policy distillation (w/o Distill), and compare them with the full model (Ours) on both datasets, as reported in Table 6.

Table 6. Performance Comparison of Key Component Ablations (Mean, N = 5)

Dataset	Variant	Latency (ms) ↓	Δ (%)	Energy (kJ) ↓	Δ (%)	Completion Rate (%) ↑	Δ (%)
Alibaba (Static)	w/o Mask	59.3	+35.4%	138.4	+13.9%	90.3	−6.3%
	w/o GNN	48.2	+10.0%	125.7	+3.5%	88.5	−8.2%
	w/o Distill	52.8	+20.5%	132.1	+8.7%	93.5	−3.0%
	Ours	43.8	–	121.5	–	96.4	–
Shanghai (Dynamic)	w/o Mask	78.7	+35.2%	167.8	+18.8%	86.1	−9.1%
	w/o GNN	64.0	+10.0%	146.5	+3.8%	86.4	−8.8%
	w/o Distill	70.6	+21.3%	155.3	+10.0%	90.4	−4.5%
	Ours	58.2	–	141.2	–	94.7	–

First, the resource affinity mask serves as the cornerstone for achieving low latency. Removing this component (w/o Mask) causes the average latency on the Alibaba dataset to increase by 35.4% (from 43.8 ms to 59.3 ms). This degradation stems from extensive early-stage exploration of invalid actions, such as assigning high-fidelity rendering tasks to CPU-only nodes, which leads to resource mismatches, queue congestion, and rendering frame-rate collapse. These results validate the necessity of physical-layer heterogeneity awareness in digital twin platforms. Second, the spatio-temporal GNN embedding is critical to maintaining DAG service-chain integrity. Without this module (w/o GNN), the task completion rate on the Shanghai dataset drops by 8.8% (from 94.7% to 86.4%), as simple MLP-based representations fail to capture complex task dependencies. Consequently, scheduling decisions violate global temporal order, resulting in virtual-physical desynchronization and control failures. Finally, federated policy distillation substantially improves system energy efficiency. Removing this mechanism (w/o Distill) increases energy consumption on the Shanghai dataset by 10.0%, since weaker nodes are unable to inherit decision policies from stronger ones, leading to suboptimal actions and increased retransmissions. This component effectively bridges structural model heterogeneity and enables knowledge transfer across cloud-edge-end deployments.

In summary, the three components play complementary roles: the resource affinity mask prevents rendering mismatches, the GNN embedding preserves service-chain integrity, and policy distillation supports heterogeneous collaboration. Rather than a simple aggregation of techniques, the proposed framework represents an organic integration tailored to the heterogeneous constraints, dynamic topology, and non-stationary workloads of urban renewal digital twin platforms, making it particularly suitable for millisecond-level real-time virtual–physical interaction.

5. Discussion

This study reveals a fundamental tension in edge-deployed urban renewal digital twins: simultaneously supporting high-fidelity rendering and millisecond-level control. Our resource affinity mask prevents a 35.4% latency increase, directly resolving the hardware matching challenge under bi-modal workloads. Unlike traditional MADRL's homogeneous assumptions, which cause rendering misallocations and disrupt visual continuity, our mechanism encodes hardware topology as hard constraints. This ensures computation-intensive services map to accelerated nodes, reducing cybersickness risks and safeguarding system availability.

The spatio-temporal GNN embedding proves highly adaptable to urban spatial entropy fluctuations. Inherently non-stationary urban renewal events frequently reconfigure sensor-simulation-rendering service chains. While purely sequential baselines like DAG-DRL suffer performance drops (completion rates falling to 87.4%), our GAT attention mechanism enables "soft dependency" modeling. This allows dynamic reweighting under unstable links, supporting graceful degradation and prioritizing real-time control loops. Nevertheless, the failure case in Figure 8 indicates that extreme anomalous events exceeding training data coverage still challenge historical aggregation-based GNNs, exposing an inherent limitation of data-driven methods in handling long-tail risks.

Federated policy distillation effectively addresses adaptation challenges in hierarchical cloud-edge-end architectures. The high variance (>0.15) in SCOF confirms that traditional parameter averaging fails across heterogeneous tiers. By transferring decision logits instead of parameters, our distillation enables edge nodes to inherit global policies while maintaining a low 2.1 ms inference latency. This helps maintain more consistent emergency-control scheduling during network disruptions.

Limitations primarily stem from the non-ergodic nature of urban systems; current datasets lack coverage of extreme anomalies like topology collapse. Although our framework scales efficiently up to 500 nodes, generalizing to city-scale deployments (thousands of nodes) introduces new concurrency and energy scaling challenges, where communication overhead may become a bottleneck. To address these limitations, future work will proceed along two directly related directions. For the topology-disruption failure shown in Figure 8, event-triggered graph updating, dynamic graph pruning, and adaptive attention thresholding will be introduced to reduce the delay between physical link disruption and scheduling adaptation. For city-scale deployments with thousands of nodes, hierarchical edge clustering and event-triggered distillation will be explored to reduce communication overhead and network-wide energy consumption.

6. Conclusion

To address the long-standing challenge of jointly supporting high-fidelity rendering and low-latency control in urban renewal digital twin platforms, this study proposes a real-time scheduling framework that integrates heterogeneity-aware masking, spatio-temporal topology embedding, and federated policy distillation. Specifically, we establish a layered methodology characterized by physical-constraint priors, dynamic topology awareness, and logical knowledge transfer, which effectively resolves the dual challenges of edge resource mismatch and service-chain fragmentation.

First, the proposed resource affinity mask explicitly encodes hardware heterogeneity as a decision constraint, eliminating invalid exploration at the physical layer. Experimental results demonstrate that this mechanism reduces average latency to 43.8 ms on the Alibaba dataset, achieving a 25.6% improvement over MADRL baselines and validating the principle that physical constraints should take precedence over purely data-driven learning. Second, to cope with fluctuations in urban spatial entropy, the spatio-temporal GNN embedding layer successfully captures dynamic service-chain characteristics. Under long-tail workload perturbations on the Shanghai dataset, the proposed method maintains a high task completion rate of 94.7%, significantly outperforming sequential prediction – based approaches. Finally, federated policy distillation overcomes the rigidity of parameter averaging and enables logical alignment across heterogeneous agents. While maintaining an inference latency of 2.1 ms, it reduces training variance to less than one-third of that observed in baseline methods.

This work introduces a new paradigm for complex digital twin platforms operating under resource-constrained environments, shifting from computational scaling to logical adaptation. The results demonstrate that, through fine-grained heterogeneous scheduling, edge infrastructures are capable of supporting real-time digital twin services. From a practical perspective, the proposed framework empowers low-cost edge gateways to execute high-stakes emergency scenarios, such as fire evacuation simulations, without complete reliance on unstable cloud connectivity, thereby enabling smooth virtual-physical interaction and consistent user experience.

In light of the limitations discussed earlier, future work will first focus on improving scheduling adaptation under extreme topology disruptions, such as the abrupt mobility case shown in Figure 8. Specifically, event-triggered graph updating, dynamic graph pruning, and adaptive attention thresholding will be investigated to reduce the delay between physical link disruption and scheduling adaptation. Furthermore, for city-level deployments with thousands of nodes, hierarchical edge clustering and event-triggered policy distillation will be explored to reduce communication overhead while maintaining real-time scheduling performance.

References

[1] R. Xu, C. W. Park, S. Khan, W. Jin, S. J. S. Moe, and D. H. Kim. Optimized task scheduling and virtual object management based on digital twin for distributed edge computing networks. *IEEE Access*, vol. 11, p. 114790-114810, 2023. https://doi.org/10.1109/access.2023.3325475

[2] R. Zhong, Y. Feng, X. Song, B. Hu, Y. Wang, P. Li, and J. Tan. Edge Computing Empowered Digital Twin: An End-to-End Computing Task Scheduling Approach. In *2024 IEEE 20th International Conference on Automation Science and Engineering (CASE)*, p. 3547-3552, 2024. https://doi.org/10.1109/case59546.2024.10711331

[3] Y. Wang, J. Fang, Y. Cheng, H. She, Y. Guo, and Zheng G. Cooperative end-edge-cloud computing and resource allocation for digital twin enabled 6g industrial iot. *IEEE Journal of Selected Topics in Signal Processing*, vol. 18, no. 1, p. 124-137, 2023. https://doi.org/10.1109/jstsp.2023.3345154

[4] S. Yuan, Z. Zhang, Q. Li, W. Li, and Y. Zhang. Joint optimization of dnn partition and continuous task scheduling for digital twin-aided mec network with deep reinforcement learning. *IEEE Access*, vol. 11, p. 27099-27110, 2023. https://doi.org/10.1109/access.2023.3257342

[5] S. R. Jeremiah, L. T. Yang, and J. H. Park. Digital twin-assisted resource allocation framework based on edge collaboration for vehicular edge computing. *Future Generation Computer Systems*, vol. 150, p. 243-254, 2024. https://doi.org/10.1016/j.future.2023.09.001

[6] J. Li, S. Guo, W. Liang, J. Wang, Q. Chen, Y. Zeng, and X. Jia. Digital twin-enabled service provisioning in edge computing via continual learning. *IEEE Transactions on Mobile Computing*, vol. 23, no. 6, p. 7335-7350, 2023. https://doi.org/10.1109/tmc.2023.3332668

[7] H. Guo, X. Zhou, J. Wang, J. Liu, and A. Benslimane. Intelligent task offloading and resource allocation in digital twin based aerial computing networks. *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 10, p. 3095-3110, 2023. https://doi.org/10.1109/jsac.2023.3310067

[8] J. Chen, Y. Yi, S. D. Okegbile, J. Cai, and X. Shen. Networking architecture and key supporting technologies for human digital twin in personalized healthcare: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, vol. 26, no. 1, p. 706-746, 2023. https://doi.org/10.1109/comst.2023.3308717

[9] X. Jiang, M. Sun, P. Wang, H. Sun, and L. Chen. Adaptive Spatio-Temporal Feature Graph Convolutional Network Prediction Model with Edge Computing Integration. *IEEE Internet of Things Journal*, 2026. https://doi.org/10.1109/jiot.2026.3661976

[10] J. Li, S. Guo, W. Liang, J. Wang, Q. Chen, W. Xu, and X. Jia. Mobility-aware utility maximization in digital twin-enabled serverless edge computing. *IEEE Transactions on Computers*, vol. 73, no. 7, p. 1837-1851, 2024. https://doi.org/10.1109/tc.2024.3388897

[11] D. Ayepah-Mensah, G. Sun, G. O. Boateng, and G. Liu. Federated Policy Distillation for Digital Twin-Enabled Intelligent Resource Trading in 5G Network Slicing. *IEEE Transactions on Network and Service Management*, vol. 22, no. 1, p. 361-379, 2024. https://doi.org/10.1109/tnsm.2024.3476480

[12] X. Zhou, X. Zheng, X. Cui, J. Shi, W. Liang, Z. Yan, and K. I. K. Wang. Digital twin enhanced federated reinforcement learning with lightweight knowledge distillation in mobile networks. *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 10, p. 3191-3211, 2023. https://doi.org/10.1109/jsac.2023.3310046

[13] Y. Dai, J. Zhao, J. Zhang, Y. Zhang, and T. Jiang. Federated deep reinforcement learning for task offloading in digital twin edge networks. *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 3, p. 2849-2863, 2024. https://doi.org/10.1109/tnse.2024.3350710

[14] R. Zhang, Z. Xie, D. Yu, W. Liang, and X. Cheng. Digital twin-assisted federated learning service provisioning over mobile edge networks. *IEEE Transactions on Computers*, vol. 73, no. 2, p. 586-598, 2023. https://doi.org/10.1109/tc.2023.3337317

[15] J. Wang, Q. Liang, and A. Mohajer. Adaptive offloading in multi-access edge networks via hierarchical federated learning and real-time system adaptation. *International Journal of Sensor Networks*, vol. 49, no. 1, p. 1-17, 2025. https://doi.org/10.1504/ijsnet.2025.10071733

[16] C. Álvarez-López, A. González-Briones, and T. Li. Explainable AI and Multi-Agent Systems for Energy Management in IoT-Edge Environments: A State of the Art Review. *Electronics*, vol. 15, no. 2, p. 385, 2026. https://doi.org/10.3390/electronics15020385

[17] N. G. Evgenidis, N. A. Mitsiou, V. I. Koutsioumpa, S. A. Tegos, P. D. Diamantoulakis, and G. Karagiannidis. Multiple access in the era of distributed computing and edge intelligence. *Proceedings of the IEEE*, vol. 112, no. 9, p. 1497-1526, 2024. https://doi.org/10.1109/jproc.2024.3417528

[18] J. Zhou, Y. Gao, X. Wang, Y. Li, and Z. Zhao. Study on Quality of AI Service Guarantee in Digital Twin Networks for XR Scenarios. *Electronics*, vol. 15, no. 2, p. 344, 2026. https://doi.org/10.3390/electronics15020344

[19] X. Chen, J. Cao, Y. Sahni, M. Zhang, Z. Liang, and L. Yang. Mobility-aware dependent task offloading in edge computing: a digital twin-assisted reinforcement learning approach. *IEEE Transactions on Mobile Computing*, 2024. https://doi.org/10.1109/tmc.2024.3506221

[20] M. Trigka and E. Dritsas. Edge and cloud computing in smart cities. *Future Internet*, vol. 17, no. 3, p. 118, 2025.

[21] E. Faliagka, E. Christopoulou, D. Ringas, T. Politi, N. Kostis, D. Leonardos, and N. Voros. Trends in digital twin framework architectures for smart cities: A case study in smart mobility. *Sensors*, vol. 24, no. 5, p. 1665, 2024. https://doi.org/10.51219/urforum.2024.christos-antonopoulos

[22] J. Yang, Q. Yuan, S. Chen, H. He, X. Jiang, and X. Tan. Cooperative task offloading for mobile edge computing based on multi-agent deep reinforcement learning. *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, p. 3205-3219, 2023.

[23] R. Guo, L. Zhou, L. Li, Y. Song, and X. Xie. Dependent Task Graph Offloading Model Based on Deep Reinforcement Learning in Mobile Edge Computing. *Electronics*, vol. 14, no. 16, p. 3184, 2025. https://doi.org/10.3390/electronics14163184

[24] H. M. Kamdjou, D. Baudry, V. Havard, and S. Ouchani. Resource-constrained extended reality operated with digital twin in industrial internet of things. *IEEE Open Journal of the Communications Society*, vol. 5, p. 928-950, 2024. https://doi.org/10.1109/ojcoms.2024.3356508

[25] K. Peng, P. Xiao, Wang S., and V. C. Leung. SCOF: Security-aware computation offloading using federated reinforcement learning in industrial internet of things with edge computing. *IEEE Transactions on Services Computing*, vol. 17, no. 4, p. 1780-1792, 2024. https://doi.org/10.1109/tsc.2024.3377899